imgaug Documentation

Release 0.4.0

Alexander Jung

Feb 18, 2020

Contents:

1 Installation **3**
- 1.1 Installation in Anaconda . . . 3
- 1.2 Installation in pip . . . 3
- 1.3 Uninstall . . . 4

2 Examples: Basics **5**
- 2.1 A standard use case . . . 5
- 2.2 A simple and common augmentation sequence . . . 5
- 2.3 Heavy Augmentations . . . 6

3 Examples: Keypoints **11**
- 3.1 Notebook . . . 11
- 3.2 A simple example . . . 11

4 Examples: Bounding Boxes **13**
- 4.1 Notebook . . . 13
- 4.2 A simple example . . . 13
- 4.3 Dealing with bounding boxes outside of the image . . . 14
- 4.4 Shifting/Moving Bounding Boxes . . . 17
- 4.5 Projection of BBs Onto Rescaled Images . . . 17
- 4.6 Computing Intersections, Unions and IoUs . . . 19

5 Examples: Heatmaps **23**
- 5.1 Notebook . . . 23
- 5.2 A simple example . . . 24
- 5.3 Multiple sub-heatmaps per heatmaps object . . . 25
- 5.4 Accessing the heatmap array . . . 27
- 5.5 Resizing heatmaps . . . 28
- 5.6 Padding heatmaps . . . 31

6 Examples: Segmentation Maps and Masks **33**
- 6.1 Notebook . . . 33
- 6.2 A simple example . . . 34
- 6.3 Using boolean masks . . . 35
- 6.4 Accessing the segmentation map array . . . 35
- 6.5 Resizing and padding . . . 38

7 Stochastic Parameters **39**
7.1 Introduction . . . 39
7.2 Continuous Probability Distributions . . . 40
7.3 Discrete Probability Distributions . . . 48
7.4 Arithmetic . . . 51
7.5 Special Parameters . . . 54
7.6 Noise Parameters . . . 60

8 Blending/Overlaying images **61**
8.1 Introduction . . . 61
8.2 Imagewise Constant Alphas Values . . . 62
8.3 BlendAlphaSimplexNoise . . . 66
8.4 FrequencyNoiseAlpha . . . 68
8.5 IterativeNoiseAggregator . . . 71
8.6 Sigmoid . . . 72

9 Overview of Augmenters **75**
9.1 augmenters.meta . . . 75
9.2 augmenters.arithmetic . . . 87
9.3 augmenters.artistic . . . 122
9.4 augmenters.blend . . . 122
9.5 augmenters.blur . . . 148
9.6 augmenters.collections . . . 157
9.7 augmenters.color . . . 161
9.8 augmenters.contrast . . . 183
9.9 augmenters.convolutional . . . 196
9.10 augmenters.debug . . . 201
9.11 augmenters.edges . . . 201
9.12 augmenters.flip . . . 205
9.13 augmenters.geometric . . . 206
9.14 augmenters.imgcorruptlike . . . 227
9.15 augmenters.pillike . . . 251
9.16 augmenters.pooling . . . 266
9.17 augmenters.segmentation . . . 274
9.18 augmenters.size . . . 283
9.19 augmenters.weather . . . 297

10 Performance **305**
10.1 Results Overview . . . 305
10.2 Images . . . 306
10.3 Heatmaps . . . 314
10.4 Keypoints and Bounding Boxes . . . 322

11 dtype support **327**
11.1 Legend . . . 327
11.2 imgaug helper functions . . . 328
11.3 imgaug.augmenters.meta . . . 334
11.4 imgaug.augmenters.arithmetic . . . 334
11.5 imgaug.augmenters.blend . . . 334
11.6 imgaug.augmenters.blur . . . 334
11.7 imgaug.augmenters.collections . . . 334
11.8 imgaug.augmenters.color . . . 334
11.9 imgaug.augmenters.contrast . . . 334
11.10 imgaug.augmenters.convolutional . . . 334
11.11 imgaug.augmenters.debug . . . 334

11.12 imgaug.augmenters.edges . . . 334
11.13 imgaug.augmenters.flip . . . 334
11.14 imgaug.augmenters.geometric . . . 338
11.15 imgaug.augmenters.imgcorruptlike . . . 338
11.16 imgaug.augmenters.pillike . . . 338
11.17 imgaug.augmenters.segmentation . . . 338
11.18 imgaug.augmenters.size . . . 338
11.19 imgaug.augmenters.weather . . . 338

12 Jupyter Notebooks **341**

13 API **343**
13.1 imgaug . . . 343
13.2 imgaug.parameters . . . 362
13.3 imgaug.multicore . . . 389
13.4 imgaug.dtypes . . . 393
13.5 imgaug.random . . . 394
13.6 imgaug.validation . . . 408
13.7 imgaug.augmentables.base . . . 409
13.8 imgaug.augmentables.batches . . . 409
13.9 imgaug.augmentables.bbs . . . 416
13.10 imgaug.augmentables.heatmaps . . . 434
13.11 imgaug.augmentables.kps . . . 439
13.12 imgaug.augmentables.lines . . . 452
13.13 imgaug.augmentables.normalization . . . 471
13.14 imgaug.augmentables.polys . . . 472
13.15 imgaug.augmentables.segmaps . . . 491
13.16 imgaug.augmentables.utils . . . 494
13.17 imgaug.augmenters.arithmetic . . . 497
13.18 imgaug.augmenters.artistic . . . 568
13.19 imgaug.augmenters.base . . . 572
13.20 imgaug.augmenters.blend . . . 572
13.21 imgaug.augmenters.blur . . . 626
13.22 imgaug.augmenters.collections . . . 645
13.23 imgaug.augmenters.color . . . 648
13.24 imgaug.augmenters.contrast . . . 703
13.25 imgaug.augmenters.convolutional . . . 729
13.26 imgaug.augmenters.debug . . . 741
13.27 imgaug.augmenters.edges . . . 745
13.28 imgaug.augmenters.flip . . . 750
13.29 imgaug.augmenters.geometric . . . 756
13.30 imgaug.augmenters.imgcorruptlike . . . 810
13.31 imgaug.augmenters.meta . . . 855
13.32 imgaug.augmenters.pillike . . . 912
13.33 imgaug.augmenters.pooling . . . 962
13.34 imgaug.augmenters.segmentation . . . 974
13.35 imgaug.augmenters.size . . . 998
13.36 imgaug.augmenters.weather . . . 1065

14 Indices and tables **1093**

Python Module Index **1095**

Index **1097**

imgaug is a library for image augmentation in machine learning experiments. It supports a wide range of augmentation techniques, allows to easily combine these and to execute them in random order or on multiple CPU cores, has a simple yet powerful stochastic interface and can not only augment images, but also keypoints/landmarks, bounding boxes, heatmaps and segmentation maps.

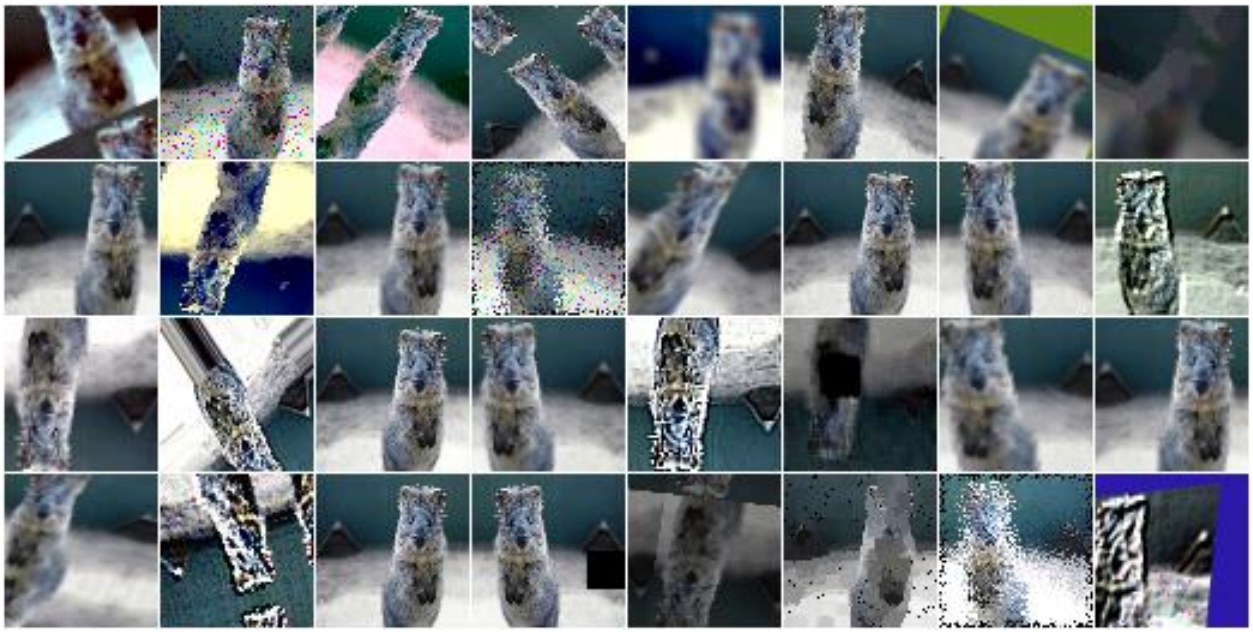

Fig. 1: Example augmentations of a single input image.

CHAPTER 1

Installation

The library uses python, which must be installed. Python 2.7, 3.4, 3.5, 3.6, 3.7 and 3.8 are supported.

The below sections explain how to install the library in anaconda and pip. If you don't know what anaconda (aka conda) is, simply use pip instead.

1.1 Installation in Anaconda

To install in anaconda simply perform the following commands

```
conda config --add channels conda-forge
conda install imgaug
```

Note that you may also use the pip-based installation commands described below. They work with and without anaconda.

To also be able to use the augmenters in `imgaug.augmenters.imgcorruptlike`, you have to manually install the `imagecorruptions` package:

```
pip install imagecorruptions
```

1.2 Installation in pip

To install the library via `pip`, simply execute:

```
pip install imgaug
```

This installs the latest version from pypi.

If you encounter any problems with `Shapely`, try the following:

```
pip install six numpy scipy Pillow matplotlib scikit-image opencv-python imageio
pip install --no-dependencies imgaug
```

The first command installs manually all dependencies except `Shapely`, the second only the library. Note that Shapely is required for some operations, mainly when operating with line strings or polygons.

The version installed above is the latest official release from pypi. That release often lags behind the latest version from github by a few months. To instead install the very latest version of imgaug use

```
pip install git+https://github.com/aleju/imgaug.git
```

Independent of whether you install from pypi or github, in order to be able to use the augmenters in `imgaug.augmenters.imgcorruptlike`, you have to manually install the `imagecorruptions` package:

```
pip install imagecorruptions
```

1.3 Uninstall

To deinstall the library use

```
conda remove imgaug
```

on anaconda and

```
pip uninstall imgaug
```

otherwise.

CHAPTER 2

Examples: Basics

2.1 A standard use case

The following example shows a standard use case. An augmentation sequence (crop + horizontal flips + gaussian blur) is defined once at the start of the script. Then many batches are loaded and augmented before being used for training.

```
from imgaug import augmenters as iaa

seq = iaa.Sequential([
    iaa.Crop(px=(0, 16)), # crop images from each side by 0 to 16px (randomly chosen)
    iaa.Fliplr(0.5), # horizontally flip 50% of the images
    iaa.GaussianBlur(sigma=(0, 3.0)) # blur images with a sigma of 0 to 3.0
])

for batch_idx in range(1000):
    # 'images' should be either a 4D numpy array of shape (N, height, width, channels)
    # or a list of 3D numpy arrays, each having shape (height, width, channels).
    # Grayscale images must have shape (height, width, 1) each.
    # All images must have numpy's dtype uint8. Values are expected to be in
    # range 0-255.
    images = load_batch(batch_idx)
    images_aug = seq(images=images)
    train_on_images(images_aug)
```

2.2 A simple and common augmentation sequence

The following example shows an augmentation sequence that might be useful for many common experiments. It applies crops and affine transformations to images, flips some of the images horizontally, adds a bit of noise and blur and also changes the contrast as well as brightness.

```
import numpy as np
import imgaug as ia
import imgaug.augmenters as iaa


ia.seed(1)

# Example batch of images.
# The array has shape (32, 64, 64, 3) and dtype uint8.
images = np.array(
    [ia.quokka(size=(64, 64)) for _ in range(32)],
    dtype=np.uint8
)

seq = iaa.Sequential([
    iaa.Fliplr(0.5), # horizontal flips
    iaa.Crop(percent=(0, 0.1)), # random crops
    # Small gaussian blur with random sigma between 0 and 0.5.
    # But we only blur about 50% of all images.
    iaa.Sometimes(
        0.5,
        iaa.GaussianBlur(sigma=(0, 0.5))
    ),
    # Strengthen or weaken the contrast in each image.
    iaa.LinearContrast((0.75, 1.5)),
    # Add gaussian noise.
    # For 50% of all images, we sample the noise once per pixel.
    # For the other 50% of all images, we sample the noise per pixel AND
    # channel. This can change the color (not only brightness) of the
    # pixels.
    iaa.AdditiveGaussianNoise(loc=0, scale=(0.0, 0.05*255), per_channel=0.5),
    # Make some images brighter and some darker.
    # In 20% of all cases, we sample the multiplier once per channel,
    # which can end up changing the color of the images.
    iaa.Multiply((0.8, 1.2), per_channel=0.2),
    # Apply affine transformations to each image.
    # Scale/zoom them, translate/move them, rotate them and shear them.
    iaa.Affine(
        scale={"x": (0.8, 1.2), "y": (0.8, 1.2)},
        translate_percent={"x": (-0.2, 0.2), "y": (-0.2, 0.2)},
        rotate=(-25, 25),
        shear=(-8, 8)
    )
], random_order=True) # apply augmenters in random order

images_aug = seq(images=images)
```

2.3 Heavy Augmentations

The following example shows a large augmentation sequence containing many different augmenters, leading to significant changes in the augmented images. Depending on the use case, the sequence might be too strong. Occasionally it can also break images by changing them too much. To weaken the effects you can lower the value of `iaa.SomeOf((0, 5), ...)` to e.g. `(0, 3)` or decrease the probability of some augmenters to be applied by decreasing in `sometimes = lambda aug: iaa.Sometimes(0.5, aug)` the value `0.5` to e.g. `0.3`.

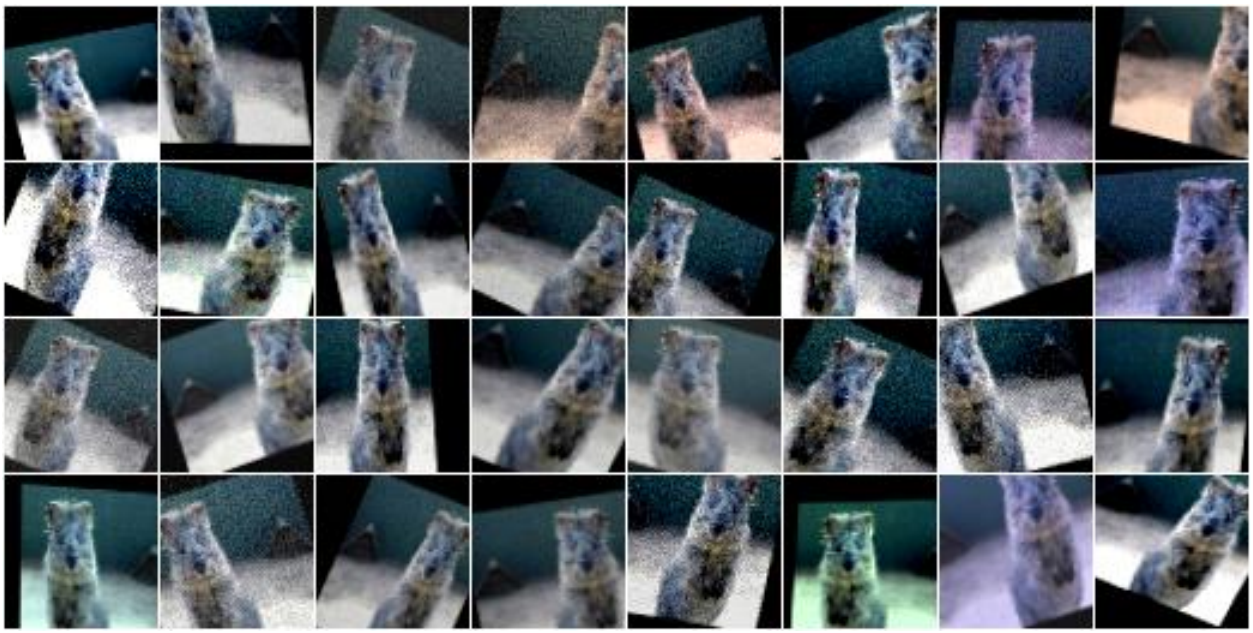

Fig. 1: Example results of the above simple augmentation sequence.

```
import numpy as np
import imgaug as ia
import imgaug.augmenters as iaa


ia.seed(1)

# Example batch of images.
# The array has shape (32, 64, 64, 3) and dtype uint8.
images = np.array(
    [ia.quokka(size=(64, 64)) for _ in range(32)],
    dtype=np.uint8
)

# Sometimes(0.5, ...) applies the given augmenter in 50% of all cases,
# e.g. Sometimes(0.5, GaussianBlur(0.3)) would blur roughly every second
# image.
sometimes = lambda aug: iaa.Sometimes(0.5, aug)

# Define our sequence of augmentation steps that will be applied to every image.
seq = iaa.Sequential(
    [
        #
        # Apply the following augmenters to most images.
        #
        iaa.Fliplr(0.5), # horizontally flip 50% of all images
        iaa.Flipud(0.2), # vertically flip 20% of all images

        # crop some of the images by 0-10% of their height/width
        sometimes(iaa.Crop(percent=(0, 0.1))),

        # Apply affine transformations to some of the images
```

(continues on next page)

(continued from previous page)

```
        # - scale to 80-120% of image height/width (each axis independently)
        # - translate by -20 to +20 relative to height/width (per axis)
        # - rotate by -45 to +45 degrees
        # - shear by -16 to +16 degrees
        # - order: use nearest neighbour or bilinear interpolation (fast)
        # - mode: use any available mode to fill newly created pixels
        #         see API or scikit-image for which modes are available
        # - cval: if the mode is constant, then use a random brightness
        #         for the newly created pixels (e.g. sometimes black,
        #         sometimes white)
        sometimes(iaa.Affine(
            scale={"x": (0.8, 1.2), "y": (0.8, 1.2)},
            translate_percent={"x": (-0.2, 0.2), "y": (-0.2, 0.2)},
            rotate=(-45, 45),
            shear=(-16, 16),
            order=[0, 1],
            cval=(0, 255),
            mode=ia.ALL
        )),

        #
        # Execute 0 to 5 of the following (less important) augmenters per
        # image. Don't execute all of them, as that would often be way too
        # strong.
        #
        iaa.SomeOf((0, 5),
            [
                # Convert some images into their superpixel representation,
                # sample between 20 and 200 superpixels per image, but do
                # not replace all superpixels with their average, only
                # some of them (p_replace).
                sometimes(
                    iaa.Superpixels(
                        p_replace=(0, 1.0),
                        n_segments=(20, 200)
                    )
                ),

                # Blur each image with varying strength using
                # gaussian blur (sigma between 0 and 3.0),
                # average/uniform blur (kernel size between 2x2 and 7x7)
                # median blur (kernel size between 3x3 and 11x11).
                iaa.OneOf([
                    iaa.GaussianBlur((0, 3.0)),
                    iaa.AverageBlur(k=(2, 7)),
                    iaa.MedianBlur(k=(3, 11)),
                ]),

                # Sharpen each image, overlay the result with the original
                # image using an alpha between 0 (no sharpening) and 1
                # (full sharpening effect).
                iaa.Sharpen(alpha=(0, 1.0), lightness=(0.75, 1.5)),

                # Same as sharpen, but for an embossing effect.
                iaa.Emboss(alpha=(0, 1.0), strength=(0, 2.0)),

                # Search in some images either for all edges or for
```

(continues on next page)

(continued from previous page)

```
            # directed edges. These edges are then marked in a black
            # and white image and overlayed with the original image
            # using an alpha of 0 to 0.7.
            sometimes(iaa.OneOf([
                iaa.EdgeDetect(alpha=(0, 0.7)),
                iaa.DirectedEdgeDetect(
                    alpha=(0, 0.7), direction=(0.0, 1.0)
                ),
            ])),

            # Add gaussian noise to some images.
            # In 50% of these cases, the noise is randomly sampled per
            # channel and pixel.
            # In the other 50% of all cases it is sampled once per
            # pixel (i.e. brightness change).
            iaa.AdditiveGaussianNoise(
                loc=0, scale=(0.0, 0.05*255), per_channel=0.5
            ),

            # Either drop randomly 1 to 10% of all pixels (i.e. set
            # them to black) or drop them on an image with 2-5% percent
            # of the original size, leading to large dropped
            # rectangles.
            iaa.OneOf([
                iaa.Dropout((0.01, 0.1), per_channel=0.5),
                iaa.CoarseDropout(
                    (0.03, 0.15), size_percent=(0.02, 0.05),
                    per_channel=0.2
                ),
            ]),

            # Invert each image's channel with 5% probability.
            # This sets each pixel value v to 255-v.
            iaa.Invert(0.05, per_channel=True), # invert color channels

            # Add a value of -10 to 10 to each pixel.
            iaa.Add((-10, 10), per_channel=0.5),

            # Change brightness of images (50-150% of original value).
            iaa.Multiply((0.5, 1.5), per_channel=0.5),

            # Improve or worsen the contrast of images.
            iaa.LinearContrast((0.5, 2.0), per_channel=0.5),

            # Convert each image to grayscale and then overlay the
            # result with the original with random alpha. I.e. remove
            # colors with varying strengths.
            iaa.Grayscale(alpha=(0.0, 1.0)),

            # In some images move pixels locally around (with random
            # strengths).
            sometimes(
                iaa.ElasticTransformation(alpha=(0.5, 3.5), sigma=0.25)
            ),

            # In some images distort local areas with varying strength.
            sometimes(iaa.PiecewiseAffine(scale=(0.01, 0.05)))
```

(continues on next page)

(continued from previous page)

```
            ],
            # do all of the above augmentations in random order
            random_order=True
        )
    ],
    # do all of the above augmentations in random order
    random_order=True
)

images_aug = seq(images=images)
```

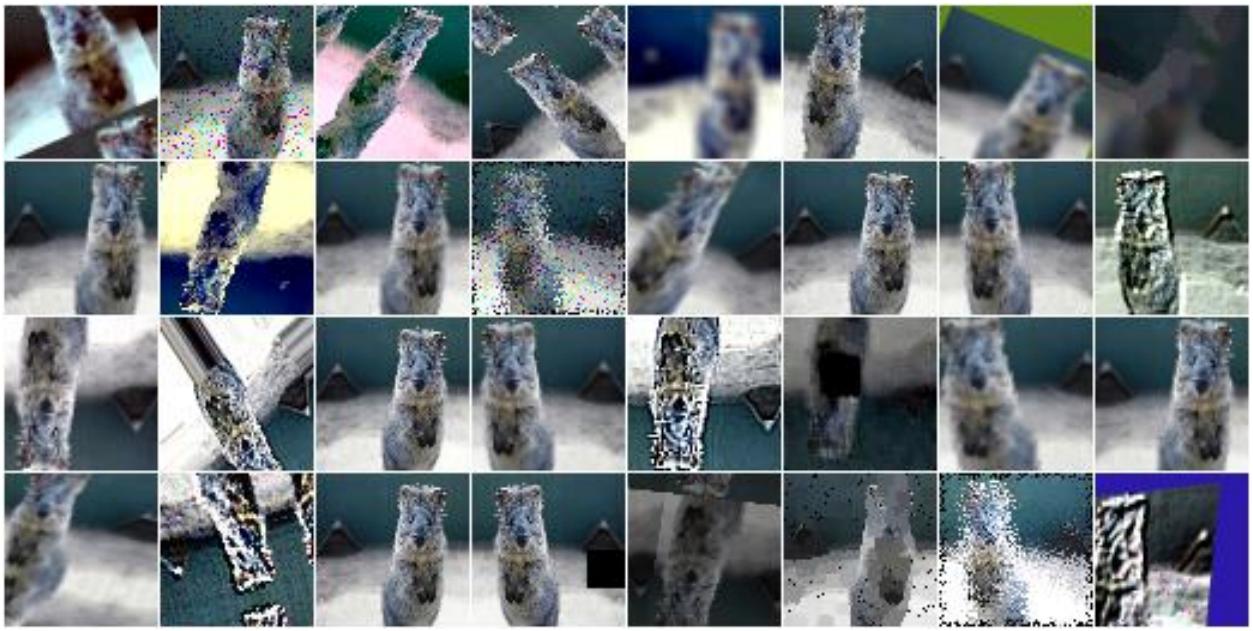

Fig. 2: Example results of the above heavy augmentation sequence.

CHAPTER 3

Examples: Keypoints

imgaug can handle not only images, but also keypoints/landmarks on these. E.g. if an image is rotated during augmentation, the library can also rotate all landmarks correspondingly.

3.1 Notebook

A jupyter notebook for keypoint augmentation is available at *Jupyter Notebooks*. The notebooks are usually more up to date and contain more examples than the ReadTheDocs documentation.

3.2 A simple example

The following example loads an image and places four keypoints on it. The image is then augmented to be brighter, slightly rotated and scaled. These augmentations are also applied to the keypoints. The image is then shown before and after augmentation (with keypoints drawn on it).

```
import imgaug as ia
import imgaug.augmenters as iaa
from imgaug.augmentables import Keypoint, KeypointsOnImage


ia.seed(1)

image = ia.quokka(size=(256, 256))
kps = KeypointsOnImage([
    Keypoint(x=65, y=100),
    Keypoint(x=75, y=200),
    Keypoint(x=100, y=100),
    Keypoint(x=200, y=80)
], shape=image.shape)

seq = iaa.Sequential([
```

(continues on next page)

(continued from previous page)

```
    iaa.Multiply((1.2, 1.5)), # change brightness, doesn't affect keypoints
    iaa.Affine(
        rotate=10,
        scale=(0.5, 0.7)
    ) # rotate by exactly 10deg and scale to 50-70%, affects keypoints
])

# Augment keypoints and images.
image_aug, kps_aug = seq(image=image, keypoints=kps)

# print coordinates before/after augmentation (see below)
# use after.x_int and after.y_int to get rounded integer coordinates
for i in range(len(kps.keypoints)):
    before = kps.keypoints[i]
    after = kps_aug.keypoints[i]
    print("Keypoint %d: (%.8f, %.8f) -> (%.8f, %.8f)" % (
        i, before.x, before.y, after.x, after.y)
    )

# image with keypoints before/after augmentation (shown below)
image_before = kps.draw_on_image(image, size=7)
image_after = kps_aug.draw_on_image(image_aug, size=7)
```

Console output of the example:

```
Keypoint 0: (65.00000000, 100.00000000) -> (97.86113503, 107.69632182)
Keypoint 1: (75.00000000, 200.00000000) -> (93.93710117, 160.01366917)
Keypoint 2: (100.00000000, 100.00000000) -> (115.85492750, 110.86911292)
Keypoint 3: (200.00000000, 80.00000000) -> (169.07878659, 109.65206321)
```

Fig. 1: Image with keypoints, before (left) and after (right) augmentation. Keypoints are shown in green and drawn in after the augmentation process.

CHAPTER 4

Examples: Bounding Boxes

imgaug offers support for bounding boxes (aka rectangles, regions of interest). E.g. if an image is rotated during augmentation, the library can also rotate all bounding boxes on it correspondingly.

Features of the library's bounding box support:

- Represent bounding boxes as objects (*imgaug.augmentables.bbs.BoundingBox*).
- Augment bounding boxes.
- Draw bounding boxes on images.
- Move/shift bounding boxes on images, project them onto other images (e.g. onto the same image after resizing), compute their intersections/unions and IoU values.

4.1 Notebook

A jupyter notebook for bounding box augmentation is available at *Jupyter Notebooks*. The notebooks are usually more up to date and contain more examples than the ReadTheDocs documentation.

4.2 A simple example

The following example loads an image and places two bounding boxes on it. The image is then augmented to be brighter, slightly rotated and scaled. These augmentations are also applied to the bounding boxes. The image is then shown before and after augmentation (with bounding boxes drawn on it).

```
import imgaug as ia
import imgaug.augmenters as iaa
from imgaug.augmentables.bbs import BoundingBox, BoundingBoxesOnImage


ia.seed(1)
```

(continues on next page)

(continued from previous page)

```
image = ia.quokka(size=(256, 256))
bbs = BoundingBoxesOnImage([
    BoundingBox(x1=65, y1=100, x2=200, y2=150),
    BoundingBox(x1=150, y1=80, x2=200, y2=130)
], shape=image.shape)

seq = iaa.Sequential([
    iaa.Multiply((1.2, 1.5)), # change brightness, doesn't affect BBs
    iaa.Affine(
        translate_px={"x": 40, "y": 60},
        scale=(0.5, 0.7)
    ) # translate by 40/60px on x/y axis, and scale to 50-70%, affects BBs
])

# Augment BBs and images.
image_aug, bbs_aug = seq(image=image, bounding_boxes=bbs)

# print coordinates before/after augmentation (see below)
# use .x1_int, .y_int, ... to get integer coordinates
for i in range(len(bbs.bounding_boxes)):
    before = bbs.bounding_boxes[i]
    after = bbs_aug.bounding_boxes[i]
    print("BB %d: (%.4f, %.4f, %.4f, %.4f) -> (%.4f, %.4f, %.4f, %.4f)" % (
        i,
        before.x1, before.y1, before.x2, before.y2,
        after.x1, after.y1, after.x2, after.y2)
    )

# image with BBs before/after augmentation (shown below)
image_before = bbs.draw_on_image(image, size=2)
image_after = bbs_aug.draw_on_image(image_aug, size=2, color=[0, 0, 255])
```

Console output of the example:

```
BB 0: (65.0000, 100.0000, 200.0000, 150.0000) -> (130.7524, 171.3311, 210.1272, 200.
↪7291)
BB 1: (150.0000, 80.0000, 200.0000, 130.0000) -> (180.7291, 159.5718, 210.1272, 188.
↪9699)
```

Note that the bounding box augmentation works by augmenting each box's edge coordinates and then drawing a bounding box around these augmented coordinates. Each of these new bounding boxes is therefore axis-aligned. This can sometimes lead to oversized new bounding boxes, especially in the case of rotation. The following image shows the result of the same code as in the example above, but *Affine* was replaced by *Affine(rotate=45)*:

4.3 Dealing with bounding boxes outside of the image

When augmenting images and their respective bounding boxes, the boxes can end up fully or partially outside of the image plane. By default, the library still returns these boxes, even though that may not be desired. The following example shows how to (a) remove bounding boxes that are fully/partially outside of the image and (b) how to clip bounding boxes that are partially outside of the image so that their are fully inside.

```
import numpy as np
import imgaug as ia
```

(continues on next page)

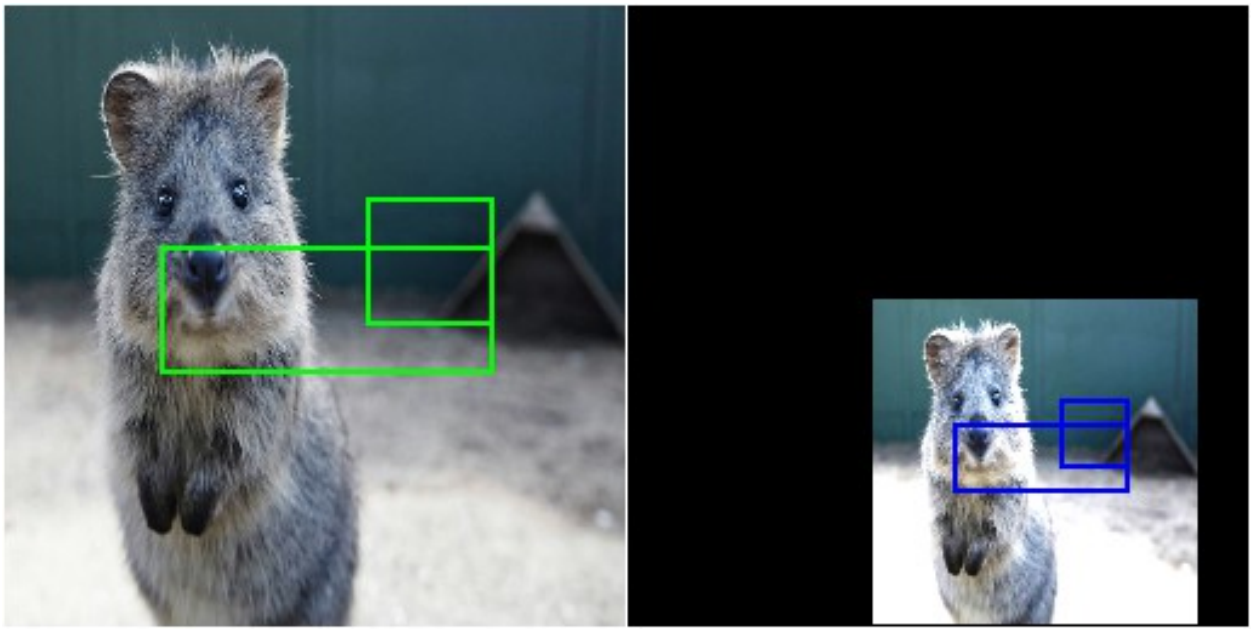

Fig. 1: Image with bounding boxes, before (left) and after (right) augmentation. Bounding boxes are shown in green (before augmentation) and blue (after augmentation).

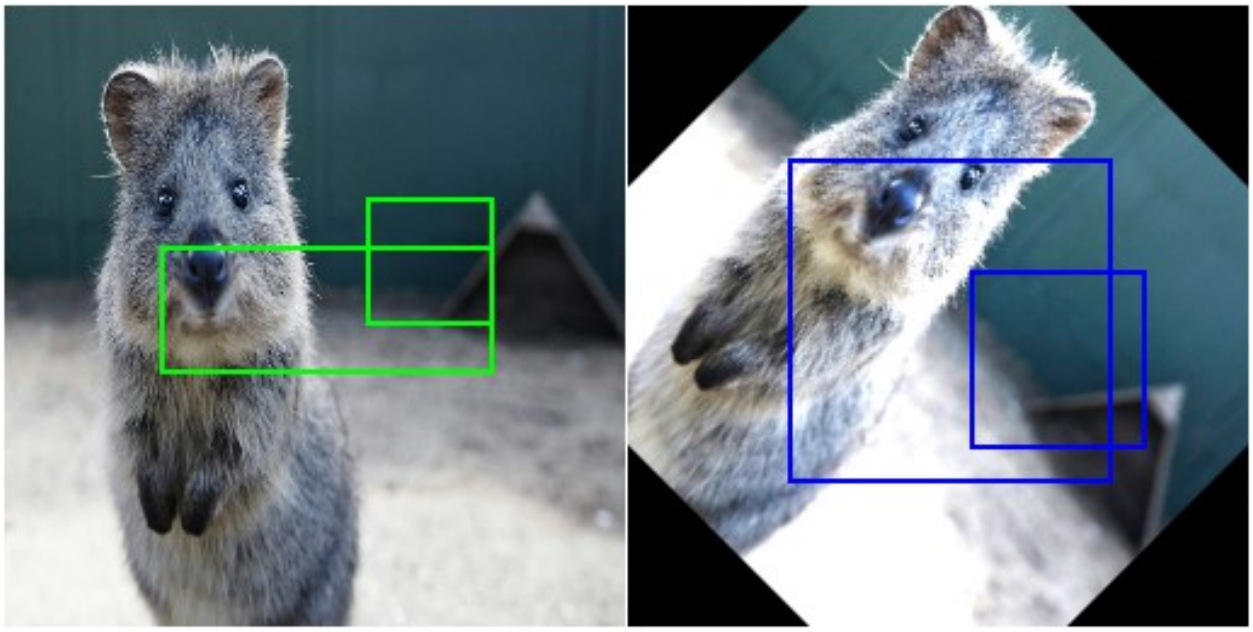

Fig. 2: Image with bounding boxes, before (left) and after (right) augmentation. The image was augmentated by rotating it by 45 degrees. The axis-aligned bounding box around the augmented keypoints ends up being oversized.

(continued from previous page)

```
import imgaug.augmenters as iaa
from imgaug.augmentables.bbs import BoundingBox, BoundingBoxesOnImage


ia.seed(1)

GREEN = [0, 255, 0]
ORANGE = [255, 140, 0]
RED = [255, 0, 0]

# Pad image with a 1px white and (BY-1)px black border
def pad(image, by):
    image_border1 = ia.pad(image, top=1, right=1, bottom=1, left=1,
                           mode="constant", cval=255)
    image_border2 = ia.pad(image_border1, top=by-1, right=by-1,
                           bottom=by-1, left=by-1,
                           mode="constant", cval=0)
    return image_border2

# Draw BBs on an image
# and before doing that, extend the image plane by BORDER pixels.
# Mark BBs inside the image plane with green color, those partially inside
# with orange and those fully outside with red.
def draw_bbs(image, bbs, border):
    image_border = pad(image, border)
    for bb in bbs.bounding_boxes:
        if bb.is_fully_within_image(image.shape):
            color = GREEN
        elif bb.is_partly_within_image(image.shape):
            color = ORANGE
        else:
            color = RED
        image_border = bb.shift(left=border, top=border)\
                         .draw_on_image(image_border, size=2, color=color)

    return image_border

# Define example image with three small square BBs next to each other.
# Augment these BBs by shifting them to the right.
image = ia.quokka(size=(256, 256))
bbs = BoundingBoxesOnImage([
    BoundingBox(x1=25, x2=75, y1=25, y2=75),
    BoundingBox(x1=100, x2=150, y1=25, y2=75),
    BoundingBox(x1=175, x2=225, y1=25, y2=75)
], shape=image.shape)

seq = iaa.Affine(translate_px={"x": 120})
image_aug, bbs_aug = seq(image=image, bounding_boxes=bbs)

# Draw the BBs (a) in their original form, (b) after augmentation,
# (c) after augmentation and removing those fully outside the image,
# (d) after augmentation and removing those fully outside the image and
# clipping those partially inside the image so that they are fully inside.
image_before = draw_bbs(image, bbs, 100)
image_after1 = draw_bbs(image_aug, bbs_aug, 100)
image_after2 = draw_bbs(image_aug, bbs_aug.remove_out_of_image(), 100)
image_after3 = draw_bbs(image_aug, bbs_aug.remove_out_of_image().clip_out_of_image(),␣
↪100)
```

(continues on next page)

(continued from previous page)

4.4 Shifting/Moving Bounding Boxes

The function *shift(top=<int>, right=<int>, bottom=<int>, left=<int>)* can be used to change the x/y position of all or specific bounding boxes.

```
import imgaug as ia
from imgaug.augmentables.bbs import BoundingBox, BoundingBoxesOnImage


ia.seed(1)

# Define image and two bounding boxes
image = ia.quokka(size=(256, 256))
bbs = BoundingBoxesOnImage([
    BoundingBox(x1=25, x2=75, y1=25, y2=75),
    BoundingBox(x1=100, x2=150, y1=25, y2=75)
], shape=image.shape)

# Move both BBs 25px to the right and the second BB 25px down
bbs_shifted = bbs.shift(left=25)
bbs_shifted.bounding_boxes[1] = bbs_shifted.bounding_boxes[1].shift(top=25)

# Draw images before/after moving BBs
image = bbs.draw_on_image(image, color=[0, 255, 0], size=2, alpha=0.75)
image = bbs_shifted.draw_on_image(image, color=[0, 0, 255], size=2, alpha=0.75)
```

4.5 Projection of BBs Onto Rescaled Images

Bounding boxes can easily be projected onto rescaled versions of the same image using the function *.on(image)*. This changes the coordinates of the bounding boxes. E.g. if the top left coordinate of the bounding box was before at x=10% and y=15%, it will still be at x/y 10%/15% on the new image, though the absolute pixel values will change depending on the height/width of the new image.

```
import imgaug as ia
from imgaug.augmentables.bbs import BoundingBox, BoundingBoxesOnImage


ia.seed(1)

# Define image with two bounding boxes
image = ia.quokka(size=(256, 256))
bbs = BoundingBoxesOnImage([
    BoundingBox(x1=25, x2=75, y1=25, y2=75),
    BoundingBox(x1=100, x2=150, y1=25, y2=75)
], shape=image.shape)

# Rescale image and bounding boxes
image_rescaled = ia.imresize_single_image(image, (512, 512))
bbs_rescaled = bbs.on(image_rescaled)
```

(continues on next page)

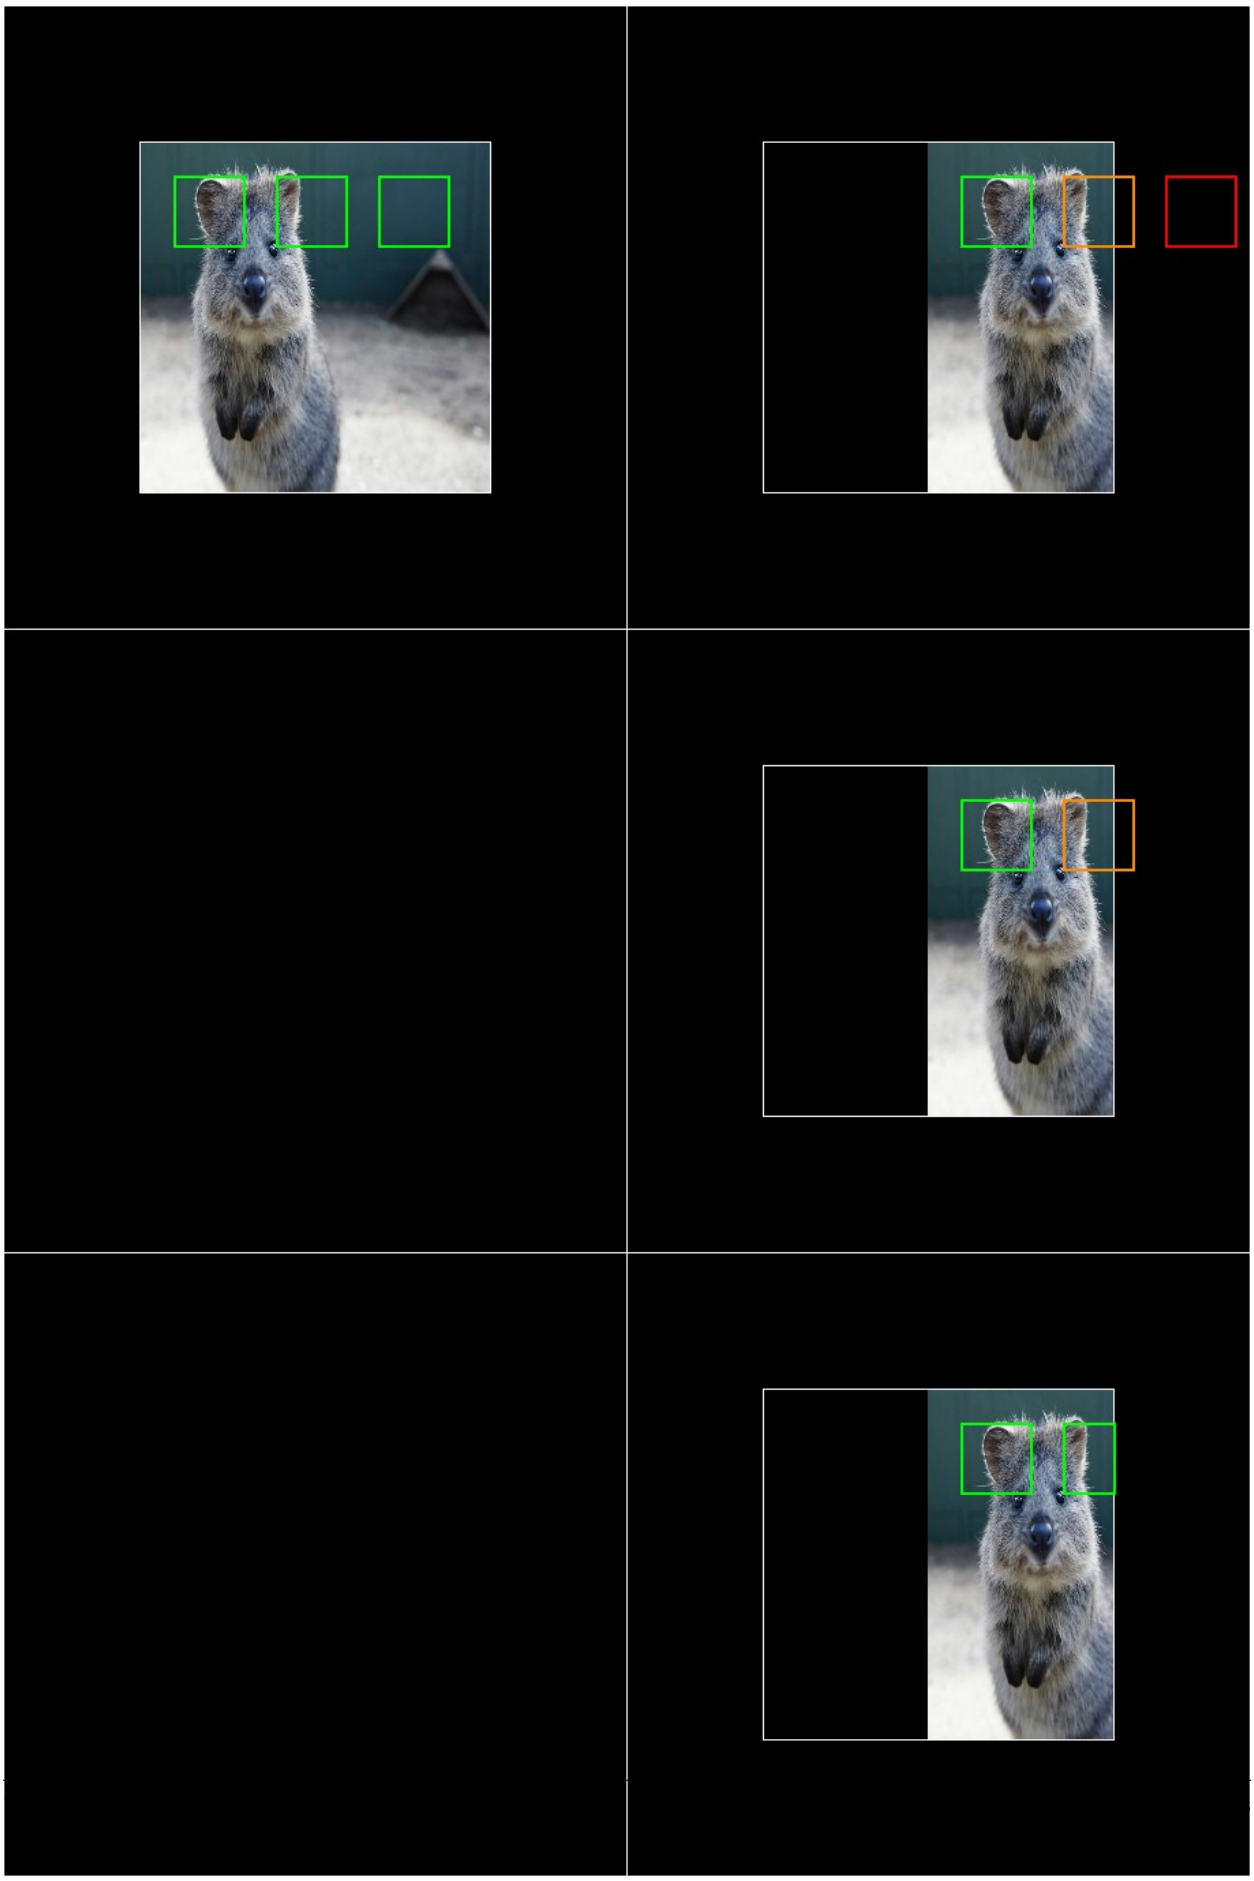

Fig. 3: Results of the above example code. Top left: Original/unaugmented image with bounding boxes (here visual

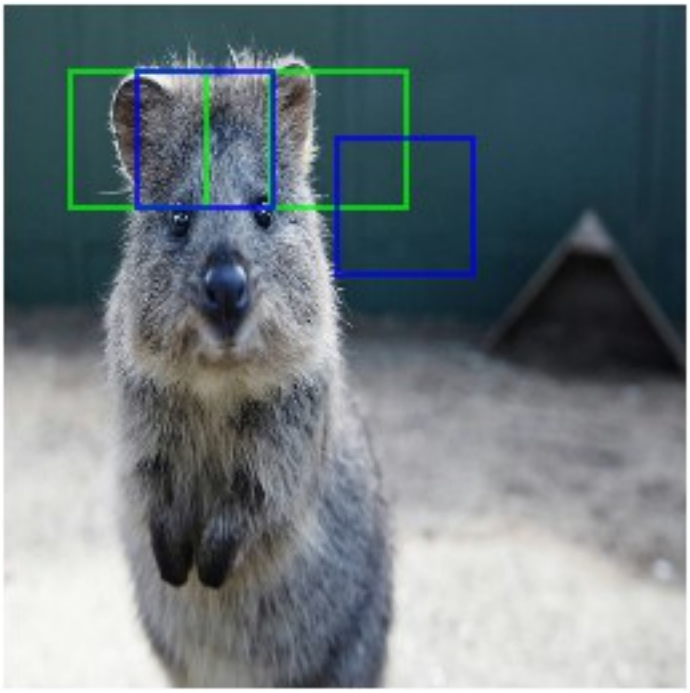

Fig. 4: Using *shift()* to move bounding boxes around (green: original BBs, blue: shifted/moved BBs).

(continued from previous page)

```
# Draw image before/after rescaling and with rescaled bounding boxes
image_bbs = bbs.draw_on_image(image, size=2)
image_rescaled_bbs = bbs_rescaled.draw_on_image(image_rescaled, size=2)
```

4.6 Computing Intersections, Unions and IoUs

Computing intersections, unions and especially IoU values (intersection over union) is common for many machine learning experiments. The library offers easy functions for that.

```
import numpy as np
import imgaug as ia
from imgaug.augmentables.bbs import BoundingBox


ia.seed(1)

# Define image with two bounding boxes.
image = ia.quokka(size=(256, 256))
bb1 = BoundingBox(x1=50, x2=100, y1=25, y2=75)
bb2 = BoundingBox(x1=75, x2=125, y1=50, y2=100)

# Compute intersection, union and IoU value
# Intersection and union are both bounding boxes. They are here
# decreased/increased in size purely for better visualization.
bb_inters = bb1.intersection(bb2).extend(all_sides=-1)
```

(continues on next page)

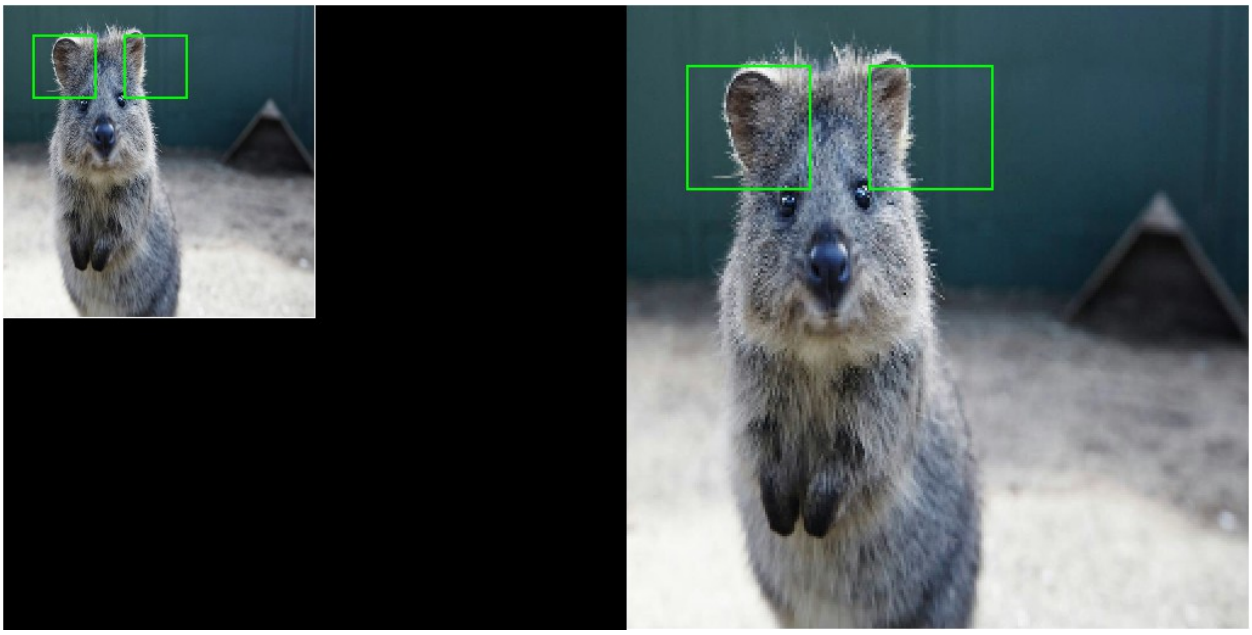

Fig. 5: Using *on()* to project bounding boxes from one image to the other, here onto an image of 2x the original size. New coordinates are determined based on their relative positions on the old image.

(continued from previous page)

```
bb_union = bb1.union(bb2).extend(all_sides=2)
iou = bb1.iou(bb2)

# Draw bounding boxes, intersection, union and IoU value on image.
image_bbs = np.copy(image)
image_bbs = bb1.draw_on_image(image_bbs, size=2, color=[0, 255, 0])
image_bbs = bb2.draw_on_image(image_bbs, size=2, color=[0, 255, 0])
image_bbs = bb_inters.draw_on_image(image_bbs, size=2, color=[255, 0, 0])
image_bbs = bb_union.draw_on_image(image_bbs, size=2, color=[0, 0, 255])
image_bbs = ia.draw_text(
    image_bbs, text="IoU=%.2f" % (iou,),
    x=bb_union.x2+10, y=bb_union.y1+bb_union.height//2,
    color=[255, 255, 255], size=13
)
```

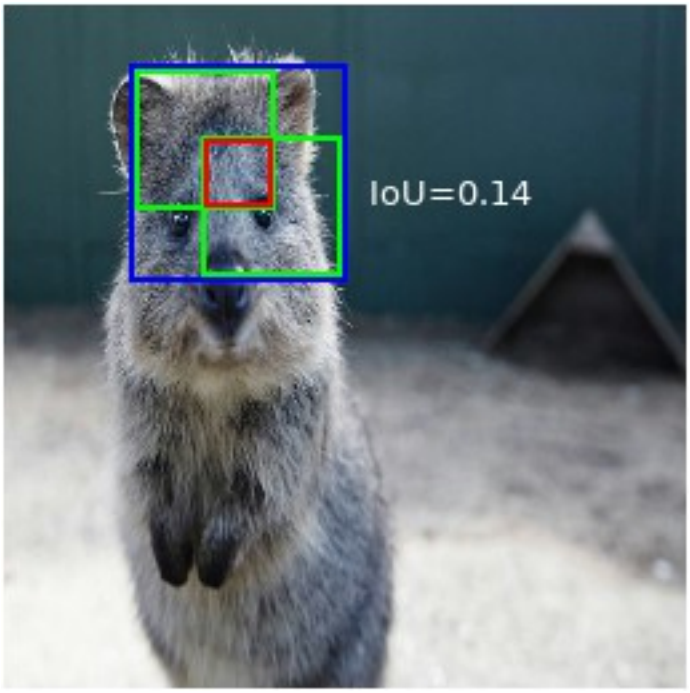

Fig. 6: Two bounding boxes on an image (green), their intersection (red, slightly shrunk), their union (blue, slightly extended) and their IoU value (white).

CHAPTER 5

Examples: Heatmaps

`imgaug` offers support for heatmap-like data. This can be used e.g. for depth map or keypoint/landmark localization maps. Heatmaps can be augmented correspondingly to images, e.g. if an image is rotated by 45°, the corresponding heatmap for that image will also be rotated by 45°.

Note:

- Heatmaps have to be bounded within value ranges, e.g. `0.0` to `1.0` for keypoint localization maps or something like `0.0` to `200.0` (meters) for depth maps. Choosing arbitrarily low/high min/max values for unbounded heatmaps is not recommended as it could lead to numerical inaccuracies.
- All augmentation functions for heatmaps are implemented under the assumption of augmenting **ground truth** data. As such, heatmaps will be affected by augmentations that change the geometry of images (e.g. affine transformations, cropping, resizing), but not by other augmentations (e.g. gaussian noise, saturation changes, grayscaling, dropout, ...).

Features of the library's heatmap support:

- Represent heatmaps as objects (`imgaug.augmentables.heatmaps.HeatmapsOnImage`).
- Augment heatmaps (only geometry-affecting augmentations, e.g. affine transformations, cropping, ...).
- Use different resolutions for heatmaps than for images (e.g. `32x32` heatmaps for `256x256` images).
- Draw heatmaps – on their own or on images (`HeatmapsOnImage.draw()`, `HeatmapsOnImage.draw_on_image()`).
- Resize, average pool or max pool heatmaps (`HeatmapsOnImage.scale()`, `HeatmapsOnImage.avg_pool()`, `HeatmapsOnImage.max_pool()`).
- Pad heatmaps by pixel amounts or to desired aspect ratios (`HeatmapsOnImage.pad()`, `HeatmapsOnImage.pad_to_aspect_ratio()`).

5.1 Notebook

A jupyter notebook for heatmap augmentation is available at *Jupyter Notebooks*. The notebooks are usually more up to date and contain more examples than the ReadTheDocs documentation.

5.2 A simple example

The following example loads a standard image and a generates a corresponding heatmap. The heatmap is supposed to be a depth map, i.e. is supposed to resemble the depth of objects in the image, where higher values indicate that objects are further away. (For simplicity we just use a simple gradient as a depth map with a cross in the center, so there is no real correspondence between the image and the depth values.)

This example shows:

- Creating heatmaps via `HeatmapsOnImage(heatmap_array, shape=image_shape)`.
- Using value ranges outside of simple `0.0` to `1.0` (here `0.0` to `50.0`) by setting `min_value` and `max_value` in the `HeatmapsOnImage` contructor.
- Resizing heatmaps, here via `HeatmapsOnImage.avg_pool(kernel_size)` (i.e. average pooling).
- Augmenting heatmaps via `Augmenter.__call__()`, which is equivalent to `Augmenter.augment()`.
- Drawing heatmaps as overlays over images `HeatmapsOnImage.draw_on_image(image)`.
- Drawing heatmaps on their own via `HeatmapsOnImage.draw()` in jet color map or via `HeatmapsOnImage.draw(cmap=None)` as intensity maps.

```
import imageio
import numpy as np
import imgaug as ia
import imgaug.augmenters as iaa
from imgaug.augmentables.heatmaps import HeatmapsOnImage

ia.seed(1)

# Load an example image (uint8, 128x128x3).
image = ia.quokka(size=(128, 128), extract="square")

# Create an example depth map (float32, 128x128).
# Here, we use a simple gradient that has low values (around 0.0)
# towards the left of the image and high values (around 50.0)
# towards the right. This is obviously a very unrealistic depth
# map, but makes the example easier.
depth = np.linspace(0, 50, 128).astype(np.float32)  # 128 values from 0.0 to 50.0
depth = np.tile(depth.reshape(1, 128), (128, 1))    # change to a horizontal gradient

# We add a cross to the center of the depth map, so that we can more
# easily see the effects of augmentations.
depth[64-2:64+2, 16:128-16] = 0.75 * 50.0  # line from left to right
depth[16:128-16, 64-2:64+2] = 1.0 * 50.0   # line from top to bottom

# Convert our numpy array depth map to a heatmap object.
# We have to add the shape of the underlying image, as that is necessary
# for some augmentations.
depth = HeatmapsOnImage(
    depth, shape=image.shape, min_value=0.0, max_value=50.0)

# To save some computation time, we want our models to perform downscaling
# and hence need the ground truth depth maps to be at a resolution of
# 64x64 instead of the 128x128 of the input image.
# Here, we use simple average pooling to perform the downscaling.
depth = depth.avg_pool(2)
```

(continues on next page)

(continued from previous page)

```
# Define our augmentation pipeline.
seq = iaa.Sequential([
    iaa.Dropout([0.05, 0.2]),      # drop 5% or 20% of all pixels
    iaa.Sharpen((0.0, 1.0)),       # sharpen the image
    iaa.Affine(rotate=(-45, 45)),  # rotate by -45 to 45 degrees (affects heatmaps)
    iaa.ElasticTransformation(alpha=50, sigma=5)  # apply water effect (affects␣
↪heatmaps)
], random_order=True)

# Augment images and heatmaps.
images_aug = []
heatmaps_aug = []
for _ in range(5):
    images_aug_i, heatmaps_aug_i = seq(image=image, heatmaps=depth)
    images_aug.append(images_aug_i)
    heatmaps_aug.append(heatmaps_aug_i)

# We want to generate an image of original input images and heatmaps
# before/after augmentation.
# It is supposed to have five columns:
# (1) original image,
# (2) augmented image,
# (3) augmented heatmap on top of augmented image,
# (4) augmented heatmap on its own in jet color map,
# (5) augmented heatmap on its own in intensity colormap.
# We now generate the cells of these columns.
#
# Note that we add a [0] after each heatmap draw command. That's because
# the heatmaps object can contain many sub-heatmaps and hence we draw
# command returns a list of drawn sub-heatmaps.
# We only used one sub-heatmap, so our lists always have one entry.
cells = []
for image_aug, heatmap_aug in zip(images_aug, heatmaps_aug):
    cells.append(image)                                                    # column 1
    cells.append(image_aug)                                                # column 2
    cells.append(heatmap_aug.draw_on_image(image_aug)[0])                  # column 3
    cells.append(heatmap_aug.draw(size=image_aug.shape[:2])[0])            # column 4
    cells.append(heatmap_aug.draw(size=image_aug.shape[:2], cmap=None)[0])  # column 5

# Convert cells to grid image and save.
grid_image = ia.draw_grid(cells, cols=5)
imageio.imwrite("example_heatmaps.jpg", grid_image)
```

5.3 Multiple sub-heatmaps per heatmaps object

The above example augmented a single heatmap with shape `(H, W)` for the example image. If you want to augment more heatmaps per image, you can simply extend the heatmap array's shape to `(H, W, C)`, where `C` is the number of heatmaps. The following example instantiates one heatmap object containing three sub-heatmaps and draws them onto the image. Heatmap augmentation would be done in the exactly same way as in the previous example.

```
import imageio
import numpy as np
import imgaug as ia
from imgaug.augmentables.heatmaps import HeatmapsOnImage
```

(continues on next page)

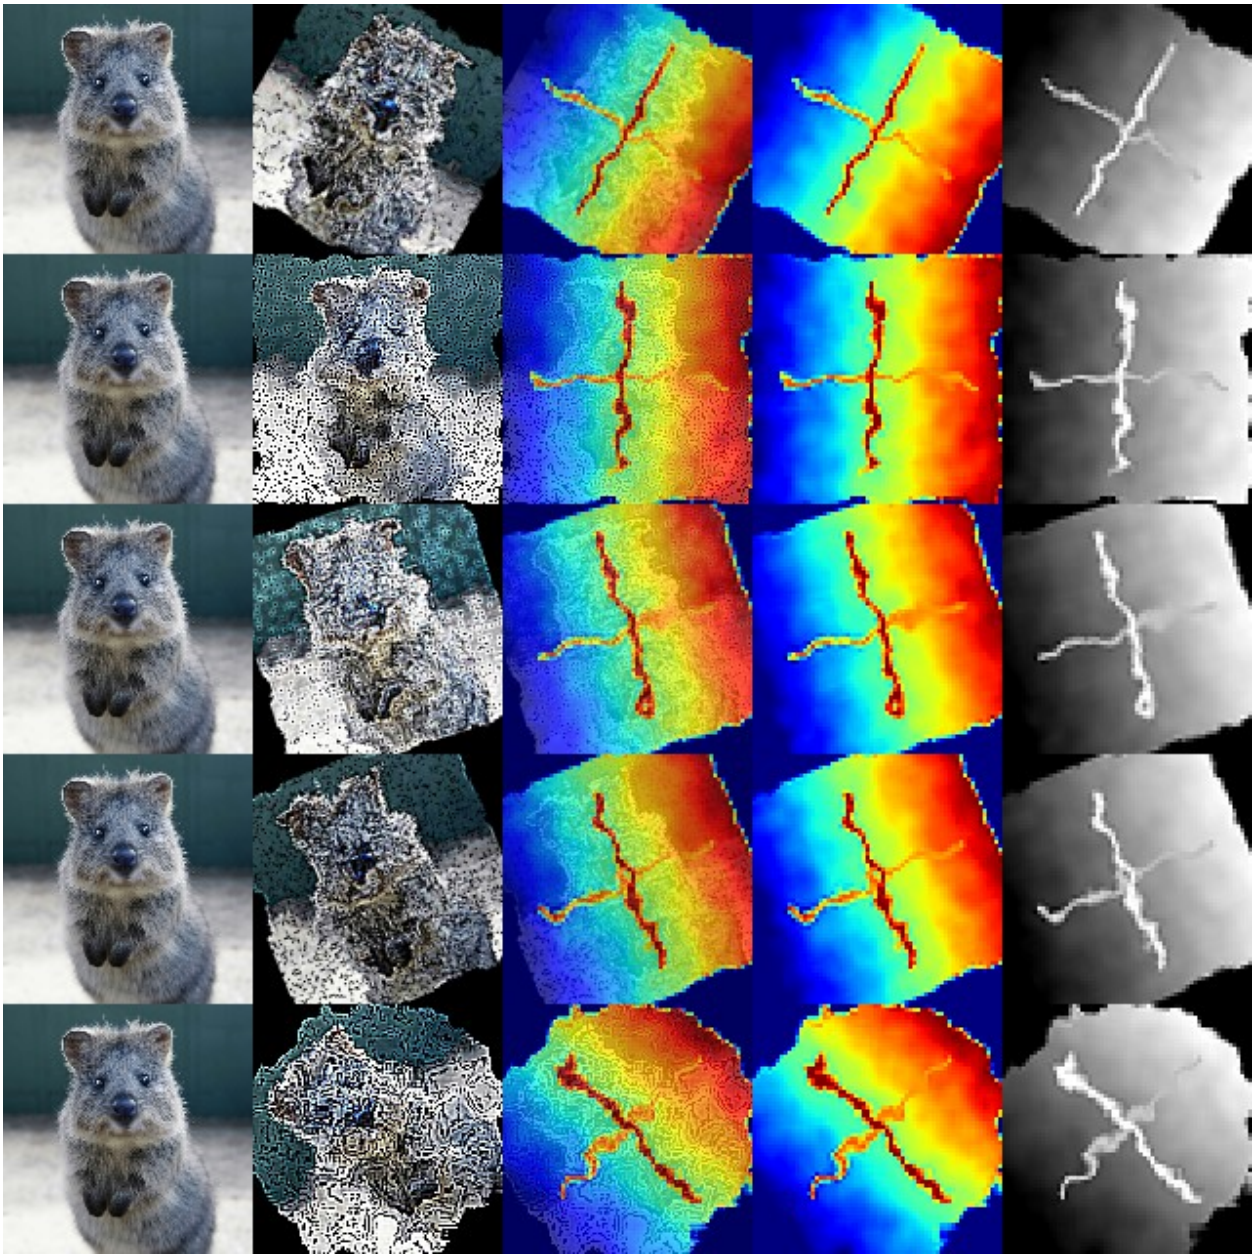

Fig. 1: Results of the above example code. Columns show: (1) Original image, (2) augmented image, (3) augmented heatmap overlayed with augmented image, (4) augmented heatmap alone in jet color map, (5) augmented heatmap alone as intensity map.

(continued from previous page)

```
# Load an image and generate a heatmap array with three sub-heatmaps.
# Each sub-heatmap contains just three horizontal lines, with one of them
# having a higher value (1.0) than the other two (0.2).
image = ia.quokka(size=(128, 128), extract="square")
heatmap = np.zeros((128, 128, 3), dtype=np.float32)
for i in range(3):
    heatmap[1*30-5:1*30+5, 10:-10, i] = 1.0 if i == 0 else 0.5
    heatmap[2*30-5:2*30+5, 10:-10, i] = 1.0 if i == 1 else 0.5
    heatmap[3*30-5:3*30+5, 10:-10, i] = 1.0 if i == 2 else 0.5
heatmap = HeatmapsOnImage(heatmap, shape=image.shape)

# Draw image and the three sub-heatmaps on it.
# We draw four columns: (1) image, (2-4) heatmaps one to three drawn on
# top of the image.
subheatmaps_drawn = heatmap.draw_on_image(image)
cells = [image, subheatmaps_drawn[0], subheatmaps_drawn[1],
         subheatmaps_drawn[2]]
grid_image = np.hstack(cells)  # Horizontally stack the images
imageio.imwrite("example_multiple_heatmaps.jpg", grid_image)
```

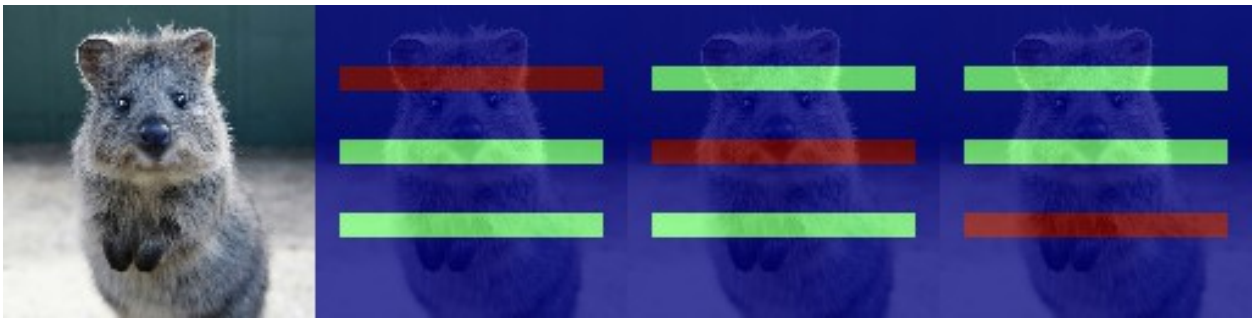

Fig. 2: Results of the above example code. It shows the original image with three heatmaps. The three heatmaps were combined in one `HeatmapsOnImage` object.

5.4 Accessing the heatmap array

After augmentation you probably want to access the heatmap's numpy array. This is done using the function `HeatmapsOnImage.get_arr()`. That functions output shape will match your original heatmap array's shape, i.e. either `(H, W)` or `(H, W, C)`. The below code shows an example, where that function's result is changed and then used to instantiate a new `HeatmapsOnImage` object.

Alternatively you could also change the heatmap object's internal array, saved as `HeatmapsOnImage.arr_0to1`. As the name indicates, it is always normalized to the range `0.0` to `1.0`, while `get_arr()` reverses that normalization. It has also always shape `(H, W, C)`, with `C>=1`.

```
import imageio
import numpy as np
import imgaug as ia
from imgaug.augmentables.heatmaps import HeatmapsOnImage

# Load an image and generate a heatmap array containing one horizontal line.
```

(continues on next page)

(continued from previous page)

```
image = ia.quokka(size=(128, 128), extract="square")
heatmap = np.zeros((128, 128, 1), dtype=np.float32)
heatmap[64-4:64+4, 10:-10, 0] = 1.0
heatmap1 = HeatmapsOnImage(heatmap, shape=image.shape)

# Extract the heatmap array from the heatmap object, change it and create
# a second heatmap.
arr = heatmap1.get_arr()
arr[10:-10, 64-4:64+4] = 0.5
heatmap2 = HeatmapsOnImage(arr, shape=image.shape)

# Draw image and heatmaps before/after changing the array.
# We draw three columns:
# (1) original image,
# (2) heatmap drawn on image,
# (3) heatmap drawn on image, with some changes made to the heatmap array.
cells = [image,
         heatmap1.draw_on_image(image)[0],
         heatmap2.draw_on_image(image)[0]]
grid_image = np.hstack(cells)  # Horizontally stack the images
imageio.imwrite("example_heatmaps_arr.jpg", grid_image)
```

Fig. 3: Results of the above example code. It shows the original image, a corresponding heatmap and again the same heatmap after its array was read out and changed.

5.5 Resizing heatmaps

When working with heatmaps it is common that the size of the input images and the heatmap sizes don't match or are supposed to not match (e.g. because predicted network output are of low resolution). `HeatmapsOnImage` offers several functions to deal with such situations: `HeatmapsOnImage.avg_pool(kernel_size)` applies average pooling to images, `HeatmapsOnImage.max_pool(kernel_size)` analogously max pooling and `HeatmapsOnImage.resize(size, [interpolation])` performs resizing. For the pooling functions the kernel size is expected to be a single integer or a tuple of two/three entries (size along each dimension). For `resize`, the size is expected to be a `(height, width)` tuple and `interpolation` can be one of the strings `nearest` (nearest neighbour interpolation), `linear`, `cubic` (default) or `area`.

The below code shows an example. It instantiates a simple `128x128` heatmap with two horizontal lines (one of which is blurred) and a small square in the center. It then applies average pooling, max pooling and resizing to heatmap sizes `64x64`, `32x32` and `16x16`. Then, an output image is generated with six rows: The first three show the results of average/max pooling and resizing, while the rows three to six show the same results after again resizing them to `128x128` using nearest neighbour upscaling.

```
import imageio
import numpy as np
import imgaug as ia
import imgaug.augmenters as iaa
from imgaug.augmentables.heatmaps import HeatmapsOnImage


def pad_by(image, amount):
    return ia.pad(image,
                  top=amount, right=amount, bottom=amount, left=amount)

def draw_heatmaps(heatmaps, upscale=False):
    drawn = []
    for heatmap in heatmaps:
        if upscale:
            drawn.append(
                heatmap.resize((128, 128), interpolation="nearest")
                       .draw()[0]
            )
        else:
            size = heatmap.get_arr().shape[0]
            pad_amount = (128-size)//2
            drawn.append(pad_by(heatmap.draw()[0], pad_amount))
    return drawn

# Generate an example heatmap with two horizontal lines (first one blurry,
# second not) and a small square.
heatmap = np.zeros((128, 128, 1), dtype=np.float32)
heatmap[32-4:32+4, 10:-10, 0] = 1.0
heatmap = iaa.GaussianBlur(3.0).augment_image(heatmap)
heatmap[96-4:96+4, 10:-10, 0] = 1.0
heatmap[64-2:64+2, 64-2:64+2, 0] = 1.0
heatmap = HeatmapsOnImage(heatmap, shape=(128, 128, 1))

# Scale the heatmaps using average pooling, max pooling and resizing with
# default interpolation (cubic).
avg_pooled = [heatmap, heatmap.avg_pool(2), heatmap.avg_pool(4),
              heatmap.avg_pool(8)]
max_pooled = [heatmap, heatmap.max_pool(2), heatmap.max_pool(4),
              heatmap.max_pool(8)]
resized = [heatmap, heatmap.resize((64, 64)), heatmap.resize((32, 32)),
           heatmap.resize((16, 16))]

# Draw an image of all scaled heatmaps.
cells = draw_heatmaps(avg_pooled)\
    + draw_heatmaps(max_pooled)\
    + draw_heatmaps(resized)\
    + draw_heatmaps(avg_pooled, upscale=True)\
    + draw_heatmaps(max_pooled, upscale=True)\
    + draw_heatmaps(resized, upscale=True)
grid_image = ia.draw_grid(cells, cols=4)
imageio.imwrite("example_heatmaps_scaling.jpg", grid_image)
```

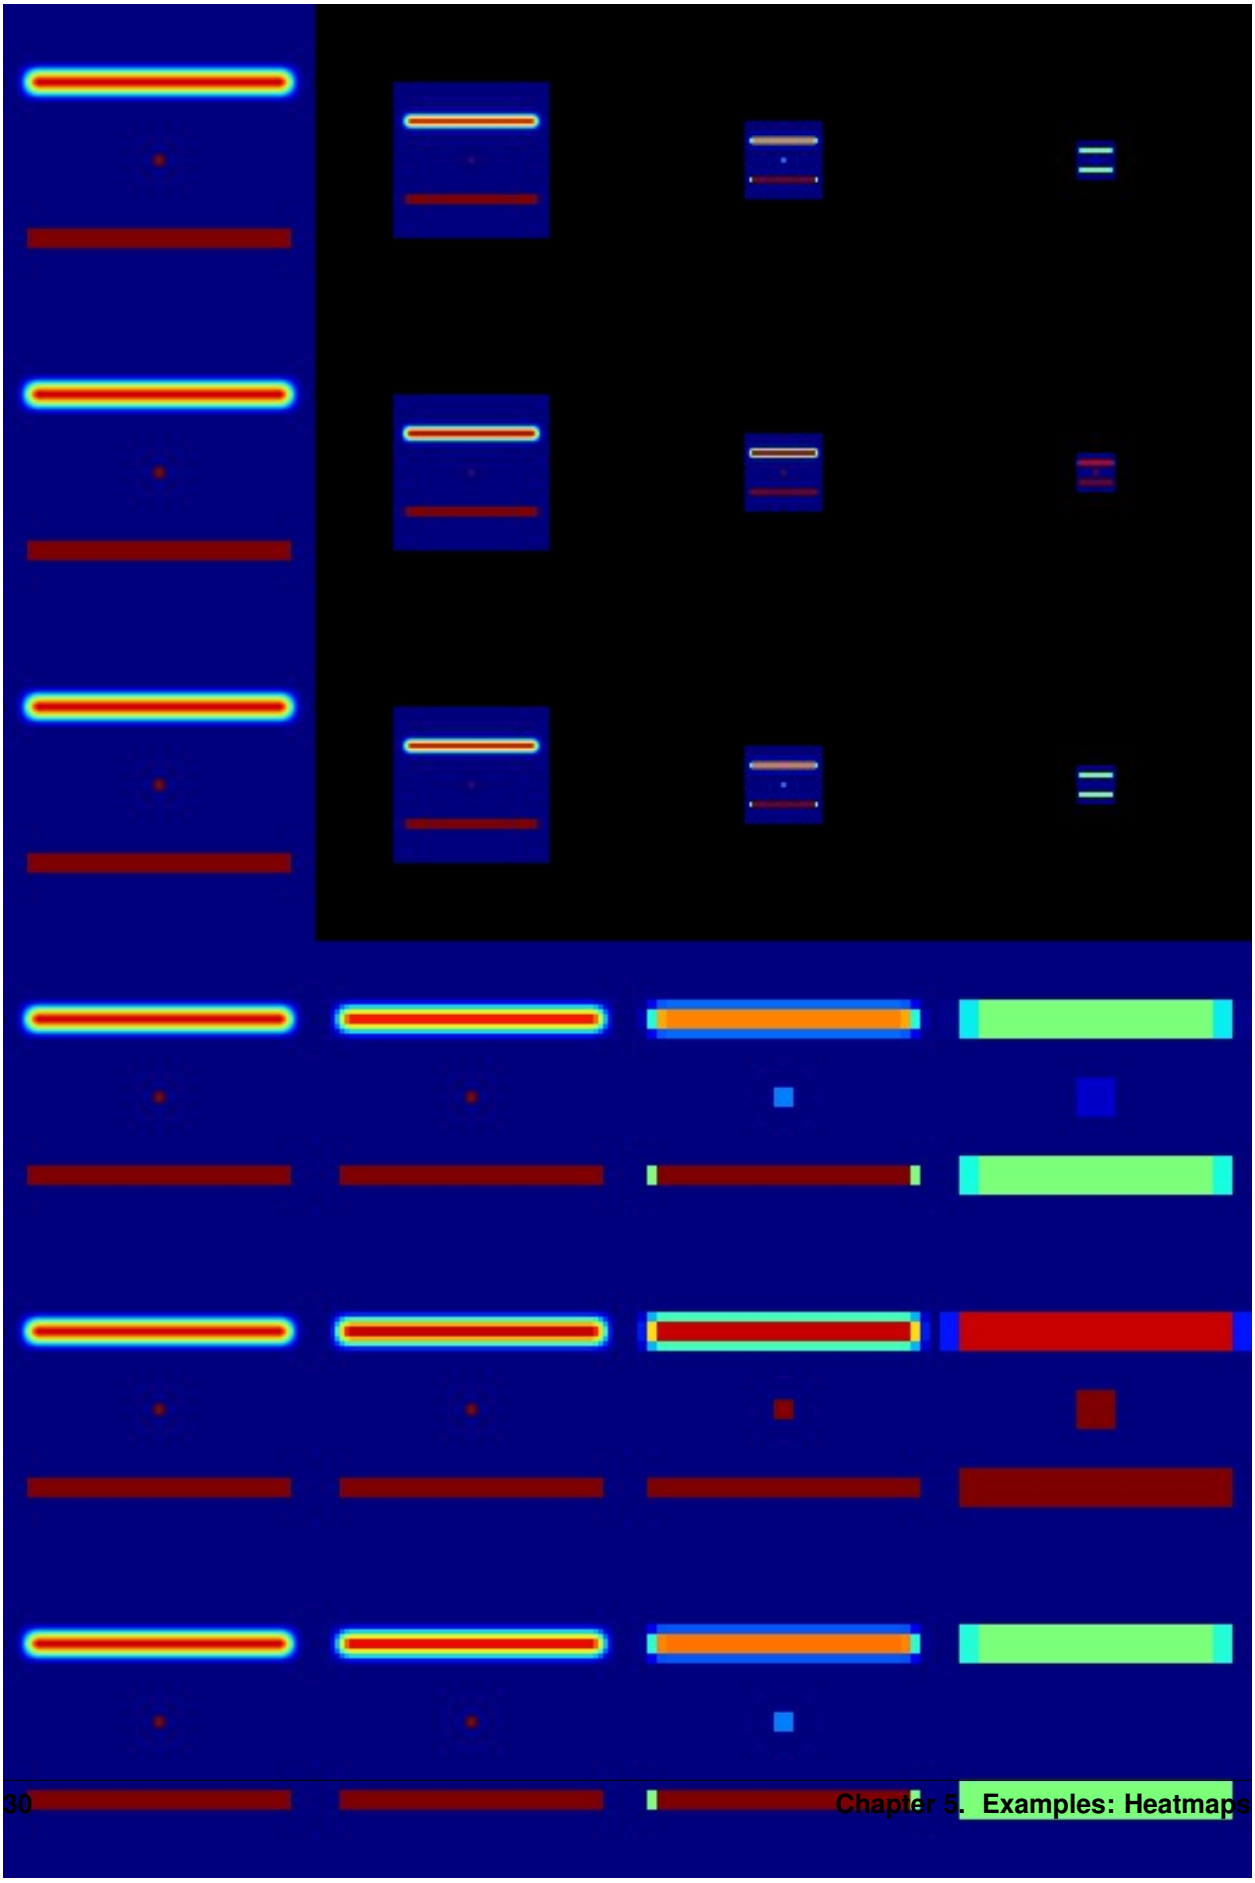

Fig. 4: Results of the above example code. It shows six rows: (Rows 1-3) scaling via average pooling, max pooling and

5.6 Padding heatmaps

Another common operation is padding of images and heatmaps, especially to squared sizes. This is done for images using `imgaug.pad(image, [top], [right], [bottom], [left], [mode], [cval])` and `imgaug.pad_to_aspect_ratio(image, aspect_ratio, [mode], [cval], [return_pad_amounts])`. For heatmaps it is done using `HeatmapsOnImage.pad([top], [right], [bottom], [left], [mode], [cval])` and `HeatmapsOnImage.pad_to_aspect_ratio(aspect_ratio, [mode], [cval], [return_pad_amounts])`. In both cases, `pad()` expects pixel amounts (i.e. integers) and `pad_to_aspect_ratio()` the target aspect ratio, given as a float denoting `ratio = width / height` (i.e. a value of `1.0` would lead to a squared image/heatmap, while `2.0` would lead to a fairly wide image/heatmap).

The below code shows an example for padding. It starts with a squared sized image and heatmap, cuts both so that they are more wide than high and then zero-pads both back to squared size.

```
import imageio
import numpy as np
import imgaug as ia
from imgaug.augmentables.heatmaps import HeatmapsOnImage

# Load example image and generate example heatmap with one horizontal line
image = ia.quokka((128, 128), extract="square")
heatmap = np.zeros((128, 128, 1), dtype=np.float32)
heatmap[64-4:64+4, 10:-10, 0] = 1.0

# Cut image and heatmap so that they are no longer squared
image = image[32:-32, :, :]
heatmap = heatmap[32:-32, :, :]

heatmap = HeatmapsOnImage(heatmap, shape=(128, 128, 1))

# Pad images and heatmaps by pixel amounts or to aspect ratios
# We pad both back to squared size of 128x128
images_padded = [
    ia.pad(image, top=32, bottom=32),
    ia.pad_to_aspect_ratio(image, 1.0)
]
heatmaps_padded = [
    heatmap.pad(top=32, bottom=32),
    heatmap.pad_to_aspect_ratio(1.0)
]

# Draw an image of all padded images and heatmaps
cells = [
    images_padded[0],
    heatmaps_padded[0].draw_on_image(images_padded[0])[0],
    images_padded[1],
    heatmaps_padded[1].draw_on_image(images_padded[1])[0]
]

grid_image = ia.draw_grid(cells, cols=2)
imageio.imwrite("example_heatmaps_padding.jpg", grid_image)
```

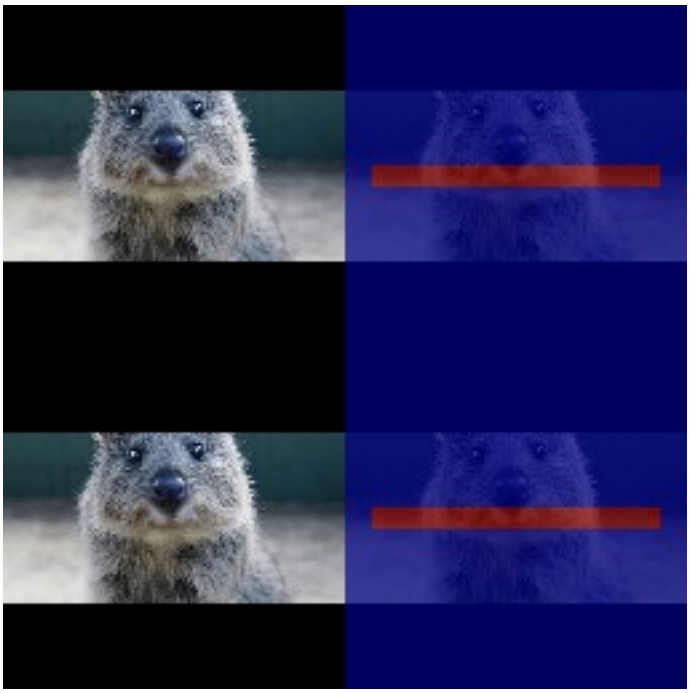

Fig. 5: Results of the above example code. It shows an input image and a heatmap that were both first cut to `64x128` and then padded back to squared size of `128x128`. First row uses `pad()`, second uses `pad_to_aspect_ratio()`.

CHAPTER 6

Examples: Segmentation Maps and Masks

`imgaug` offers support for segmentation map data, such as semantic segmentation maps, instance segmentation maps or ordinary masks. Segmentation maps can be augmented correspondingly to images. E.g. if an image is rotated by 45°, the corresponding segmentation map for that image will also be rotated by 45°.

Note: All augmentation functions for segmentation maps are implemented under the assumption of augmenting **ground truth** data. As such, segmentation maps will be affected by augmentations that change the geometry of images (e.g. affine transformations, cropping, resizing), but not by other augmentations (e.g. gaussian noise, saturation changes, grayscaling, dropout, ...).

Features of the library's segmentation map support:

- Represent segmentation maps as objects (`imgaug.augmentables.segmaps.SegmentationMapsOnImage`).
- Support integer maps (integer dtypes, usually `int32`) and boolean masks (dtype `numpy.bool_`).
- Augment segmentation maps (only geometry-affecting augmentations, e.g. affine transformations, cropping, ...).
- Use different resolutions for segmentation maps and images (e.g. `32x32` segmentation maps and `256x256` for the corresponding images).
- Draw segmentation maps – on their own or on images (`SegmentationMapsOnImage.draw()`, `SegmentationMapsOnImage.draw_on_image()`).
- Resize segmentation maps (`SegmentationMapsOnImage.resize()`).
- Pad segmentation maps by pixel amounts or to desired aspect ratios (`SegmentationMapsOnImage.pad()`, `SegmentationMapsOnImage.pad_to_aspect_ratio()`).

6.1 Notebook

A jupyter notebook for segmentation map augmentation is available at *Jupyter Notebooks*. The notebooks are usually more up to date and contain more examples than the ReadTheDocs documentation.

6.2 A simple example

The following example loads a standard image and defines a corresponding `int32` segmentation map. The image and segmentation map are augmented in the same way and the results are visualized.

```
import imageio
import numpy as np
import imgaug as ia
import imgaug.augmenters as iaa
from imgaug.augmentables.segmaps import SegmentationMapsOnImage

ia.seed(1)

# Load an example image (uint8, 128x128x3).
image = ia.quokka(size=(128, 128), extract="square")

# Define an example segmentation map (int32, 128x128).
# Here, we arbitrarily place some squares on the image.
# Class 0 is our intended background class.
segmap = np.zeros((128, 128, 1), dtype=np.int32)
segmap[28:71, 35:85, 0] = 1
segmap[10:25, 30:45, 0] = 2
segmap[10:25, 70:85, 0] = 3
segmap[10:110, 5:10, 0] = 4
segmap[118:123, 10:110, 0] = 5
segmap = SegmentationMapsOnImage(segmap, shape=image.shape)

# Define our augmentation pipeline.
seq = iaa.Sequential([
    iaa.Dropout([0.05, 0.2]),      # drop 5% or 20% of all pixels
    iaa.Sharpen((0.0, 1.0)),       # sharpen the image
    iaa.Affine(rotate=(-45, 45)),  # rotate by -45 to 45 degrees (affects segmaps)
    iaa.ElasticTransformation(alpha=50, sigma=5)  # apply water effect (affects␣
↪segmaps)
], random_order=True)

# Augment images and segmaps.
images_aug = []
segmaps_aug = []
for _ in range(5):
    images_aug_i, segmaps_aug_i = seq(image=image, segmentation_maps=segmap)
    images_aug.append(images_aug_i)
    segmaps_aug.append(segmaps_aug_i)

# We want to generate an image containing the original input image and
# segmentation maps before/after augmentation. (Both multiple times for
# multiple augmentations.)
#
# The whole image is supposed to have five columns:
# (1) original image,
# (2) original image with segmap,
# (3) augmented image,
# (4) augmented segmap on augmented image,
# (5) augmented segmap on its own in.
#
# We now generate the cells of these columns.
#
```

(continues on next page)

(continued from previous page)

```
# Note that draw_on_image() and draw() both return lists of drawn
# images. Assuming that the segmentation map array has shape (H,W,C),
# the list contains C items.
cells = []
for image_aug, segmap_aug in zip(images_aug, segmaps_aug):
    cells.append(image)                                       # column 1
    cells.append(segmap.draw_on_image(image)[0])              # column 2
    cells.append(image_aug)                                   # column 3
    cells.append(segmap_aug.draw_on_image(image_aug)[0])      # column 4
    cells.append(segmap_aug.draw(size=image_aug.shape[:2])[0])  # column 5

# Convert cells to a grid image and save.
grid_image = ia.draw_grid(cells, cols=5)
imageio.imwrite("example_segmaps.jpg", grid_image)
```

6.3 Using boolean masks

In order to augment masks, you can simply use boolean arrays. Everything else is identical to `int32` maps. The below code shows an example and is very similar to the previous code for `int32` maps. It noteably changes `np.zeros((128, 128, 1), dtype=np.int32)` to `np.zeros((128, 128, 1), dtype=bool)`.

```
import imageio
import numpy as np
import imgaug as ia
from imgaug.augmentables.segmaps import SegmentationMapsOnImage

# Load an example image (uint8, 128x128x3).
image = ia.quokka(size=(128, 128), extract="square")

# Create an example mask (bool, 128x128).
# Here, we arbitrarily place a square on the image.
segmap = np.zeros((128, 128, 1), dtype=bool)
segmap[28:71, 35:85, 0] = True
segmap = SegmentationMapsOnImage(segmap, shape=image.shape)

# Draw three columns: (1) original image,
# (2) original image with mask on top, (3) only mask
cells = [
    image,
    segmap.draw_on_image(image)[0],
    segmap.draw(size=image.shape[:2])[0]
]

# Convert cells to a grid image and save.
grid_image = ia.draw_grid(cells, cols=3)
imageio.imwrite("example_segmaps_bool.jpg", grid_image)
```

6.4 Accessing the segmentation map array

After augmentation it is often desired to re-access the segmentation map array. This can be done using `SegmentationMapsOnImage.get_arr()`, which returns a segmentation map array with the same shape and dtype as was originally provided as `arr` to `SegmentationMapsOnImage(arr, ...)`.

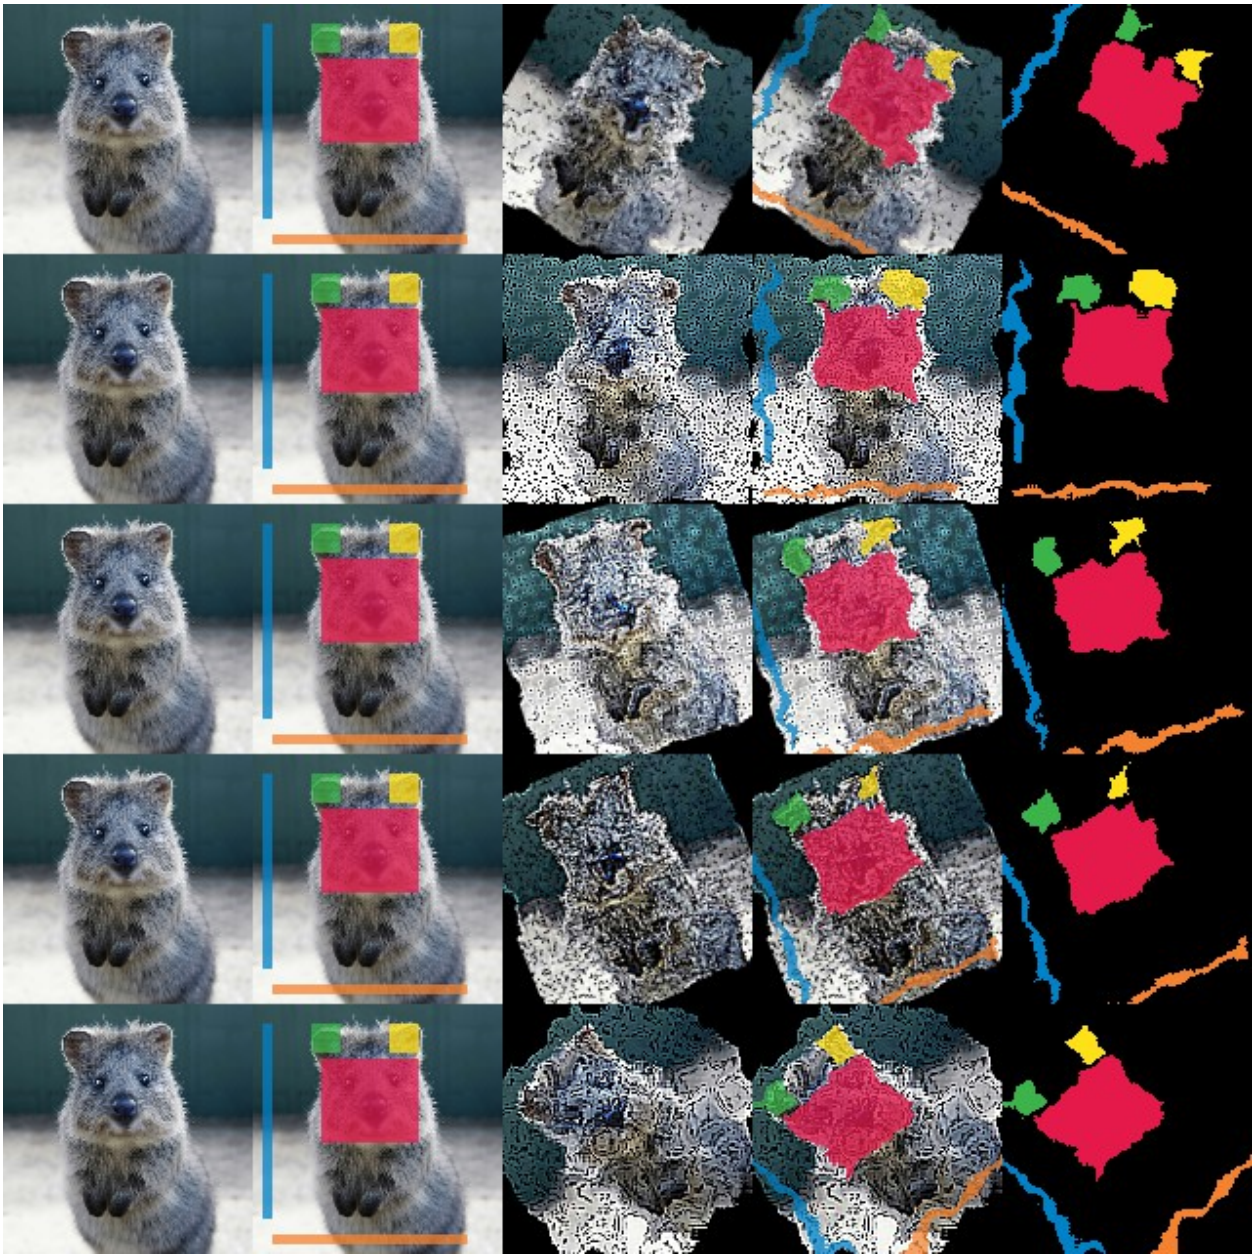

Fig. 1: Results of the above example code. Columns show: (1) Original image, (2) original segmentation map drawn on original image, (3) augmented image, (4) augmented segmentation map drawn on augmented image, (5) augmented segmentation map drawn on its own.

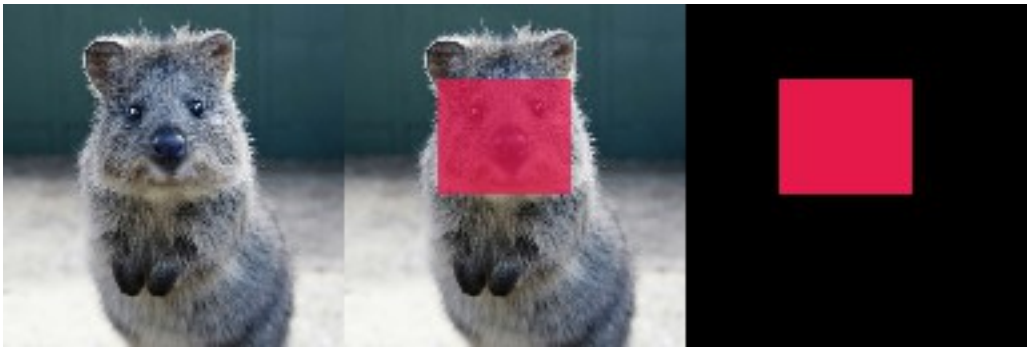

Fig. 2: Results of the above example code. Columns show: (1) Original image, (2) boolean segmentation map (i.e. mask) drawn on image, (3) boolean segmentation map drawn on its own.

The below code shows an example that accesses and changes the array.

```
import imageio
import numpy as np
import imgaug as ia
from imgaug.augmentables.segmaps import SegmentationMapsOnImage

# Load an example image (uint8, 128x128x3).
image = ia.quokka(size=(128, 128), extract="square")

# Create an example segmentation map (int32, 128x128).
# Here, we arbitrarily place some squares on the image.
# Class 0 is the background class.
segmap = np.zeros((128, 128, 1), dtype=np.int32)
segmap[28:71, 35:85, 0] = 1
segmap[10:25, 30:45, 0] = 2
segmap[10:25, 70:85, 0] = 3
segmap[10:110, 5:10, 0] = 4
segmap[118:123, 10:110, 0] = 5
segmap1 = SegmentationMapsOnImage(segmap, shape=image.shape)

# Read out the segmentation map's array, change it and create a new
# segmentation map
arr = segmap1.get_arr()
arr[10:110, 5:10, 0] = 5
segmap2 = ia.SegmentationMapsOnImage(arr, shape=image.shape)

# Draw three columns: (1) original image, (2) original image with
# unaltered segmentation map on top, (3) original image with altered
# segmentation map on top
cells = [
    image,
    segmap1.draw_on_image(image)[0],
    segmap2.draw_on_image(image)[0]
]

# Convert cells to grid image and save.
grid_image = ia.draw_grid(cells, cols=3)
imageio.imwrite("example_segmaps_array.jpg", grid_image)
```

Fig. 3: Results of the above example code. Columns show: (1) Original image, (2) original segmentation map drawn on original image, (3) segmentation map with modified array drawn on image.

6.5 Resizing and padding

Segmentation maps can easily be resized and padded. The methods are identical to the ones used for heatmaps (see :doc:`examples_heatmaps`), though segmentation maps do not offer resizing via average or max pooling. The `resize()` method also defaults to nearest neighbour interpolation (instead of `cubic` interpolation) and it is recommended to not change that.

The functions for resizing and padding are:

- `SegmentationMapsOnImage.resize(sizes, interpolation="nearest")`: Resizes to `sizes` given as a tuple `(height, width)`. Interpolation can be `nearest`, `linear`, `cubic` and `area`, but only `nearest` is actually recommended.
- `SegmentationMapsOnImage.pad(top=0, right=0, bottom=0, left=0, mode="constant", cval=0)`: Pads the segmentation map by given pixel amounts. Uses by default constant value padding with value `0`, i.e. zero-padding. Possible padding modes are the same as for `numpy.pad()`, i.e. `constant`, `edge`, `linear_ramp`, `maximum`, `mean`, `median`, `minimum`, `reflect`, `symmetric` and `wrap`.
- `SegmentationMapsOnImage.pad_to_aspect_ratio(aspect_ratio, mode="constant", cval=0, return_pad_amounts=False)`: Same as `pad()`, but pads an image towards a desired aspect ratio (`ratio = width / height`). E.g. use `1.0` for squared segmentation maps or `2.0` for maps that are twice as wide as they are high.

CHAPTER 7

Stochastic Parameters

7.1 Introduction

When augmenting images during experiments, usually one wants to augment each image in different ways. E.g. when rotating images, not every image is supposed to be rotated by 10 degrees. Instead, only some are supposed to be rotated by 10 degrees, while others should be rotated by 17 degrees or 5 degrees or -12 degrees - and so on. This can be achieved using random functions, but reimplementing these, making sure that they generate the expected values and getting them to work with *determinism* is cumbersome. To avoid all of this work, the library uses *Stochastic Parameters*. These are usually abstract representations of probability distributions, e.g. the normal distribution *N(0, 1.0)* or the uniform range *[0.0, 10.0]*. Basically all augmenters accept these stochastic parameters, making it easy to control value ranges. They are all adapted to work with *determinism* out of the box.

The below code shows their usage:

```
from imgaug import augmenters as iaa
from imgaug import parameters as iap

seq = iaa.Sequential([
    iaa.GaussianBlur(
        sigma=iap.Uniform(0.0, 1.0)
    ),
    iaa.ContrastNormalization(
        iap.Choice(
            [1.0, 1.5, 3.0],
            p=[0.5, 0.3, 0.2]
        )
    ),
    iaa.Affine(
        rotate=iap.Normal(0.0, 30),
        translate_px=iap.RandomSign(iap.Poisson(3))
    ),
    iaa.AddElementwise(
        iap.Discretize(
            (iap.Beta(0.5, 0.5) * 2 - 1.0) * 64
```

(continues on next page)

(continued from previous page)

```
        )
    ),
    iaa.Multiply(
        iap.Positive(iap.Normal(0.0, 0.1)) + 1.0
    )
])
```

The example does the following:

- Blur each image by *sigma*, where *sigma* is sampled from the uniform range *[0.0, 1.0)*. Example values: *0.053, 0.414, 0.389, 0.277, 0.981*.
- Increase the contrast either to 100% (50% chance of being chosen) or by 150% (30% chance of being chosen) or 300% (20% chance of being chosen).
- Rotate each image by a random amount of degrees, where the degree is sampled from the normal distribution *N(0, 30)*. Most of the values will be in the range *-60* to *60*.
- Translate each image by n pixels, where n is sampled from a poisson distribution with alpha=3 (pick should be around x=3). As we cant translate by a fraction of a pixel, we pick a discrete distribution here, which poisson is. However, we do not just want to translate towards the right/top (only positive values). So we randomly flip the sign sometimes to get negative pixel amounts too.
- Add to each pixel a random value, sampled from the beta distribution *Beta(0.5, 0.5)*. This distribution has its peaks around 0.0 and 1.0. We multiply this with 2 and subtract 1 to get it into the range *[-1, 1]*. Then we multiply by 64 to get the range *[-64, 64]*. As we beta distribution is continuous, we convert it to a discrete distribution. The result is that a lot of pixel intensities are shifted by -64 or 64 (or a value very close to these two). Some other pixel intensities are kept (mostly) at their old values.
- We use Multiply to make each image brighter. The brightness increase is sampled from a normal distribution, converted to have only positive values. So most values are expected to be in the range *0.0* to *0.2*. We add *1.0* to set the brightness to *1.0* (100%) to *1.2* (120%).

7.2 Continuous Probability Distributions

The following continuous probability distributions are available:

- *Normal(loc, scale)*: The popular normal distribution with mean *loc* and standard deviation *scale*. Example:

```
from imgaug import parameters as iap
params = [
    iap.Normal(0, 1),
    iap.Normal(5, 3),
    iap.Normal(iap.Choice([-3, 3]), 1),
    iap.Normal(iap.Uniform(-3, 3), 1)
]
iap.show_distributions_grid(params)
```

- *Laplace(loc, scale)*: Similarly shaped to a normal distribution. Has its peak at *loc* and width *scale*. Example:

```
from imgaug import parameters as iap
params = [
    iap.Laplace(0, 1),
    iap.Laplace(5, 3),
    iap.Laplace(iap.Choice([-3, 3]), 1),
    iap.Laplace(iap.Uniform(-3, 3), 1)
```

(continues on next page)

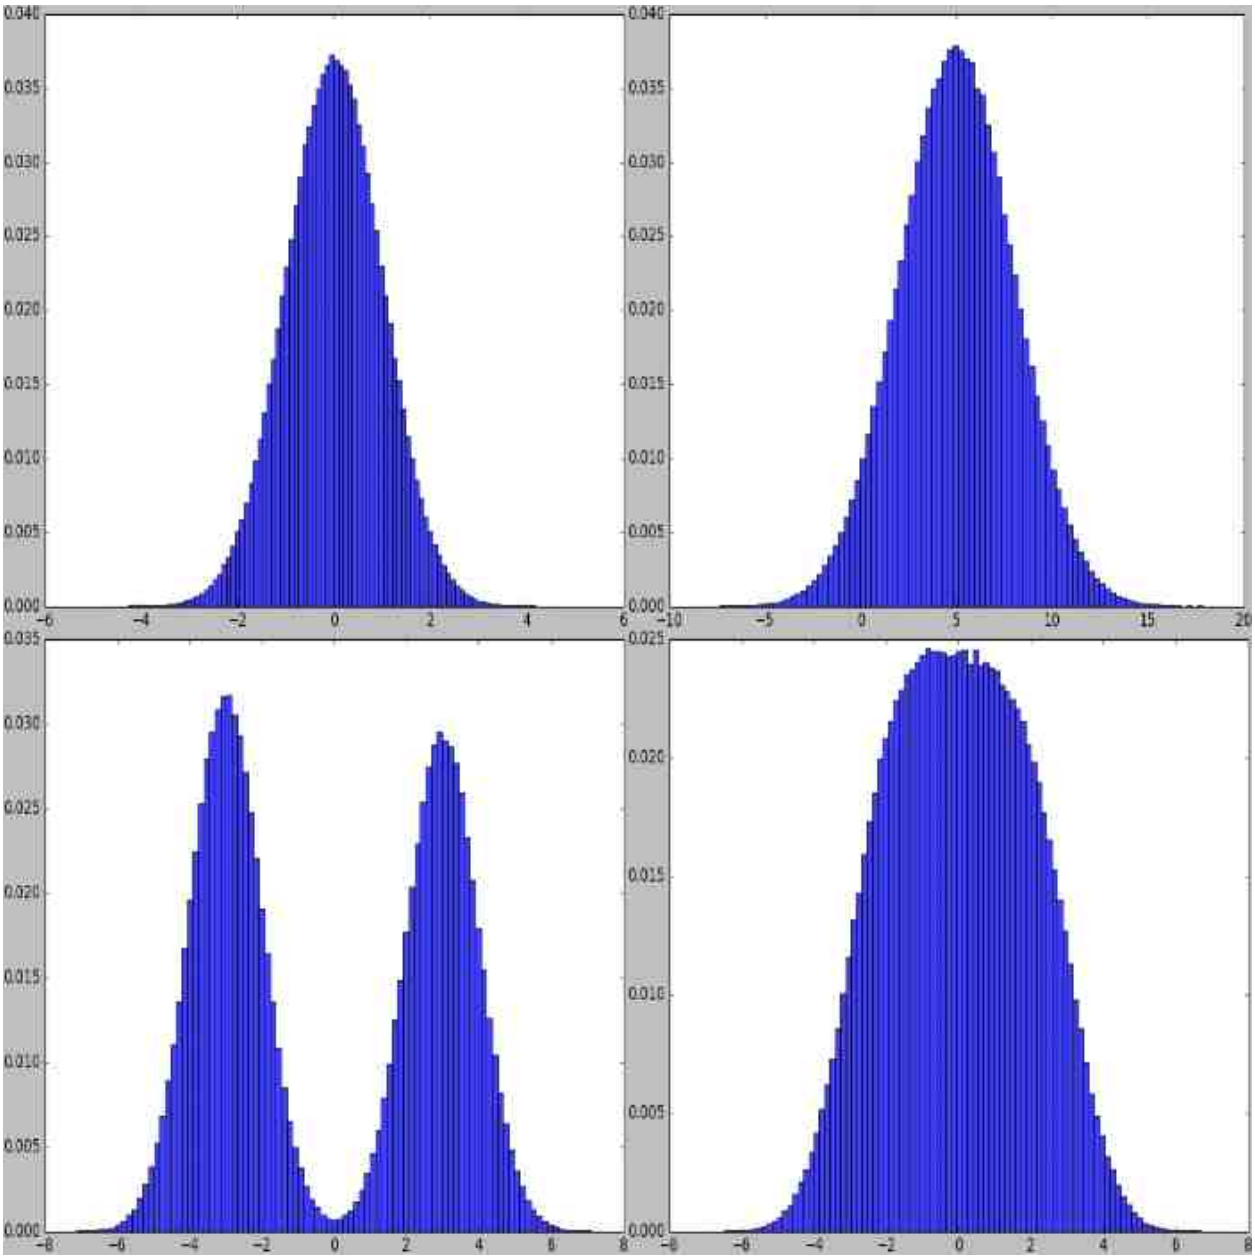

(continued from previous page)

```
]
iap.show_distributions_grid(params)
```

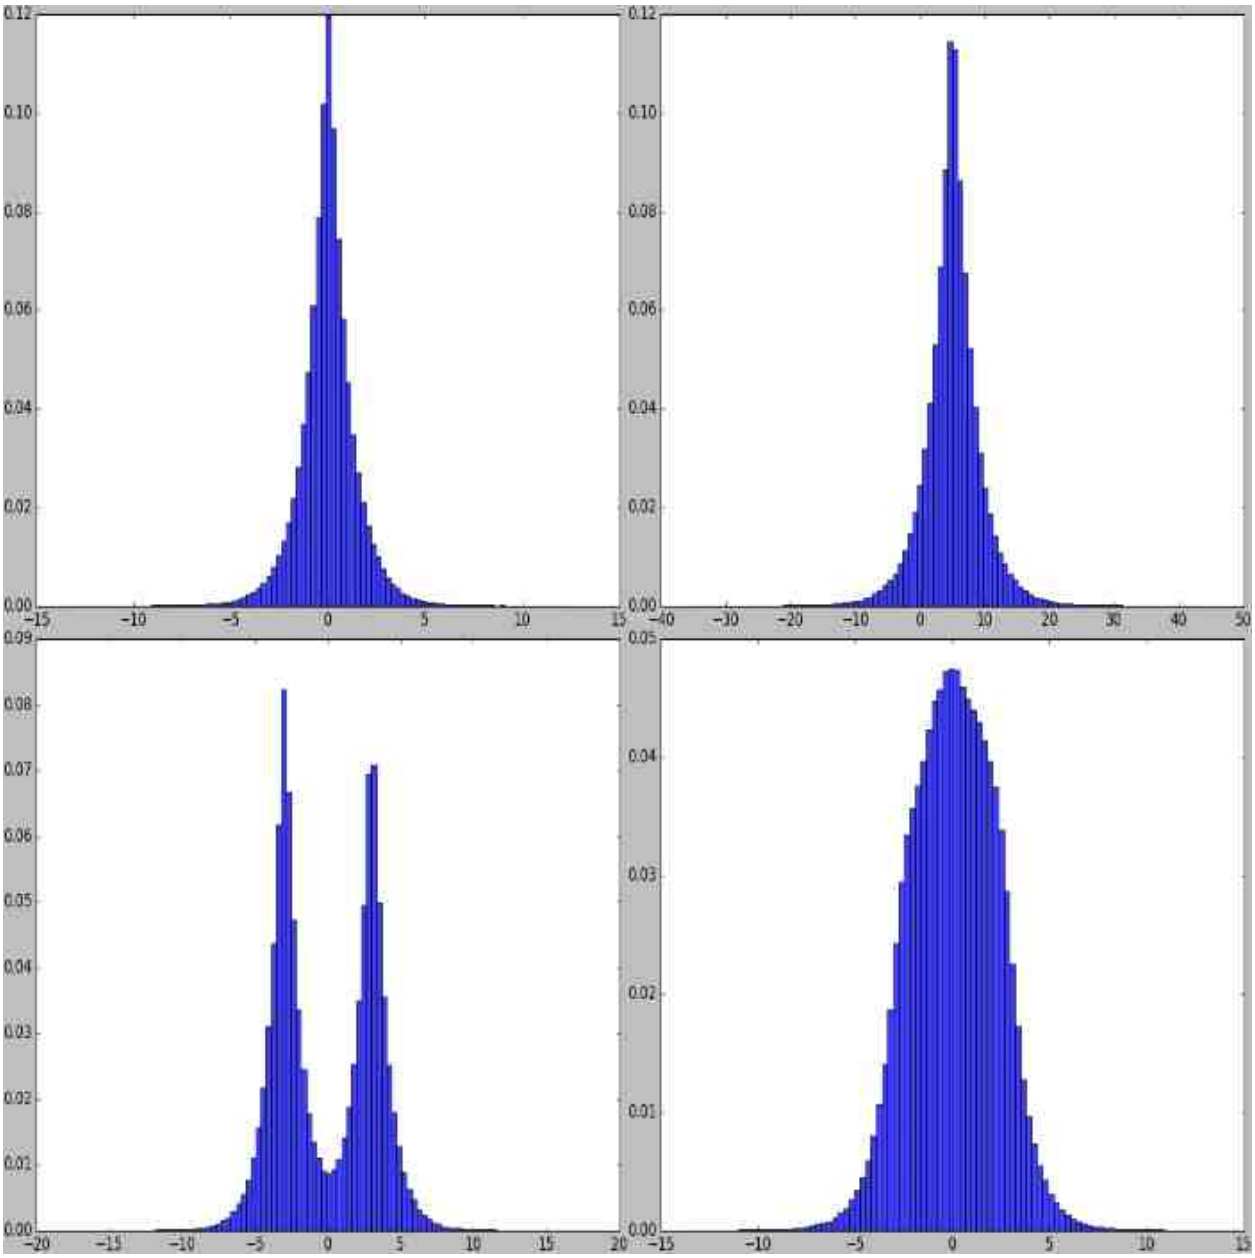

- ***ChiSquare(df)*: The chi-square ("X^2") distribution with *df* degrees of freedom.** Roughly similar to a continuous version of the poisson distribution. Has its peak at *df* and no negative values, only positive ones. Example:

```
from imgaug import parameters as iap
params = [
    iap.ChiSquare(1),
    iap.ChiSquare(3),
    iap.ChiSquare(iap.Choice([1, 5])),
```

(continues on next page)

(continued from previous page)

```
    iap.RandomSign(iap.ChiSquare(3))
]
iap.show_distributions_grid(params)
```

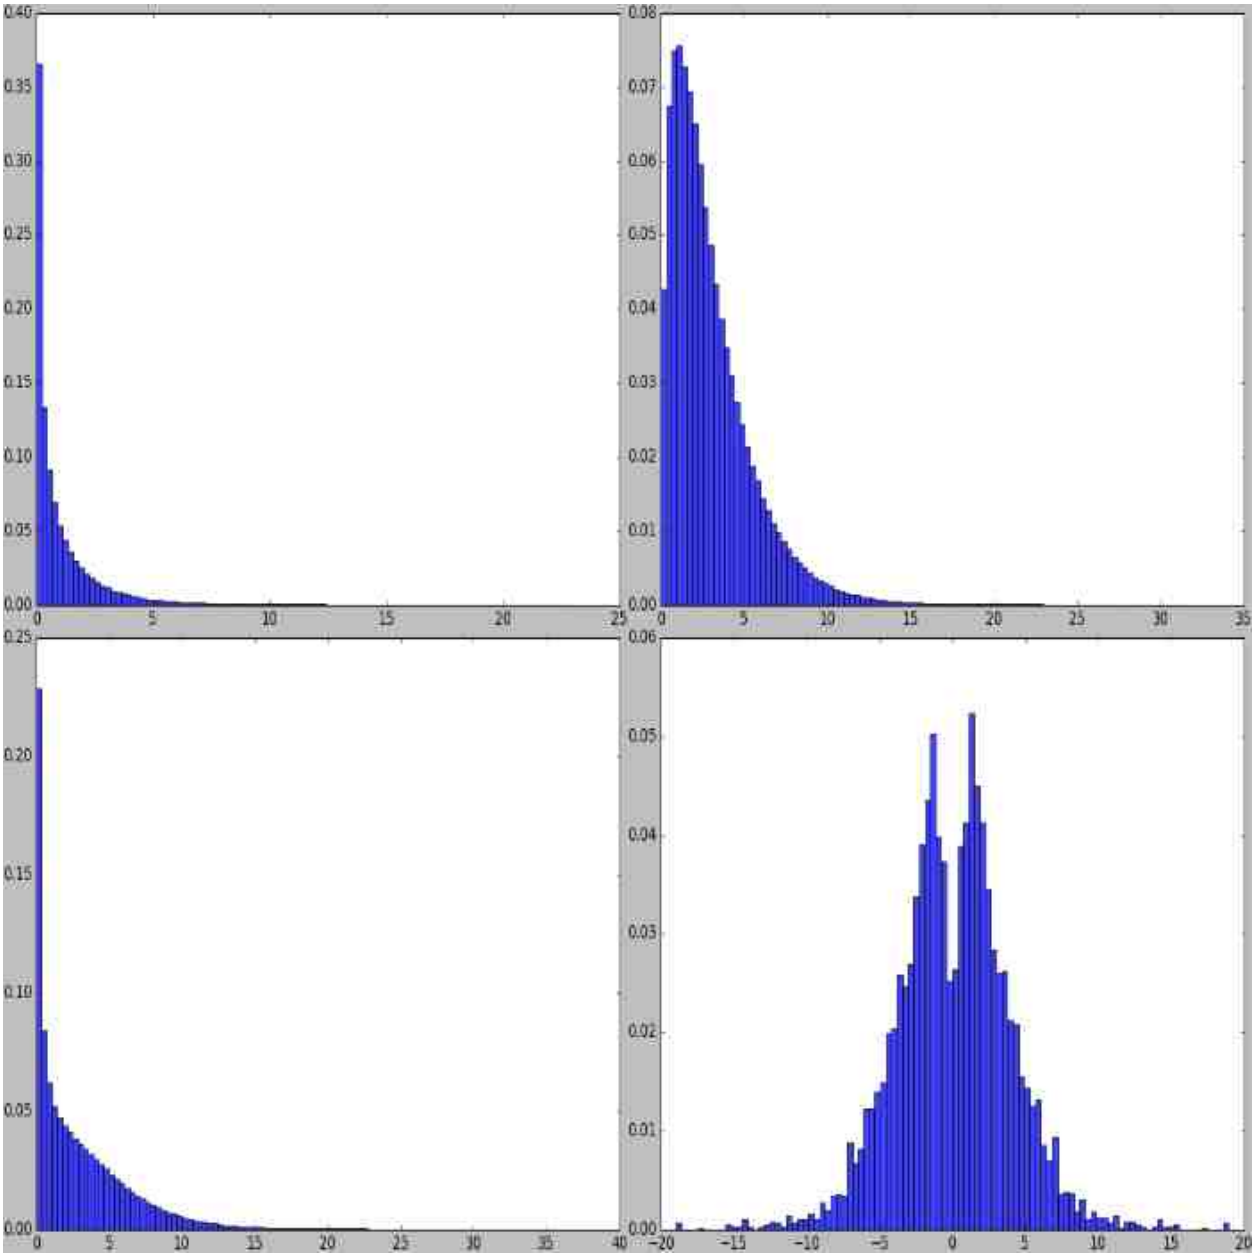

- *Weibull(a)*: Weibull distribution with shape *a*. Example:

```
from imgaug import parameters as iap
params = [
    iap.Weibull(0.5),
    iap.Weibull(1),
    iap.Weibull(1.5),
    iap.Weibull((0.5, 1.5))
]
```

(continues on next page)

(continued from previous page)

```
iap.show_distributions_grid(params)
```

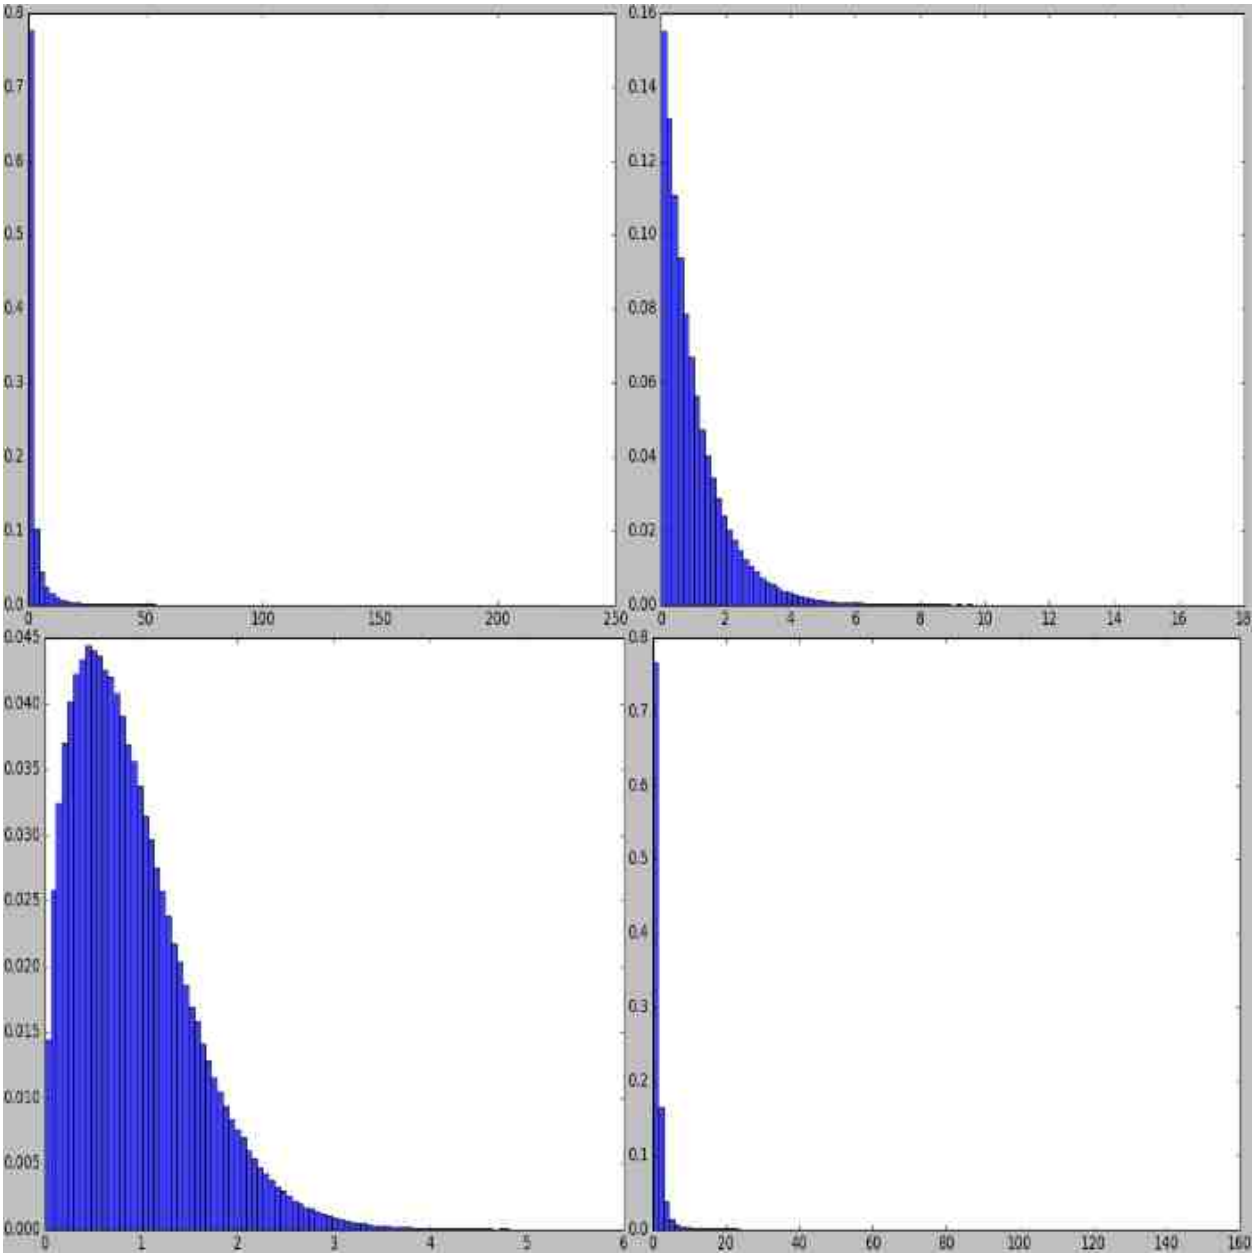

- *Uniform(a, b)*: Uniform distribution in the range *[a, b)*. Example:

```
from imgaug import parameters as iap
params = [
    iap.Uniform(0, 1),
    iap.Uniform(iap.Normal(-3, 1), iap.Normal(3, 1)),
    iap.Uniform([-1, 0], 1),
    iap.Uniform((-1, 0), 1)
]
iap.show_distributions_grid(params)
```

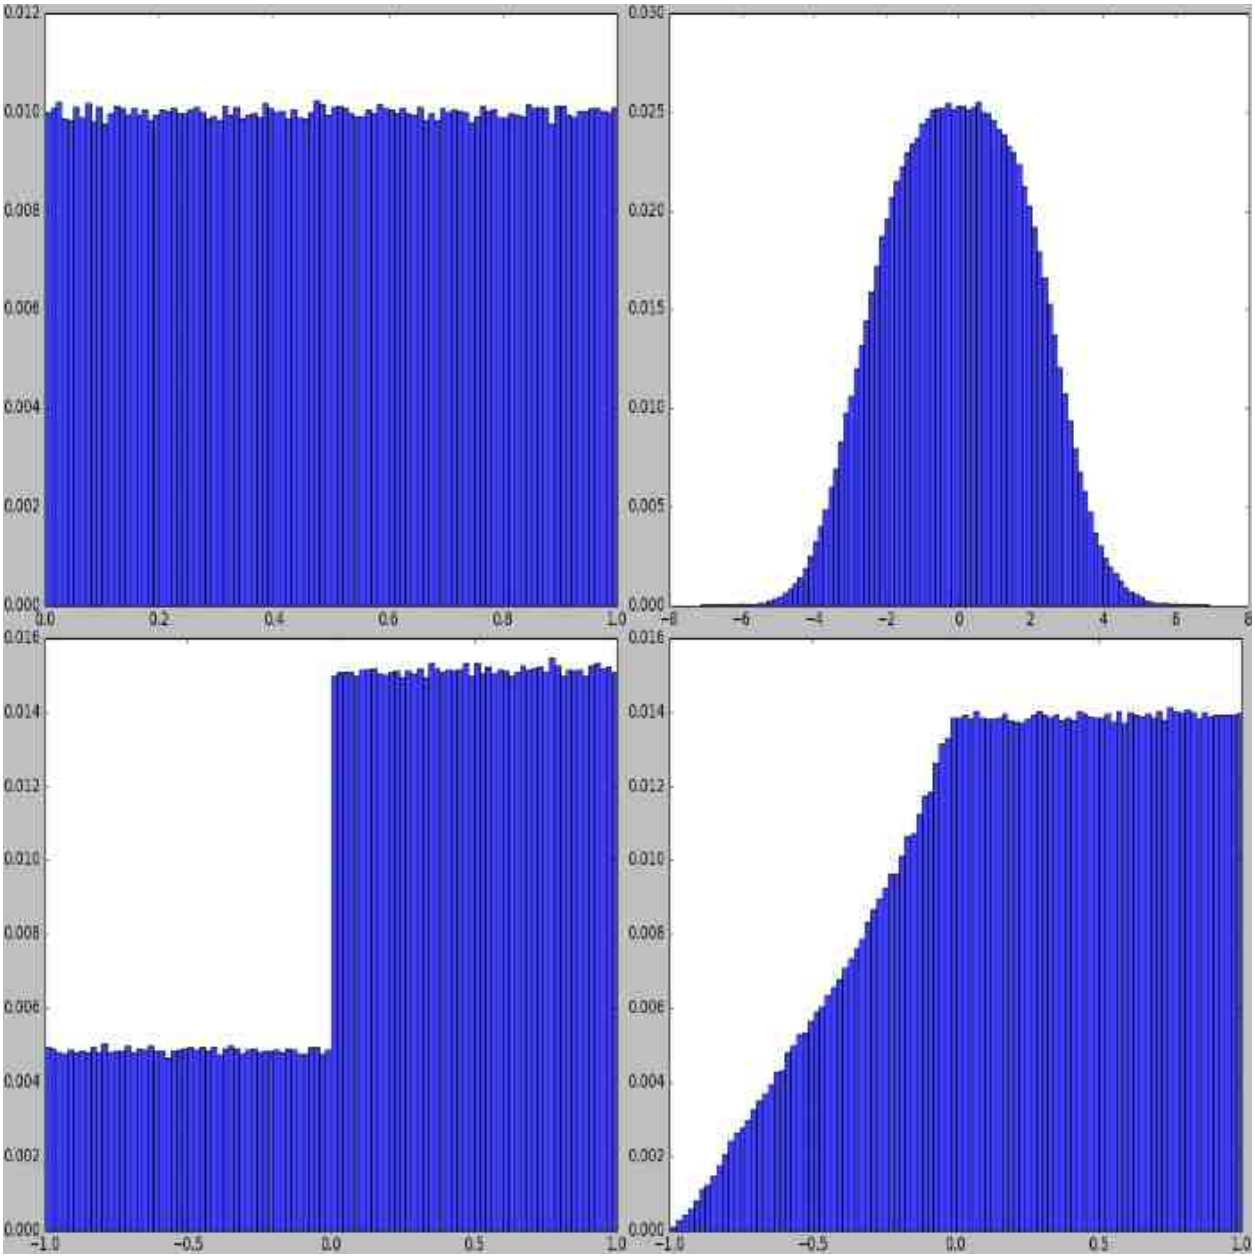

- *Beta(alpha, beta)*: Beta distribution with parameters *alpha* and *beta*. Example:

```
from imgaug import parameters as iap
params = [
    iap.Beta(0.5, 0.5),
    iap.Beta(2.0, 2.0),
    iap.Beta(1.0, 0.5),
    iap.Beta(0.5, 1.0)
]
iap.show_distributions_grid(params)
```

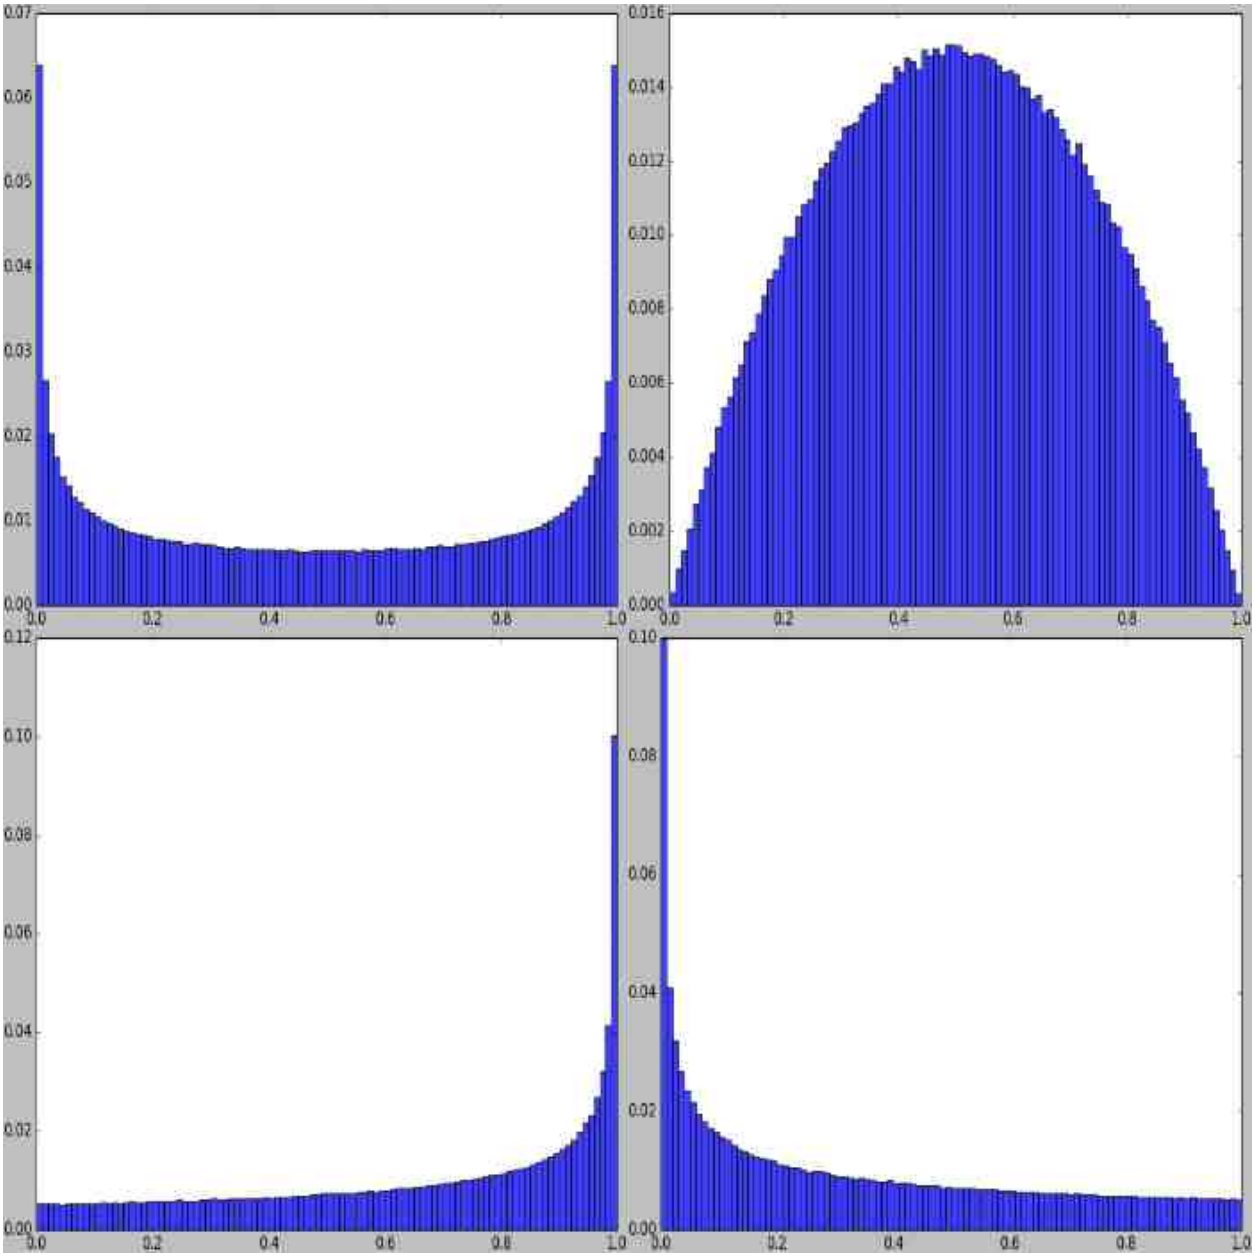

7.3 Discrete Probability Distributions

The following discrete probability distributions are available:

- *Binomial(p)*: The common binomial distribution with probability *p*. Useful to simulate coinflips. Example:

```
from imgaug import parameters as iap
params = [
    iap.Binomial(0.5),
    iap.Binomial(0.9)
]
iap.show_distributions_grid(params)
```

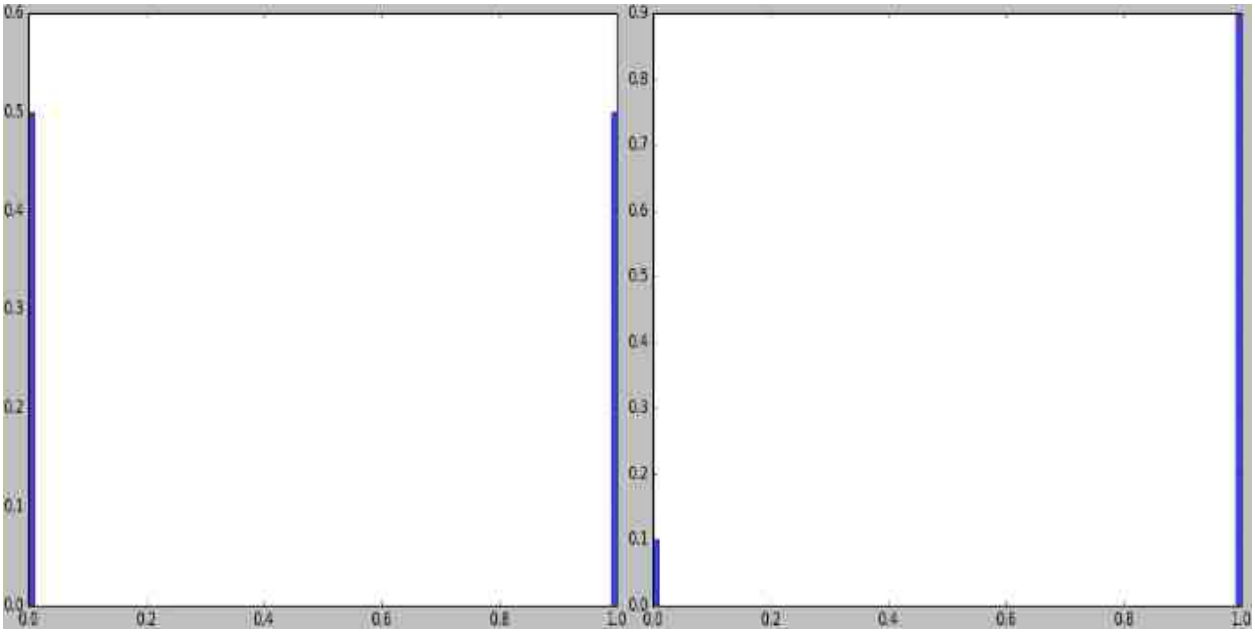

- *DiscreteUniform(a, b)*: The discrete uniform distribution in the range *[a..b]*. Example:

```
from imgaug import parameters as iap
params = [
    iap.DiscreteUniform(0, 10),
    iap.DiscreteUniform(-10, 10),
    iap.DiscreteUniform([-10, -9, -8, -7], 10),
    iap.DiscreteUniform((-10, -7), 10)
]
iap.show_distributions_grid(params)
```

- *Poisson(lam)*: Poisson distribution with shape *lam*. Generates no negative values. Example:

```
from imgaug import parameters as iap
params = [
    iap.Poisson(1),
    iap.Poisson(2.5),
    iap.Poisson((1, 2.5)),
    iap.RandomSign(iap.Poisson(2.5))
]
iap.show_distributions_grid(params)
```

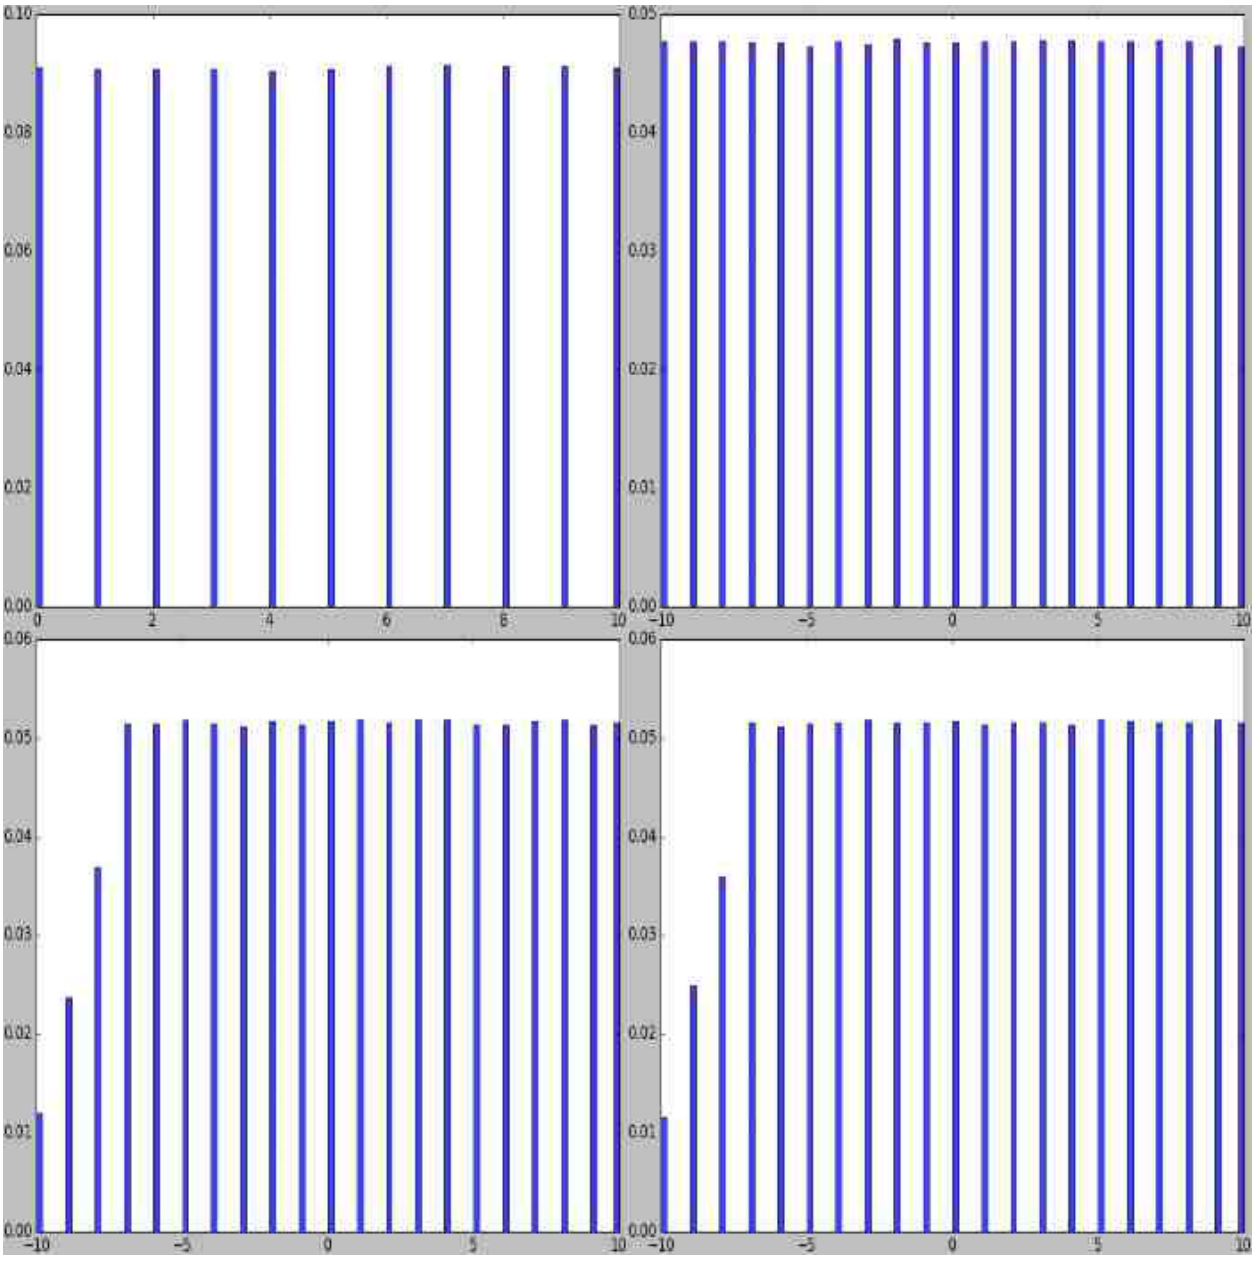

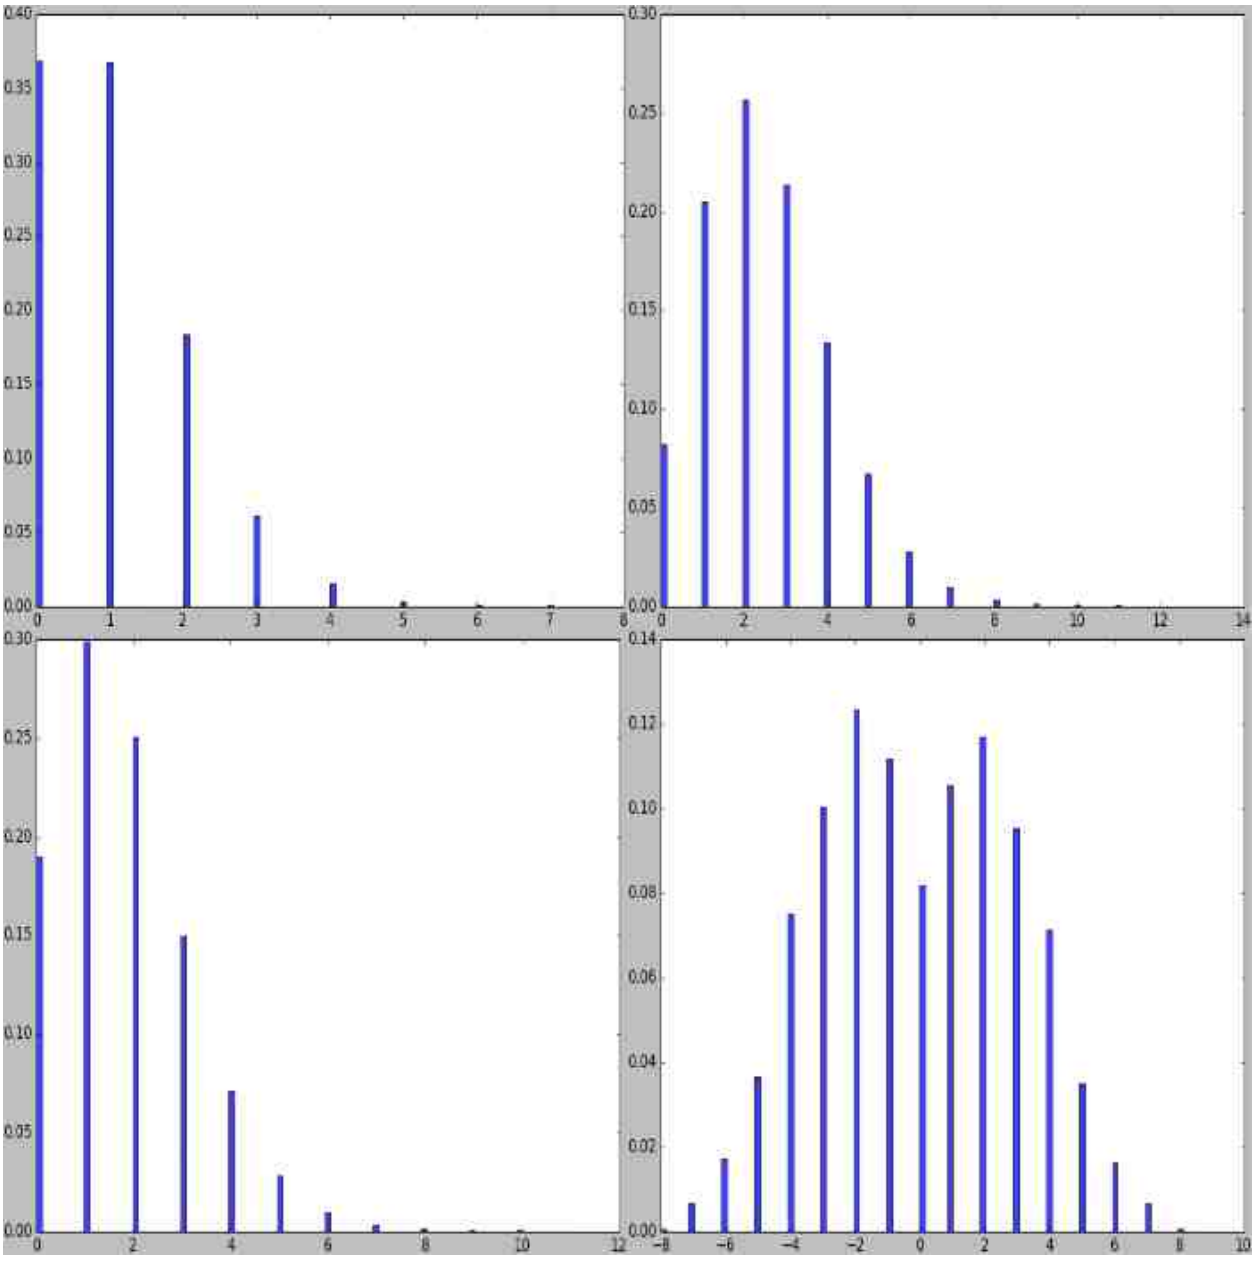

7.4 Arithmetic

The library supports arithmetic operations on stochastic parameters. This allows to modify values sampled from distributions or combine several distributions with each other.

- *Add(param, val, elementwise)*: Add *val* to the values sampled from *param*. The shortcut is +, e.g. *Uniform(...) + 1*. *val* can be a stochastic parameter itself. Usually, only one value is sampled from *val* per sampling run and added to all samples generated by *param*. Alternatively, *elementwise* can be set to *True* in order to generate as many samples from *val* as from *param* and add them elementwise. Note that *Add* merely adds to the results of *param* and does *not* combine probability density functions (see e.g. example image 3 and 4). Example:

```
from imgaug import parameters as iap
params = [
    iap.Uniform(0, 1) + 1, # identical to: Add(Uniform(0, 1), 1)
    iap.Add(iap.Uniform(0, 1), iap.Choice([0, 1], p=[0.7, 0.3])),
    iap.Normal(0, 1) + iap.Uniform(-5.5, -5) + iap.Uniform(5, 5.5),
    iap.Normal(0, 1) + iap.Uniform(-7, 5) + iap.Poisson(3),
    iap.Add(iap.Normal(-3, 1), iap.Normal(3, 1)),
    iap.Add(iap.Normal(-3, 1), iap.Normal(3, 1), elementwise=True)
]
iap.show_distributions_grid(
    params,
    rows=2,
    sample_sizes=[ # (iterations, samples per iteration)
        (1000, 1000), (1000, 1000), (1000, 1000),
        (1000, 1000), (1, 100000), (1, 100000)
    ]
)
```

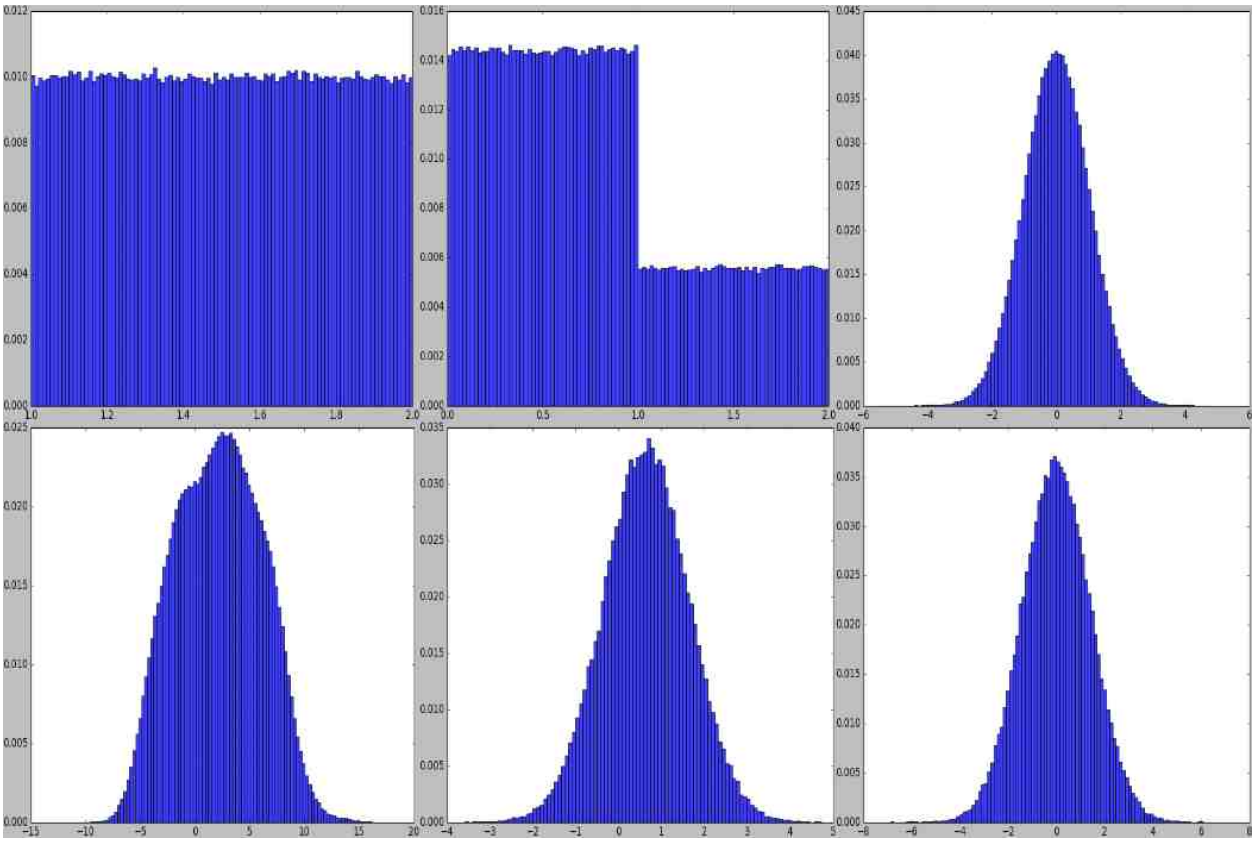

- *Subtract(param, val, elementwise)*: Same as *Add*, but subtracts *val* from the results of *param*. The shortcut is -, e.g. *Uniform(...) - 1*.
- *Multiply(param, val, elementwise)*: Same as *Add*, but multiplies *val* with the results of *param*. The shortcut is *, e.g. *Uniform(...) * 2*. Example:

```
from imgaug import parameters as iap
params = [
    iap.Uniform(0, 1) * 2, # identical to: Multiply(Uniform(0, 1), 2)
    iap.Multiply(iap.Uniform(0, 1), iap.Choice([0, 1], p=[0.7, 0.3])),
    (iap.Normal(0, 1) * iap.Uniform(-5.5, -5)) * iap.Uniform(5, 5.5),
    (iap.Normal(0, 1) * iap.Uniform(-7, 5)) * iap.Poisson(3),
    iap.Multiply(iap.Normal(-3, 1), iap.Normal(3, 1)),
    iap.Multiply(iap.Normal(-3, 1), iap.Normal(3, 1), elementwise=True)
]
iap.show_distributions_grid(
    params,
    rows=2,
    sample_sizes=[ # (iterations, samples per iteration)
        (1000, 1000), (1000, 1000), (1000, 1000),
        (1000, 1000), (1, 100000), (1, 100000)
    ]
)
```

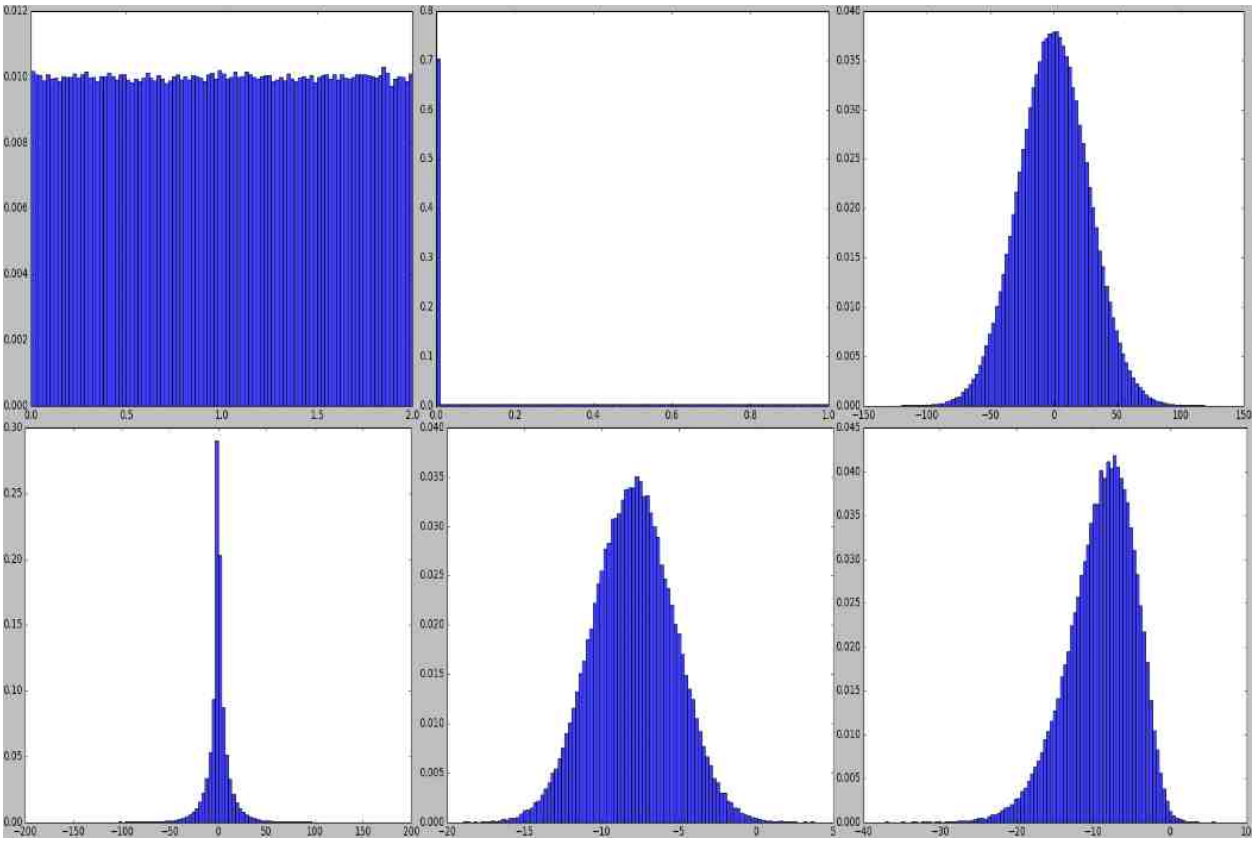

- *Divide(param, val, elementwise)*: Same as *Multiply*, but divides by *val*. The shortcut is */*, e.g. *Uniform(...) / 2*. Division by zero is automatically prevented (zeros are replaced by ones). Example:

```
from imgaug import parameters as iap
```

(continues on next page)

(continued from previous page)

```
params = [
    iap.Uniform(0, 1) / 2, # identical to: Divide(Uniform(0, 1), 2)
    iap.Divide(iap.Uniform(0, 1), iap.Choice([0, 2], p=[0.7, 0.3])),
    (iap.Normal(0, 1) / iap.Uniform(-5.5, -5)) / iap.Uniform(5, 5.5),
    (iap.Normal(0, 1) * iap.Uniform(-7, 5)) / iap.Poisson(3),
    iap.Divide(iap.Normal(-3, 1), iap.Normal(3, 1)),
    iap.Divide(iap.Normal(-3, 1), iap.Normal(3, 1), elementwise=True)
]
iap.show_distributions_grid(
    params,
    rows=2,
    sample_sizes=[ # (iterations, samples per iteration)
        (1000, 1000), (1000, 1000), (1000, 1000),
        (1000, 1000), (1, 100000), (1, 100000)
    ]
)
```

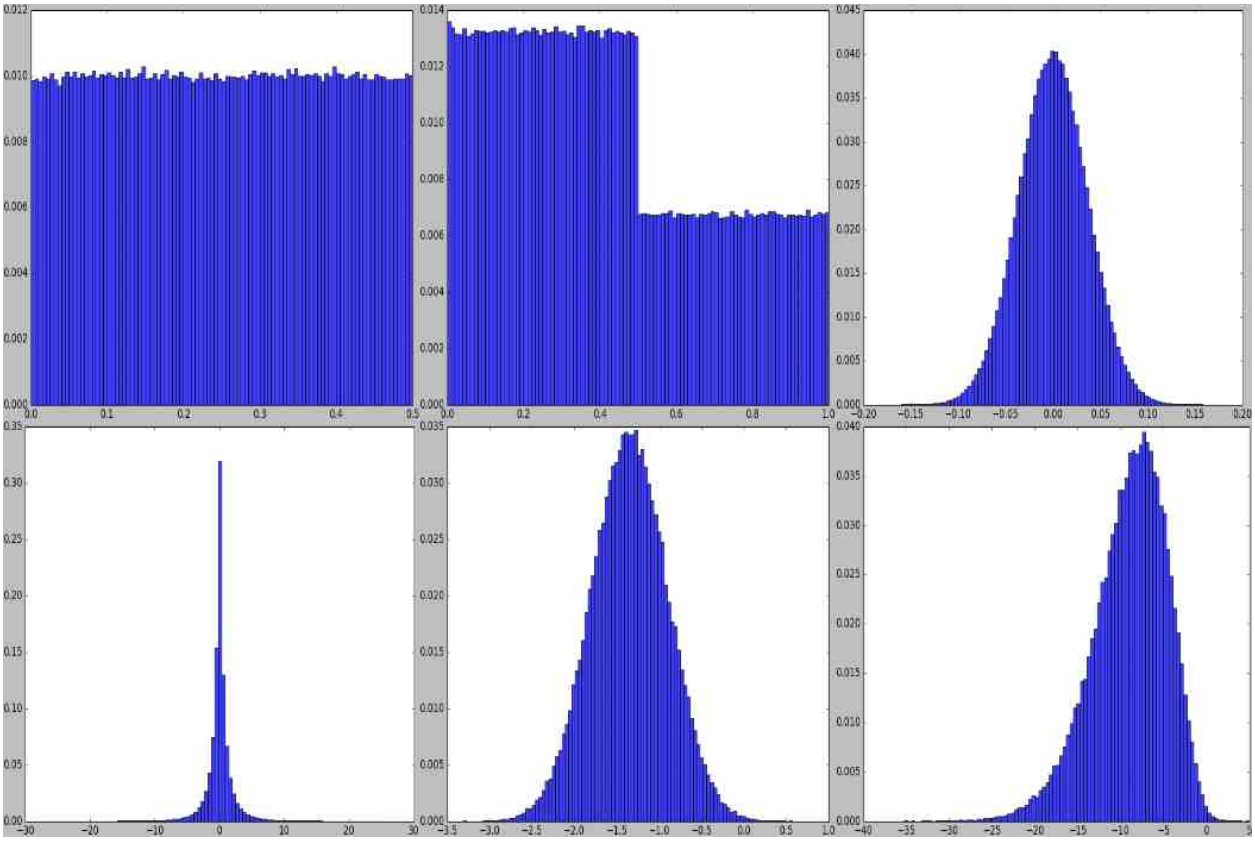

- *Power(param, val, elementwise)*: Same as *Add*, but raises sampled values to the exponent *val*. The shortcut is **. Example:

```
from imgaug import parameters as iap
params = [
    iap.Uniform(0, 1) ** 2, # identical to: Power(Uniform(0, 1), 2)
    iap.Clip(iap.Uniform(-1, 1) ** iap.Normal(0, 1), -4, 4)
]
iap.show_distributions_grid(params, rows=1)
```

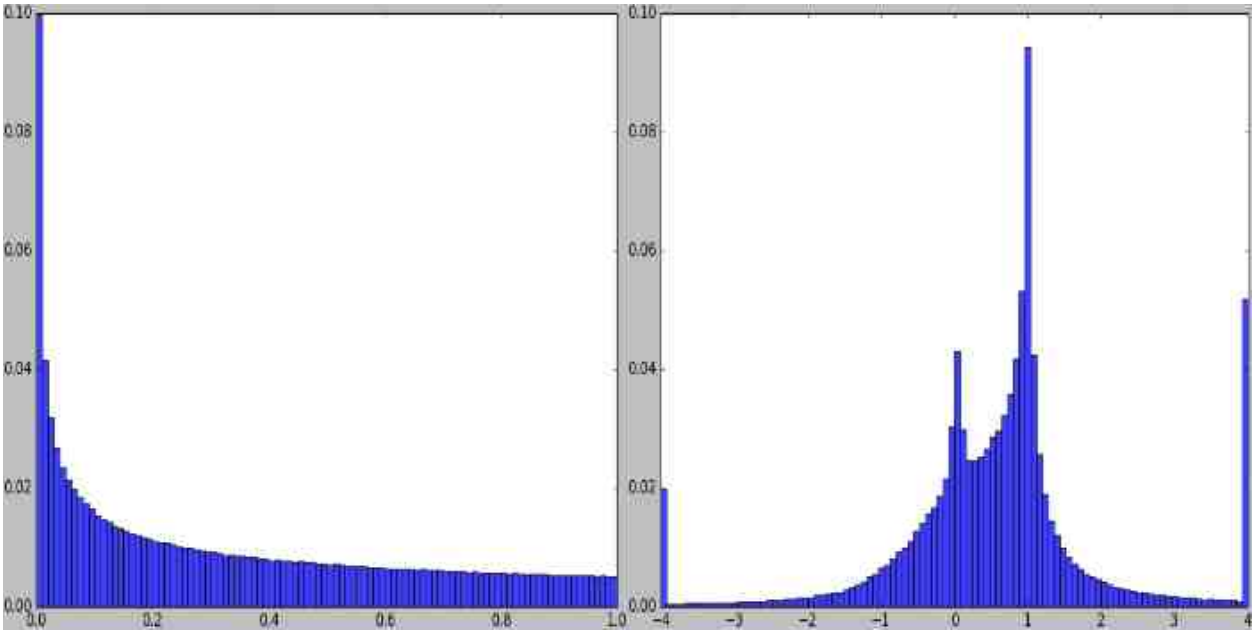

7.5 Special Parameters

- *Deterministic(v)*: A constant. Upon sampling, this always returns *v*.
- *Choice(values, replace=True, p=None)*: Upon sampling, this parameter picks randomly elements from a list *values*. If *replace* is set to *True* (default), the picking happens with replacement. By default, all elements have the same probability of being picked. This can be modified using *p*. Note that *values* may also contain strings and other stochastic parameters. In the latter case, each picked parameter will be replaced by a sample from that parameter. This allows merging of probability mass functions, but is a rather slow process. All elements in *values* should have the same datatype (except for stochastic parameters). Example:

```
from imgaug import parameters as iap
params = [
    iap.Choice([0, 1, 2]),
    iap.Choice([0, 1, 2], p=[0.15, 0.5, 0.35]),
    iap.Choice([iap.Normal(-3, 1), iap.Normal(3, 1)]),
    iap.Choice([iap.Normal(-3, 1), iap.Poisson(3)])
]
iap.show_distributions_grid(params)
```

- *Clip(param, minval=None, maxval=None)*: Clips the values sampled from *param* to the range *[minval, maxval]*. *minval* and *maxval* may be *None*. In that case, only minimum or maximum clipping is applied (depending on what is *None*). Example:

```
from imgaug import parameters as iap
params = [
    iap.Clip(iap.Normal(0, 1), -2, 2),
    iap.Clip(iap.Normal(0, 1), -2, None)
]
iap.show_distributions_grid(params, rows=1)
```

- *Discretize(param)*: Converts a continuous parameter *param* into a discrete one (using rounding). Discrete parameters are not changed. Example:

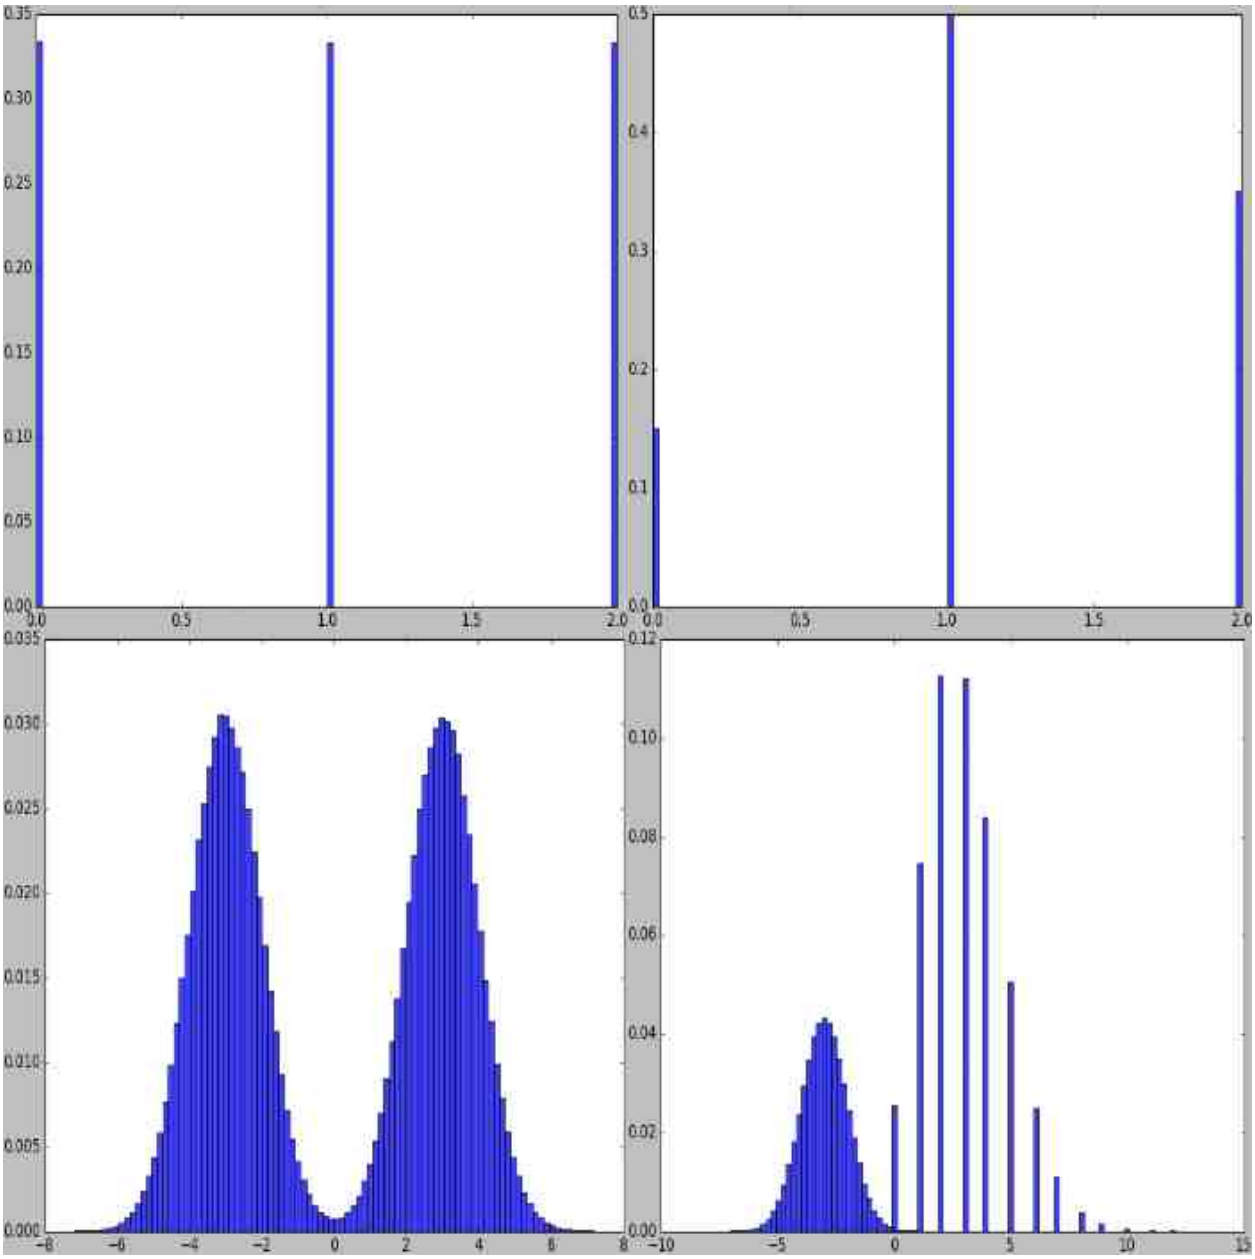

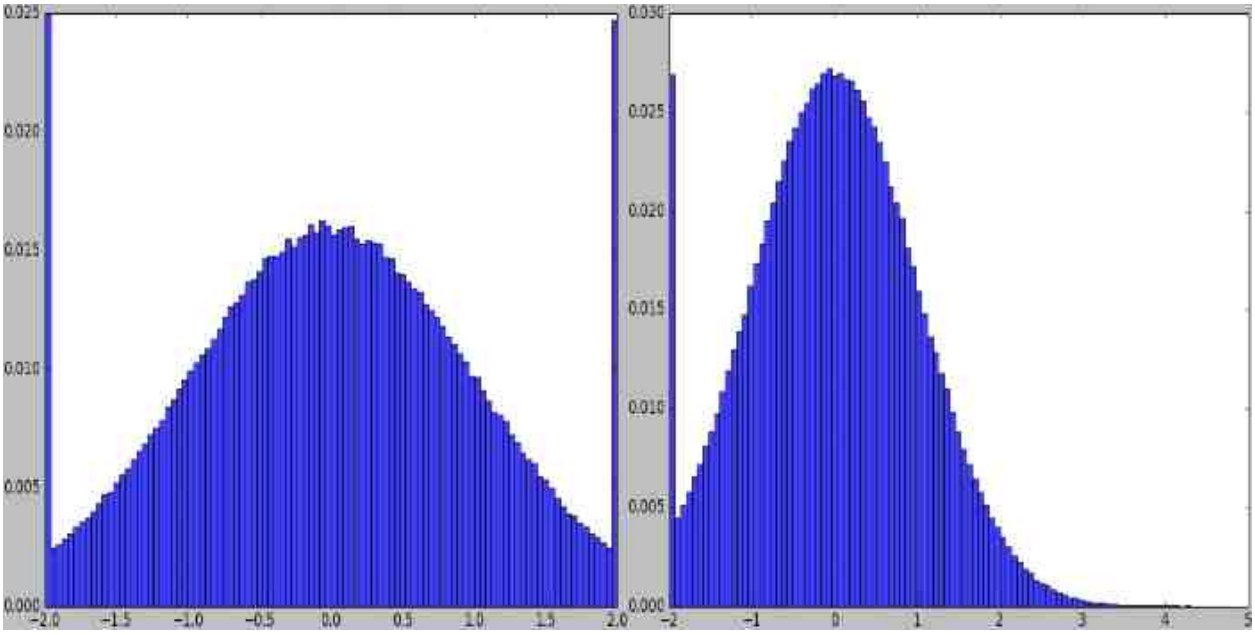

```
from imgaug import parameters as iap
params = [
    iap.Discretize(iap.Normal(0, 1)),
    iap.Discretize(iap.ChiSquare(3))
]
iap.show_distributions_grid(params, rows=1)
```

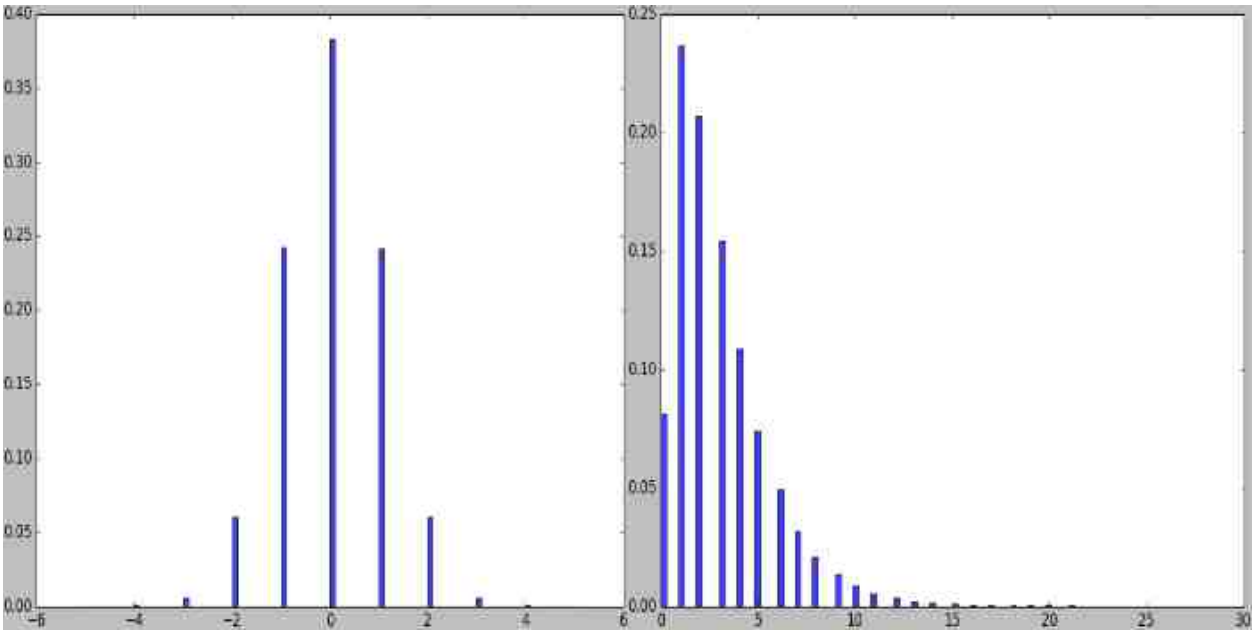

- *Absolute(param)*: Applies an absolute function to each value sampled from *param*, turning them to positive ones. Example:

```
from imgaug import parameters as iap
```

(continues on next page)

(continued from previous page)

```
params = [
    iap.Absolute(iap.Normal(0, 1)),
    iap.Absolute(iap.Laplace(0, 1))
]
iap.show_distributions_grid(params, rows=1)
```

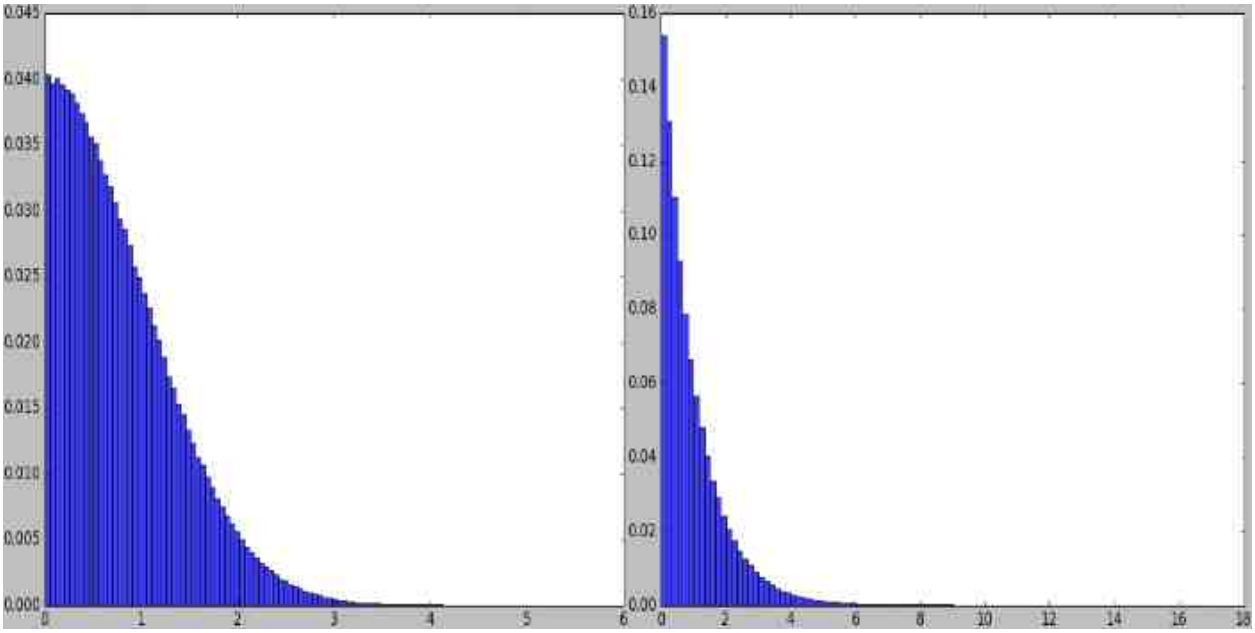

- *RandomSign(param, p_positive=0.5)*: Randomly flips the signs of values sampled from *param*. Optionally, the probability of flipping a value's sign towards positive can be set. Example:

```
from imgaug import parameters as iap
params = [
    iap.ChiSquare(3),
    iap.RandomSign(iap.ChiSquare(3)),
    iap.RandomSign(iap.ChiSquare(3), p_positive=0.75),
    iap.RandomSign(iap.ChiSquare(3), p_positive=0.9)
]
iap.show_distributions_grid(params)
```

- *ForceSign(param, positive, mode="invert", reroll_count_max=2)*: Converts all values sampled from *param* to positive or negative ones. Signs of positive/negative values may simply be flipped (*mode="invert"*) or re-sampled from param (mode="reroll"). When rerolling, the number of iterations is limited to *reroll_count_max* (afterwards mode="invert" is used). Example:

```
from imgaug import parameters as iap
params = [
    iap.ForceSign(iap.Normal(0, 1), positive=True),
    iap.ChiSquare(3) - 3.0,
    iap.ForceSign(iap.ChiSquare(3) - 3.0, positive=True, mode="invert"),
    iap.ForceSign(iap.ChiSquare(3) - 3.0, positive=True, mode="reroll")
]
iap.show_distributions_grid(params)
```

- *Positive(other_param, mode="invert", reroll_count_max=2)*: Shortcut for *ForceSign* with *positive=True*. E.g. *Positive(Normal(0, 1))* restricts a normal distribution to only positive values.

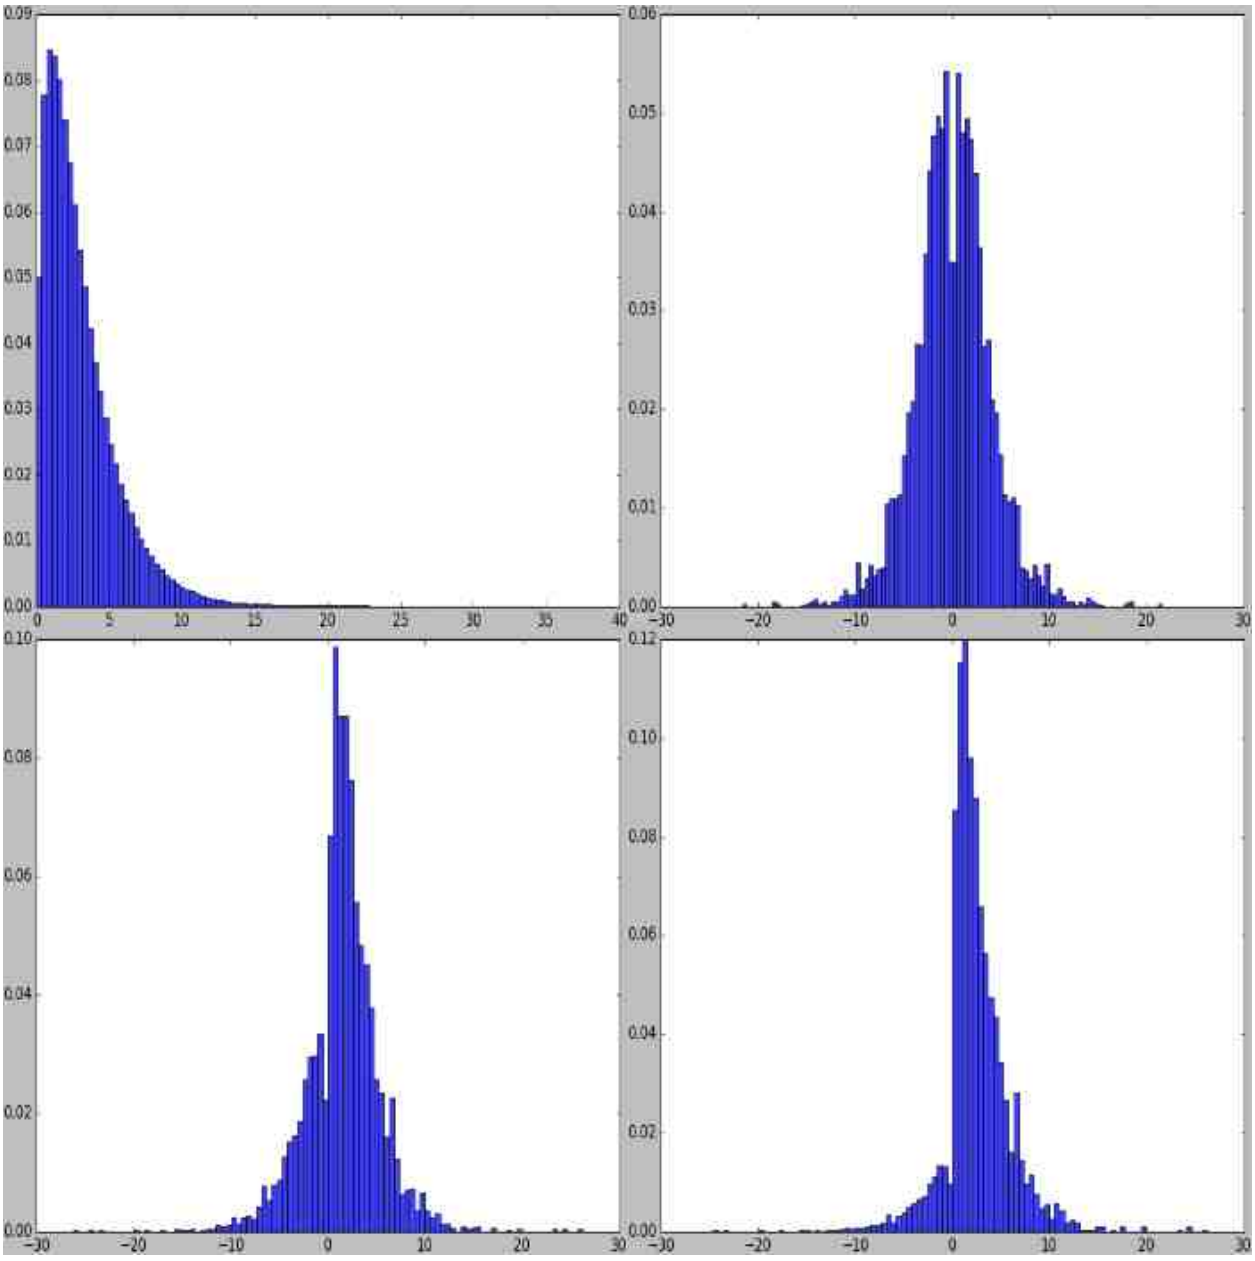

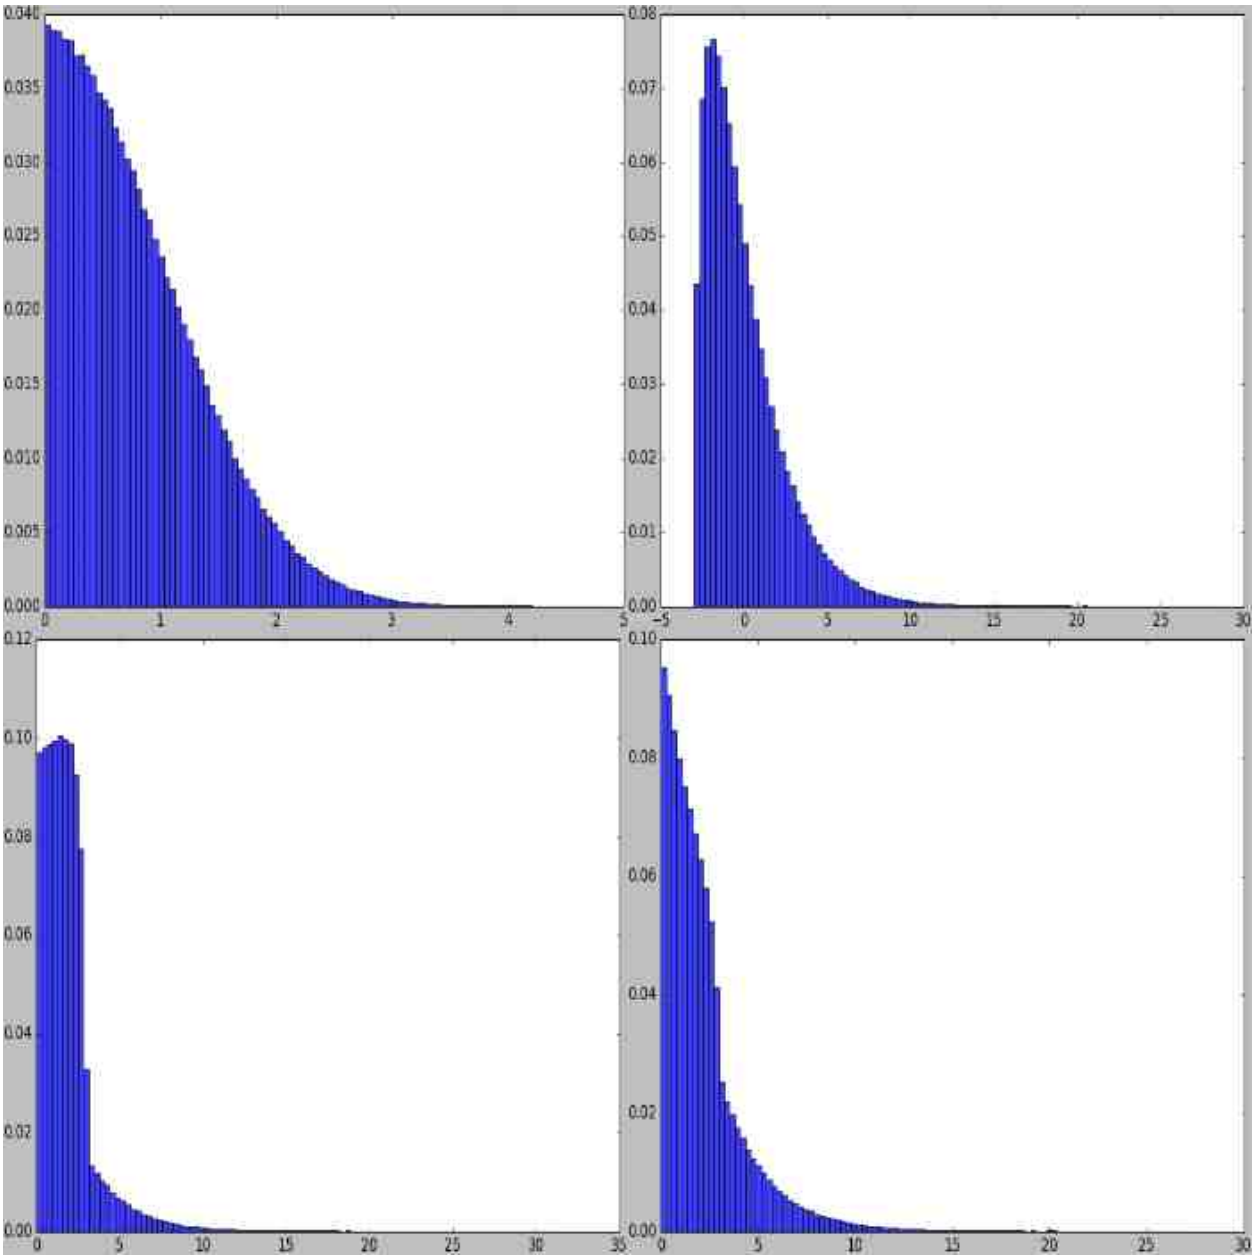

- *Negative(other_param, mode="invert", reroll_count_max=2)*: Shortcut for *ForceSign* with *positive=False*. E.g. *Negative(Normal(0, 1))* restricts a normal distribution to only negative values.
- *FromLowerResolution(other_param, size_percent=None, size_px=None, method="nearest", min_size=1)*: Intended for 2d-sampling processes, e.g. for masks. Samples these in a lower resolution space. E.g. instead of sampling a mask at 100x100, this allows to sample it at 10x10 and then upsample to 100x100. One advantage is, that this can be faster. Another possible use is, that the upsampling may result in large, correlated blobs (linear interpolation) or rectangles (nearest neighbour interpolation).

7.6 Noise Parameters

TODO

CHAPTER 8

Blending/Overlaying images

8.1 Introduction

Most augmenters in the library affect images in uniform ways per image. Sometimes one might not want that and instead desires more localized effects (e.g. change the color of some image regions, while keeping the others unchanged) or wants to keep a fraction of the old image (e.g. blur the image and mix in a bit of the unblurred image). Blending augmenters are intended for these use cases. They either mix two images using a constant alpha factor or using a pixel-wise mask. Below image shows examples.

```
# First row
iaa.BlendAlpha(
    (0.0, 1.0),
    foreground=iaa.MedianBlur(11),
    per_channel=True
)

# Second row
iaa.BlendAlphaSimplexNoise(
    foreground=iaa.EdgeDetect(1.0),
    per_channel=False
)

# Third row
iaa.BlendAlphaSimplexNoise(
    foreground=iaa.EdgeDetect(1.0),
    background=iaa.LinearContrast((0.5, 2.0)),
    per_channel=0.5
)

# Forth row
iaa.BlendAlphaFrequencyNoise(
    foreground=iaa.Affine(
        rotate=(-10, 10),
        translate_px={"x": (-4, 4), "y": (-4, 4)}
```

(continues on next page)

(continued from previous page)

```
    ),
    background=iaa.AddToHueAndSaturation((-40, 40)),
    per_channel=0.5
)

# Fifth row
iaa.BlendAlphaSimplexNoise(
    foreground=iaa.BlendAlphaSimplexNoise(
        foreground=iaa.EdgeDetect(1.0),
        background=iaa.LinearContrast((0.5, 2.0)),
        per_channel=True
    ),
    background=iaa.BlendAlphaFrequencyNoise(
        exponent=(-2.5, -1.0),
        foreground=iaa.Affine(
            rotate=(-10, 10),
            translate_px={"x": (-4, 4), "y": (-4, 4)}
        ),
        background=iaa.AddToHueAndSaturation((-40, 40)),
        per_channel=True
    ),
    per_channel=True,
    aggregation_method="max",
    sigmoid=False
)
```

8.2 Imagewise Constant Alphas Values

The augmenter *imgaug.augmenters.blend.BlendAlpha* allows to mix the results of two augmentation branches using an alpha factor that is constant throughout the whole image, i.e. it follows roughly `I_blend = alpha * I_fg + (1 - alpha) * I_bg` per image, where `I_fg` is the image from the foreground branch and `I_bg` is the image from the background branch. Often, the first branch will be an augmented version of the image and the second branch will be the identity function, leading to a blend of augmented and unaugmented image. The background branch can also contain non-identity augmenters, leading to a blend of two distinct augmentation effects.

imgaug.augmenters.blend.BlendAlpha is already built into some augmenters as a parameter, e.g. into *imgaug.augmenters.convolutional.EdgeDetect*.

The below example code generates images that are a blend between *imgaug.augmenters.convolutional.Sharpen* and *imgaug.augmenters.arithmetic.CoarseDropout*. Notice how the sharpening does not affect the black rectangles from dropout, as the two augmenters are both applied to the original images and merely blended.

```
import imgaug as ia
from imgaug import augmenters as iaa

ia.seed(1)

# Example batch of images.
# The array has shape (8, 128, 128, 3) and dtype uint8.
images = np.array(
    [ia.quokka(size=(128, 128)) for _ in range(8)],
    dtype=np.uint8
)
```

(continues on next page)

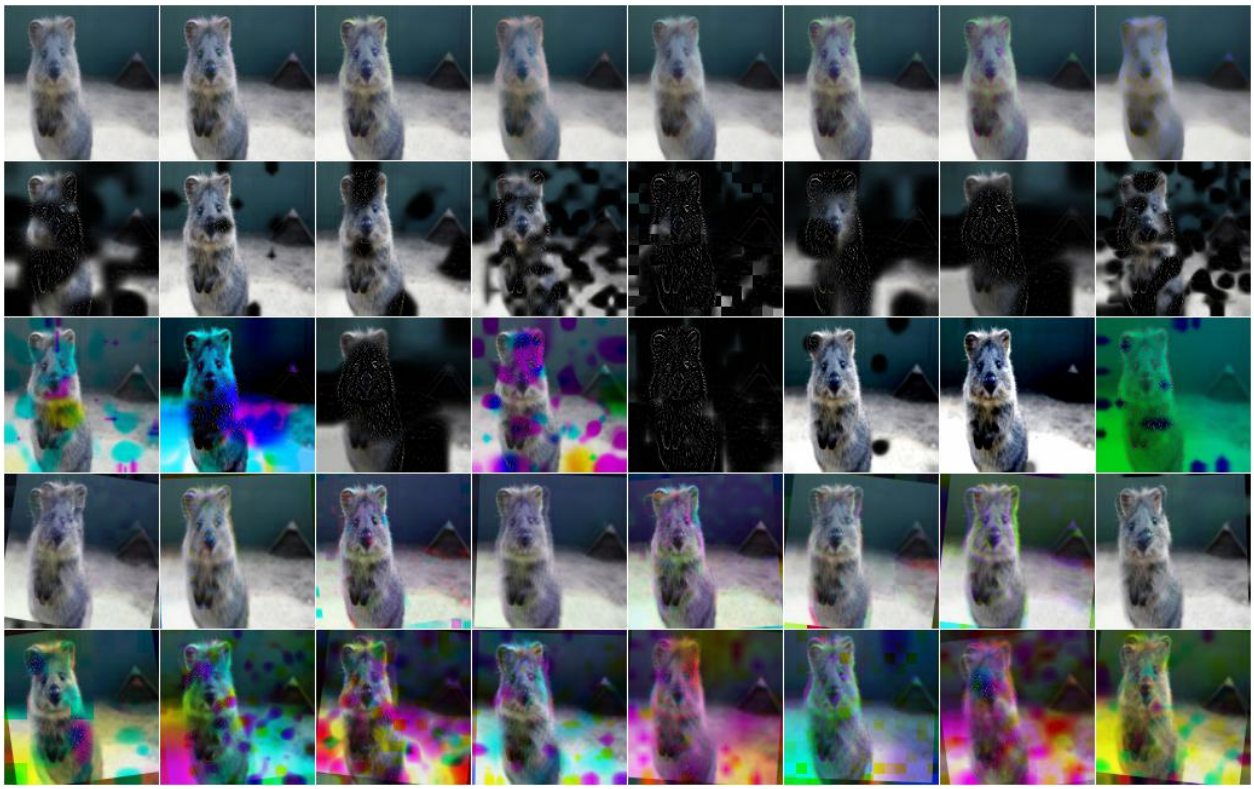

Fig. 1: Various effects of combining alpha-augmenters with other augmenters. First row shows *imgaug.augmenters.blend.BlendAlpha* with *imgaug.augmenters.blur.MedianBlur*, second *imgaug.augmenters.blend.BlendAlphaSimplexNoise* with *imgaug.augmenters.convolutional.EdgeDetect*, third *imgaug.augmenters.blend.BlendAlphaSimplexNoise* with *imgaug.augmenters.convolutional.EdgeDetect* and imgaug.augmenters.contrast.ContrastNormalization, third shows *imgaug.augmenters.blend.BlendAlphaFrequencyNoise* with *imgaug.augmenters.geometric.Affine* and *imgaug.augmenters.color.AddToHueAndSaturation* and forth row shows a mixture *imgaug.augmenters.blend.BlendAlphaSimplexNoise* and *imgaug.augmenters.blend.BlendAlphaFrequencyNoise*.

(continued from previous page)

```
seq = iaa.BlendAlpha(
    factor=(0.2, 0.8),
    foreground=iaa.Sharpen(1.0, lightness=2),
    background=iaa.CoarseDropout(p=0.1, size_px=8)
)

images_aug = seq(images=images)
```

Fig. 2: Mixing *imgaug.augmenters.convolutional.Sharpen* and *imgaug.augmenters.arithmetic.CoarseDropout* via *imgaug.augmenters.blend.BlendAlpha*. The resulting effect is very different from executing them in sequence.

Similar to other augmenters, *imgaug.augmenters.blend.BlendAlpha* supports a `per_channel` mode, in which it samples blending strengths for each channel independently. As a result, some channels may show more from the foreground (or background) branch's outputs than other channels. This can lead to visible color effects. The following example is the same as the one above, only `per_channel` was activated.

```
iaa.BlendAlpha(..., per_channel=True)
```

imgaug.augmenters.blend.BlendAlpha can also be used with augmenters that change the position of pixels, leading to "ghost" images. (This should not be done when also augmenting keypoints, as their position becomes unclear.)

```
seq = iaa.BlendAlpha(
    factor=(0.2, 0.8),
    foreground=iaa.Affine(rotate=(-20, 20)),
    per_channel=True
)
```

Fig. 3: Mixing *imgaug.augmenters.convolutional.Sharpen* and *imgaug.augmenters.arithmetic.CoarseDropout* via *imgaug.augmenters.blend.BlendAlpha* and `per_channel` set to `True`.

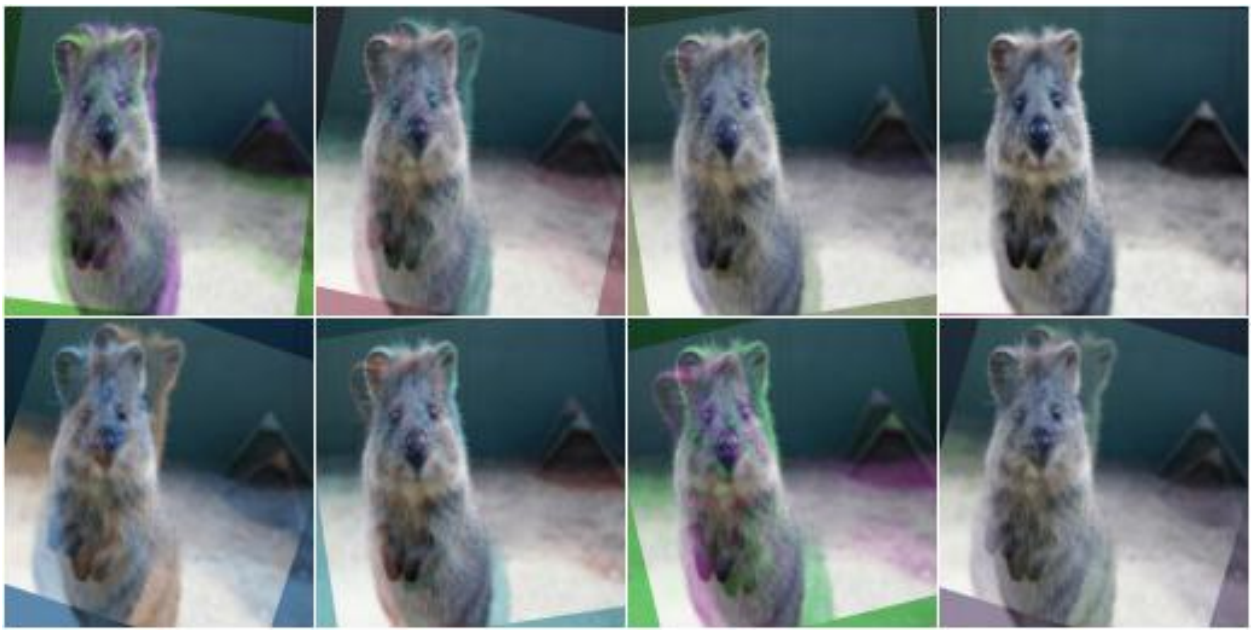

Fig. 4: Mixing original images with their rotated version. Some channels are more visibly rotated than others.

8.3 BlendAlphaSimplexNoise

imgaug.augmenters.blend.BlendAlpha uses a constant blending factor per image (or per channel). This limits its possibilities. Often, a more localized factor is desired to create unusual patterns. *imgaug.augmenters.blend.BlendAlphaSimplexNoise* is an augmenter that does that. It generates continuous masks following simplex noise and uses them to perform local blending. The following example shows a combination of *imgaug.augmenters.blend.BlendAlphaSimplexNoise* and *imgaug.augmenters.arithmetic.Multiply* (with `per_channel=True`) that creates blobs of various colors in the image.

```
import imgaug as ia
from imgaug import augmenters as iaa

ia.seed(1)

# Example batch of images.
# The array has shape (8, 128, 128, 3) and dtype uint8.
images = np.array(
    [ia.quokka(size=(128, 128)) for _ in range(8)],
    dtype=np.uint8
)

seq = iaa.SimplexNoiseAlpha(
    foreground=iaa.Multiply(iap.Choice([0.5, 1.5]), per_channel=True)
)

images_aug = seq(images=images)
```

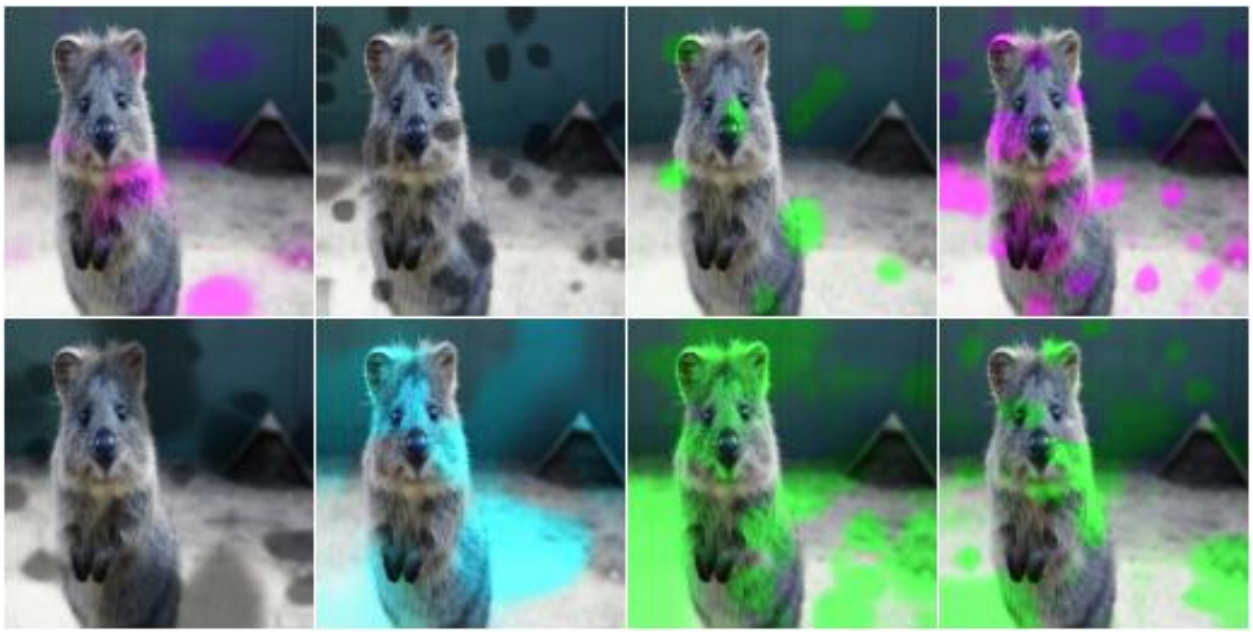

Fig. 5: Mixing original images with their versions modified by *imgaug.augmenters.arithmetic.Multiply* (with `per_channel` set to `True`). Simplex noise masks are used for the blending process, leading to blobby patterns.

imgaug.augmenters.blend.BlendAlphaSimplexNoise also supports `per_channel=True`, leading to unique noise masks sampled per channel. The following example shows the combination of *imgaug.augmenters.blend.BlendAlphaSimplexNoise* (with `per_channel=True`) and *imgaug.*

augmenters.convolutional.EdgeDetect. Even though *imgaug.augmenters.convolutional.EdgeDetect* usually generates black and white images (white=edges, black=everything else), here the combination leads to strong color effects as the channel-wise noise masks only blend EdgeDetect's result for some channels.

```
seq = iaa.BlendAlphaSimplexNoise(
    foreground=iaa.EdgeDetect(1.0),
    per_channel=True
)
```

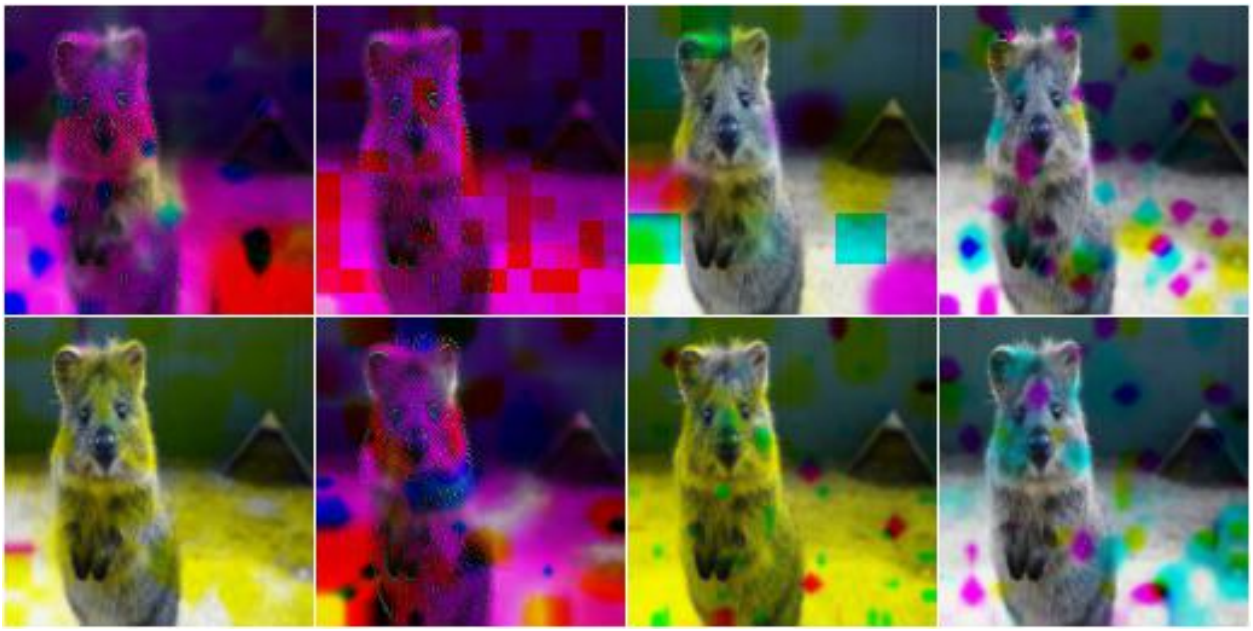

Fig. 6: Blending images via simplex noise can lead to unexpected but diverse patterns when `per_channel` is set to `True`. Here, a mixture of original images with `EdgeDetect(1.0)` is used.

imgaug.augmenters.blend.BlendAlphaSimplexNoise uses continuous noise masks (2d arrays with values in the range [0.0, 1.0]) to blend images. The below image shows examples of 64x64 noise masks generated by *imgaug.augmenters.blend.BlendAlphaSimplexNoise* with default settings. Values close to `1.0` (white) indicate that pixel colors will be taken from the first image source, while `0.0` (black) values indicate that pixel colors will be taken from the second image source. (Often only one image source will be given in the form of augmenters and the second will fall back to the original images fed into *imgaug.augmenters.blend.BlendAlphaSimplexNoise*.)

Fig. 7: Examples of noise masks generated by *imgaug.augmenters.blend.BlendAlphaSimplexNoise* using default settings.

imgaug.augmenters.blend.BlendAlphaSimplexNoise generates its noise masks in low resolution images and then upscales the masks to the size of the input images. During upscaling it usually uses nearest neighbour interpolation (`nearest`), linear interpolation (`linear`) or cubic interpolation (`cubic`). Nearest neighbour interpolation leads to noise maps with rectangular blobs. The below example shows noise maps generated when only using nearest neighbour interpolation.

```
seq = iaa.BlendAlphaSimplexNoise(
    ...,
    upscale_method="nearest"
)
```

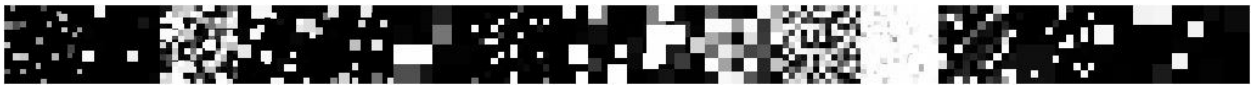

Fig. 8: Examples of noise masks generated by *imgaug.augmenters.blend.BlendAlphaSimplexNoise* when restricting the upscaling method to `nearest`.

Similarly, the following example shows noise maps generated when only using linear interpolation.

```
seq = iaa.BlendAlphaSimplexNoise(
    ...,
    upscale_method="linear"
)
```

Fig. 9: Examples of noise masks generated by *imgaug.augmenters.blend.BlendAlphaSimplexNoise* when restricting the upscaling method to `linear`.

8.4 FrequencyNoiseAlpha

imgaug.augmenters.blend.BlendAlphaFrequencyNoise is mostly identical to *imgaug.augmenters.blend.BlendAlphaSimplexNoise*. In contrast to *imgaug.augmenters.blend.BlendAlphaSimplexNoise* it uses a different sampling process to generate the blend masks. The process is based on starting with random frequencies, weighting them with a random exponent and then transforming from frequency domain to spatial domain. When using a low exponent value this leads to large, smooth blobs. Slightly higher exponents lead to cloudy patterns. High exponent values lead to recurring, small patterns. The below example shows the usage of *imgaug.augmenters.blend.BlendAlphaFrequencyNoise*.

```
import imgaug as ia
from imgaug import augmenters as iaa
from imgaug import parameters as iap

ia.seed(1)

# Example batch of images.
# The array has shape (8, 64, 64, 3) and dtype uint8.
images = np.array(
    [ia.quokka(size=(128, 128)) for _ in range(8)],
    dtype=np.uint8
)

seq = iaa.BlendAlphaFrequencyNoise(
    foreground=iaa.Multiply(iap.Choice([0.5, 1.5]), per_channel=True)
)

images_aug = seq.augment_images(images)
```

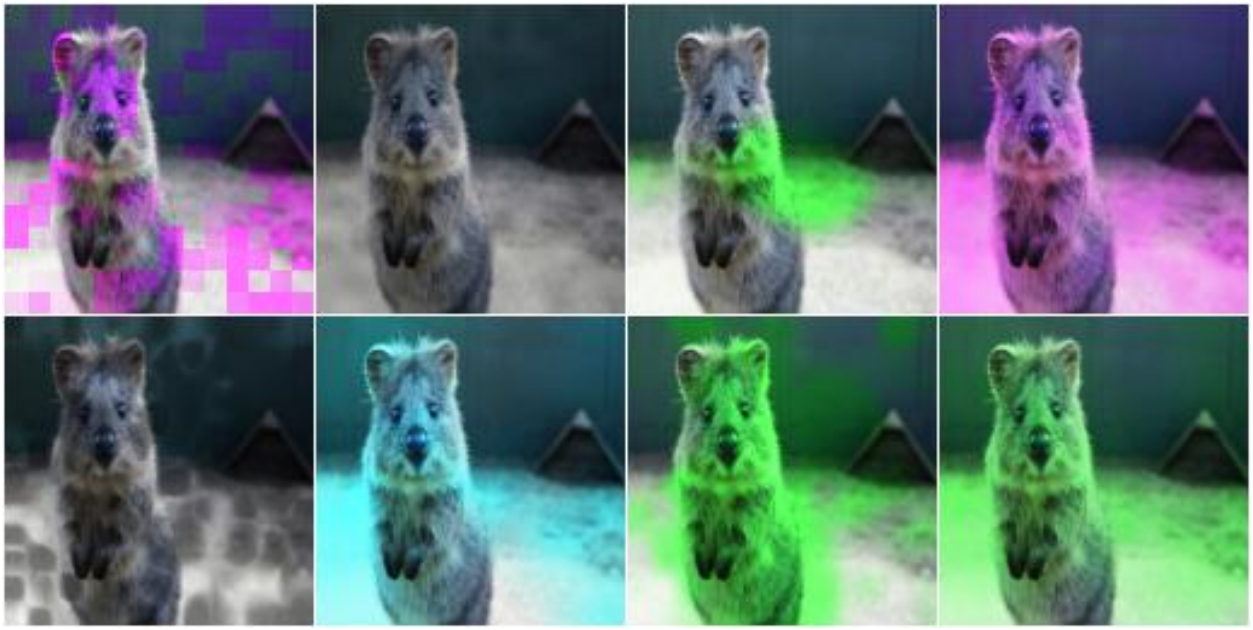

Fig. 10: Mixing original images with their versions modified by *imgaug.augmenters.arithmetic.Multiply* (with `per_channel` set to `True`). Frequency noise masks are used for the blending process, leading to blobby patterns.

Similarly to simplex noise, *imgaug.augmenters.blend.BlendAlphaFrequencyNoise* also supports `per_channel=True`, leading to different noise maps per image channel.

```
seq = iaa.BlendAlphaFrequencyNoise(
    foreground=iaa.EdgeDetect(1.0),
    per_channel=True
)
```

The below image shows random example noise masks generated by *imgaug.augmenters.blend.BlendAlphaFrequencyNoise* with default settings.

The following image shows the effects of varying `exponent` between `-4.0` and `4.0`. To show these effects more clearly, a few features of *imgaug.augmenters.blend.BlendAlphaFrequencyNoise* were deactivated (e.g. multiple iterations). In the code, `E` is the value of the exponent (e.g. `E=-2.0`).

```
seq = iaa.BlendAlphaFrequencyNoise(
    exponent=E,
    foreground=iaa.Multiply(iap.Choice([0.5, 1.5]), per_channel=True),
    size_px_max=32,
    upscale_method="linear",
    iterations=1,
    sigmoid=False
)
```

Similarly to *imgaug.augmenters.blend.BlendAlphaSimplexNoise*, *imgaug.augmenters.blend.BlendAlphaFrequencyNoise* also generates the noise masks as low resolution versions and then upscales them to the full image size. The following images show the usage of nearest neighbour interpolation (`upscale_method="nearest"`) and linear interpolation (`upscale_method="linear"`).

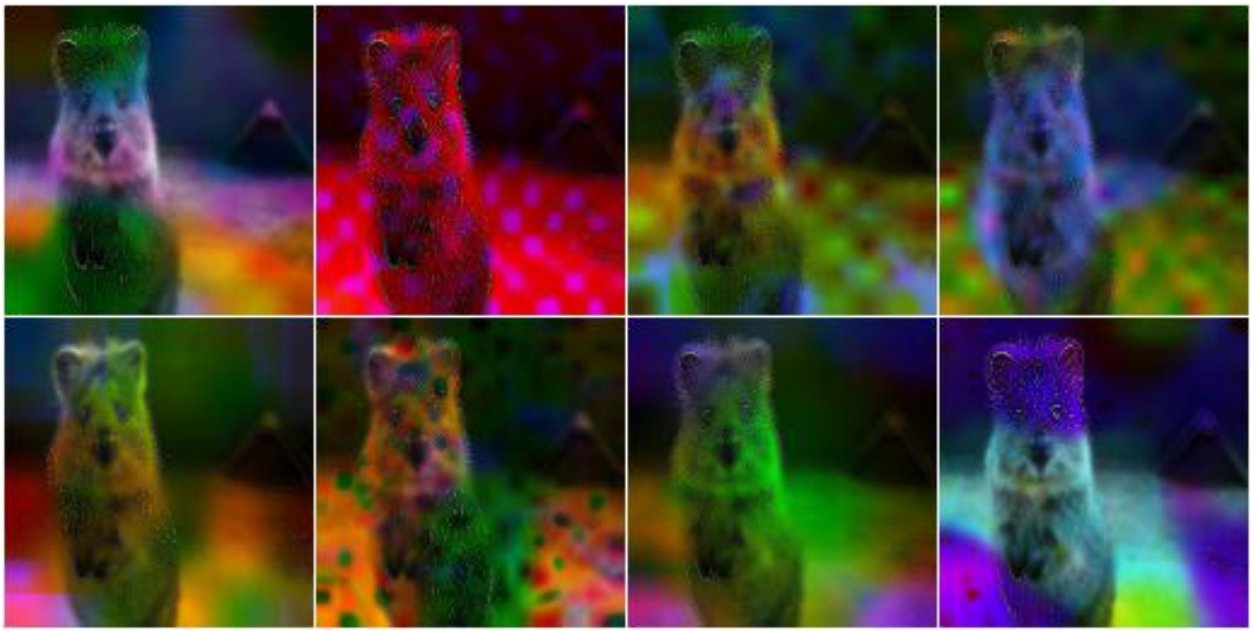

Fig. 11: Blending images via frequency noise can lead to unexpected but diverse patterns when `per_channel` is set to `True`. Here, a mixture of original images with `imgaug.augmenters.convolutional.EdgeDetect(1.0)` is used.

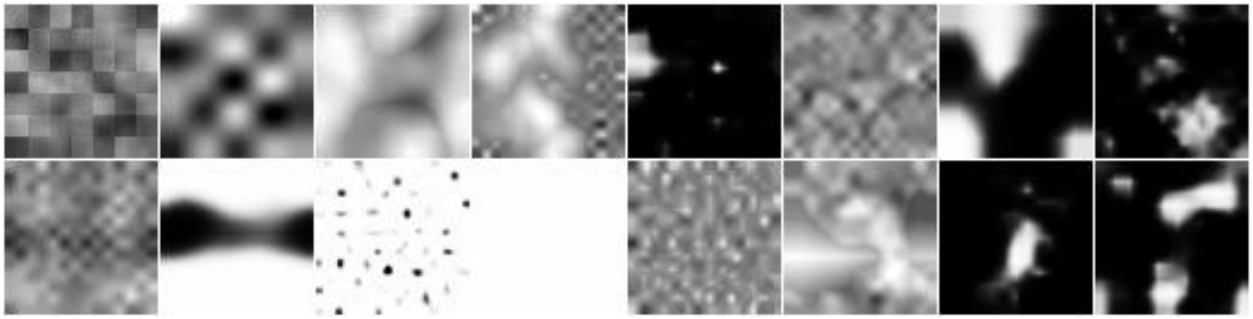

Fig. 12: Examples of noise masks generated by *imgaug.augmenters.blend.FrequencyNoiseAlpha* using default settings.

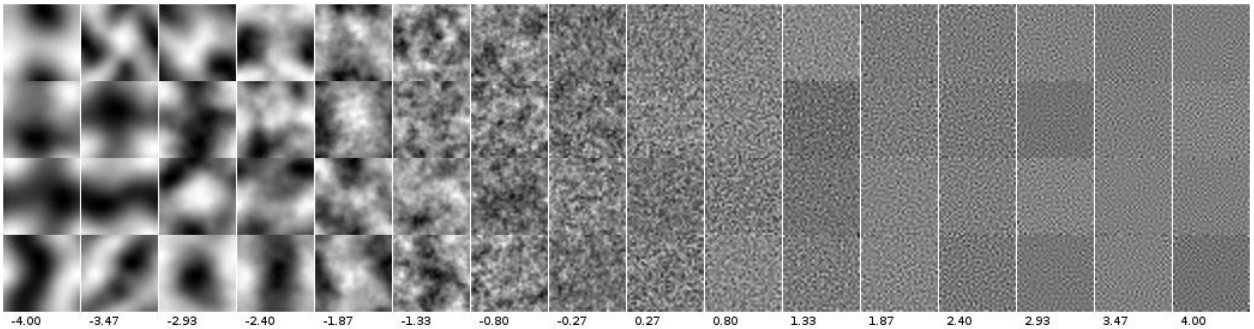

Fig. 13: Examples of noise masks generated by *imgaug.augmenters.blend.BlendAlphaFrequencyNoise* using default settings with varying exponents.

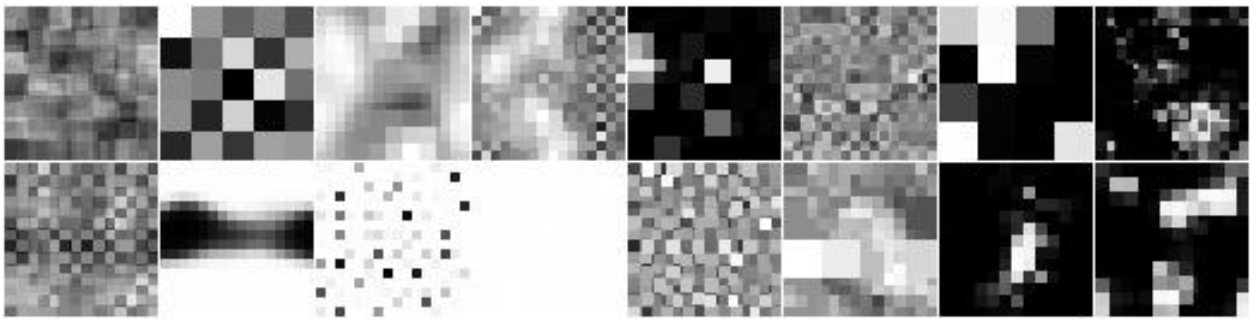

Fig. 14: Examples of noise masks generated by *imgaug.augmenters.blend.BlendAlphaFrequencyNoise* when restricting the upscaling method to `nearest`.

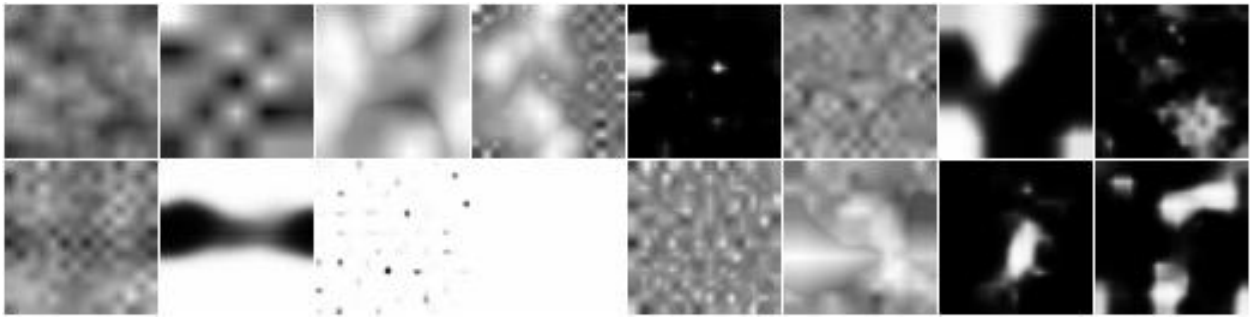

Fig. 15: Examples of noise masks generated by *imgaug.augmenters.blend.BlendAlphaFrequencyNoise* when restricting the upscaling method to `linear`.

8.5 IterativeNoiseAggregator

Both *imgaug.augmenters.blend.BlendAlphaSimplexNoise* and *imgaug.augmenters.blend.BlendAlphaFrequencyNoise* wrap around *imgaug.parameters.IterativeNoiseAggregator*, a component to generate noise masks in multiple iterations. It has parameters for the number of iterations (1 to N) and for the aggregation methods, which controls how the noise masks from the different iterations are to be combined. Valid aggregation methods are `"min"`, `"avg"` and `"max"`, where `min` takes the minimum over all iteration's masks, `max` the maxmimum and `avg` the average. As a result, masks generated with method `min` tend to be close to 0.0 (mostly black values), those generated with `max` close to `1.0` and `avg` converges towards `0.5`. (`0.0` means that the results of the second image dominate the final image, so in many cases the original images before the augmenter). The following image shows the effects of changing the number of iterations when combining *imgaug.parameters.FrequencyNoise* with *imgaug.parameters.IterativeNoiseAggregator*.

```
# This is how the iterations would be changed for BlendAlphaFrequencyNoise.
# (Same for BlendAlphaSimplexNoise.)
seq = iaa.BlendAlphaFrequencyNoise(
    ...,
    iterations=N
)
```

The following image shows the effects of changing the aggregation mode (with varying iterations).

```
# This is how the iterations and aggregation method would be changed for
# BlendAlphaFrequencyNoise. (Same for BlendAlphaSimplexNoise.)
seq = iaa.BlendAlphaFrequencyNoise(
```

(continues on next page)

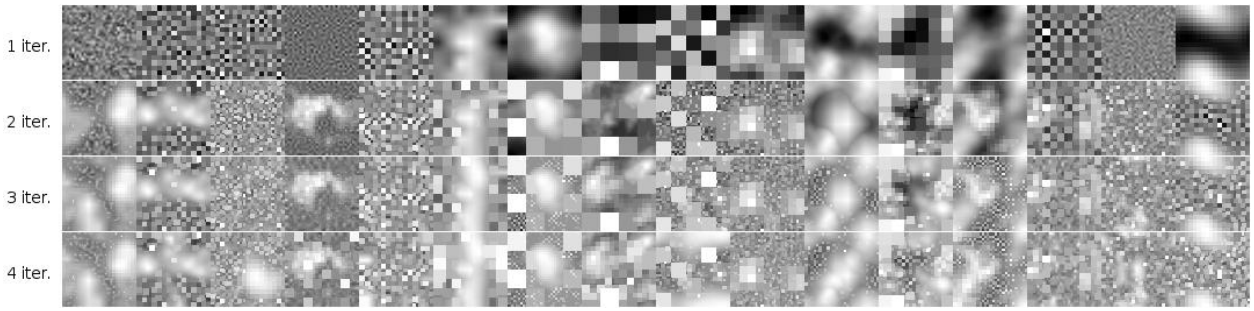

Fig. 16: Examples of varying the number of iterations in *imgaug.parameters.IterativeNoiseAggregator* (here in combination with *imgaug.parameters.FrequencyNoise*).

(continued from previous page)

```
    ...,
    iterations=N,
    aggregation_method=M
)
```

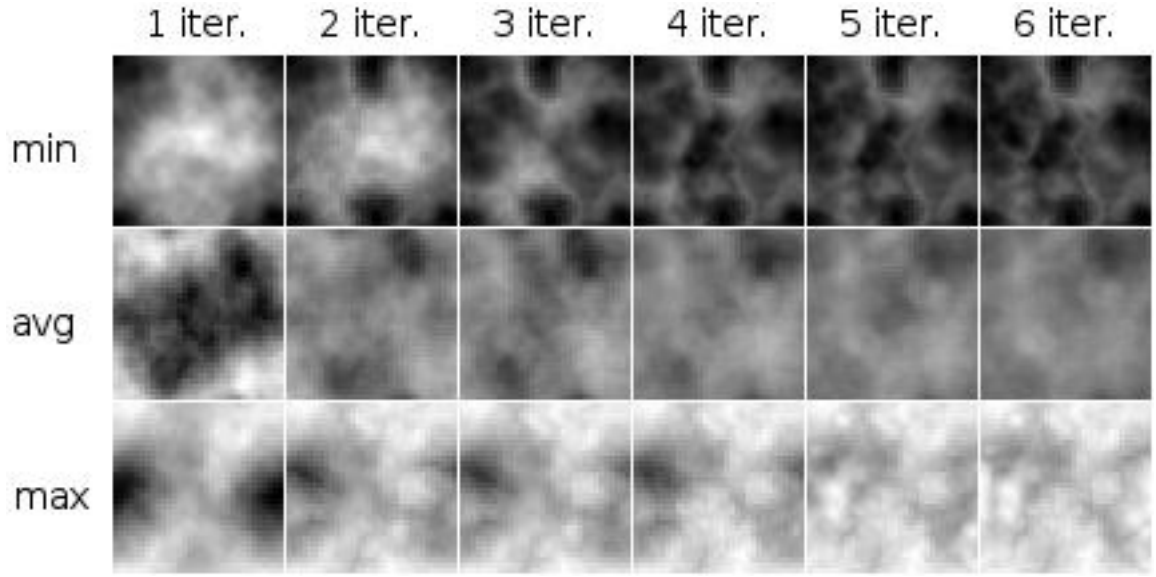

Fig. 17: Examples of varying the aggregation method and iterations in *imgaug.parameters.IterativeNoiseAggregator* (here in combination with *imgaug.parameters.FrequencyNoise*).

8.6 Sigmoid

Generated noise masks can often end up having many values around 0.5, especially when running *imgaug.parameters.IterativeNoiseAggregator* with many iterations and aggregation method `avg`. This can be undesired. *imgaug.parameters.Sigmoid* is a method to compensate that. It applies a sigmoid function to the noise masks, forcing the values to mostly lie close to 0.0 or 1.0 and only rarely in between. This can lead to blobs of values close to 1.0 ("use only colors from images coming from source A"), surrounded by blobs with values close to 0.0 ("use only colors from images coming from source B"). This is similar to taking *either* from one image source (per pixel) or the other, but usually not both. Sigmoid is inte-

grated into both class:*imgaug.augmenters.blend.BlendAlphaSimplexNoise* and *imgaug.augmenters.blend.BlendAlphaFrequencyNoise*. It can be dynamically activated/deactivated and has a threshold parameter that controls how aggressive and pushes the noise values towards 1.0.

```
# This is how the Sigmoid would be activated/deactivated for
# BlendAlphaFrequencyNoise (same for BlendAlphaSimplexNoise). P is the
# probability of the Sigmoid being activated (can be True/False), T is the
# threshold (sane values are usually around -10 to +10, can be a
# tuple, e.g. sigmoid_thresh=(-10, 10), to indicate a uniform range).
seq = iaa.BlendAlphaFrequencyNoise(
    ...,
    sigmoid=P,
    sigmoid_thresh=T
)
```

The below image shows the effects of applying *imgaug.parameters.Sigmoid* to noise masks generated by *imgaug.parameters.FrequencyNoise*.

Fig. 18: Examples of noise maps without and with activated *imgaug.parameters.Sigmoid* (noise maps here from *imgaug.parameters.FrequencyNoise*).

The below image shows the effects of varying the sigmoid's threshold. Lower values place the threshold further to the "left" (lower x values), leading to more x-values being above the threshold values, leading to more 1.0s in the noise masks.

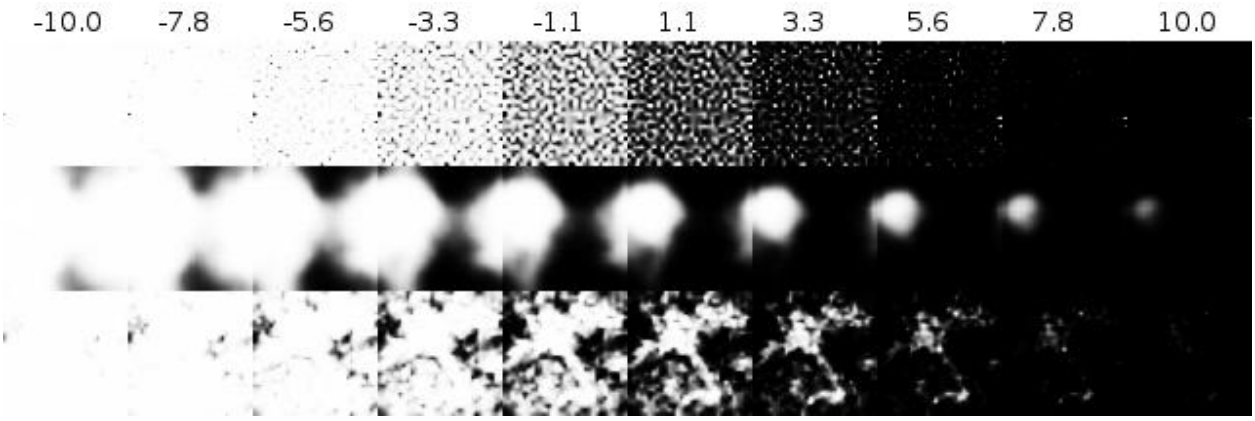

Fig. 19: Examples of varying the *imgaug.parameters.Sigmoid* threshold from `-10.0` to `10.0`.

CHAPTER 9

Overview of Augmenters

9.1 augmenters.meta

9.1.1 Sequential

List augmenter that may contain other augmenters to apply in sequence or random order.

API link: *Sequential*

Example. Apply in predefined order:

```
import imgaug.augmenters as iaa
aug = iaa.Sequential([
    iaa.Affine(translate_px={"x":-40}),
    iaa.AdditiveGaussianNoise(scale=0.1*255)
])
```

Example. Apply in random order (note that the order is sampled once per batch and then the same for all images within the batch):

```
aug = iaa.Sequential([
      iaa.Affine(translate_px={"x":-40}),
      iaa.AdditiveGaussianNoise(scale=0.1*255)
], random_order=True)
```

9.1.2 SomeOf

List augmenter that applies only some of its children to images.

API link: *SomeOf*

Example. Apply two of four given augmenters:

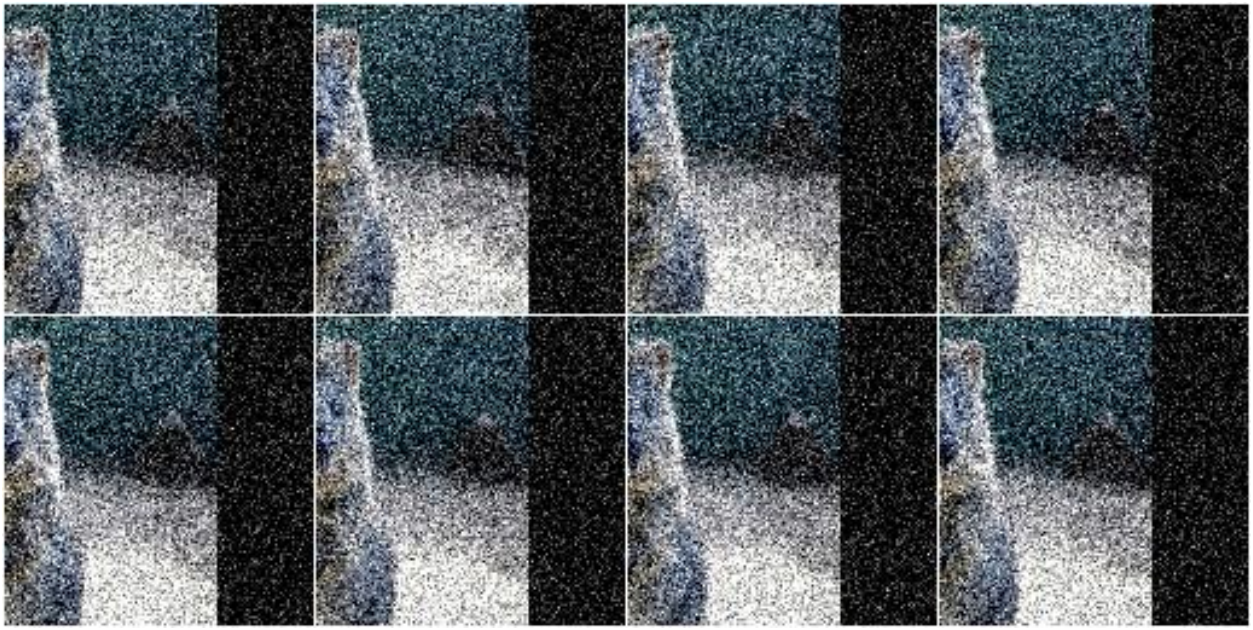

```
import imgaug.augmenters as iaa
aug = iaa.SomeOf(2, [
    iaa.Affine(rotate=45),
    iaa.AdditiveGaussianNoise(scale=0.2*255),
    iaa.Add(50, per_channel=True),
    iaa.Sharpen(alpha=0.5)
])
```

Example. Apply `0` to `<max>` given augmenters (where `<max>` is automatically replaced with the number of children):

```
aug = iaa.SomeOf((0, None), [
    iaa.Affine(rotate=45),
    iaa.AdditiveGaussianNoise(scale=0.2*255),
    iaa.Add(50, per_channel=True),
    iaa.Sharpen(alpha=0.5)
])
```

Example. Pick two of four given augmenters and apply them in random order:

```
aug = iaa.SomeOf(2, [
    iaa.Affine(rotate=45),
    iaa.AdditiveGaussianNoise(scale=0.2*255),
    iaa.Add(50, per_channel=True),
    iaa.Sharpen(alpha=0.5)
], random_order=True)
```

9.1.3 OneOf

Augmenter that always executes exactly one of its children.

API link: *OneOf()*

Example. Apply one of four augmenters to each image:

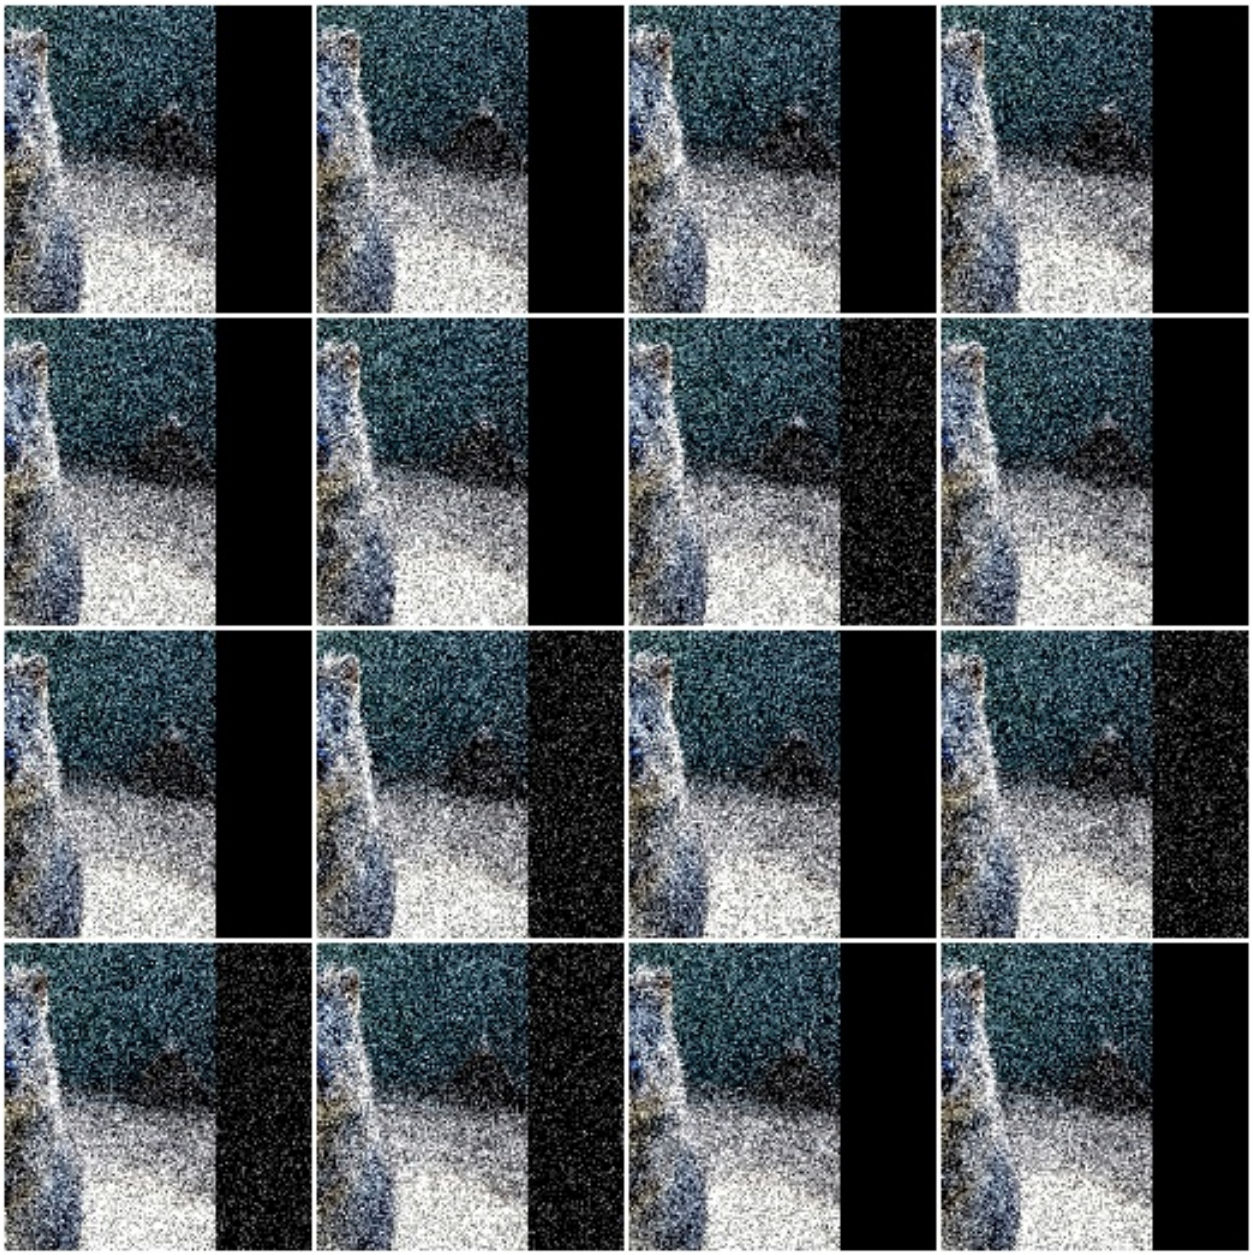

```
import imgaug.augmenters as iaa
aug = iaa.OneOf([
    iaa.Affine(rotate=45),
    iaa.AdditiveGaussianNoise(scale=0.2*255),
    iaa.Add(50, per_channel=True),
    iaa.Sharpen(alpha=0.5)
])
```

9.1.4 Sometimes

Augment only p percent of all images with one or more augmenters.

API link: *Sometimes*

Example. Apply gaussian blur to about 50% of all images:

```
import imgaug.augmenters as iaa
aug = iaa.Sometimes(0.5, iaa.GaussianBlur(sigma=2.0))
```

Example. Apply gaussian blur to about 50% of all images. Apply a mixture of affine rotations and sharpening to the other 50%.

```
aug = iaa.Sometimes(
      0.5,
      iaa.GaussianBlur(sigma=2.0),
      iaa.Sequential([iaa.Affine(rotate=45), iaa.Sharpen(alpha=1.0)])
  )
```

9.1.5 WithChannels

Apply child augmenters to specific channels.

API link: *WithChannels*

Example. Increase each pixel's R-value (redness) by `10` to `100`:

```
import imgaug.augmenters as iaa
aug = iaa.WithChannels(0, iaa.Add((10, 100)))
```

Example. Rotate each image's red channel by `0` to `45` degrees:

```
aug = iaa.WithChannels(0, iaa.Affine(rotate=(0, 45)))
```

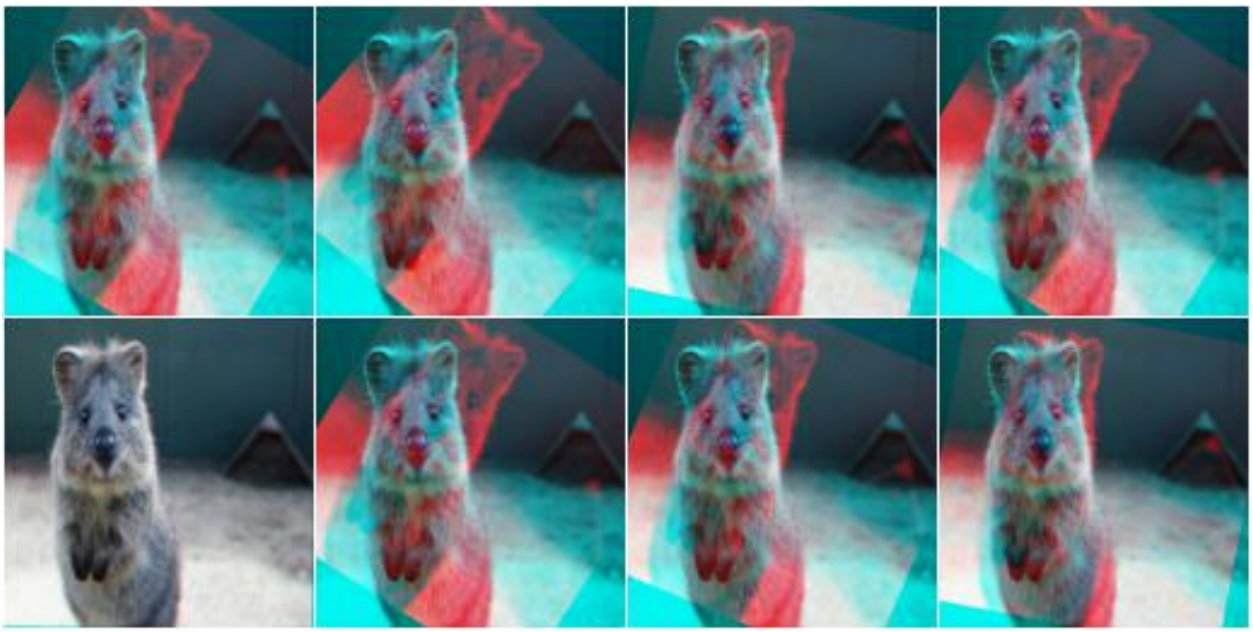

9.1.6 Identity

Augmenter that does not change the input data.

This augmenter is useful e.g. during validation/testing as it allows to re-use the training code without actually performing any augmentation.

API link: *Identity*

Example. Create an augmenter that does not change inputs:

```
import imgaug.augmenters as iaa
aug = iaa.Identity()
```

9.1.7 Noop

Alias for augmenter `Identity`.

It is recommended to now use `Identity`. `Noop` might be deprecated in the future.

API link: *Noop*

Example. Create an augmenter that does nothing:

```
import imgaug.augmenters as iaa
aug = iaa.Noop()
```

9.1.8 Lambda

Augmenter that calls a lambda function for each batch of input image.

API link: *Lambda*

Example. Replace in every image each fourth row with black pixels:

```
import imgaug.augmenters as iaa

def img_func(images, random_state, parents, hooks):
    for img in images:
        img[::4] = 0
    return images

def keypoint_func(keypoints_on_images, random_state, parents, hooks):
    return keypoints_on_images

aug = iaa.Lambda(img_func, keypoint_func)
```

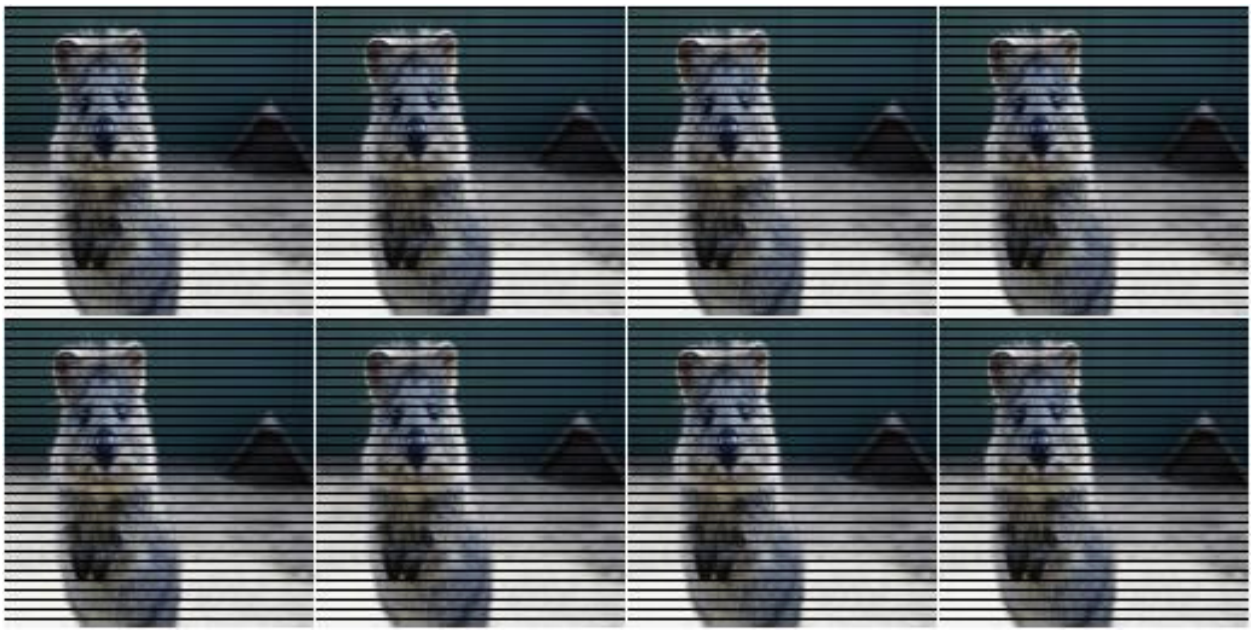

9.1.9 AssertLambda

Augmenter that runs an assert on each batch of input images using a lambda function as condition.

API link: *AssertLambda*

TODO examples

9.1.10 AssertShape

Augmenter to make assumptions about the shape of input image(s) and keypoints.

API link: *AssertShape*

Example. Check if each image in a batch has shape `32x32x3`, otherwise raise an exception:

```
import imgaug.augmenters as iaa
seq = iaa.Sequential([
    iaa.AssertShape((None, 32, 32, 3)),
    iaa.Fliplr(0.5) # only executed if shape matches
])
```

Example. Check if each image in a batch has a height in the range `32<=x<64`, a width of exactly `64` and either `1` or `3` channels:

```
seq = iaa.Sequential([
    iaa.AssertShape((None, (32, 64), 32, [1, 3])),
    iaa.Fliplr(0.5)
])
```

9.1.11 ChannelShuffle

Randomize the order of channels in input images.

API link: *ChannelShuffle*

Example. Shuffle all channels of 35% of all images:

```
import imgaug.augmenters as iaa
aug = iaa.ChannelShuffle(0.35)
```

Example. Shuffle only channels `0` and `1` of 35% of all images. As the new channel orders `0, 1` and `1, 0` are both valid outcomes of the shuffling, it means that for `0.35 * 0.5 = 0.175` or 17.5% of all images the order of channels `0` and `1` is inverted.

```
aug = iaa.ChannelShuffle(0.35, channels=[0, 1])
```

9.1.12 RemoveCBAsByOutOfImageFraction

Remove coordinate-based augmentables exceeding an out of image fraction.

This augmenter inspects all coordinate-based augmentables (e.g. bounding boxes, line strings) within a given batch and removes any such augmentable which's out of image fraction is exactly a given value or greater than that. The out of image fraction denotes the fraction of the augmentable's area that is outside of the image, e.g. for a bounding box that has half of its area outside of the image it would be `0.5`.

API link: *RemoveCBAsByOutOfImageFraction*

Example. Translate all inputs by `-100` to `100` pixels on the x-axis, then remove any coordinate-based augmentable (e.g. bounding boxes) which has at least `50%` of its area outside of the image plane:

```
import imgaug.augmenters as iaa
aug = iaa.Sequential([
    iaa.Affine(translate_px={"x": (-100, 100)}),
    iaa.RemoveCBAsByOutOfImageFraction(0.5)
])
```

Example. Create a bounding box on an example image, then translate the image so that `50%` of the bounding box's area is outside of the image and compare the effects and using `RemoveCBAsByOutOfImageFraction` with not using it.

```
import imgaug as ia
import imgaug.augmenters as iaa
image = ia.quokka_square((100, 100))
bb = ia.BoundingBox(x1=50-25, y1=0, x2=50+25, y2=100)
bbsoi = ia.BoundingBoxesOnImage([bb], shape=image.shape)
aug_without = iaa.Affine(translate_px={"x": 51})
aug_with = iaa.Sequential([
    iaa.Affine(translate_px={"x": 51}),
    iaa.RemoveCBAsByOutOfImageFraction(0.5)
```

(continues on next page)

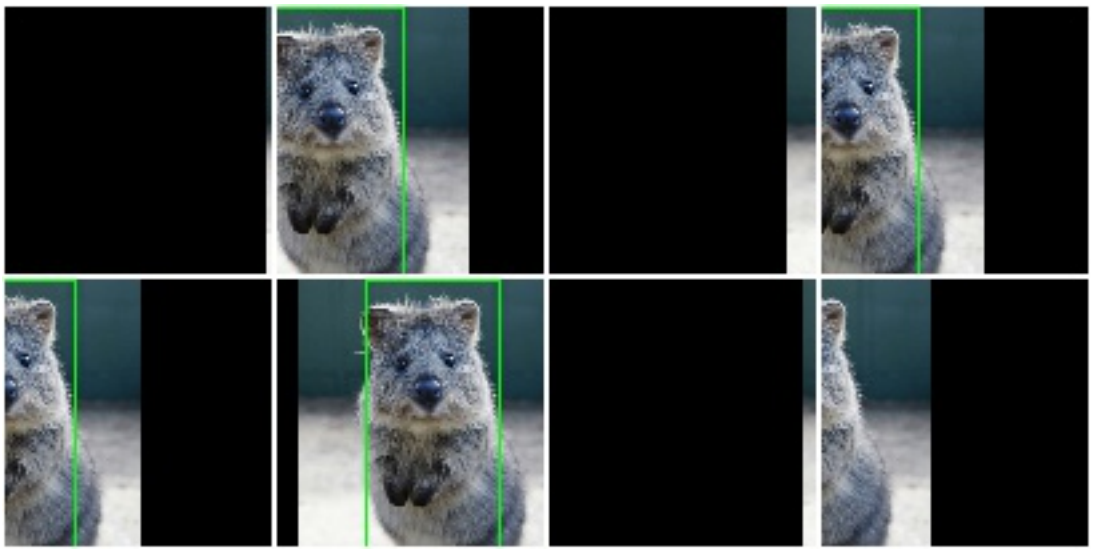

(continued from previous page)

```
])

image_without, bbsoi_without = aug_without(
    image=image, bounding_boxes=bbsoi)
image_with, bbsoi_with = aug_with(
    image=image, bounding_boxes=bbsoi)

assert len(bbsoi_without.bounding_boxes) == 1
assert len(bbsoi_with.bounding_boxes) == 0
```

9.1.13 ClipCBAsToImagePlanes

Clip coordinate-based augmentables to areas within the image plane.

This augmenter inspects all coordinate-based augmentables (e.g. bounding boxes, line strings) within a given batch and from each of them parts that are outside of the image plane. Parts within the image plane will be retained. This may e.g. shrink down bounding boxes. For keypoints, it removes any single points outside of the image plane. Any augmentable that is completely outside of the image plane will be removed.

API link: *ClipCBAsToImagePlanes*

Example. Translate input data on the x-axis by `-100` to `100` pixels, then cut all coordinate-based augmentables (e.g. bounding boxes) down to areas that are within the image planes of their corresponding images:

```
import imgaug.augmenters as iaa
aug = iaa.Sequential([
    iaa.Affine(translate_px={"x": (-100, 100)}),
    iaa.ClipCBAsToImagePlanes()
])
```

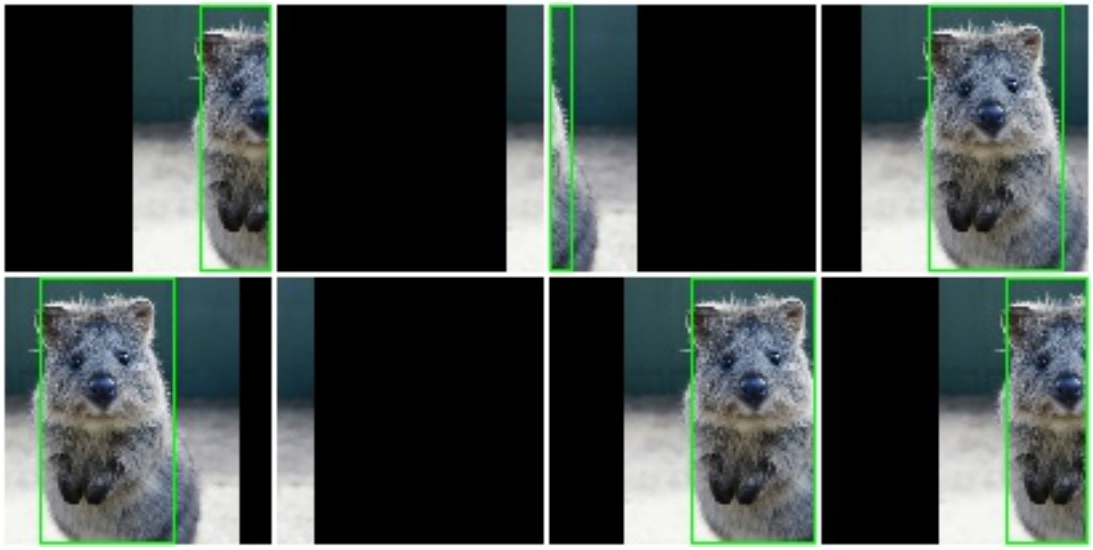

9.2 augmenters.arithmetic

9.2.1 Add

Add a value to all pixels in an image.

API link: *Add*

Example. Add random values between -40 and 40 to images, with each value being sampled once per image and then being the same for all pixels:

```
import imgaug.augmenters as iaa
aug = iaa.Add((-40, 40))
```

Example. Add random values between -40 and 40 to images. In 50% of all images the values differ per channel (3 sampled value). In the other 50% of all images the value is the same for all channels:

```
aug = iaa.Add((-40, 40), per_channel=0.5)
```

9.2.2 AddElementwise

Add values to the pixels of images with possibly different values for neighbouring pixels.

API link: *AddElementwise*

Example. Add random values between -40 and 40 to images, with each value being sampled per pixel:

```
import imgaug.augmenters as iaa
aug = iaa.AddElementwise((-40, 40))
```

Example. Add random values between -40 and 40 to images. In 50% of all images the values differ per channel (3 sampled values per pixel). In the other 50% of all images the value is the same for all channels per pixel:

```
aug = iaa.AddElementwise((-40, 40), per_channel=0.5)
```

9.2.3 AdditiveGaussianNoise

Add noise sampled from gaussian distributions elementwise to images.

API link: *AdditiveGaussianNoise()*

Example. Add gaussian noise to an image, sampled once per pixel from a normal distribution `N(0, s)`, where `s` is sampled per image and varies between `0` and `0.2*255`:

```
import imgaug.augmenters as iaa
aug = iaa.AdditiveGaussianNoise(scale=(0, 0.2*255))
```

Example. Add gaussian noise to an image, sampled once per pixel from a normal distribution `N(0, 0.05*255)`:

```
aug = iaa.AdditiveGaussianNoise(scale=0.2*255)
```

Example. Add gaussian noise to an image, sampled channelwise from `N(0, 0.2*255)` (i.e. three independent samples per pixel):

```
aug = iaa.AdditiveGaussianNoise(scale=0.2*255, per_channel=True)
```

9.2.4 AdditiveLaplaceNoise

Add noise sampled from laplace distributions elementwise to images.

The laplace distribution is similar to the gaussian distribution, but puts more weight on the long tail. Hence, this noise will add more outliers (very high/low values). It is somewhere between gaussian noise and salt and pepper noise.

API link: *AdditiveLaplaceNoise()*

Example. Add laplace noise to an image, sampled once per pixel from `Laplace(0, s)`, where `s` is sampled per image and varies between `0` and `0.2*255`:

```
import imgaug.augmenters as iaa
aug = iaa.AdditiveLaplaceNoise(scale=(0, 0.2*255))
```

Example. Add laplace noise to an image, sampled once per pixel from `Laplace(0, 0.2*255)`:

```
aug = iaa.AdditiveLaplaceNoise(scale=0.2*255)
```

Example. Add laplace noise to an image, sampled channelwise from `Laplace(0, 0.2*255)` (i.e. three independent samples per pixel):

```
aug = iaa.AdditiveLaplaceNoise(scale=0.2*255, per_channel=True)
```

9.2.5 AdditivePoissonNoise

Add noise sampled from poisson distributions elementwise to images.

Poisson noise is comparable to gaussian noise, as e.g. generated via `AdditiveGaussianNoise`. As poisson distributions produce only positive numbers, the sign of the sampled values are here randomly flipped.

Values of around `20.0` for `lam` lead to visible noise (for `uint8`). Values of around `40.0` for `lam` lead to very visible noise (for `uint8`). It is recommended to usually set `per_channel` to `True`.

API link: *AdditivePoissonNoise()*

Example. Add poisson noise to an image, sampled once per pixel from `Poisson(lam)`, where `lam` is sampled per image and varies between `0` and `40`:

```
import imgaug.augmenters as iaa
aug = iaa.AdditivePoissonNoise(scale=(0, 40))
```

Example. Add poisson noise to an image, sampled once per pixel from `Poisson(40)`:

```
aug = iaa.AdditivePoissonNoise(40)
```

Example. Add poisson noise to an image, sampled channelwise from `Poisson(40)` (i.e. three independent samples per pixel):

```
aug = iaa.AdditivePoissonNoise(scale=40, per_channel=True)
```

9.2.6 Multiply

Multiply all pixels in an image with a specific value, thereby making the image darker or brighter.

API link: *Multiply*

Example. Multiply each image with a random value between 0.5 and 1.5:

```
import imgaug.augmenters as iaa
aug = iaa.Multiply((0.5, 1.5))
```

Example. Multiply 50% of all images with a random value between 0.5 and 1.5 and multiply the remaining 50% channel-wise, i.e. sample one multiplier independently per channel:

```
aug = iaa.Multiply((0.5, 1.5), per_channel=0.5)
```

9.2.7 MultiplyElementwise

Multiply values of pixels with possibly different values for neighbouring pixels, making each pixel darker or brighter.

API link: *MultiplyElementwise*

Example. Multiply each pixel with a random value between 0.5 and 1.5:

```
import imgaug.augmenters as iaa
aug = iaa.MultiplyElementwise((0.5, 1.5))
```

Example. Multiply in 50% of all images each pixel with random values between 0.5 and 1.5 and multiply in the remaining 50% of all images the pixels channel-wise, i.e. sample one multiplier independently per channel and pixel:

```
aug = iaa.MultiplyElementwise((0.5, 1.5), per_channel=0.5)
```

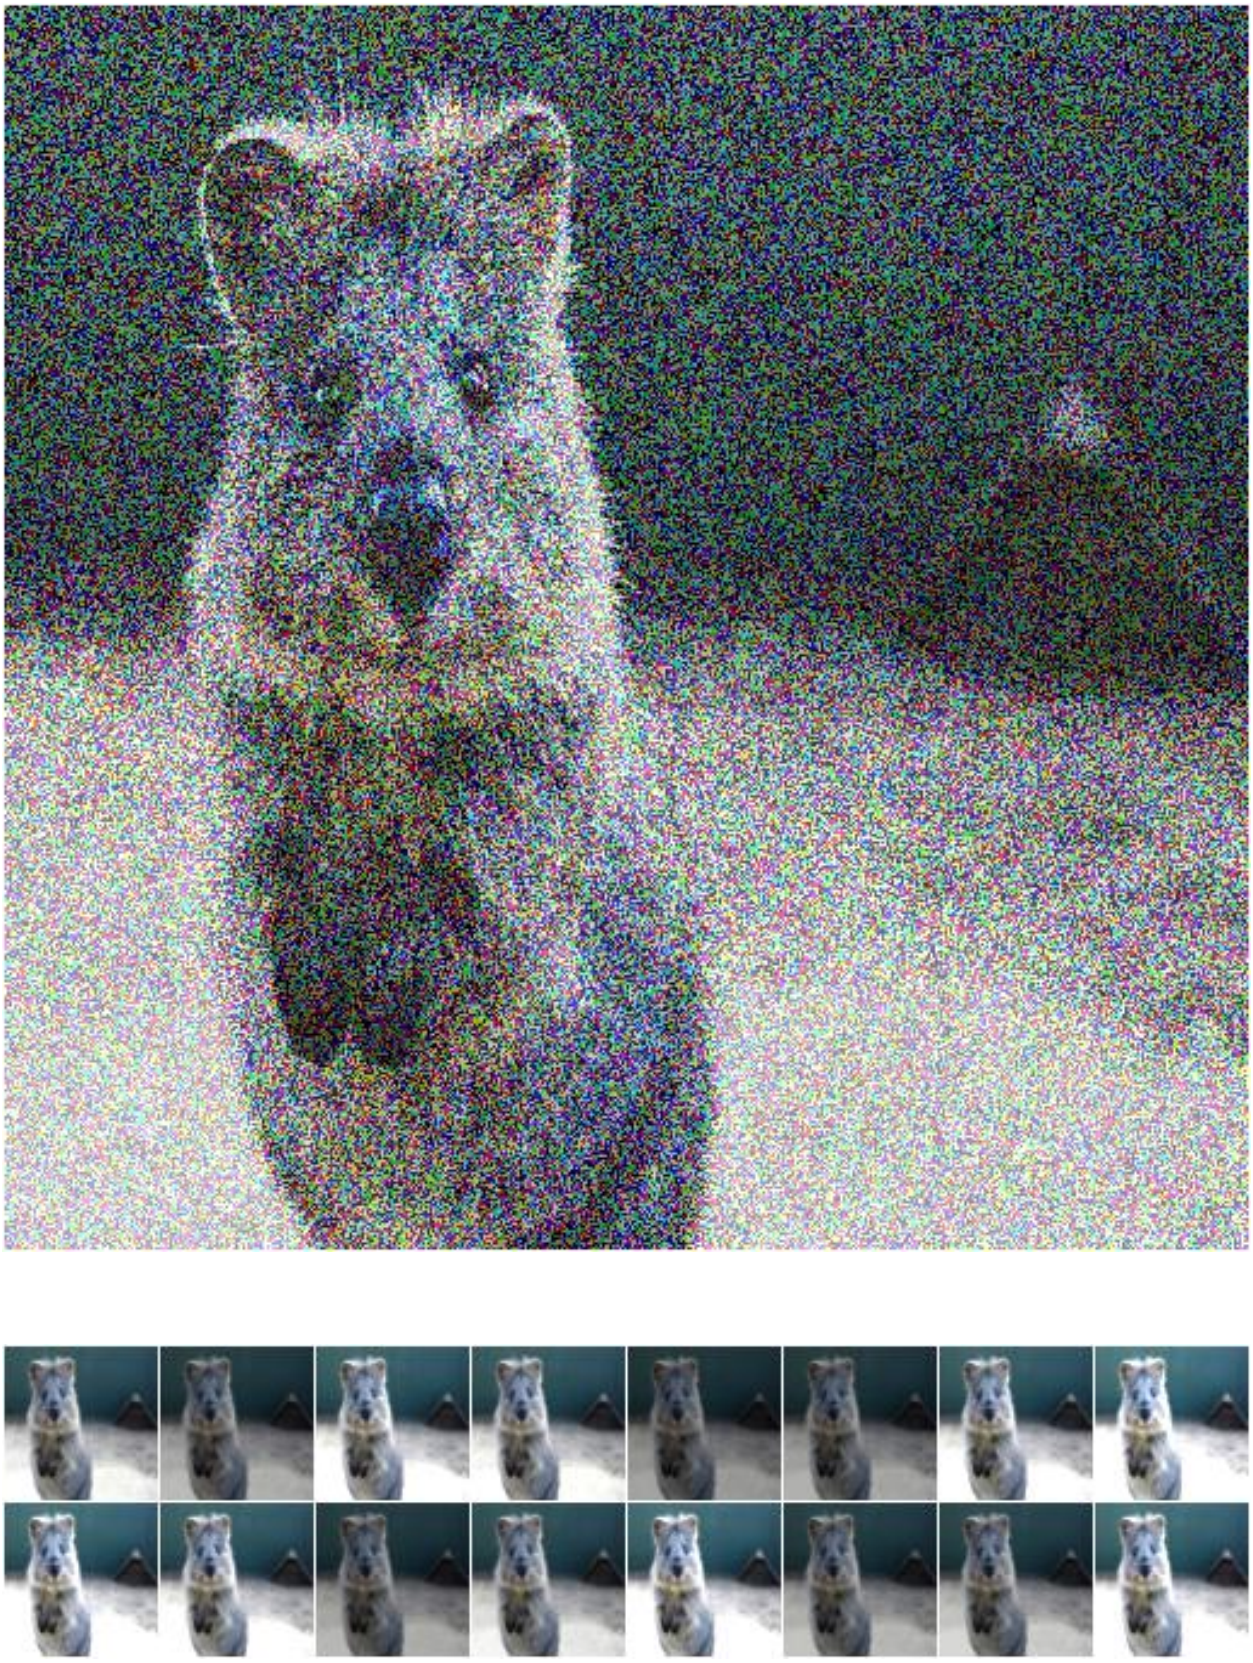

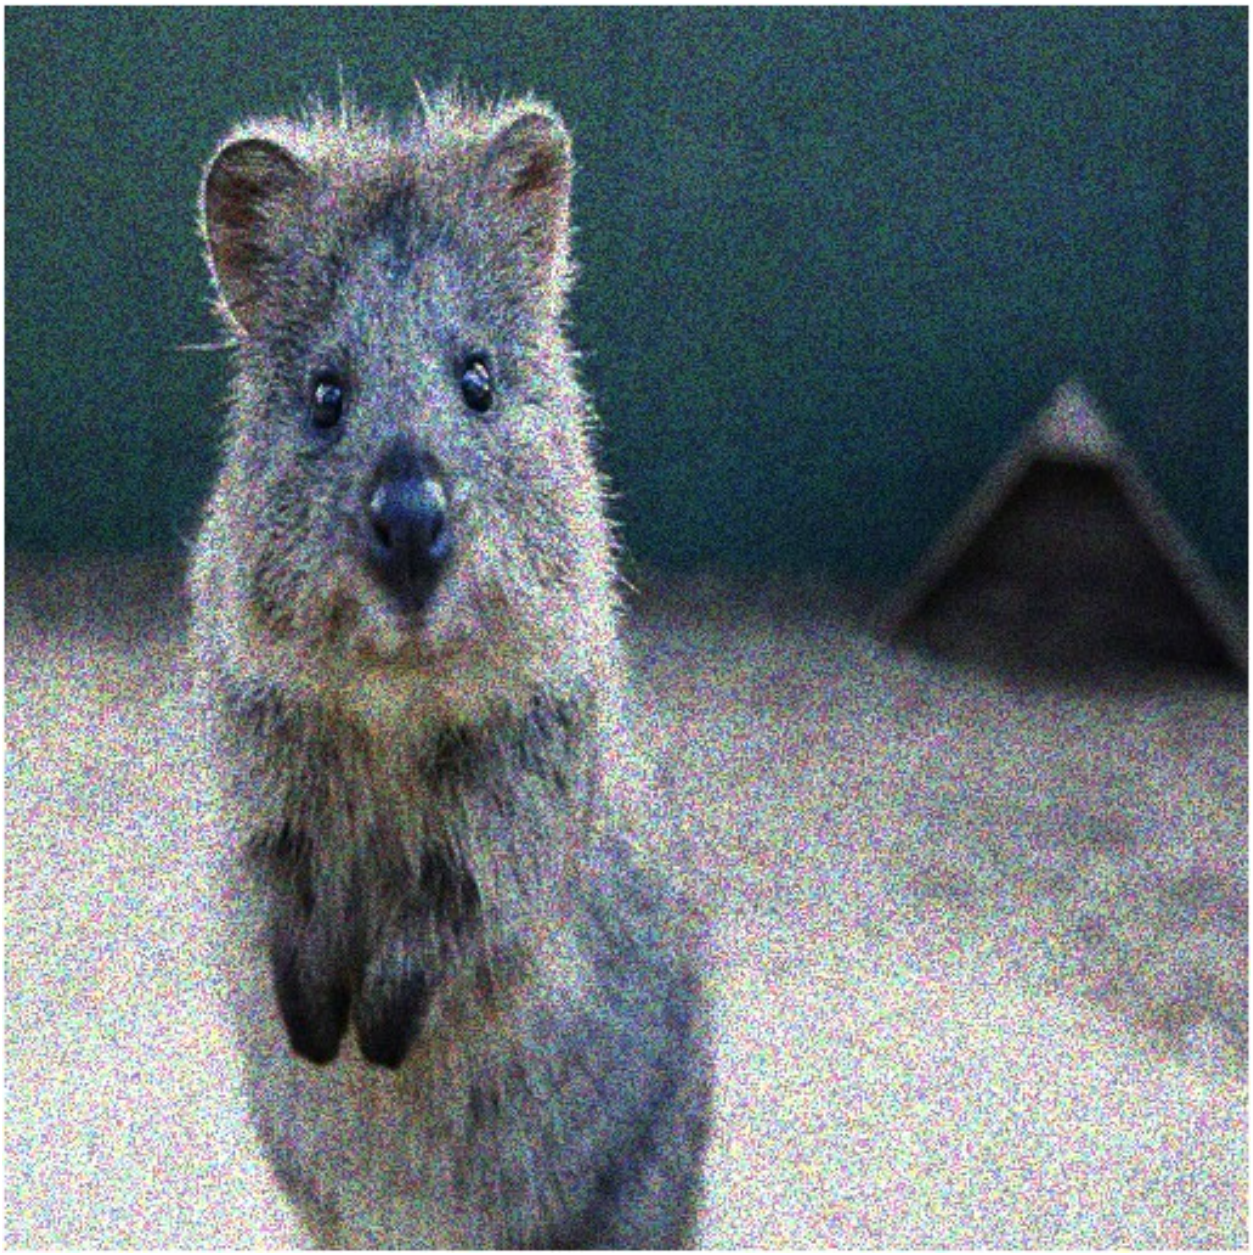

9.2.8 Cutout

Fill one or more rectangular areas in an image using a fill mode.

See paper "Improved Regularization of Convolutional Neural Networks with Cutout" by DeVries and Taylor.

In contrast to the paper, this implementation also supports replacing image sub-areas with gaussian noise, random intensities or random RGB colors. It also supports non-squared areas. While the paper uses absolute pixel values for the size and position, this implementation uses relative values, which seems more appropriate for mixed-size datasets. The position parameter furthermore allows more flexibility, e.g. gaussian distributions around the center.

Note: This augmenter affects only image data. Other datatypes (e.g. segmentation map pixels or keypoints within the filled areas) are not affected.

Note: Gaussian fill mode will assume that float input images contain values in the interval `[0.0, 1.0]` and hence sample values from a gaussian within that interval, i.e. from `N(0.5, std=0.5/3)`.

API link: *MultiplyElementwise*

Example. Fill per image two random areas, by default with grayish pixels:

```
import imgaug.augmenters as iaa
aug = iaa.Cutout(nb_iterations=2)
```

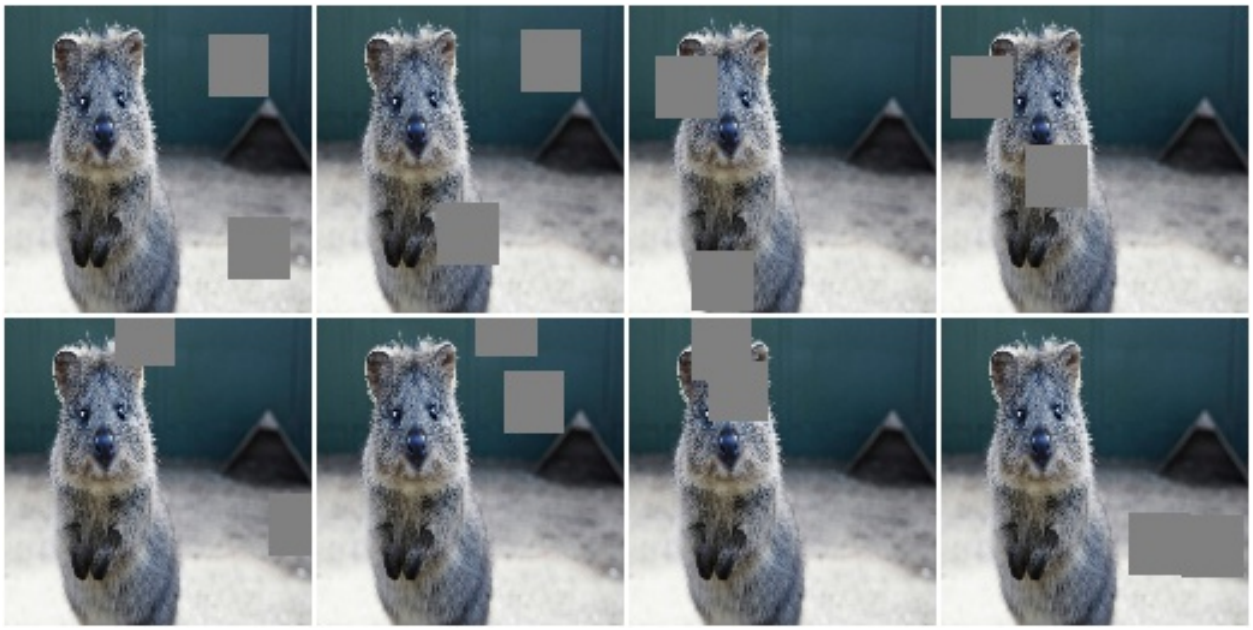

Example. Fill per image between one and five areas, each having `20%` of the corresponding size of the height and width (for non-square images this results in non-square areas to be filled).

```
aug = iaa.Cutout(nb_iterations=(1, 5), size=0.2, squared=False)
```

Example. Fill all areas with white pixels:

```
aug = iaa.Cutout(fill_mode="constant", cval=255)
```

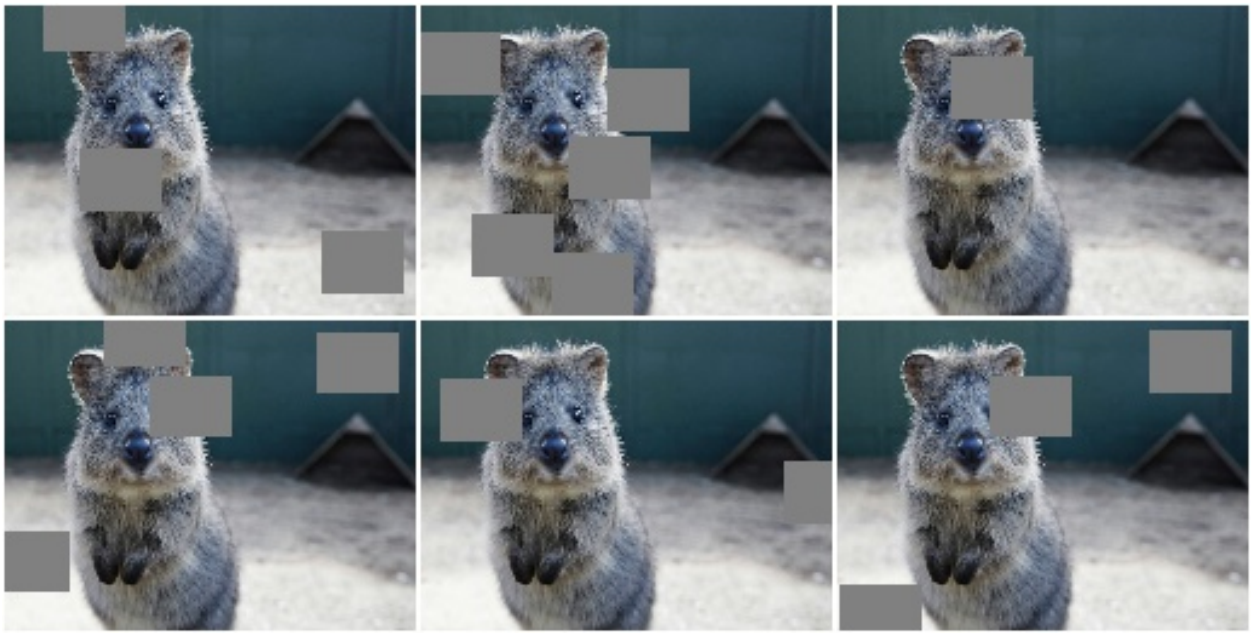

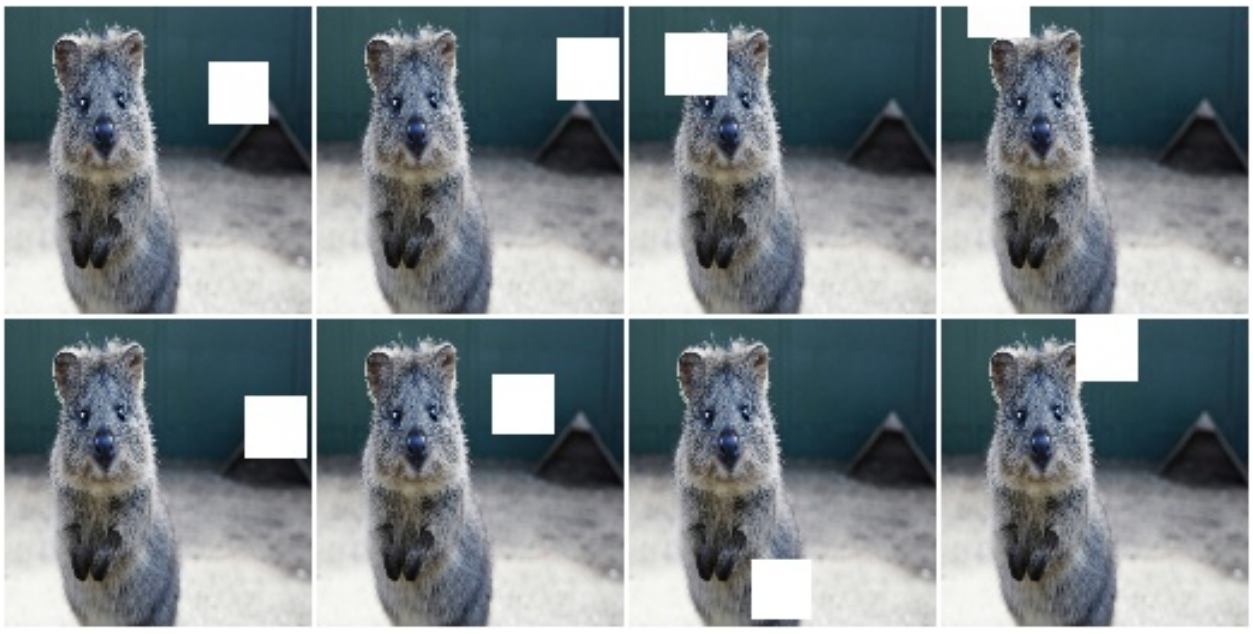

Example. Fill 50% of all areas with a random intensity value between 0 and 256. Fill the other 50% of all areas with random colors.

```
aug = iaa.Cutout(fill_mode="constant", cval=(0, 255),
                 fill_per_channel=0.5)
```

Example. Fill areas with gaussian channelwise noise (i.e. usually RGB).

```
aug = iaa.Cutout(fill_mode="gaussian", fill_per_channel=True)
```

9.2.9 Dropout

Augmenter that sets a certain fraction of pixels in images to zero.

API link: *Dropout()*

Example. Sample per image a value `p` from the range `0<=p<=0.2` and then drop `p` percent of all pixels in the image (i.e. convert them to black pixels):

```
import imgaug.augmenters as iaa
aug = iaa.Dropout(p=(0, 0.2))
```

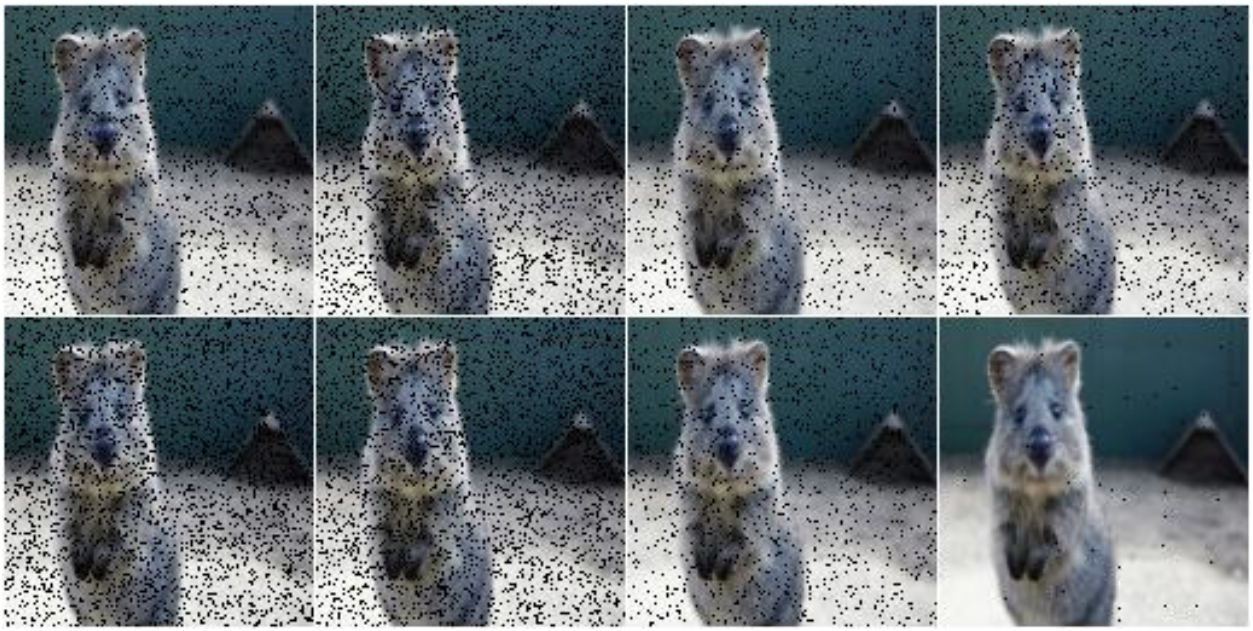

Example. Sample per image a value `p` from the range `0<=p<=0.2` and then drop `p` percent of all pixels in the image (i.e. convert them to black pixels), but do this independently per channel in 50% of all images:

```
aug = iaa.Dropout(p=(0, 0.2), per_channel=0.5)
```

9.2.10 CoarseDropout

Augmenter that sets rectangular areas within images to zero.

API link: *CoarseDropout()*

Example. Drop 2% of all pixels by converting them to black pixels, but do that on a lower-resolution version of the image that has 50% of the original size, leading to 2x2 squares being dropped:

```
import imgaug.augmenters as iaa
aug = iaa.CoarseDropout(0.02, size_percent=0.5)
```

Example. Drop 0 to 5% of all pixels by converting them to black pixels, but do that on a lower-resolution version of the image that has 5% to 50% of the original size, leading to large rectangular areas being dropped:

```
import imgaug.augmenters as iaa
aug = iaa.CoarseDropout((0.0, 0.05), size_percent=(0.02, 0.25))
```

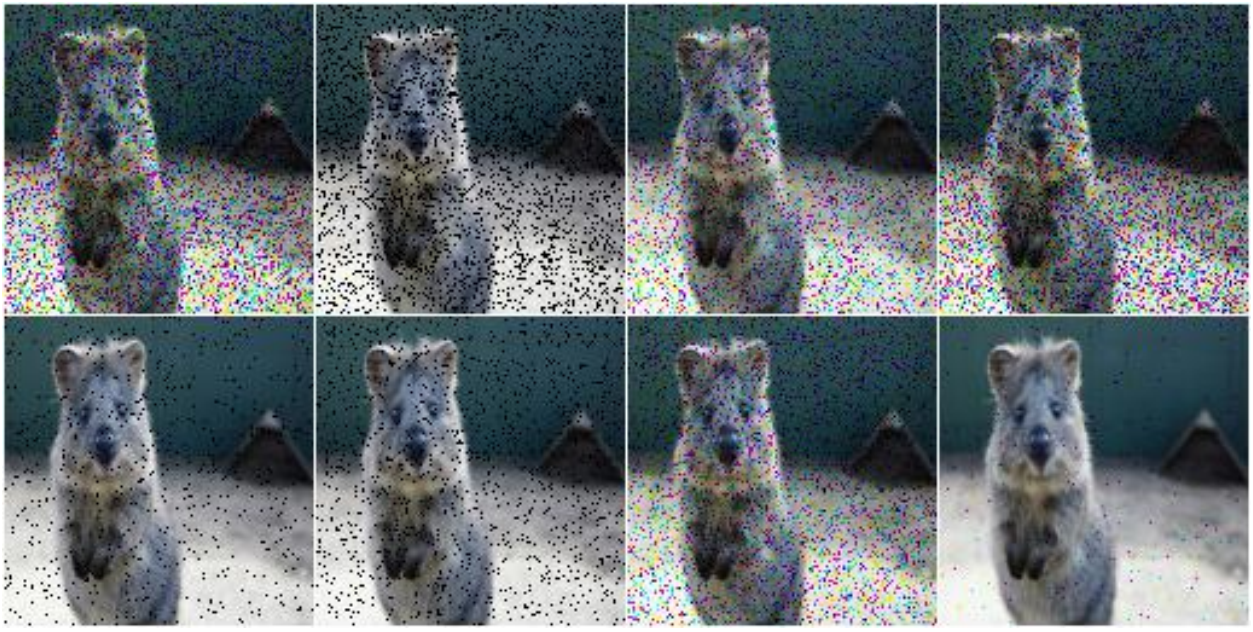

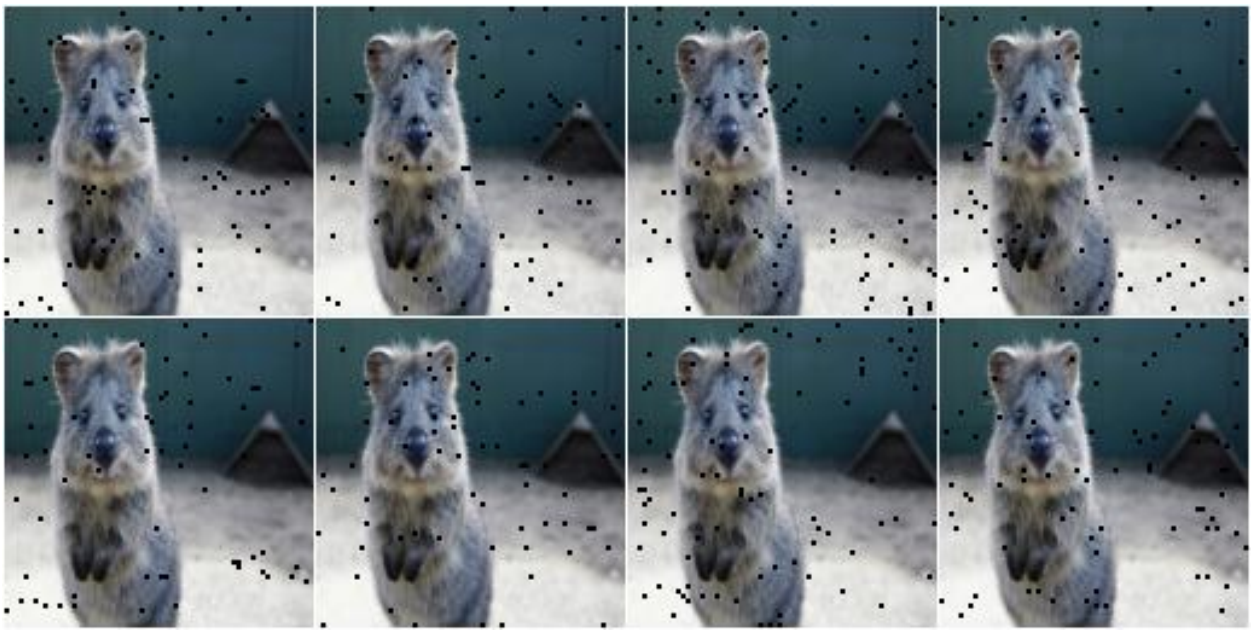

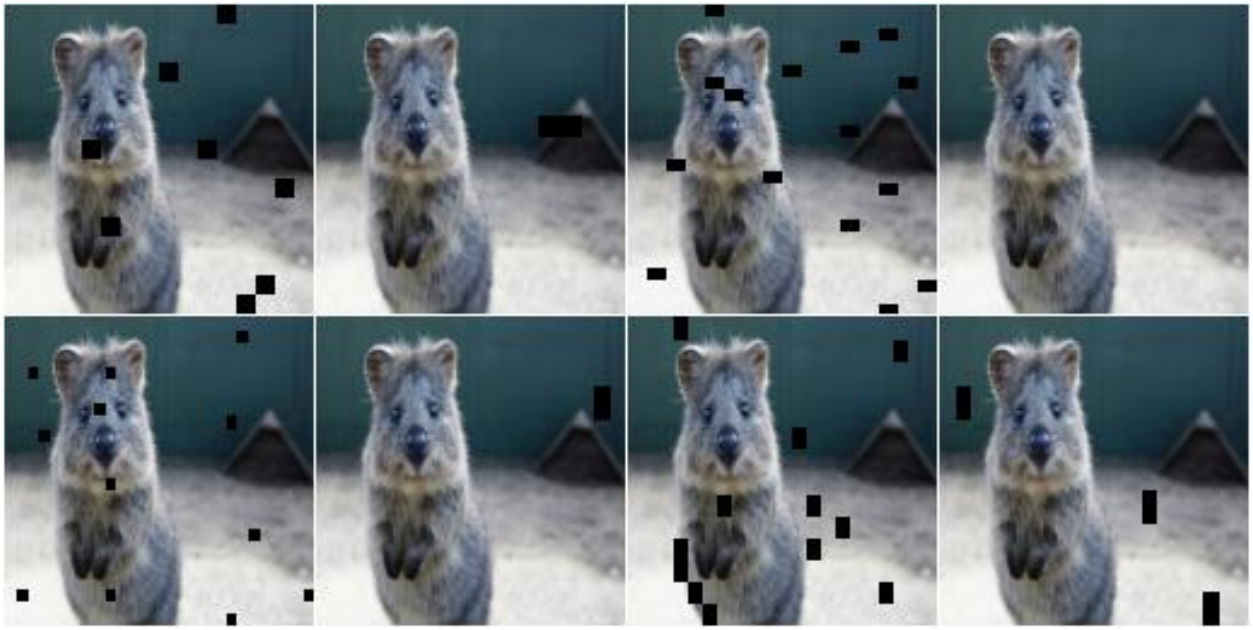

Example. Drop 2% of all pixels by converting them to black pixels, but do that on a lower-resolution version of the image that has 50% of the original size, leading to 2x2 squares being dropped. Also do this in 50% of all images channel-wise, so that only the information of some channels in set to 0 while others remain untouched:

```
aug = iaa.CoarseDropout(0.02, size_percent=0.15, per_channel=0.5)
```

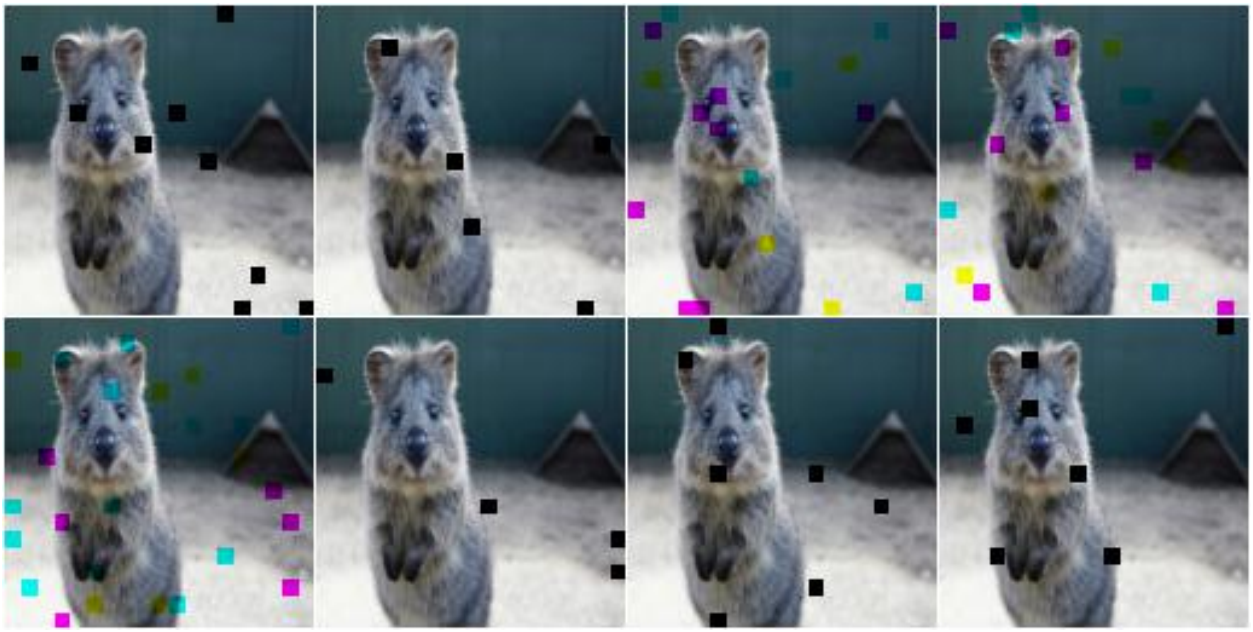

9.2.11 Dropout2D

Drop random channels from images.

For image data, dropped channels will be filled with zeros.

Note: This augmenter may also set the arrays of heatmaps and segmentation maps to zero and remove all coordinate-based data (e.g. it removes all bounding boxes on images that were filled with zeros). It does so if and only if *all* channels of an image are dropped. If `nb_keep_channels >= 1` then that never happens.

API link: *Dropout2d()*

Example. Create a dropout augmenter that drops on average half of all image channels. Dropped channels will be filled with zeros. At least one channel is kept unaltered in each image (default setting).

```
import imgaug.augmenters as iaa
aug = iaa.Dropout2d(p=0.5)
```

Example. Create a dropout augmenter that drops on average half of all image channels *and* may drop *all* channels in

an image (i.e. images may contain nothing but zeros):

```
import imgaug.augmenters as iaa
aug = iaa.Dropout2d(p=0.5, nb_keep_channels=0)
```

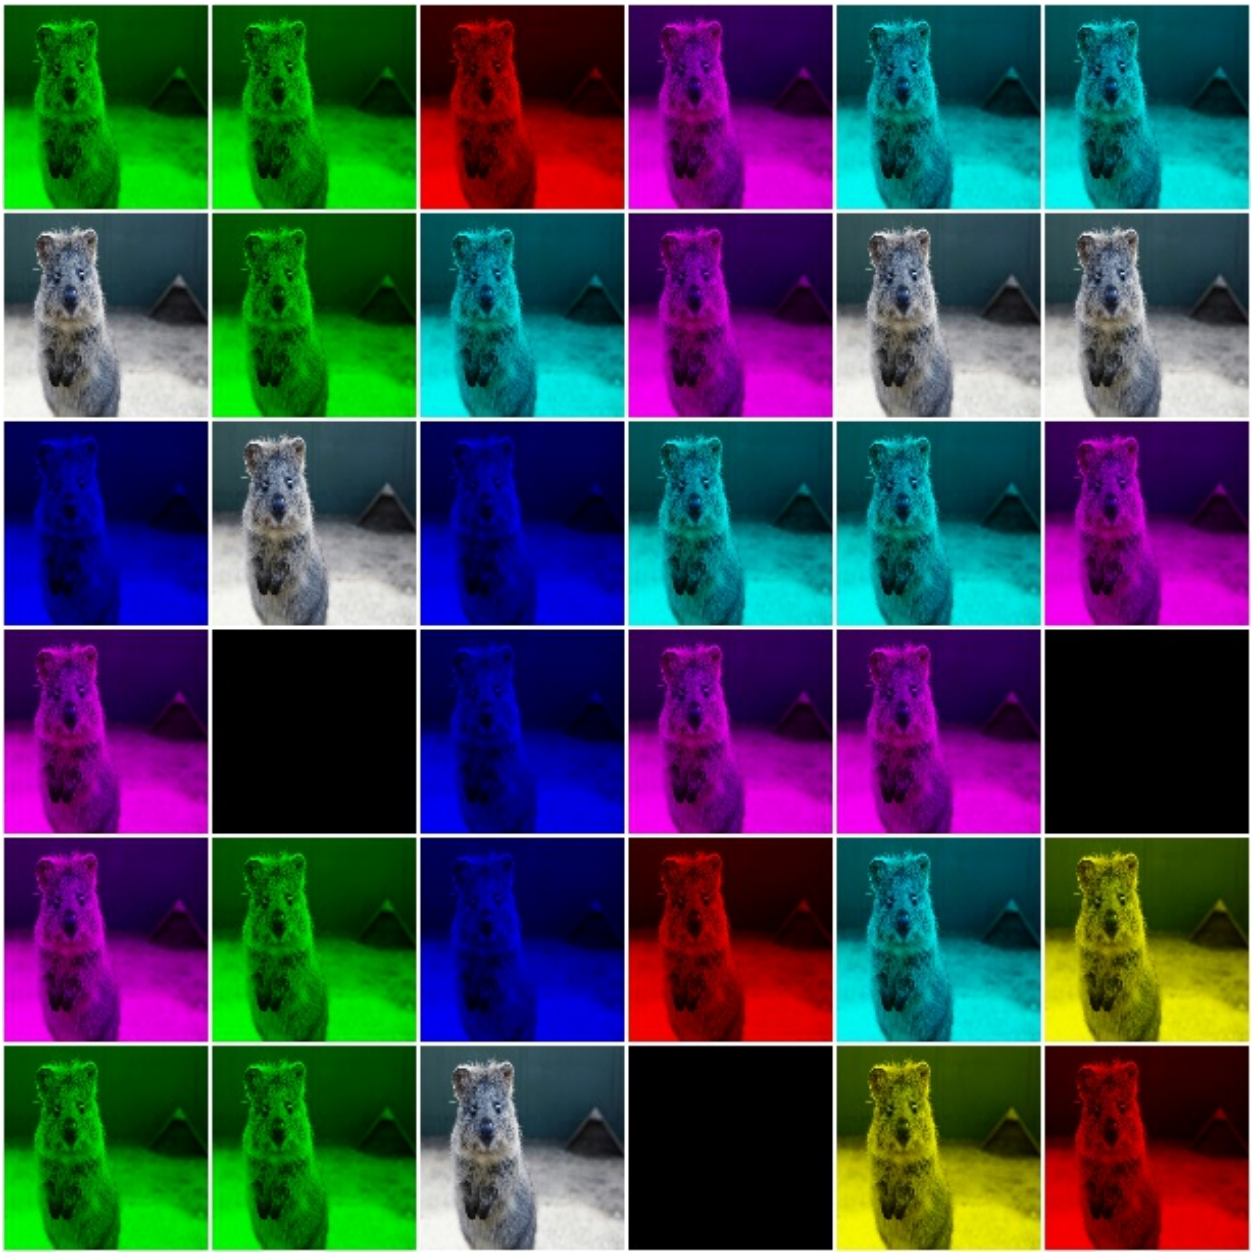

9.2.12 TotalDropout

Drop all channels of a defined fraction of all images.

For image data, all components of dropped images will be filled with zeros.

Note: This augmenter also sets the arrays of heatmaps and segmentation maps to zero and removes all coordinate-

based data (e.g. it removes all bounding boxes on images that were filled with zeros).

API link: *TotalDropout()*

Example. Create an augmenter that sets *all* components of all images to zero:

```
import imgaug.augmenters as iaa
aug = iaa.TotalDropout(1.0)
```

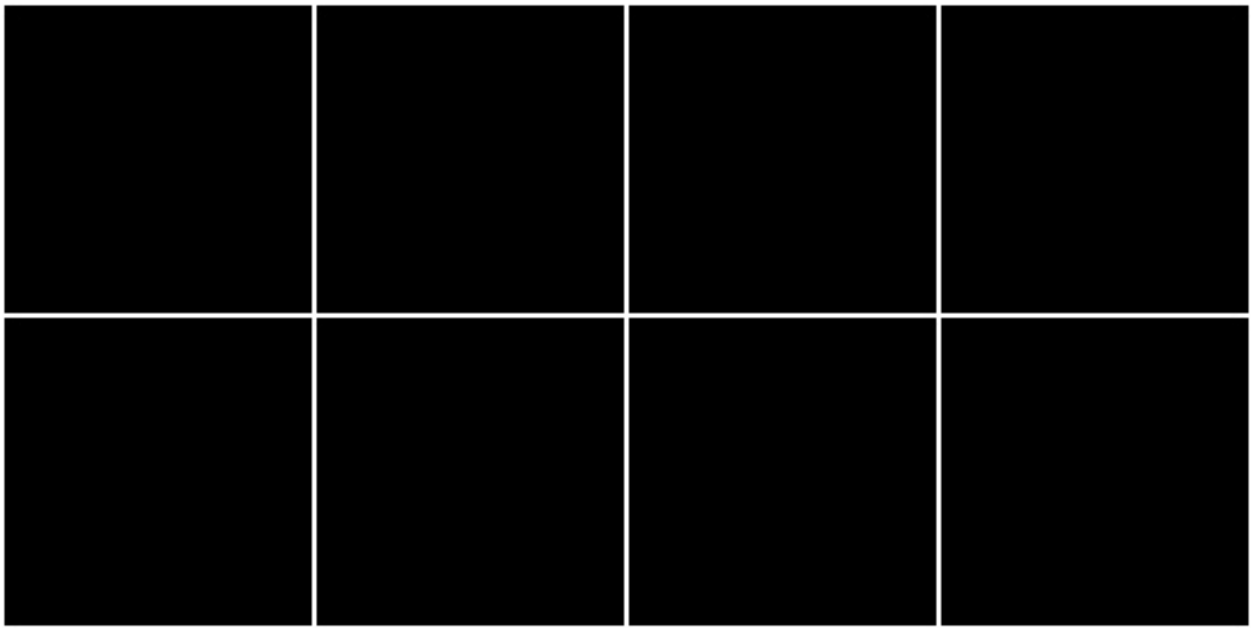

Example. Create an augmenter that sets *all* components of `50%` of all images to zero:

```
aug = iaa.TotalDropout(0.5)
```

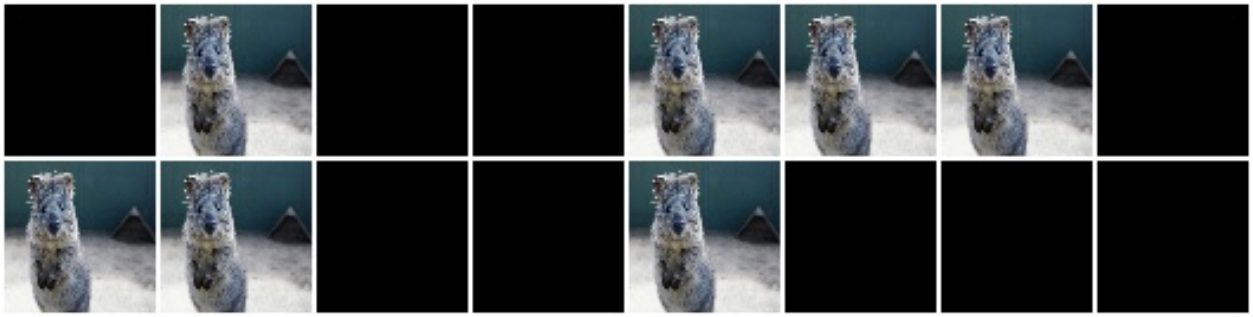

9.2.13 ReplaceElementwise

Replace pixels in an image with new values.

API link: *ReplaceElementwise*

Example. Replace `10%` of all pixels with either the value `0` or the value `255`:

```
import imgaug.augmenters as iaa
aug = ReplaceElementwise(0.1, [0, 255])
```

Example. For `50%` of all images, replace `10%` of all pixels with either the value `0` or the value `255` (same as in the previous example). For the other `50%` of all images, replace *channelwise* `10%` of all pixels with either the value `0` or the value `255`. So, it will be very rare for each pixel to have all channels replaced by `255` or `0`.

```
aug = ReplaceElementwise(0.1, [0, 255], per_channel=0.5)
```

Example. Replace `10%` of all pixels by gaussian noise centered around `128`. Both the replacement mask and the gaussian noise are sampled for `50%` of all images.

```
import imgaug.parameters as iap
aug = ReplaceElementwise(0.1, iap.Normal(128, 0.4*128), per_channel=0.5)
```

Example. Replace `10%` of all pixels by gaussian noise centered around `128`. Sample the replacement mask at a lower resolution (`8x8` pixels) and upscale it to the image size, resulting in coarse areas being replaced by gaussian noise.

```
aug = ReplaceElementwise(
    iap.FromLowerResolution(iap.Binomial(0.1), size_px=8),
    iap.Normal(128, 0.4*128),
    per_channel=0.5)
```

9.2.14 ImpulseNoise

Add impulse noise to images.

This is identical to `SaltAndPepper`, except that `per_channel` is always set to `True`.

API link: *ImpulseNoise()*

Example. Replace `10%` of all pixels with impulse noise:

```
import imgaug.augmenters as iaa
aug = iaa.ImpulseNoise(0.1)
```

9.2.15 SaltAndPepper

Replace pixels in images with salt/pepper noise (white/black-ish colors).

API link: *SaltAndPepper()*

Example. Replace `10%` of all pixels with salt and pepper noise:

```
import imgaug.augmenters as iaa
aug = iaa.SaltAndPepper(0.1)
```

Example. Replace *channelwise* `10%` of all pixels with salt and pepper noise:

```
aug = iaa.SaltAndPepper(0.1, per_channel=True)
```

9.2.16 CoarseSaltAndPepper

Replace rectangular areas in images with white/black-ish pixel noise.

API link: *CoarseSaltAndPepper()*

Example. Mark `5%` of all pixels in a mask to be replaced by salt/pepper noise. The mask has `1%` to `10%` the size of the input image. The mask is then upscaled to the input image size, leading to large rectangular areas being marked as to be replaced. These areas are then replaced in the input image by salt/pepper noise.

```
import imgaug.augmenters as iaa
aug = iaa.CoarseSaltAndPepper(0.05, size_percent=(0.01, 0.1))
```

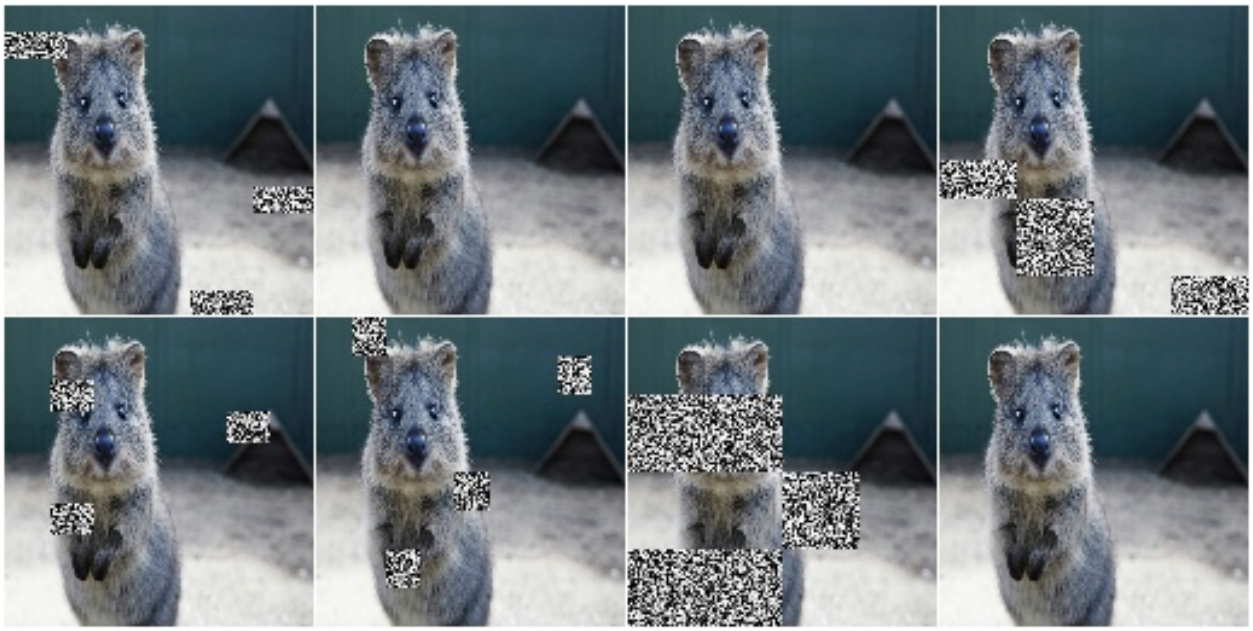

Example. Same as in the previous example, but the replacement mask before upscaling has a size between `4x4` and `16x16` pixels (the axis sizes are sampled independently, i.e. the mask may be rectangular).

```
aug = iaa.CoarseSaltAndPepper(0.05, size_px=(4, 16))
```

Example. Same as in the first example, but mask and replacement are each sampled independently per image channel.

```
aug = iaa.CoarseSaltAndPepper(
    0.05, size_percent=(0.01, 0.1), per_channel=True)
```

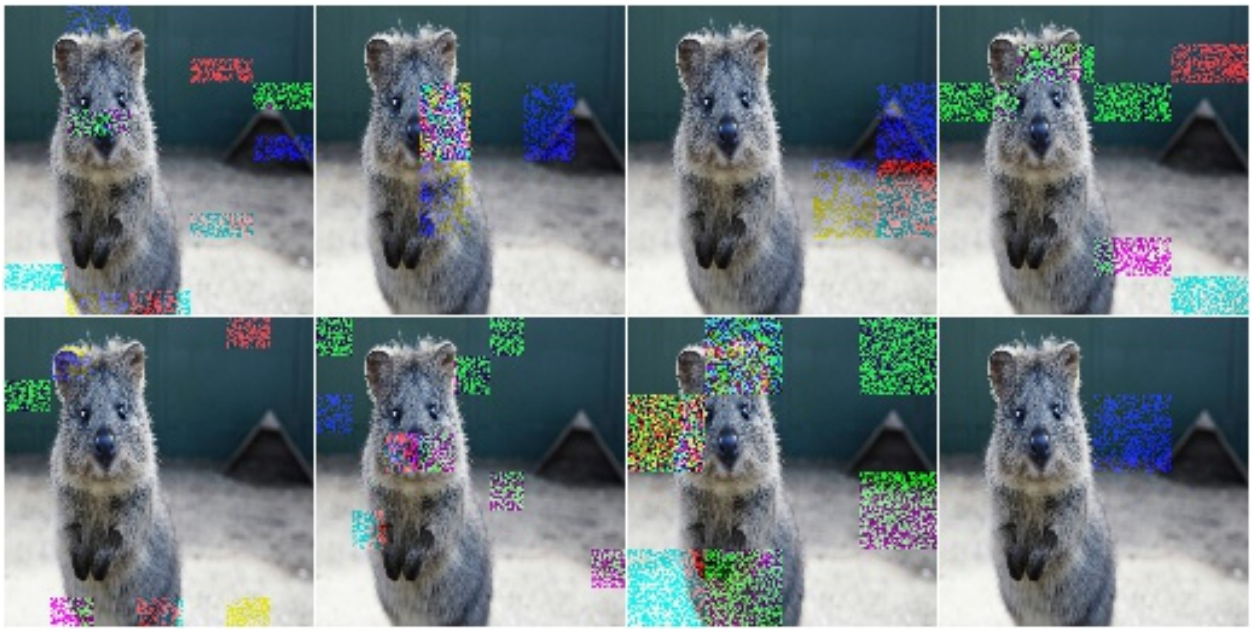

9.2.17 Salt

Replace pixels in images with salt noise, i.e. white-ish pixels.

This augmenter is similar to `SaltAndPepper`, but adds no pepper noise to images.

API link: *Salt()*

Example. Replace `10%` of all pixels with salt noise (white-ish colors):

```
import imgaug.augmenters as iaa
aug = iaa.Salt(0.1)
```

Example. Similar to `SaltAndPepper`, this augmenter also supports the `per_channel` argument, which is skipped here for brevity.

9.2.18 CoarseSalt

Replace rectangular areas in images with white-ish pixel noise.

This augmenter is similar to `CoarseSaltAndPepper`, but adds no pepper noise to images.

API link: *CoarseSalt()*

Example. Mark `5%` of all pixels in a mask to be replaced by salt noise. The mask has `1%` to `10%` the size of the input image. The mask is then upscaled to the input image size, leading to large rectangular areas being marked as to be replaced. These areas are then replaced in the input image by salt noise.

```
import imgaug.augmenters as iaa
aug = iaa.CoarseSalt(0.05, size_percent=(0.01, 0.1))
```

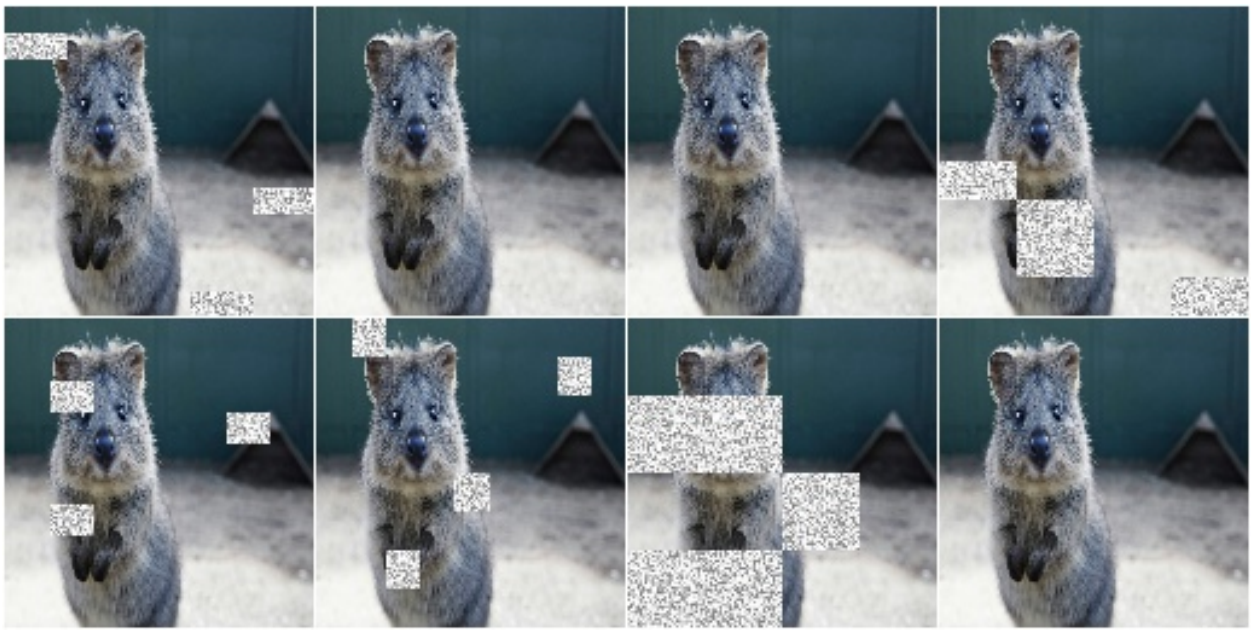

Similar to `CoarseSaltAndPepper`, this augmenter also supports the `per_channel` argument, which is skipped here for brevity

9.2.19 Pepper

Replace pixels in images with pepper noise, i.e. black-ish pixels.

This augmenter is similar to `SaltAndPepper`, but adds no salt noise to images.

This augmenter is similar to `Dropout`, but slower and the black pixels are not uniformly black.

API link: *Pepper()*

Example. Replace `10%` of all pixels with pepper noise (black-ish colors):

```
import imgaug.augmenters as iaa
aug = iaa.Pepper(0.1)
```

Similar to `SaltAndPepper`, this augmenter also supports the `per_channel` argument, which is skipped here for brevity.

9.2.20 CoarsePepper

Replace rectangular areas in images with black-ish pixel noise.

This augmenter is similar to `CoarseSaltAndPepper`, but adds no salt noise to images.

API link: *CoarsePepper()*

Example. Mark `5%` of all pixels in a mask to be replaced by pepper noise. The mask has `1%` to `10%` the size of the input image. The mask is then upscaled to the input image size, leading to large rectangular areas being marked as to be replaced. These areas are then replaced in the input image by pepper noise.

```
import imgaug.augmenters as iaa
aug = iaa.CoarsePepper(0.05, size_percent=(0.01, 0.1))
```

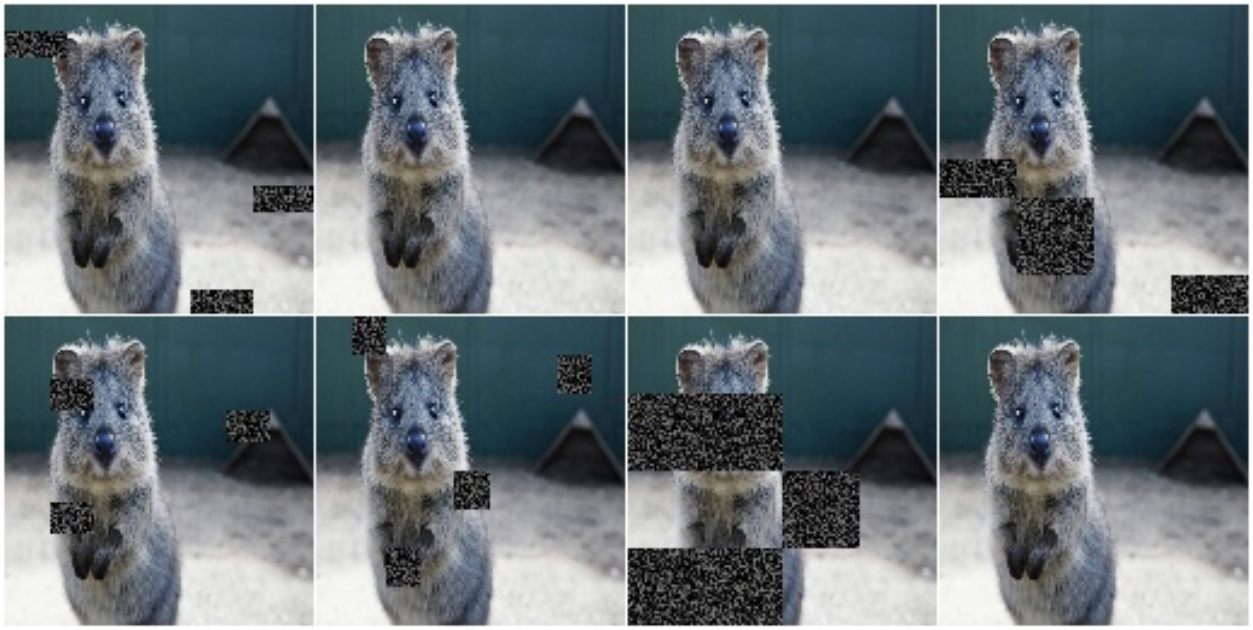

Similar to `CoarseSaltAndPepper`, this augmenter also supports the `per_channel` argument, which is skipped here for brevity

9.2.21 Invert

Augmenter that inverts all values in images, i.e. sets a pixel from value `v` to `255-v`.

API link: *Invert*

Example. Invert in 50% of all images all pixels:

```
import imgaug.augmenters as iaa
aug = iaa.Invert(0.5)
```

Example. For 50% of all images, invert all pixels in these images with 25% probability (per image). In the remaining 50% of all images, invert 25% of all channels:

```
aug = iaa.Invert(0.25, per_channel=0.5)
```

9.2.22 Solarize

Invert all values above a threshold in images.

This is the same as `Invert`, but sets a default threshold around `128` (+/- 64, decided per image) and default *invert_above_threshold* to `True` (i.e. only values above the threshold will be inverted).

API link: *Solarize*

Example. Invert the colors in `50` percent of all images for pixels with a value between `32` and `128` or more. The threshold is sampled once per image. The thresholding operation happens per channel.

```
import imgaug.augmenters as iaa
aug = iaa.Solarize(0.5, threshold=(32, 128))
```

9.2.23 JpegCompression

Degrade the quality of images by JPEG-compressing them.

API link: *JpegCompression*

Example. Remove high frequency components in images via JPEG compression with a *compression strength* between `80` and `95` (randomly and uniformly sampled per image). This corresponds to a (very low) *quality* setting of `5` to `20`.

```
import imgaug.augmenters as iaa
aug = iaa.JpegCompression(compression=(70, 99))
```

9.3 augmenters.artistic

9.3.1 Cartoon

Convert the style of images to a more cartoonish one.

This augmenter was primarily designed for images with a size of `200` to `800` pixels. Smaller or larger images may cause issues.

Note that the quality of the results can currently not compete with learned style transfer, let alone human-made images. A lack of detected edges or also too many detected edges are probably the most significant drawbacks.

API link: *`Cartoon`*

Example. Create an example image, then apply a cartoon filter to it:

```
import imgaug.augmenters as iaa
aug = iaa.Cartoon()
```

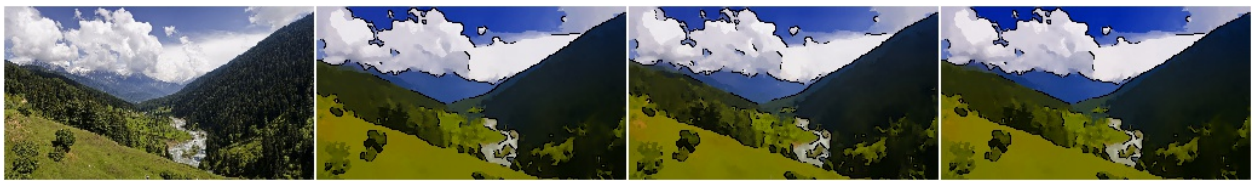

Example. Create a non-stochastic cartoon augmenter that produces decent-looking images:

```
aug = iaa.Cartoon(blur_ksize=3, segmentation_size=1.0,
                  saturation=2.0, edge_prevalence=1.0)
```

9.4 augmenters.blend

Note: It is not recommended to use blending augmenter with child augmenters that change the geometry of images (e.g. horizontal flips, affine transformations) if you *also* want to augment coordinates (e.g. keypoints, bounding boxes, polygons, ...), as it is not clear which of the two coordinate results (first or second branch) should be used as the

coordinates after augmentation. Currently, all blending augmenters try to use the augmented coordinates of the branch that makes up most of the augmented image.

9.4.1 BlendAlpha

Alpha-blend two image sources using an alpha/opacity value.

The two image sources can be imagined as branches. If a source is not given, it is automatically the same as the input. Let `FG` be the foreground branch and `BG` be the background branch. Then the result images are defined as `factor * FG + (1-factor) * BG`, where `factor` is an overlay factor.

API link: *BlendAlpha*

Example. Convert each image to pure grayscale and alpha-blend the result with the original image using an alpha of `50%`, thereby removing about `50%` of all color. This is equivalent to `iaa.Grayscale(0.5)`.

```
import imgaug.augmenters as iaa
aug = iaa.BlendAlpha(0.5, iaa.Grayscale(1.0))
```

Example. Same as in the previous example, but the alpha factor is sampled uniformly from the interval `[0.0, 1.0]` once per image, thereby removing a random fraction of all colors. This is equivalent to `iaa.Grayscale((0.0, 1.0))`.

```
aug = iaa.BlendAlpha((0.0, 1.0), iaa.Grayscale(1.0))
```

Example. First, rotate each image by a random degree sampled uniformly from the interval `[-20, 20]`. Then, alpha-blend that new image with the original one using a random factor sampled uniformly from the interval `[0.0, 1.0]`. For `50%` of all images, the blending happens channel-wise and the factor is sampled independently per channel (`per_channel=0.5`). As a result, e.g. the red channel may look visibly rotated (factor near `1.0`), while the green and blue channels may not look rotated (factors near `0.0`).

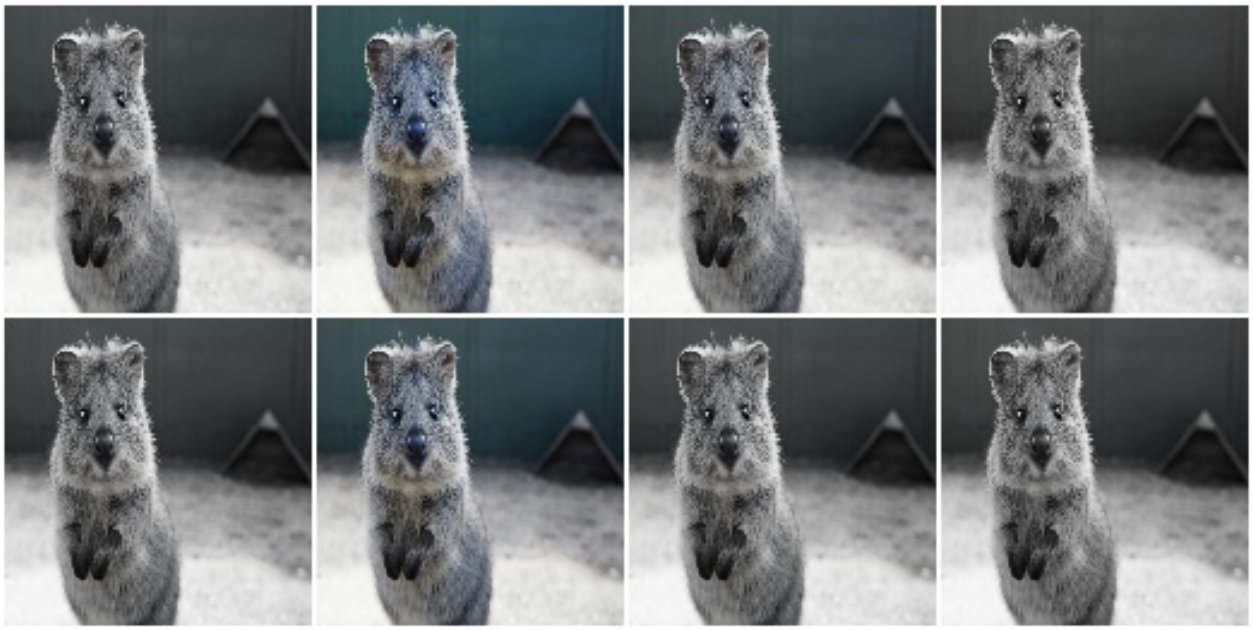

```
aug = iaa.BlendAlpha(
    (0.0, 1.0),
    iaa.Affine(rotate=(-20, 20)),
    per_channel=0.5)
```

Example. Apply two branches of augmenters – `A` and `B` – *independently* to input images and alpha-blend the results of these branches using a factor `f`. Branch `A` increases image pixel intensities by `100` and `B` multiplies the pixel intensities by `0.2`. `f` is sampled uniformly from the interval `[0.0, 1.0]` per image. The resulting images contain a bit of `A` and a bit of `B`.

```
aug = iaa.BlendAlpha(
    (0.0, 1.0),
    foreground=iaa.Add(100),
    background=iaa.Multiply(0.2))
```

Example. Apply median blur to each image and alpha-blend the result with the original image using an alpha factor of either exactly `0.25` or exactly `0.75` (sampled once per image).

```
aug = iaa.BlendAlpha([0.25, 0.75], iaa.MedianBlur(13))
```

9.4.2 BlendAlphaMask

Alpha-blend two image sources using non-binary masks generated per image.

This augmenter queries for each image a mask generator to generate a `(H,W)` or `(H,W,C)` channelwise mask `[0.0, 1.0]`, where `H` is the image height and `W` the width. The mask will then be used to alpha-blend pixel- and possibly channel-wise between a foreground branch of augmenters and a background branch. (Both branches default to the identity operation if not provided.)

See also *BlendAlpha*.

API link: *BlendAlphaMask*

Example. Create an augmenter that sometimes adds clouds at the bottom and sometimes at the top of the image:

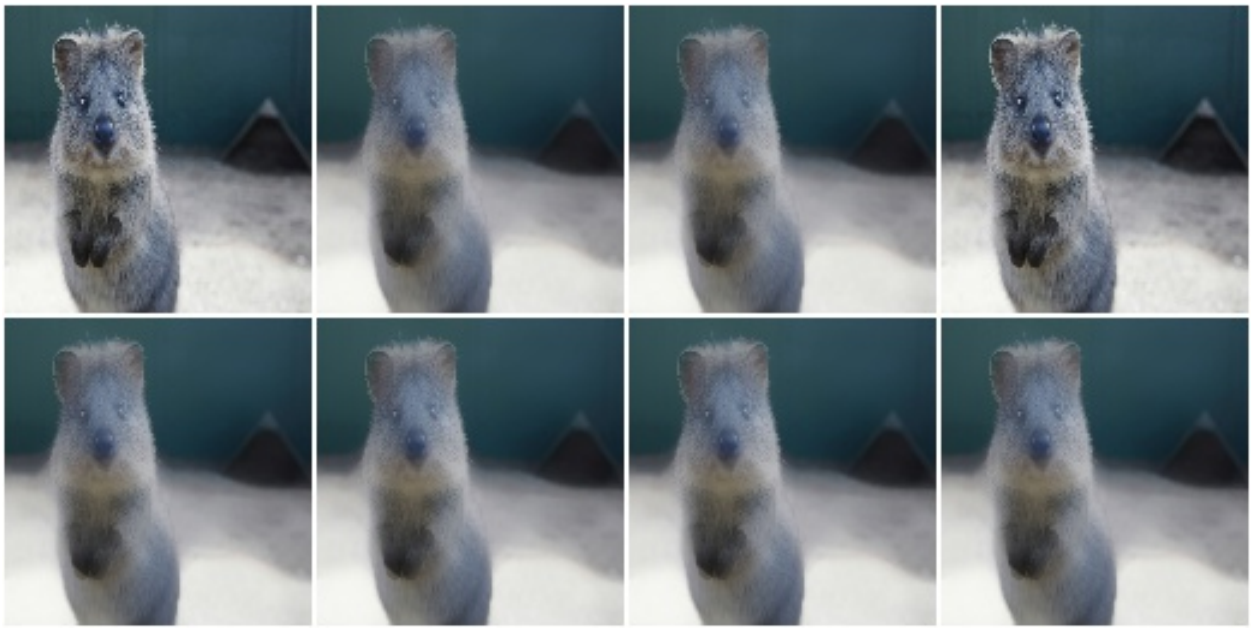

```
import imgaug.augmenters as iaa
aug = iaa.BlendAlphaMask(
    iaa.InvertMaskGen(0.5, iaa.VerticalLinearGradientMaskGen()),
    iaa.Clouds()
)
```

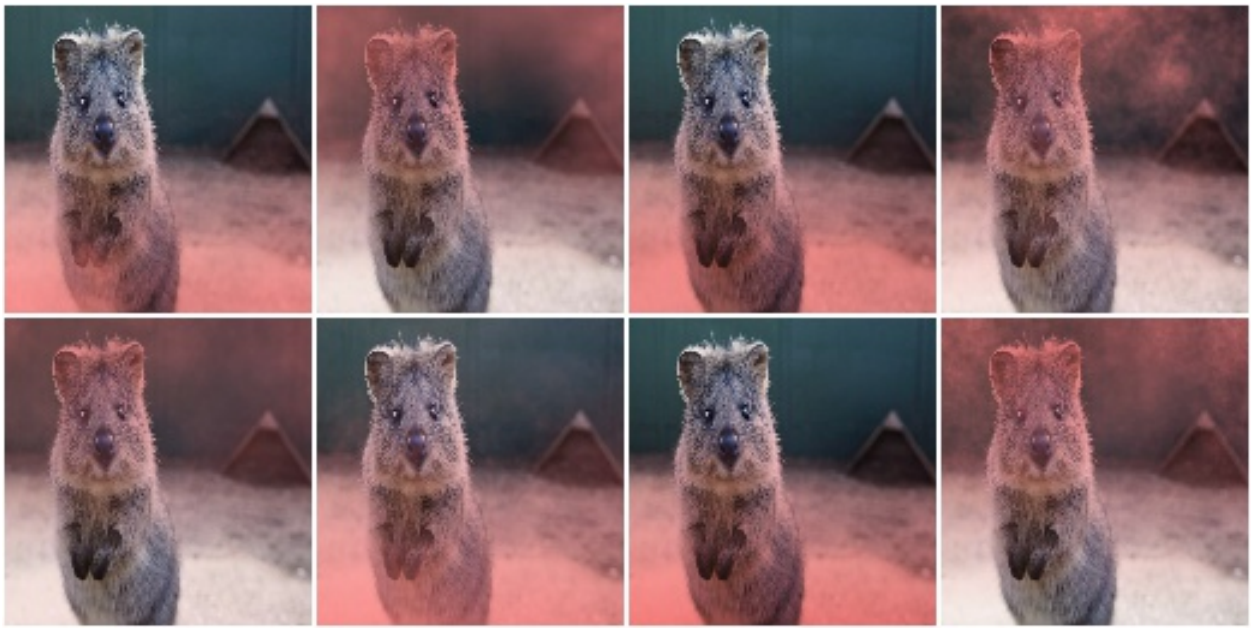

9.4.3 BlendAlphaElementwise

Alpha-blend two image sources using alpha/opacity values sampled per pixel.

This is the same as `BlendAlpha`, except that the opacity factor is sampled once per *pixel* instead of once per *image* (or a few times per image, if `BlendAlpha.per_channel` is set to `True`).

See `BlendAlpha` for more details.

This class is a wrapper around *`BlendAlphaMask`*.

API link: *`BlendAlphaElementwise`*

Example. Convert each image to pure grayscale and alpha-blend the result with the original image using an alpha of `50%` for all pixels, thereby removing about `50%` of all color. This is equivalent to `iaa.Grayscale(0.5)`. This is also equivalent to `iaa.Alpha(0.5, iaa.Grayscale(1.0))`, as the opacity has a fixed value of `0.5` and is hence identical for all pixels.

```
import imgaug.augmenters as iaa
aug = iaa.BlendAlphaElementwise(0.5, iaa.Grayscale(1.0))
```

Example. Same as in the previous example, but here with hue-shift instead of grayscaling and additionally the alpha factor is sampled uniformly from the interval `[0.0, 1.0]` once per pixel, thereby shifting the hue by a random fraction for each pixel.

```
aug = iaa.BlendAlphaElementwise((0, 1.0), iaa.AddToHue(100))
```

Example. First, rotate each image by a random degree sampled uniformly from the interval `[-20, 20]`. Then, alpha-blend that new image with the original one using a random factor sampled uniformly from the interval `[0.0,`

1.0] per pixel. For 50% of all images, the blending happens channel-wise and the factor is sampled independently per pixel *and* channel (per_channel=0.5). As a result, e.g. the red channel may look visibly rotated (factor near 1.0), while the green and blue channels may not look rotated (factors near 0.0).

```
aug = iaa.BlendAlphaElementwise(
    (0.0, 1.0),
    iaa.Affine(rotate=(-20, 20)),
    per_channel=0.5)
```

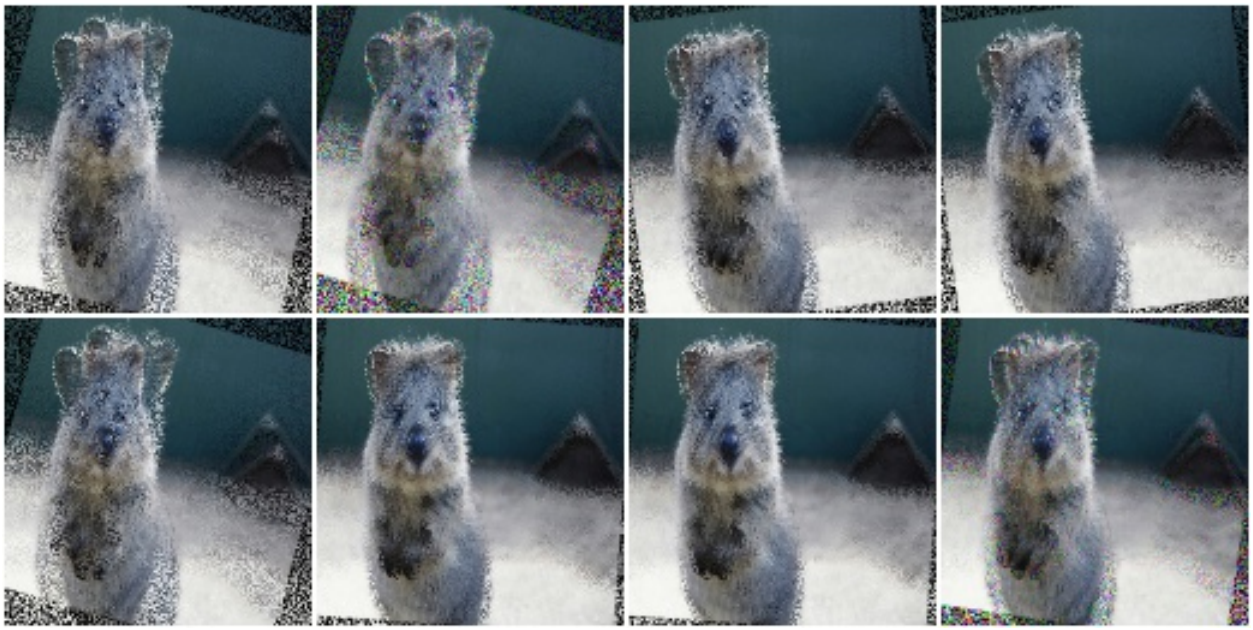

Example. Apply two branches of augmenters – A and B – *independently* to input images and alpha-blend the results of these branches using a factor f. Branch A increases image pixel intensities by 100 and B multiplies the pixel intensities by 0.2. f is sampled uniformly from the interval [0.0, 1.0] per pixel. The resulting images contain a bit of A and a bit of B.

```
aug = iaa.BlendAlphaElementwise(
    (0.0, 1.0),
    foreground=iaa.Add(100),
    background=iaa.Multiply(0.2))
```

Example. Apply median blur to each image and alpha-blend the result with the original image using an alpha factor of either exactly 0.25 or exactly 0.75 (sampled once per pixel).

```
aug = iaa.BlendAlphaElementwise([0.25, 0.75], iaa.MedianBlur(13))
```

9.4.4 BlendAlphaSimplexNoise

Alpha-blend two image sources using simplex noise alpha masks.

The alpha masks are sampled using a simplex noise method, roughly creating connected blobs of 1s surrounded by 0s. If nearest neighbour upsampling is used, these blobs can be rectangular with sharp edges.

API link: *BlendAlphaSimplexNoise*

Example. Detect per image all edges, mark them in a black and white image and then alpha-blend the result with the original image using simplex noise masks.

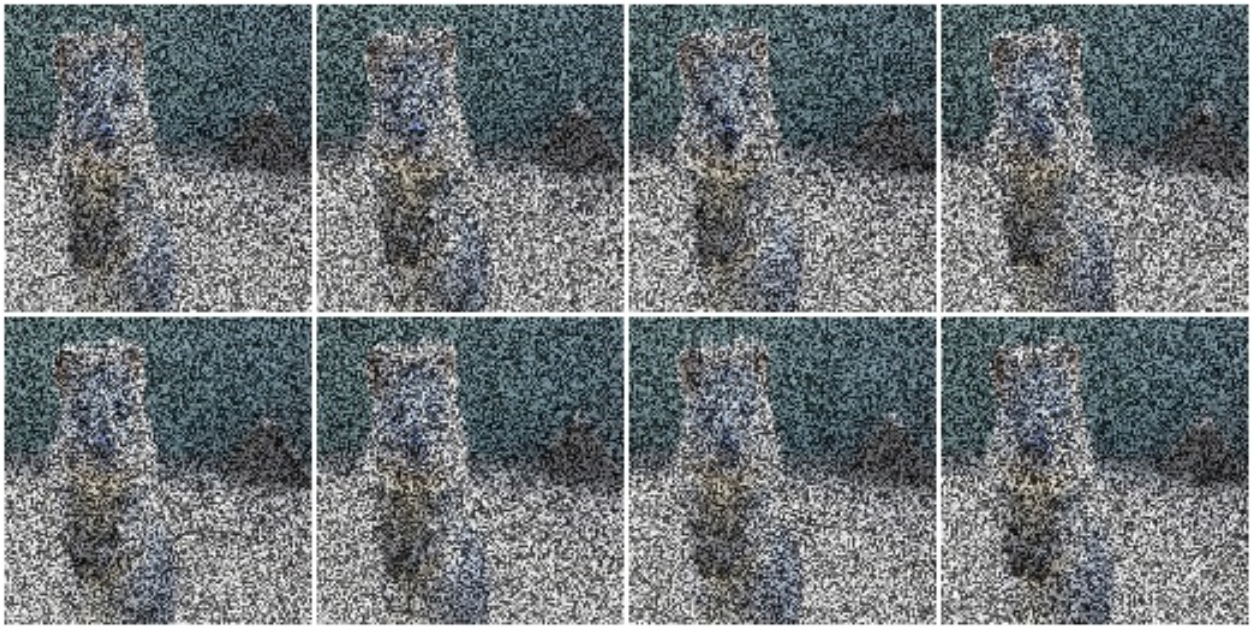

```
import imgaug.augmenters as iaa
aug = iaa.BlendAlphaSimplexNoise(iaa.EdgeDetect(1.0))
```

Example. Same as in the previous example, but using only nearest neighbour upscaling to scale the simplex noise masks to the final image sizes, i.e. no nearest linear upsampling is used. This leads to rectangles with sharp edges.

```
aug = iaa.BlendAlphaSimplexNoise(
    iaa.EdgeDetect(1.0),
    upscale_method="nearest")
```

Example. Same as in the previous example, but using only linear upscaling to scale the simplex noise masks to the final image sizes, i.e. no nearest neighbour upsampling is used. This leads to rectangles with smooth edges.

```
aug = iaa.BlendAlphaSimplexNoise(
    iaa.EdgeDetect(1.0),
    upscale_method="linear")
```

Example. Same as in the first example, but using a threshold for the sigmoid function that is further to the right. This is more conservative, i.e. the generated noise masks will be mostly black (values around `0.0`), which means that most of the original images (parameter/branch *second*) will be kept, rather than using the results of the augmentation (parameter/branch *first*).

```
import imgaug.parameters as iap
aug = iaa.BlendAlphaSimplexNoise(
    iaa.EdgeDetect(1.0),
    sigmoid_thresh=iap.Normal(10.0, 5.0))
```

9.4.5 BlendAlphaFrequencyNoise

Alpha-blend two image sources using frequency noise masks.

The alpha masks are sampled using frequency noise of varying scales, which can sometimes create large connected blobs of `1` s surrounded by `0` s and other times results in smaller patterns. If nearest neighbour upsampling is used, these blobs can be rectangular with sharp edges.

API link: *BlendAlphaFrequencyNoise*

Example. Detect per image all edges, mark them in a black and white image and then alpha-blend the result with the original image using frequency noise masks.

```
import imgaug.augmenters as iaa
aug = iaa.BlendAlphaFrequencyNoise(first=iaa.EdgeDetect(1.0))
```

Example. Same as the first example, but using only linear upscaling to scale the frequency noise masks to the final image sizes, i.e. no nearest neighbour upsampling is used. This results in smooth edges.

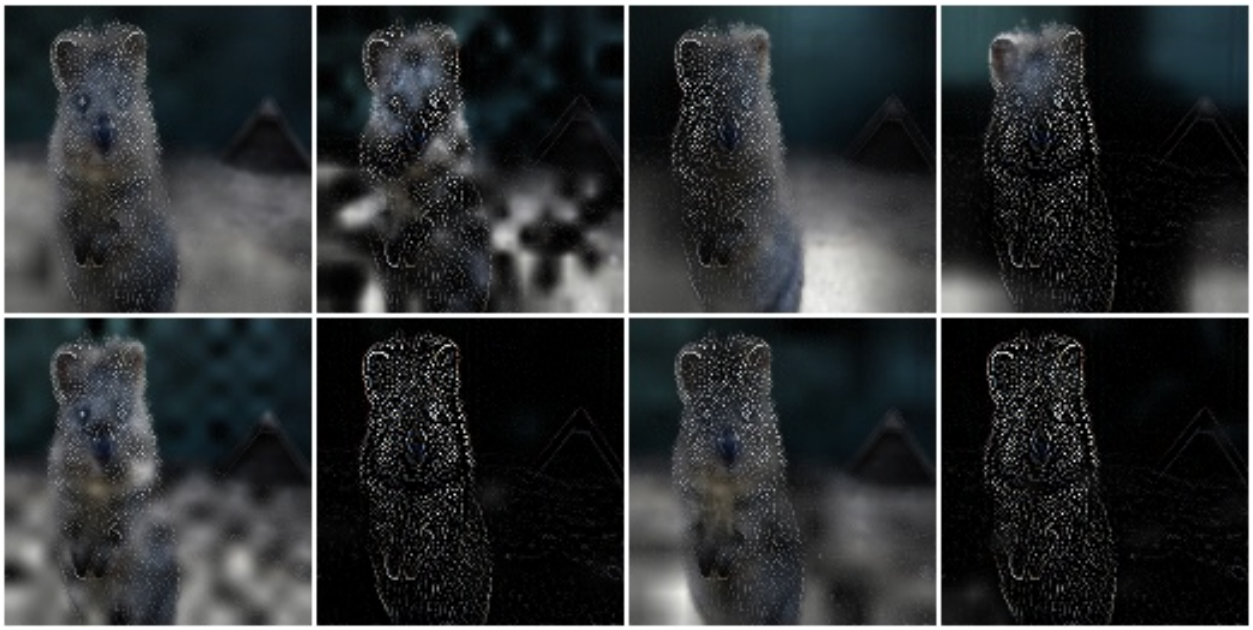

```
aug = iaa.BlendAlphaFrequencyNoise(
    first=iaa.EdgeDetect(1.0),
    upscale_method="nearest")
```

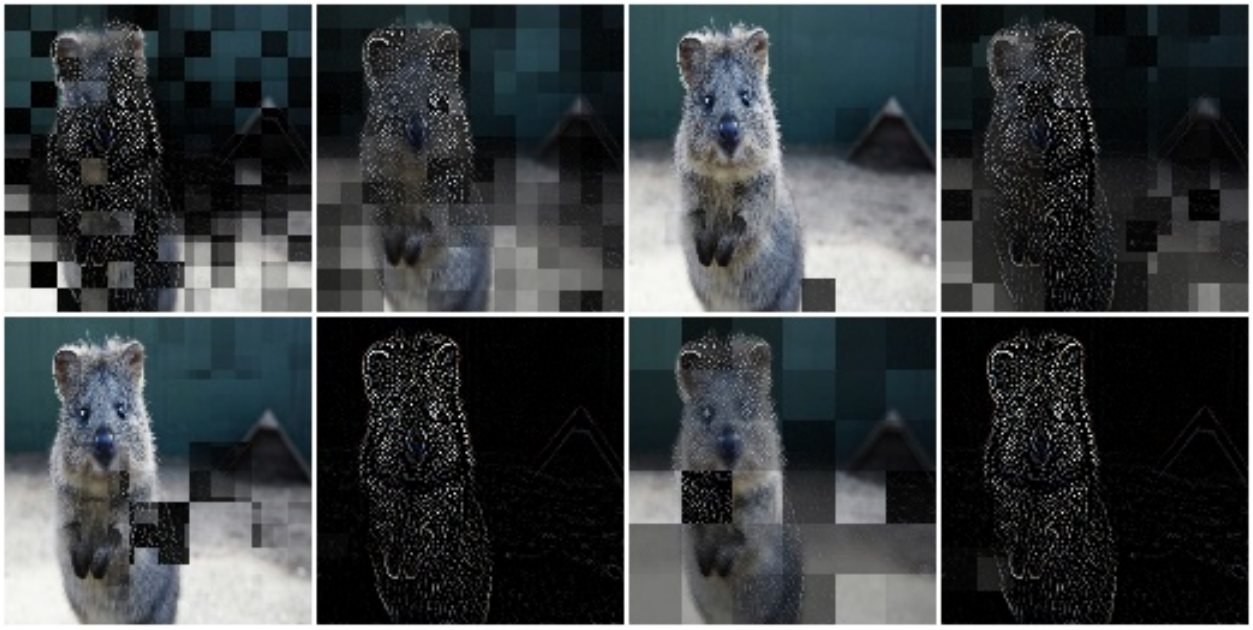

Example. Same as the first example, but using only linear upscaling to scale the frequency noise masks to the final image sizes, i.e. no nearest neighbour upsampling is used. This results in smooth edges.

```
aug = iaa.BlendAlphaFrequencyNoise(
    first=iaa.EdgeDetect(1.0),
    upscale_method="linear")
```

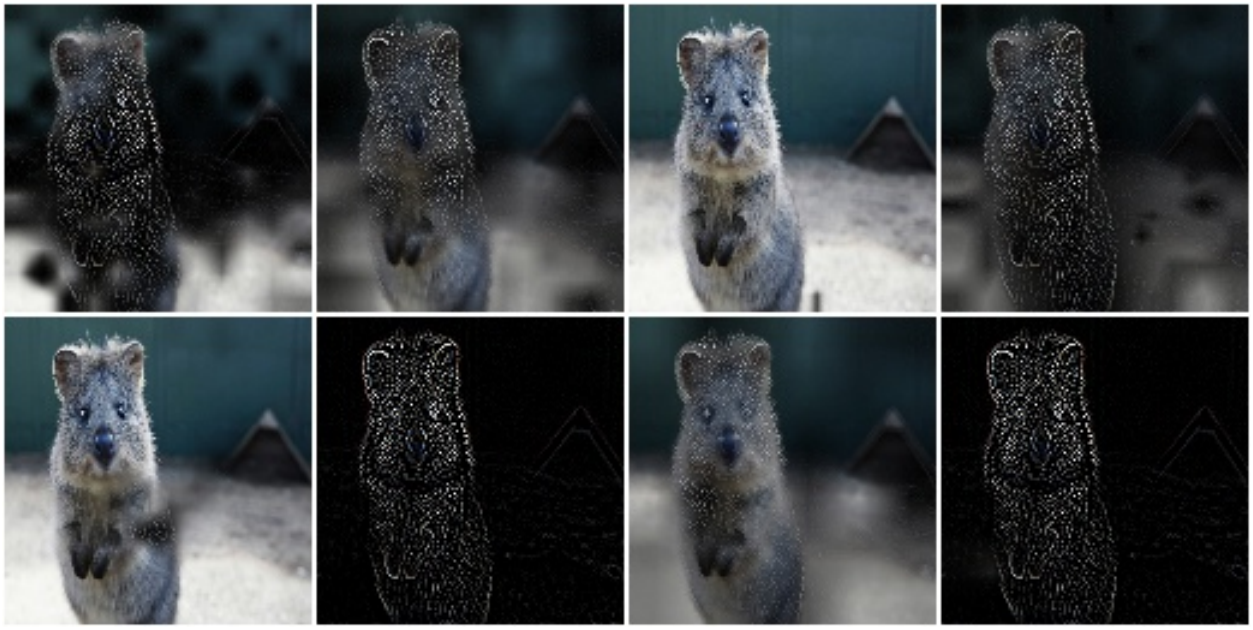

Example. Same as in the previous example, but with the exponent set to a constant -2 and the sigmoid deactivated, resulting in cloud-like patterns without sharp edges.

```
aug = iaa.BlendAlphaFrequencyNoise(
    first=iaa.EdgeDetect(1.0),
    upscale_method="linear",
    exponent=-2,
    sigmoid=False)
```

Example. Same as the first example, but using a threshold for the sigmoid function that is further to the right. This is more conservative, i.e. the generated noise masks will be mostly black (values around `0.0`), which means that most of the original images (parameter/branch *second*) will be kept, rather than using the results of the augmentation (parameter/branch *first*).

```
import imgaug.parameters as iap
aug = iaa.BlendAlphaFrequencyNoise(
    first=iaa.EdgeDetect(1.0),
    sigmoid_thresh=iap.Normal(10.0, 5.0))
```

9.4.6 BlendAlphaSomeColors

Blend images from two branches using colorwise masks.

This class generates masks that "mark" a few colors and replace the pixels within these colors with the results of the foreground branch. The remaining pixels are replaced with the results of the background branch (usually the identity function). That allows to e.g. selectively grayscale a few colors, while keeping other colors unchanged.

This class is a thin wrapper around *BlendAlphaMask* together with *SomeColorsMaskGen*.

API link: *BlendAlphaSomeColors*

Example. Create an augmenter that turns randomly removes some colors in images by grayscaling them:

```
import imgaug.augmenters as iaa
aug = iaa.BlendAlphaSomeColors(iaa.Grayscale(1.0))
```

Example. Create an augmenter that removes some colors in images by replacing them with black pixels:

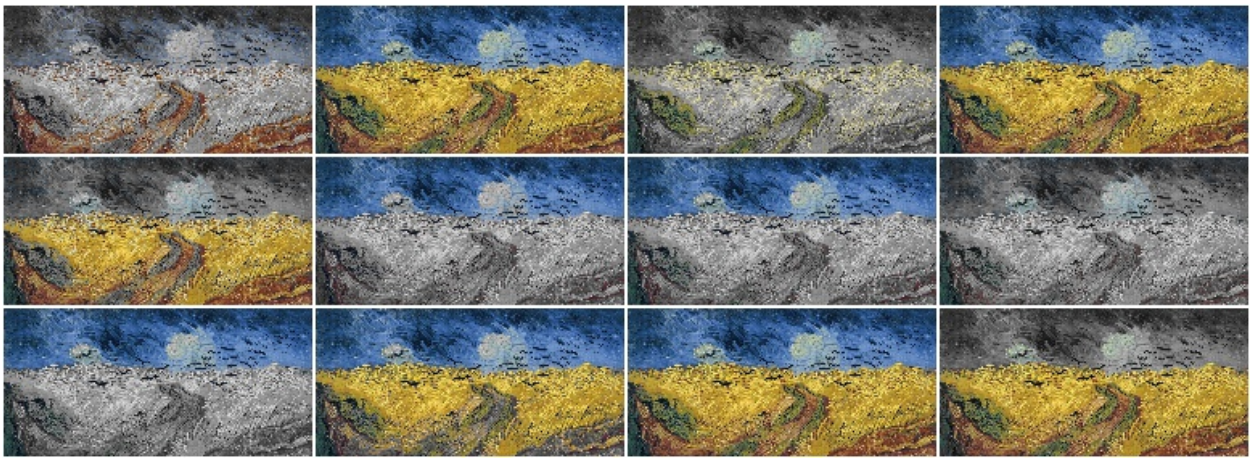

```
aug = iaa.BlendAlphaSomeColors(iaa.TotalDropout(1.0))
```

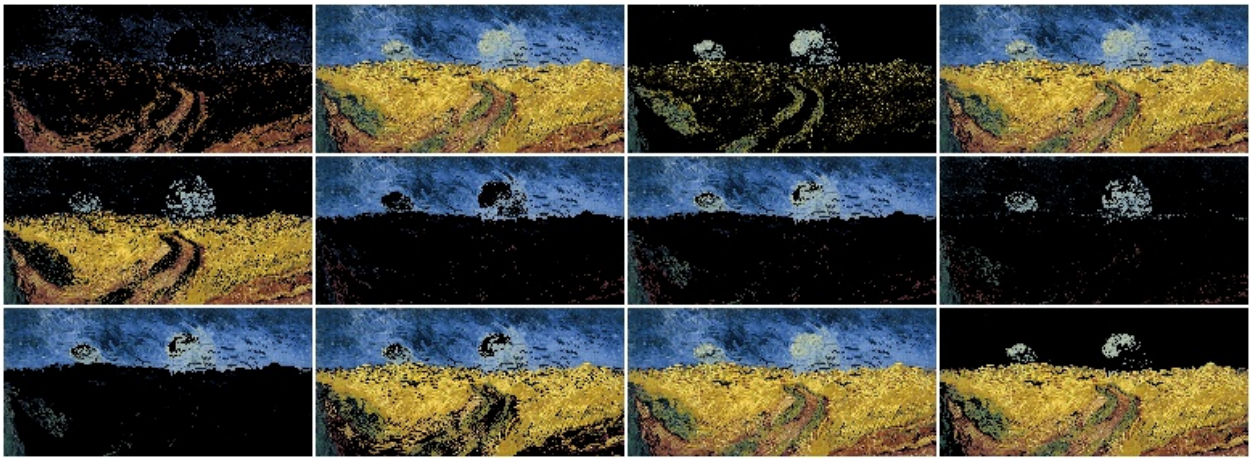

Example. Create an augmenter that desaturates some colors and increases the saturation of the remaining ones:

```
aug = iaa.BlendAlphaSomeColors(
    iaa.MultiplySaturation(0.5), iaa.MultiplySaturation(1.5))
```

Example. Create an augmenter that applies average pooling to some colors. Each color tune is either selected (alpha of `1.0`) or not selected (`0.0`). There is no gradual change between similar colors.

```
aug = iaa.BlendAlphaSomeColors(
    iaa.AveragePooling(7), alpha=[0.0, 1.0], smoothness=0.0)
```

Example. Create an augmenter that applies average pooling to some colors. Choose on average half of all colors in images for the blending operation.

```
aug = iaa.BlendAlphaSomeColors(
    iaa.AveragePooling(7), nb_bins=2, smoothness=0.0)
```

Example. Create an augmenter that applies average pooling to some colors with input images being in BGR colorspace

```
aug = iaa.BlendAlphaSomeColors(
    iaa.AveragePooling(7), from_colorspace="BGR")
```

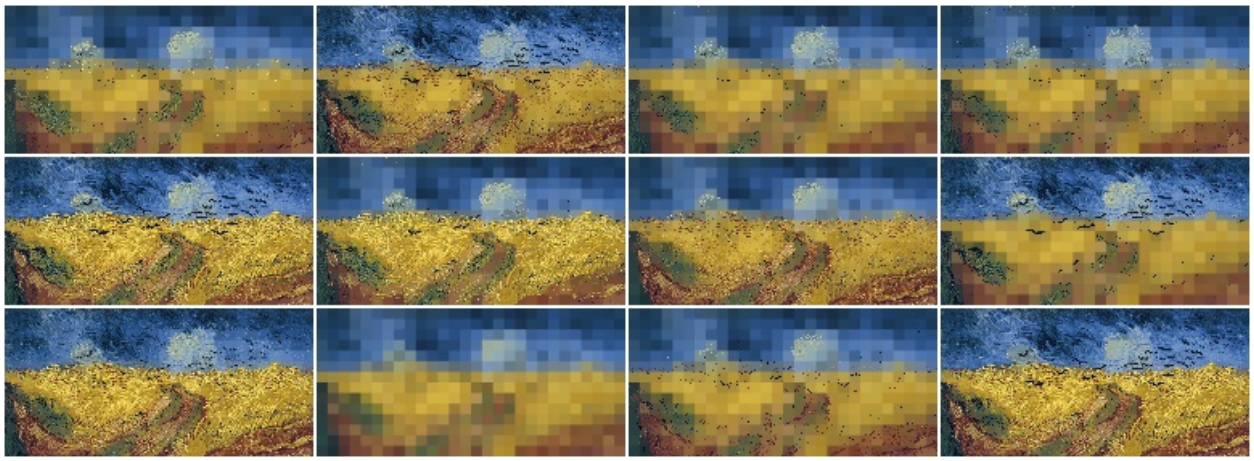

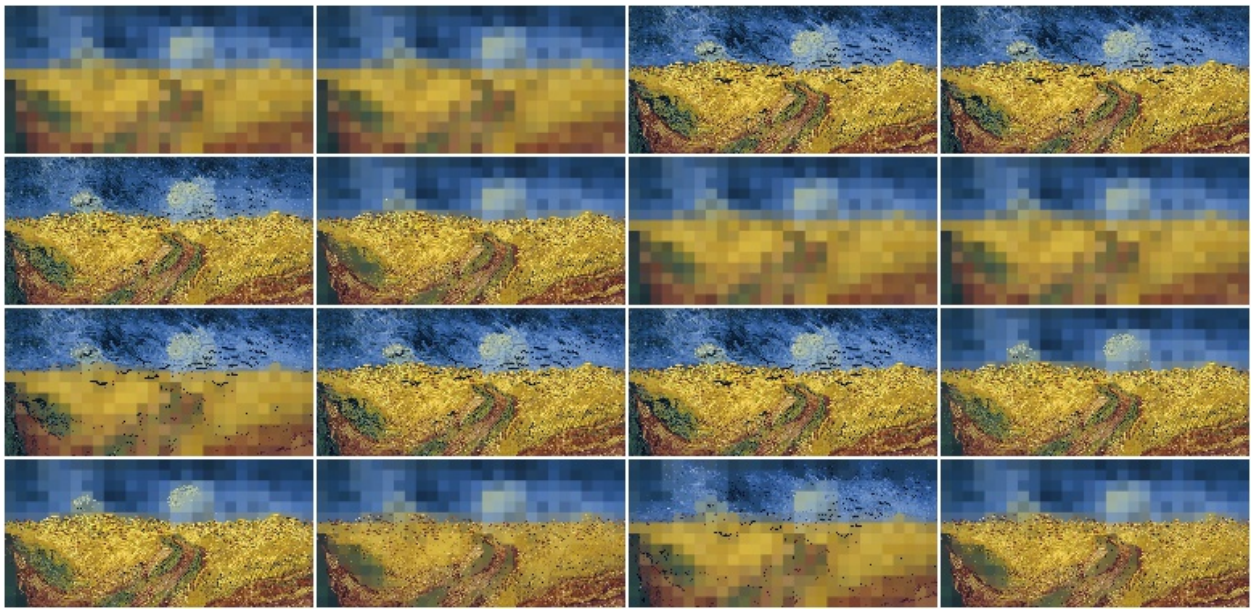

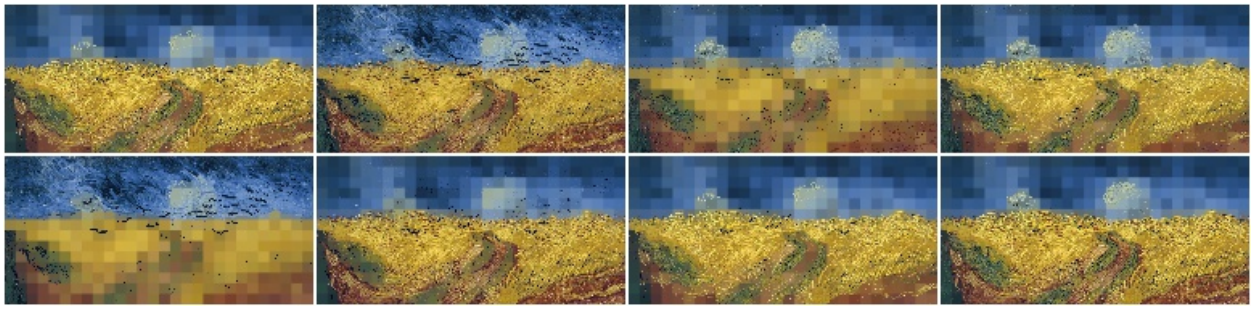

9.4.7 BlendAlphaHorizontalLinearGradient

Blend images from two branches along a horizontal linear gradient.

This class generates a horizontal linear gradient mask (i.e. usually a mask with low values on the left and high values on the right) and alphas-blends between foreground and background branch using that mask.

This class is a thin wrapper around *BlendAlphaMask* together with *HorizontalLinearGradientMaskGen*.

API link: *BlendAlphaHorizontalLinearGradient*

Example. Create an augmenter that removes more color towards the right of the image:

```
import imgaug.augmenters as iaa
aug = iaa.BlendAlphaHorizontalLinearGradient(iaa.AddToHue((-100, 100)))
```

Example. Create an augmenter that replaces pixels towards the right with darker and darker values. However it always keeps at least 20% (`1.0 - max_value`) of the original pixel value on the far right and always replaces at least 20% on the far left (`min_value=0.2`).

```
aug = iaa.BlendAlphaHorizontalLinearGradient(
    iaa.TotalDropout(1.0),
    min_value=0.2, max_value=0.8)
```

Example. Create an augmenter that blends with an average-pooled image according to a horizontal gradient that starts at a random x-coordinate and reaches its maximum at another random x-coordinate. Due to that randomness, the gradient may increase towards the left or right.

```
aug = iaa.BlendAlphaHorizontalLinearGradient(
    iaa.AveragePooling(11),
    start_at=(0.0, 1.0), end_at=(0.0, 1.0))
```

9.4.8 BlendAlphaVerticalLinearGradient

Blend images from two branches along a vertical linear gradient.

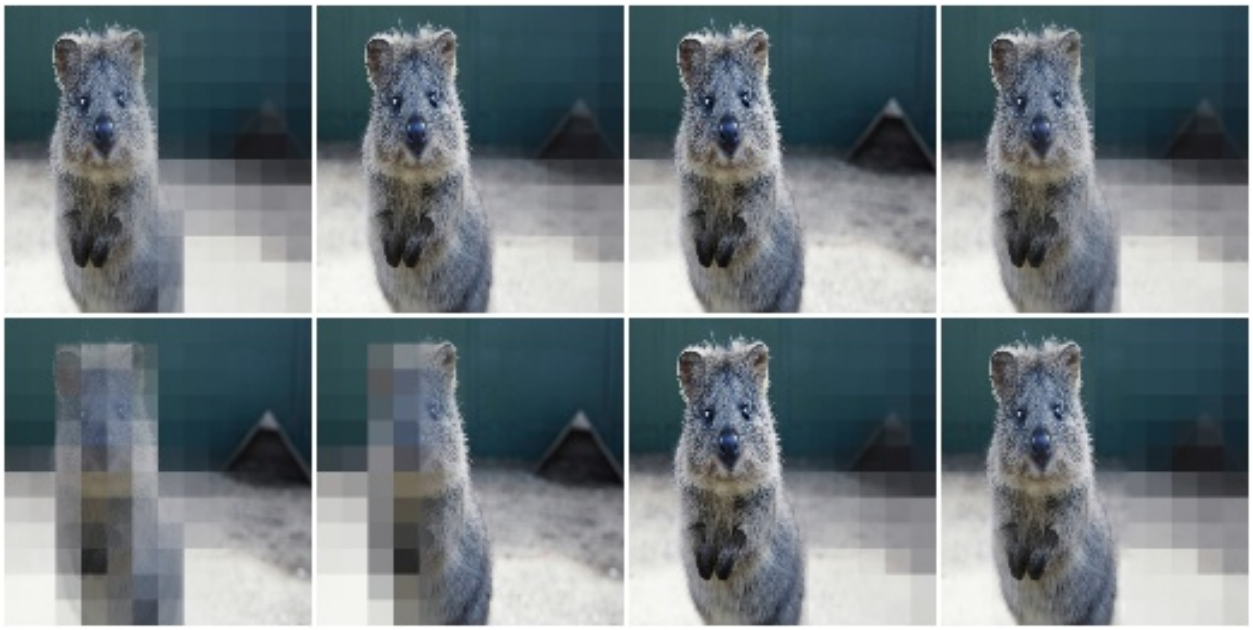

This class generates a vertical linear gradient mask (i.e. usually a mask with low values on the left and high values on the right) and alphas-blends between foreground and background branch using that mask.

This class is a thin wrapper around *BlendAlphaMask* together with *VerticalLinearGradientMaskGen*.

API link: *BlendAlphaVerticalLinearGradient*

Example. Create an augmenter that removes more color towards the bottom of the image:

```
import imgaug.augmenters as iaa
aug = iaa.BlendAlphaVerticalLinearGradient(iaa.AddToHue((-100, 100)))
```

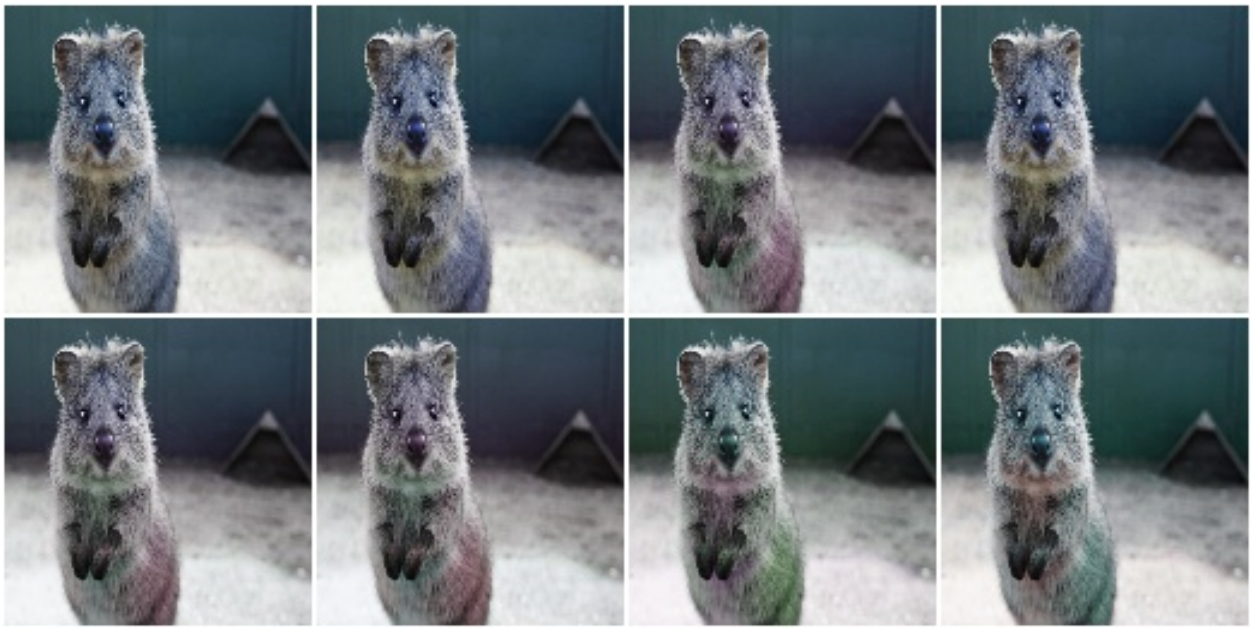

Example. Create an augmenter that replaces pixels towards the bottom with darker and darker values. However it always keeps at least 20% (`1.0 - max_value`) of the original pixel value on the far bottom and always replaces at least 20% on the far top (`min_value=0.2`).

```
aug = iaa.BlendAlphaVerticalLinearGradient(
    iaa.TotalDropout(1.0),
    min_value=0.2, max_value=0.8)
```

Example. Create an augmenter that blends with an average-pooled image according to a vertical gradient that starts at a random y-coordinate and reaches its maximum at another random y-coordinate. Due to that randomness, the gradient may increase towards the bottom or top.

```
aug = iaa.BlendAlphaVerticalLinearGradient(
    iaa.AveragePooling(11),
    start_at=(0.0, 1.0), end_at=(0.0, 1.0))
```

Example. Create an augmenter that draws clouds in roughly the top quarter of the image:

```
aug = iaa.BlendAlphaVerticalLinearGradient(
    iaa.Clouds(),
    start_at=(0.15, 0.35), end_at=0.0)
```

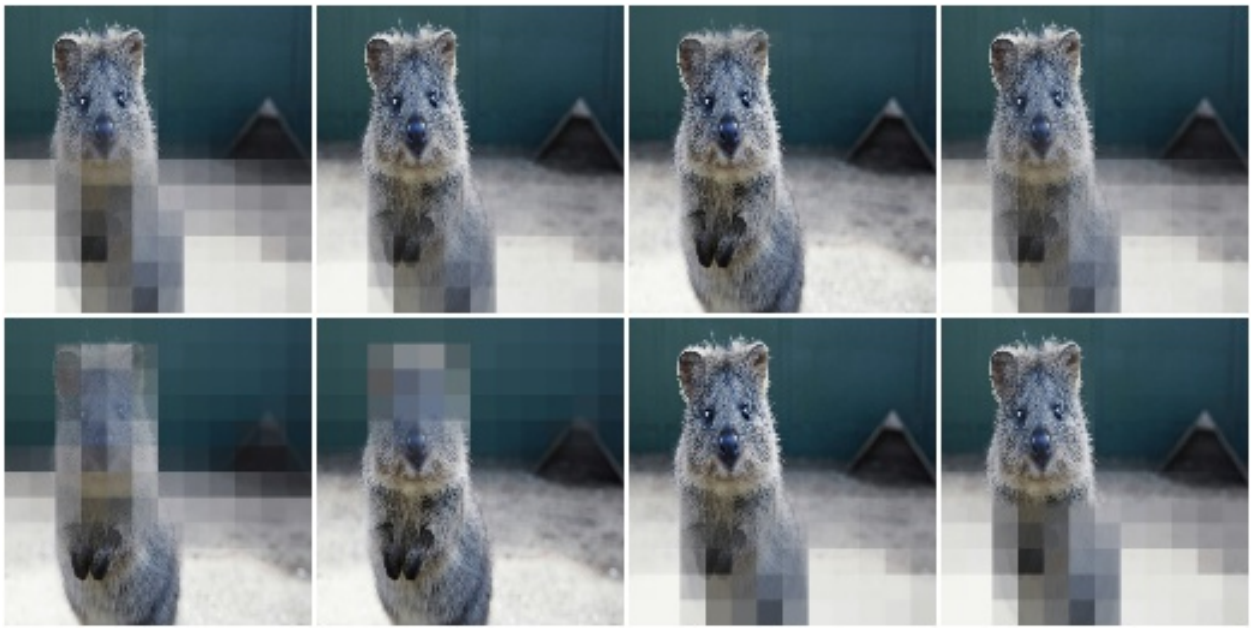

9.4.9 BlendAlphaRegularGrid

Blend images from two branches according to a regular grid.

This class generates for each image a mask that splits the image into a grid-like pattern of `H` rows and `W` columns. Each cell is then filled with an alpha value, sampled randomly per cell.

The difference to `AlphaBlendCheckerboard` is that this class samples random alpha values per grid cell, while in the checkerboard the alpha values follow a fixed pattern.

This class is a thin wrapper around *`BlendAlphaMask`* together with *`RegularGridMaskGen`*.

API link: *`BlendAlphaRegularGrid`*

Example. Create an augmenter that places a `HxW` grid on each image, where `H` (rows) is randomly and uniformly sampled from the interval `[4, 6]` and `W` is analogously sampled from the interval `[1, 4]`. Roughly half of the cells in the grid are filled with `0.0`, the remaining ones are unaltered. Which cells exactly are "dropped" is randomly decided per image. The resulting effect is similar to *`CoarseDropout`*.

```
import imgaug.augmenters as iaa
aug = iaa.BlendAlphaRegularGrid(nb_rows=(4, 6), nb_cols=(1, 4),
                                foreground=iaa.Multiply(0.0))
```

Example. Create an augmenter that always placed `2x2` cells on each image and sets about `1/3` of them to zero (foreground branch) and the remaining `2/3` to a pixelated version (background branch).

```
aug = iaa.BlendAlphaRegularGrid(nb_rows=2, nb_cols=2,
                                foreground=iaa.Multiply(0.0),
                                background=iaa.AveragePooling(8),
                                alpha=[0.0, 0.0, 1.0])
```

9.4.10 BlendAlphaCheckerboard

Blend images from two branches according to a checkerboard pattern.

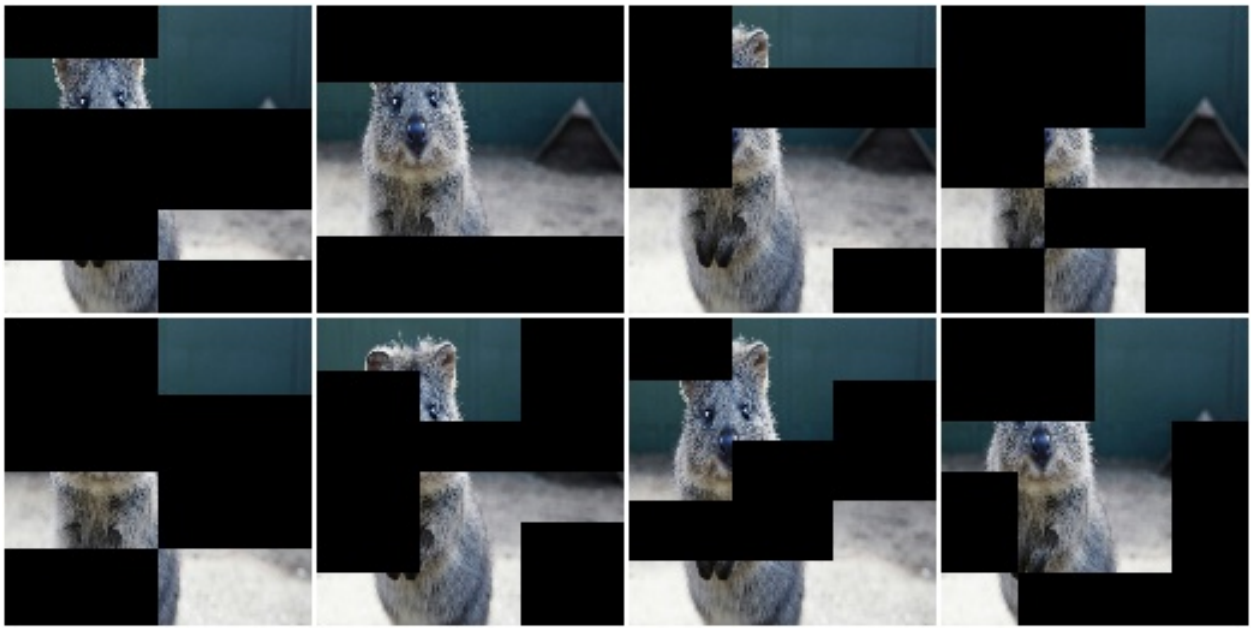

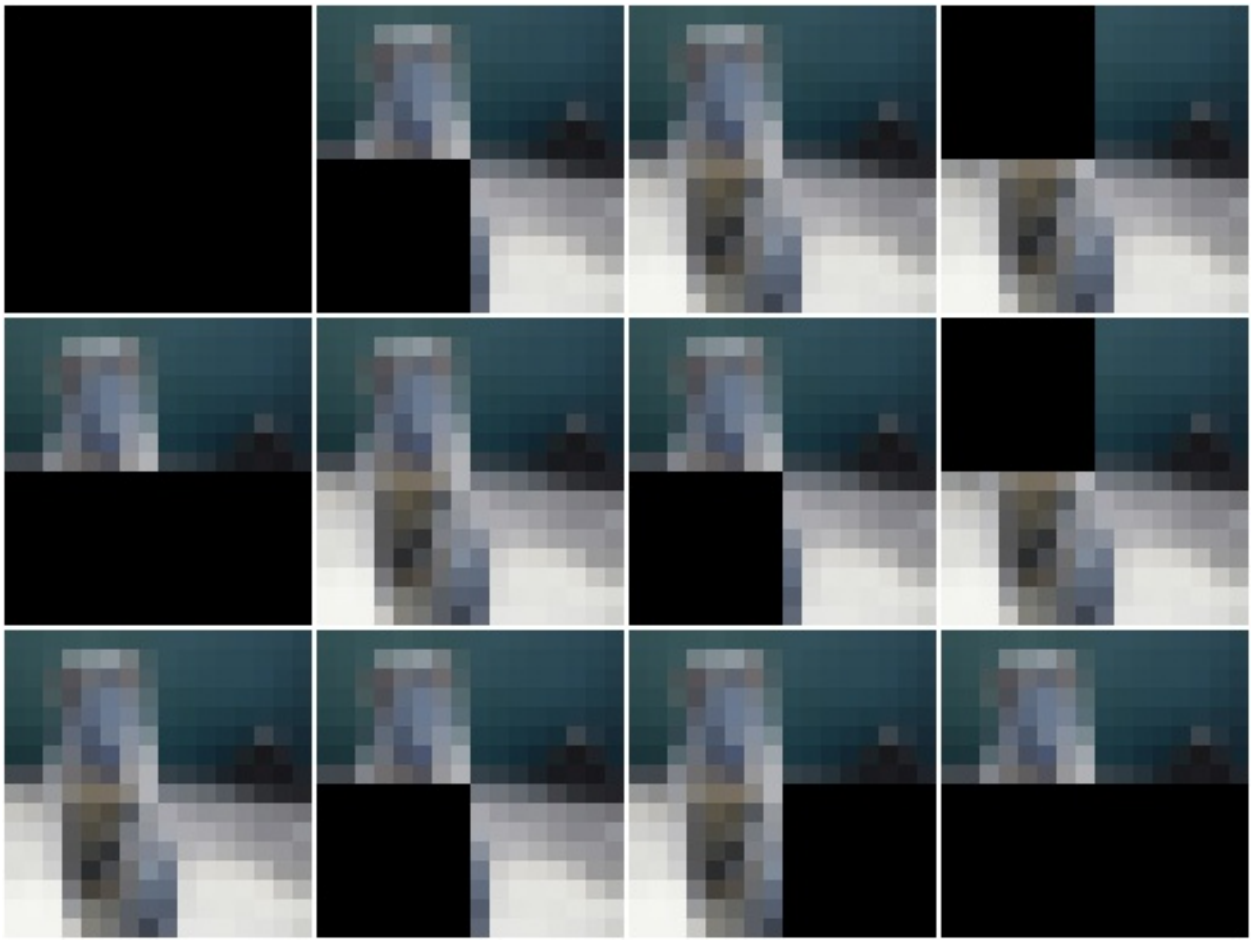

This class generates for each image a mask following a checkboard layout of `H` rows and `W` columns. Each cell is then filled with either `1.0` or `0.0`. The cell at the top-left is always `1.0`. Its right and bottom neighbour cells are `0.0`. The 4-neighbours of any cell always have a value opposite to the cell's value (`0.0` vs. `1.0`).

This class is a thin wrapper around *BlendAlphaMask* together with *CheckerboardMaskGen*.

API link: *BlendAlphaCheckerboard*

Example. Create an augmenter that places a `HxW` grid on each image, where `H` (rows) is always `2` and `W` is randomly and uniformly sampled from the interval `[1, 4]`. Half of the cells in the grid are grayscaled, the other half is unaltered.

```
import imgaug.augmenters as iaa
aug = iaa.BlendAlphaCheckerboard(nb_rows=2, nb_cols=(1, 4),
                                 foreground=iaa.AddToHue((-100, 100)))
```

9.4.11 BlendAlphaSegMapClassIds

Blend images from two branches based on segmentation map ids.

This class generates masks that are `1.0` at pixel locations covered by specific classes in segmentation maps.

This class is a thin wrapper around *BlendAlphaMask* together with *SegMapClassIdsMaskGen*.

Note: Segmentation maps can have multiple channels. If that is the case then for each position `(x, y)` it is sufficient that any class id in any channel matches one of the desired class ids.

Note: This class will produce an `AssertionError` if there are no segmentation maps in a batch.

API link: *BlendAlphaSegMapClassIds*

Example. Create an augmenter that removes color wherever the segmentation maps contain the classes `1` or `3`:

```
import imgaug.augmenters as iaa
aug = iaa.BlendAlphaSegMapClassIds(
    [1, 3],
    foreground=iaa.AddToHue((-100, 100)))
```

Example. Create an augmenter that randomly picks `2` classes from the list `[1, 2, 3, 4]` and blurs the image content wherever these classes appear in the segmentation map. Note that as the sampling of class ids happens *with replacement*, it is not guaranteed to sample two *unique* class ids.

```
aug = iaa.BlendAlphaSegMapClassIds(
    [1, 2, 3, 4],
    nb_sample_classes=2,
    foreground=iaa.GaussianBlur(3.0))
```

Example. Create an augmenter that zeros for roughly every fifth image all image pixels that do *not* belong to class id `2` (note that the *background* branch was used, not the *foreground* branch). Example use case: Human body landmark detection where both the landmarks/keypoints and the body segmentation map are known. Train the model to detect landmarks and sometimes remove all non-body information to force the model to become more independent of the background.

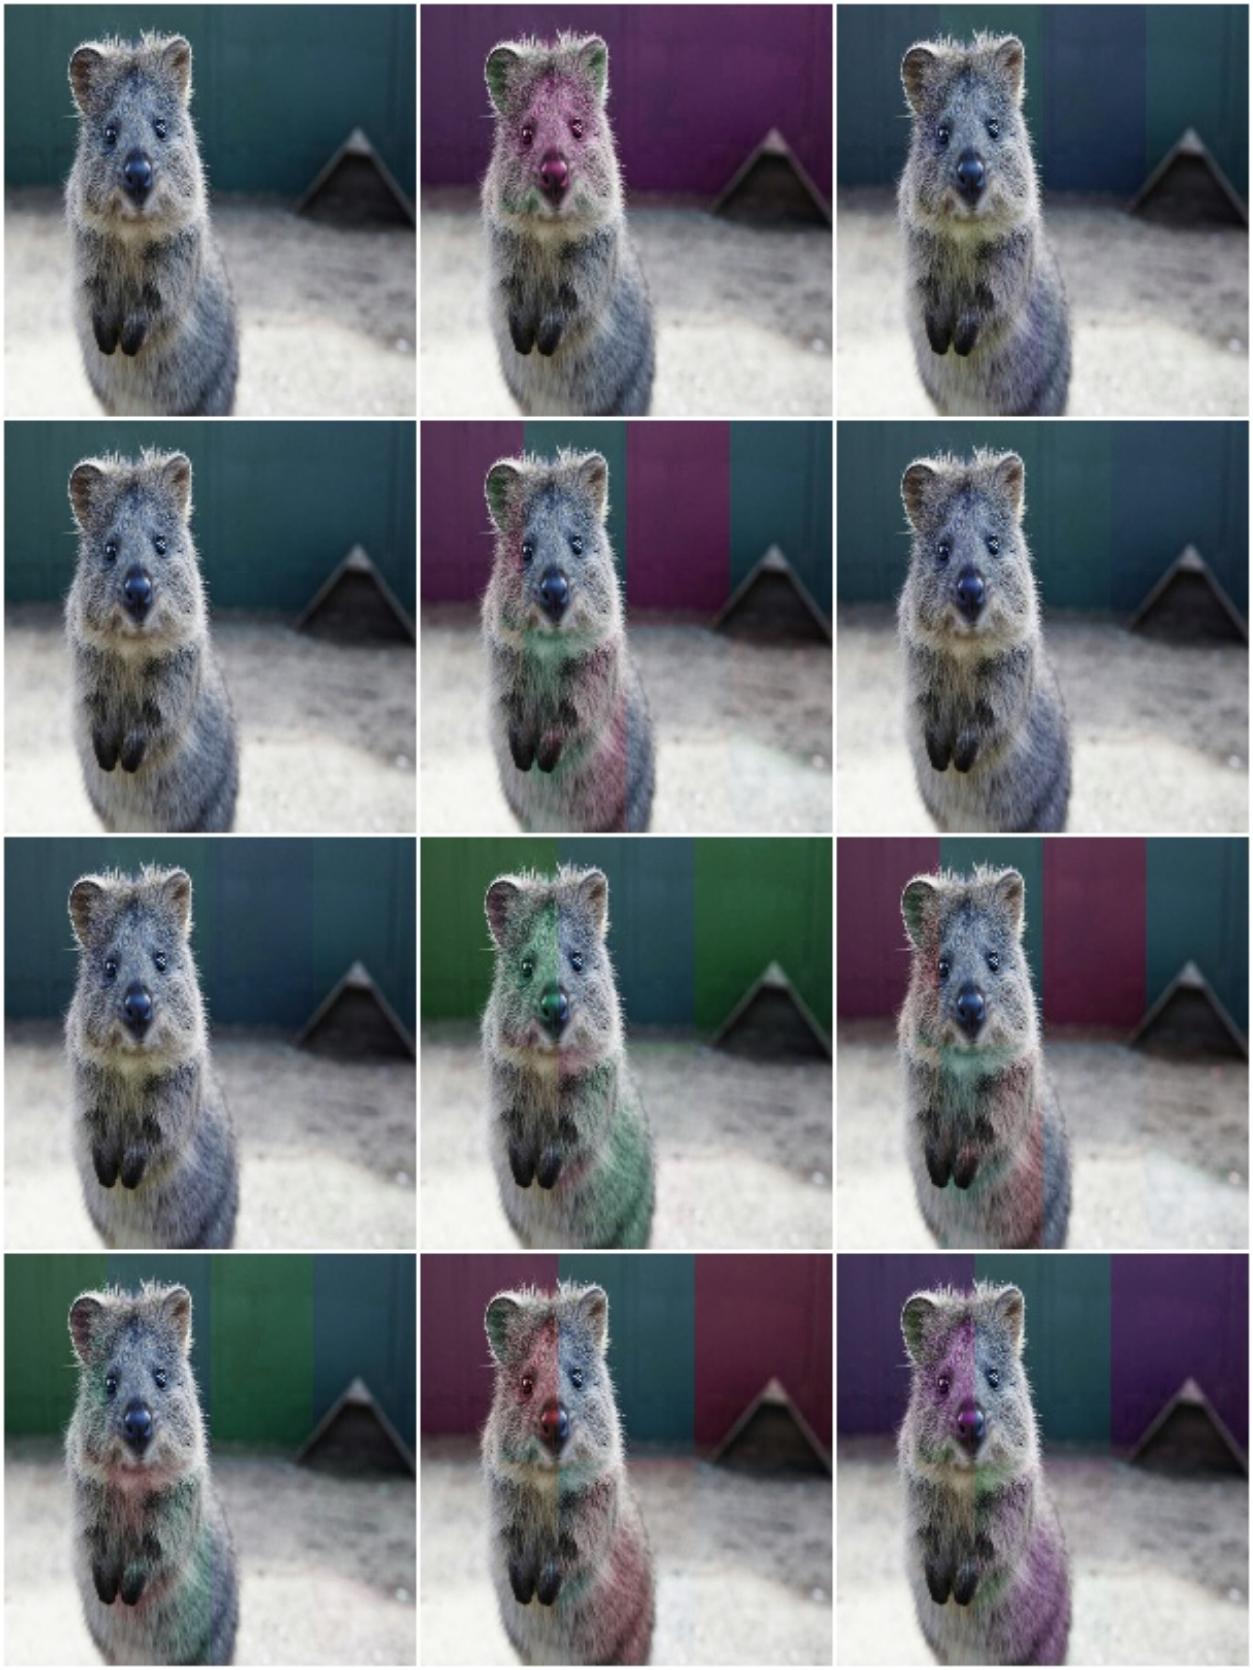

```
aug = iaa.Sometimes(0.2,
    iaa.BlendAlphaSegMapClassIds(
        2,
        background=iaa.TotalDropout(1.0)))
```

9.4.12 BlendAlphaBoundingBoxes

Blend images from two branches based on areas enclosed in bounding boxes.

This class generates masks that are `1.0` within bounding boxes of given labels. A mask pixel will be set to `1.0` if *at least* one bounding box covers the area and has one of the requested labels.

This class is a thin wrapper around *BlendAlphaMask* together with *BoundingBoxesMaskGen*.

Note: Avoid using augmenters as children that affect pixel locations (e.g. horizontal flips). See *BlendAlphaMask* for details.

Note: This class will produce an `AssertionError` if there are no bounding boxes in a batch.

API link: *BlendAlphaBoundingBoxes*

Example. Create an augmenter that removes color within bounding boxes having the label `person`:

```
import imgaug.augmenters as iaa
aug = iaa.BlendAlphaBoundingBoxes("person",
                                  foreground=iaa.Grayscale(1.0))
```

Example. Create an augmenter that randomizes the hue within bounding boxes that have the label `person` or `car`:

```
aug = iaa.BlendAlphaBoundingBoxes(["person", "car"],
                                  foreground=iaa.AddToHue((-255, 255)))
```

Example. Create an augmenter that randomizes the hue within bounding boxes that have either the label `person` or `car`. Only one label is picked per image. Note that the sampling happens with replacement, so if `nb_sample_classes` would be `>1`, it could still lead to only one *unique* label being sampled.

```
aug = iaa.BlendAlphaBoundingBoxes(["person", "car"],
                                  foreground=iaa.AddToHue((-255, 255)),
                                  nb_sample_labels=1)
```

Example. Create an augmenter that zeros all pixels (`Multiply(0.0)`) that are *not* (`background` branch) within bounding boxes of *any* (`None`) label. In other words, all pixels outside of bounding boxes become black. Note that we don't use `TotalDropout` here, because by default it will also remove all coordinate-based augmentables, which will break the blending of such inputs.

```
aug = iaa.BlendAlphaBoundingBoxes(None,
                                  background=iaa.Multiply(0.0))
```

9.5 augmenters.blur

9.5.1 GaussianBlur

Augmenter to blur images using gaussian kernels.

API link: *GaussianBlur*

Example. Blur each image with a gaussian kernel with a sigma of `3.0`:

```
import imgaug.augmenters as iaa
aug = iaa.GaussianBlur(sigma=(0.0, 3.0))
```

9.5.2 AverageBlur

Blur an image by computing simple means over neighbourhoods.

API link: *AverageBlur*

Example. Blur each image using a mean over neihbourhoods that have a random size between 2x2 and 11x11:

```
import imgaug.augmenters as iaa
aug = iaa.AverageBlur(k=(2, 11))
```

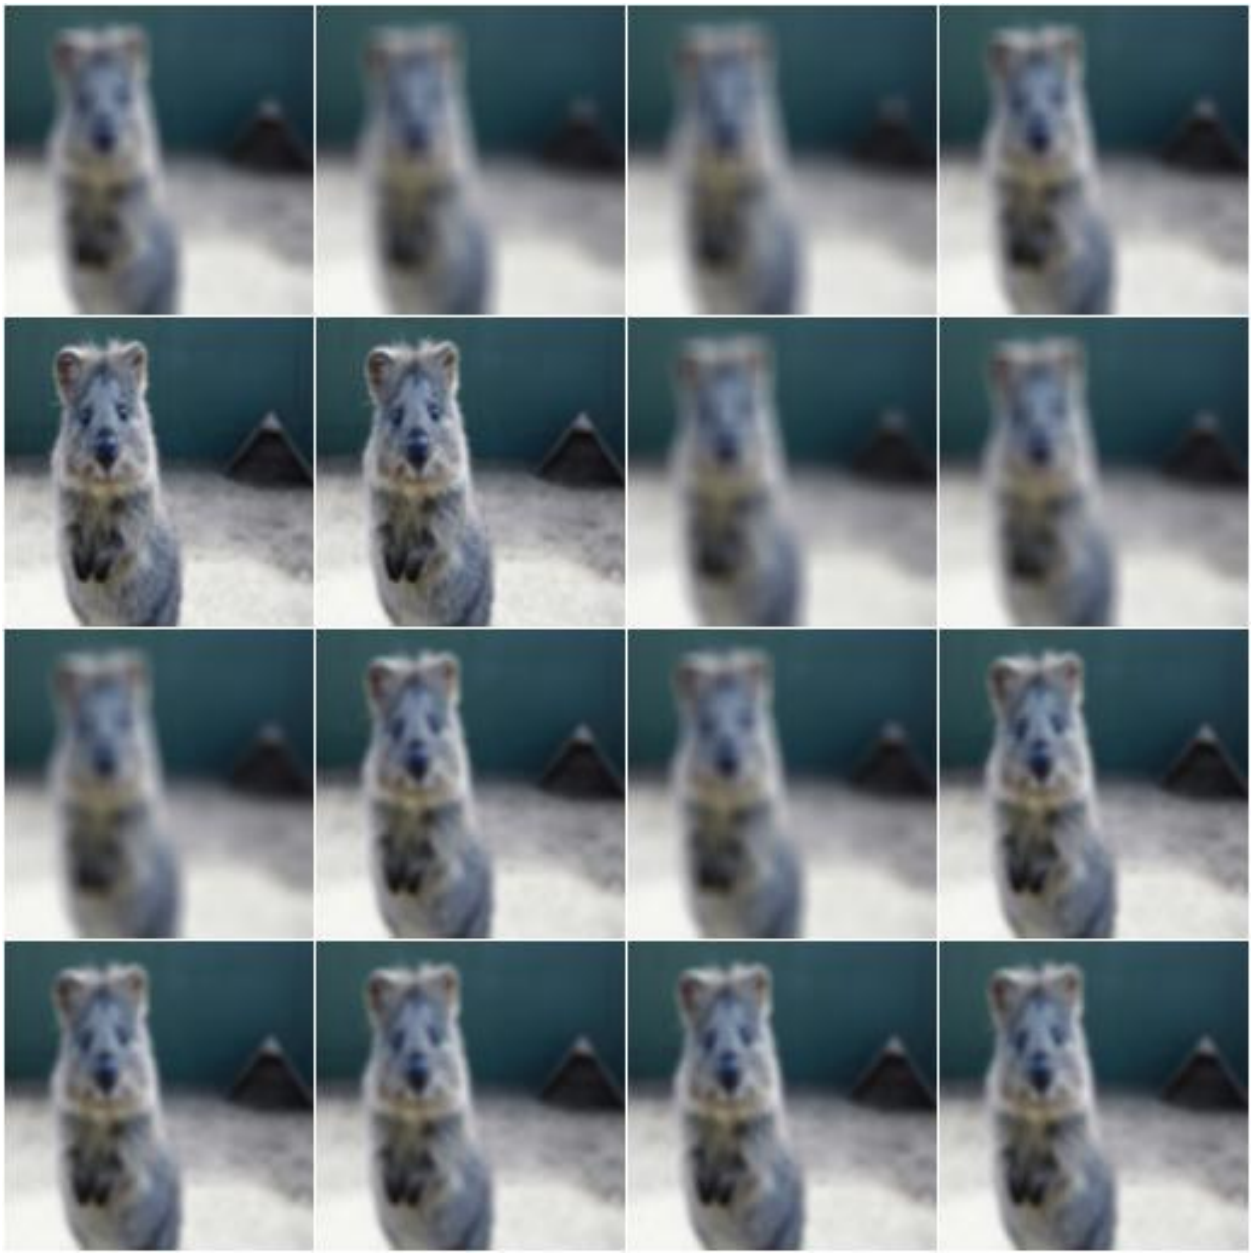

Example. Blur each image using a mean over neihbourhoods that have random sizes, which can vary between 5 and 11 in height and 1 and 3 in width:

```
aug = iaa.AverageBlur(k=((5, 11), (1, 3)))
```

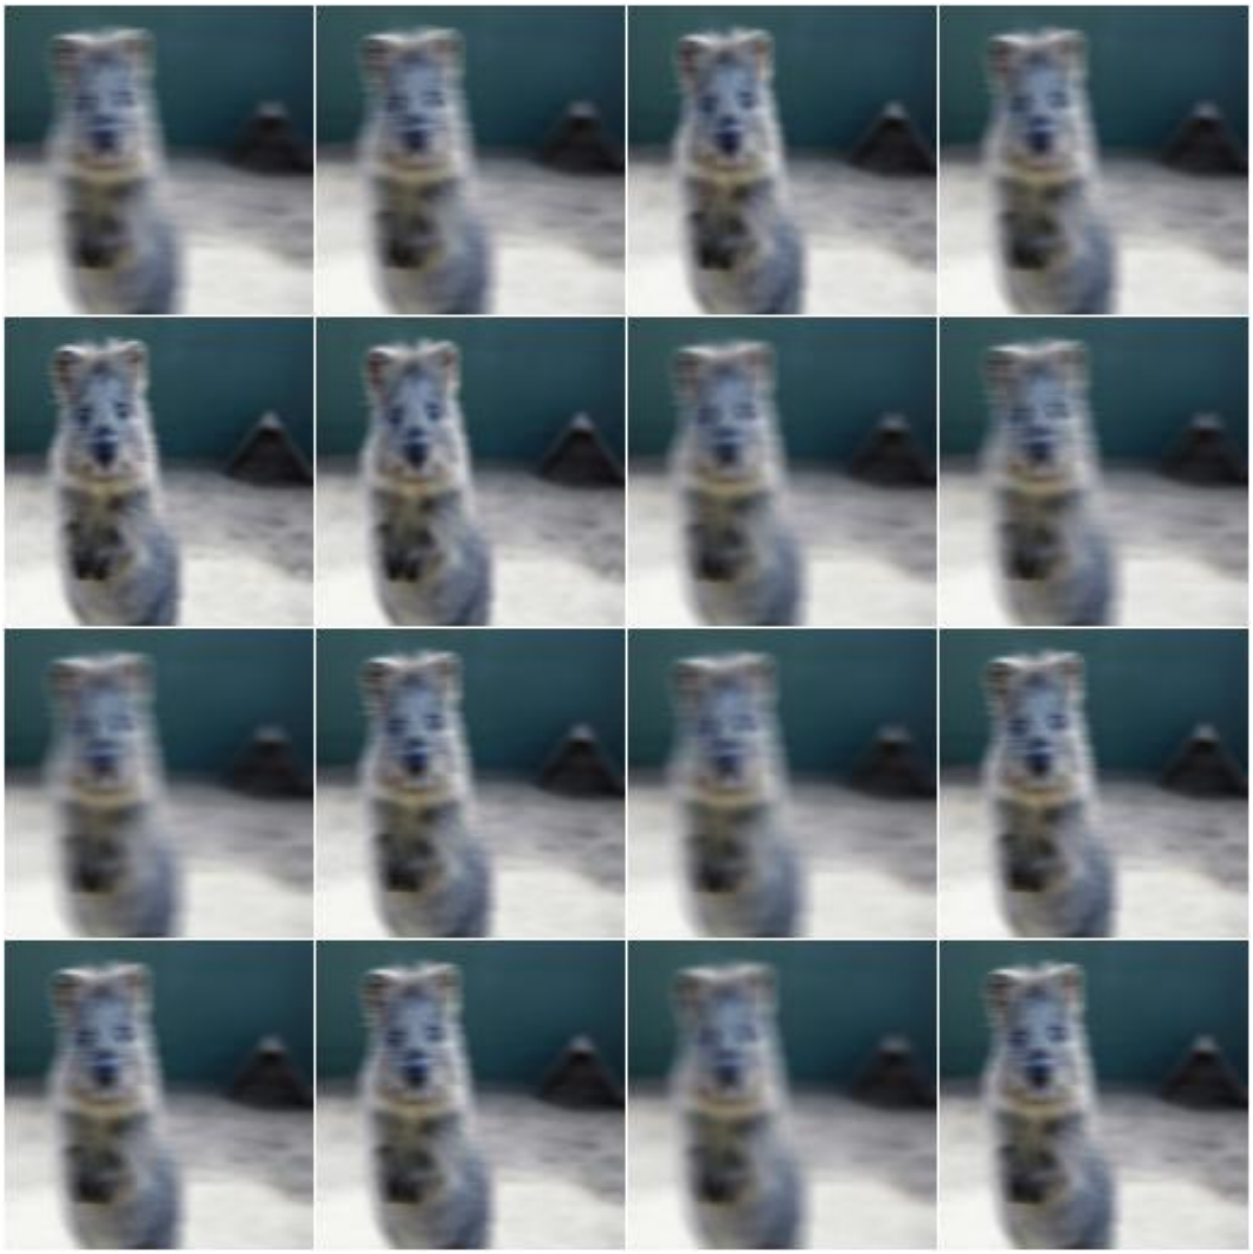

9.5.3 MedianBlur

Blur an image by computing median values over neighbourhoods.

API link: *MedianBlur*

Example. Blur each image using a median over neihbourhoods that have a random size between 3x3 and 11x11:

```
import imgaug.augmenters as iaa
aug = iaa.MedianBlur(k=(3, 11))
```

9.5.4 BilateralBlur

Blur/Denoise an image using a bilateral filter.

Bilateral filters blur homogenous and textured areas, while trying to preserve edges.

API link: *BilateralBlur*

Example. Blur all images using a bilateral filter with a *max distance* sampled uniformly from the interval `[3, 10]` and wide ranges for *sigma_color* and *sigma_space*:

```
import imgaug.augmenters as iaa
aug = iaa.BilateralBlur(
    d=(3, 10), sigma_color=(10, 250), sigma_space=(10, 250))
```

9.5.5 MotionBlur

Blur images in a way that fakes camera or object movements.

API link: *MotionBlur*

Example. Apply motion blur with a kernel size of `15x15` pixels to images:

```
import imgaug.augmenters as iaa
aug = iaa.MotionBlur(k=15)
```

Example. Apply motion blur with a kernel size of `15x15` pixels and a blur angle of either `-45` or `45` degrees (randomly picked per image):

```
aug = iaa.MotionBlur(k=15, angle=[-45, 45])
```

9.5.6 MeanShiftBlur

Apply a pyramidic mean shift filter to each image.

See also *blur_mean_shift_()* for details.

This augmenter expects input images of shape `(H,W)` or `(H,W,1)` or `(H,W,3)`.

Note: This augmenter is quite slow.

API link: *MeanShiftBlur*

Example. Create a mean shift blur augmenter:

```
import imgaug.augmenters as iaa
aug = iaa.MeanShiftBlur()
```

9.6 augmenters.collections

9.6.1 RandAugment

Apply RandAugment to inputs as described in the corresponding paper.

See paper:

```
Cubuk et al.

RandAugment: Practical automated data augmentation with a reduced
search space
```

Note: The paper contains essentially no hyperparameters for the individual augmentation techniques. The hyperparameters used here come mostly from the official code repository, which however seems to only contain code for CIFAR10 and SVHN, not for ImageNet. So some guesswork was involved and a few of the hyperparameters were also taken from https://github.com/ildoonet/pytorch-randaugment/blob/master/RandAugment/augmentations.py .

This implementation deviates from the code repository for all PIL enhance operations. In the repository these use a factor of `0.1 + M*1.8/M_max`, which would lead to a factor of `0.1` for the weakest `M` of `M=0`. For e.g. `Brightness` that would result in a basically black image. This definition is fine for AutoAugment (from where the code and hyperparameters are copied), which optimizes each transformation's `M` individually, but not for RandAugment, which uses a single fixed `M`. We hence redefine these hyperparameters to `1.0 + S * M * 0.9/M_max`, where `S` is randomly either `1` or `-1`.

We also note that it is not entirely clear which transformations were used in the ImageNet experiments. The paper lists some transformations in Figure 2, but names others in the text too (e.g. crops, flips, cutout). While Figure 2 lists the Identity function, this transformation seems to not appear in the repository (and in fact, the function `randaugment(N, M)` doesn't seem to exist in the repository either). So we also make a best guess here about what transformations might have been used.

Warning: This augmenter only works with image data, not e.g. bounding boxes. The used PIL-based affine transformations are not yet able to process non-image data. (This augmenter uses PIL-based affine transformations to ensure that outputs are as similar as possible to the paper's implementation.)

API link: *RandAugment*

Example. Create a RandAugment augmenter similar to the suggested hyperparameters in the paper:

```
import imgaug.augmenters as iaa
aug = iaa.RandAugment(n=2, m=9)
```

Example. Create a RandAugment augmenter with maximum magnitude/strength:

```
aug = iaa.RandAugment(m=30)
```

Example. Create a RandAugment augmenter that applies its transformations with a random magnitude between `0` (very weak) and `9` (recommended for ImageNet and ResNet-50). `m` is sampled per transformation:

```
aug = iaa.RandAugment(m=(0, 9))
```

Example. Create a RandAugment augmenter that applies `0` to `3` of its child transformations to images. Horizontal flips (p=50%) and crops are always applied.

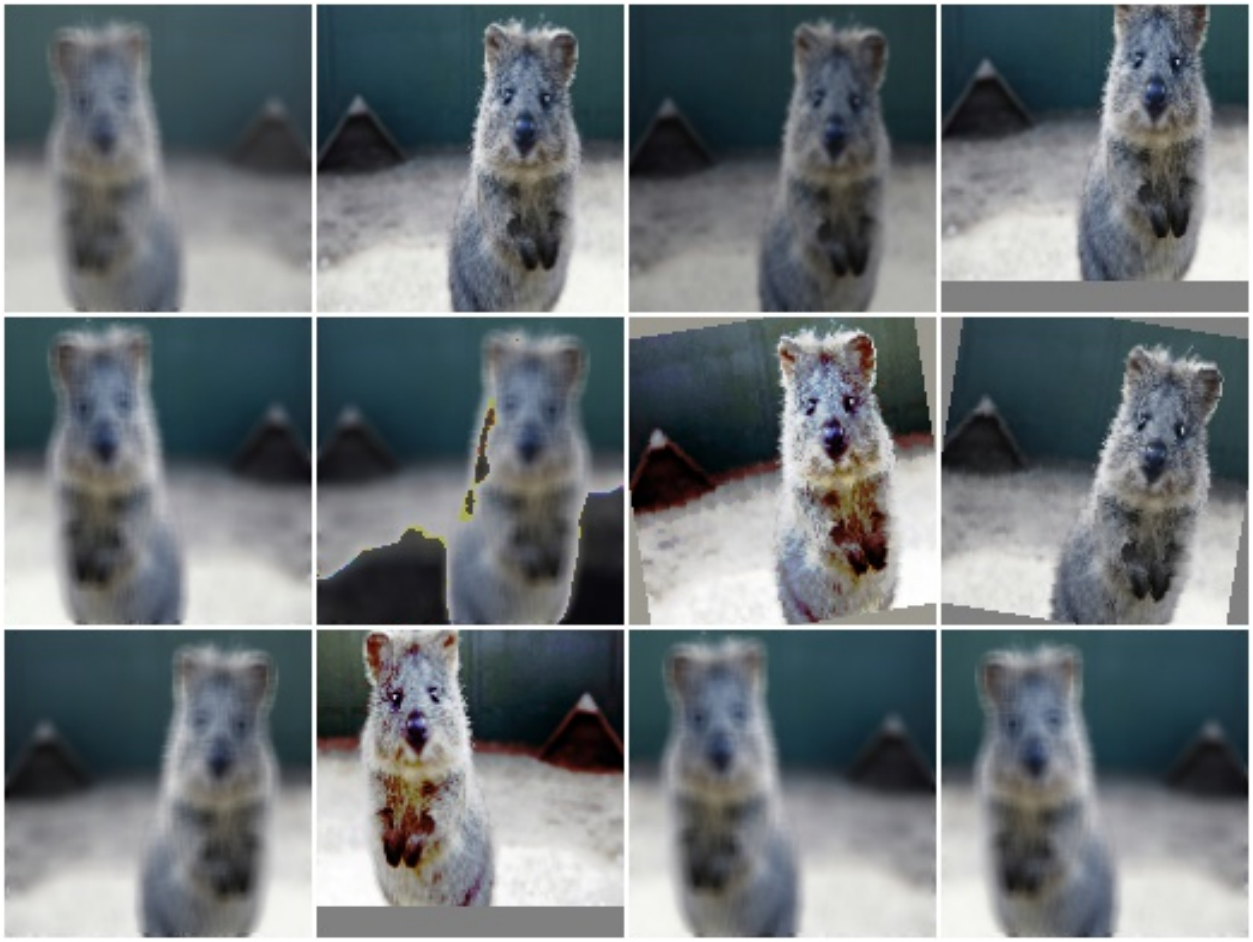

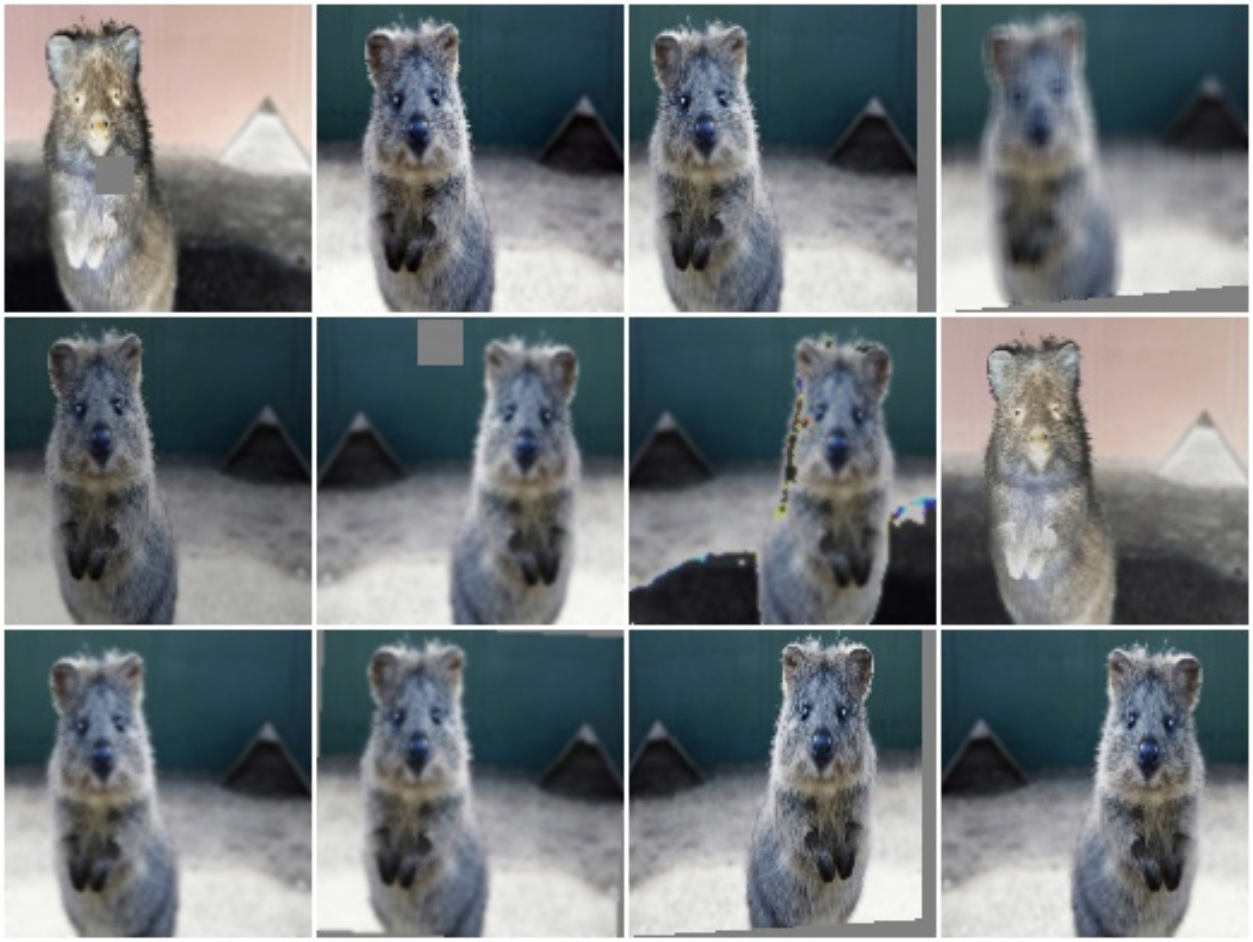

```
aug = iaa.RandAugment(n=(0, 3))
```

9.7 augmenters.color

9.7.1 WithColorspace

Apply child augmenters within a specific colorspace.

This augumenter takes a source colorspace A and a target colorspace B as well as children C. It changes images from A to B, then applies the child augmenters C and finally changes the colorspace back from B to A. See also ChangeColorspace() for more.

API link: *WithColorspace*

Example. Convert to `HSV` colorspace, add a value between `0` and `50` (uniformly sampled per image) to the Hue channel, then convert back to the input colorspace (`RGB`).

```
import imgaug.augmenters as iaa
aug = iaa.WithColorspace(
    to_colorspace="HSV",
    from_colorspace="RGB",
    children=iaa.WithChannels(
```

(continues on next page)

(continued from previous page)

```
        0,
        iaa.Add((0, 50))
    )
)
```

9.7.2 WithBrightnessChannels

Augmenter to apply child augmenters to brightness-related image channels.

This augmenter first converts an image to a random colorspace containing a brightness-related channel (e.g. `V` in `HSV`), then extracts that channel and applies its child augmenters to this one channel. Afterwards, it reintegrates the augmented channel into the full image and converts back to the input colorspace.

API link: *WithBrightnessChannels*

Example. Add `-50` to `50` to the brightness-related channels of each image:

```
import imgaug.augmenters as iaa
aug = iaa.WithBrightnessChannels(iaa.Add((-50, 50)))
```

Example. Add `-50` to `50` to the brightness-related channels of each image, but pick those brightness-related channels only from `Lab` (`L`) and `HSV` (`V`) colorspaces.

```
aug = iaa.WithBrightnessChannels(
    iaa.Add((-50, 50)), to_colorspace=[iaa.CSPACE_Lab, iaa.CSPACE_HSV])
```

Example. Add `-50` to `50` to the brightness-related channels of each image, where the images are provided in `BGR` colorspace instead of the standard `RGB`.

```
aug = iaa.WithBrightnessChannels(
    iaa.Add((-50, 50)), from_colorspace=iaa.CSPACE_BGR)
```

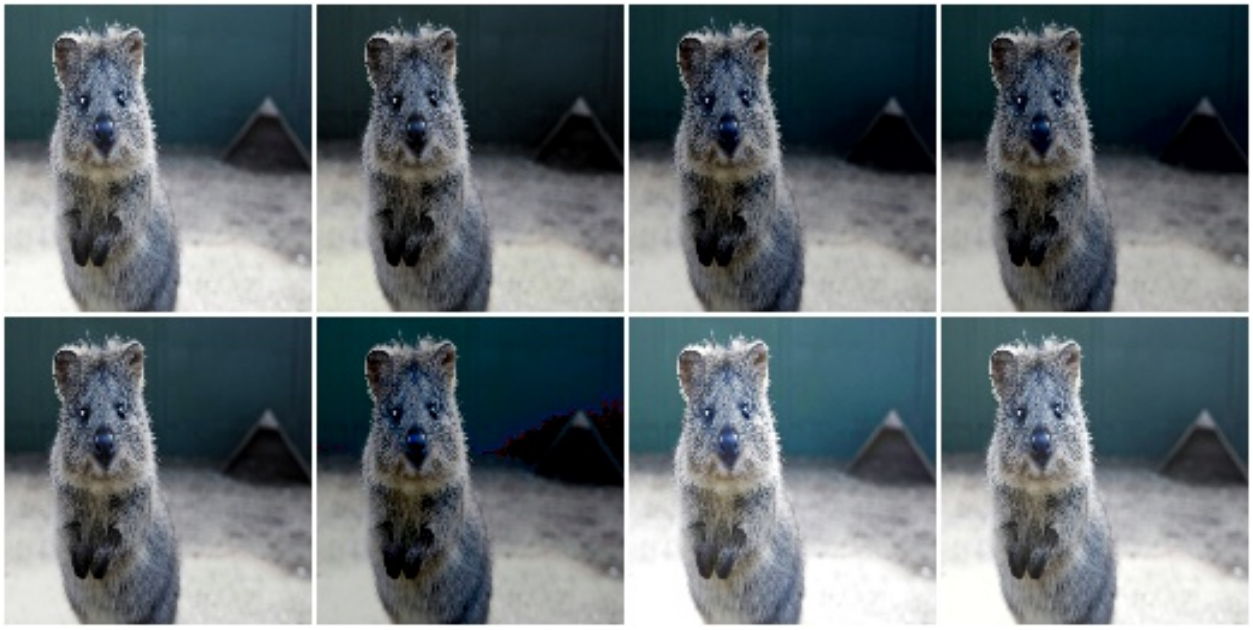

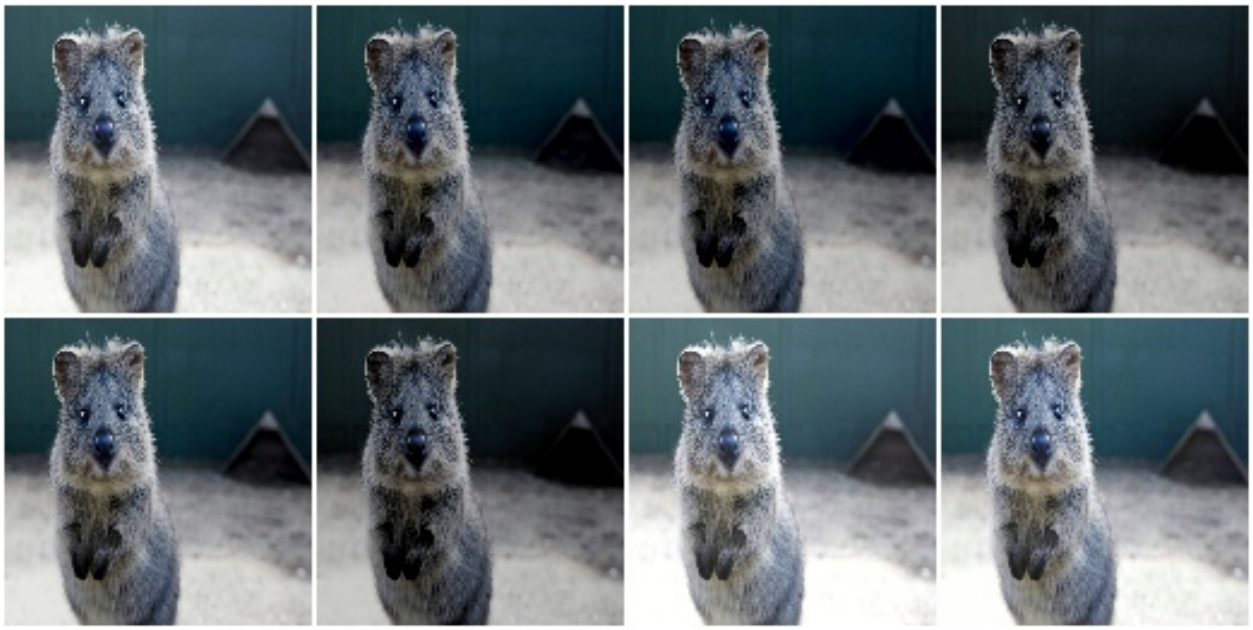

9.7.3 MultiplyAndAddToBrightness

Multiply and add to the brightness channels of input images.

This is a wrapper around *WithBrightnessChannels* and hence performs internally the same projection to random colorspaces.

API link: *MultiplyAndAddToBrightness*

Example. Convert each image to a colorspace with a brightness-related channel, extract that channel, multiply it by a factor between `0.5` and `1.5`, add a value between `-30` and `30` and convert back to the original colorspace.

```
import imgaug.augmenters as iaa
aug = iaa.MultiplyAndAddToBrightness(mul=(0.5, 1.5), add=(-30, 30))
```

9.7.4 MultiplyBrightness

Multiply the brightness channels of input images.

This is a wrapper around *WithBrightnessChannels* and hence performs internally the same projection to random colorspaces.

API link: *MultiplyBrightness*

Example. Convert each image to a colorspace with a brightness-related channel, extract that channel, multiply it by a factor between `0.5` and `1.5`, and convert back to the original colorspace.

```
import imgaug.augmenters as iaa
aug = iaa.MultiplyBrightness((0.5, 1.5))
```

9.7.5 AddToBrightness

Add to the brightness channels of input images.

This is a wrapper around *WithBrightnessChannels* and hence performs internally the same projection to random colorspaces.

API link: *AddToBrightness*

Example. Convert each image to a colorspace with a brightness-related channel, extract that channel, add between `-30` and `30` and convert back to the original colorspace:

```
import imgaug.augmenters as iaa
aug = iaa.AddToBrightness((-30, 30))
```

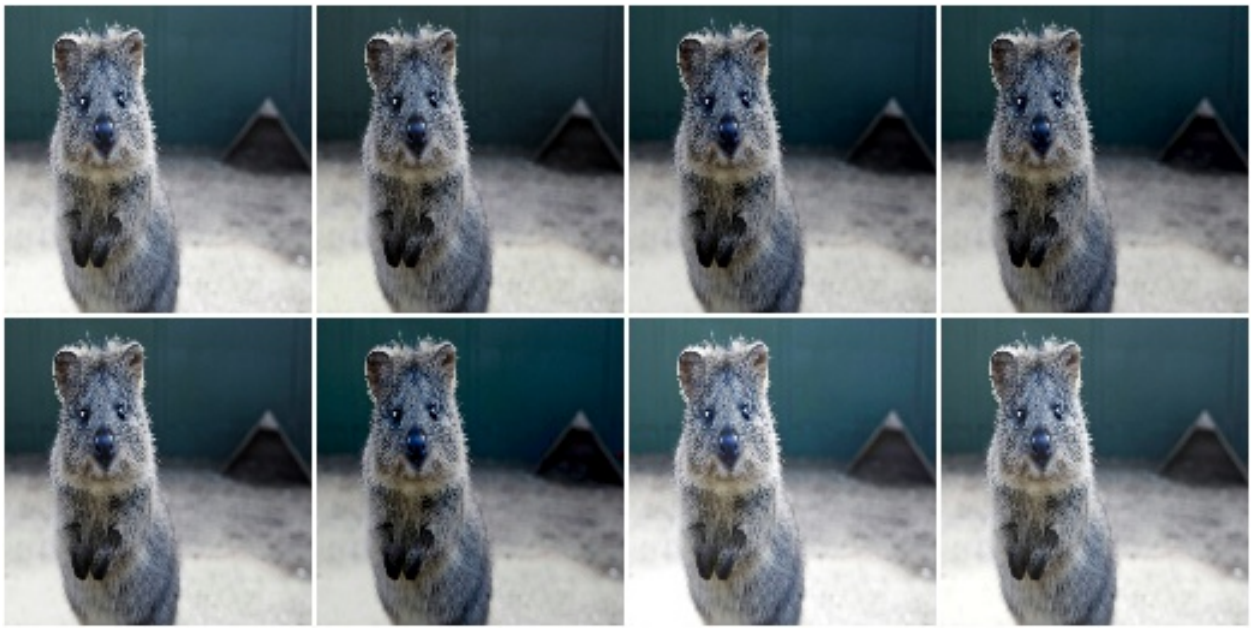

9.7.6 WithHueAndSaturation

Apply child augmenters to hue and saturation channels.

This augumenter takes an image in a source colorspace, converts it to HSV, extracts the H (hue) and S (saturation) channels, applies the provided child augmenters to these channels and finally converts back to the original colorspace.

The image array generated by this augmenter and provided to its children is in `int16` (**sic!** only augmenters that can handle `int16` arrays can be children!). The hue channel is mapped to the value range `[0, 255]`. Before converting back to the source colorspace, the saturation channel's values are clipped to `[0, 255]`. A modulo operation is applied to the hue channel's values, followed by a mapping from `[0, 255]` to `[0, 180]` (and finally the colorspace conversion).

API link: *WithHueAndSaturation*

Example. Create an augmenter that will add a random value between `0` and `50` (uniformly sampled per image) hue channel in HSV colorspace. It automatically accounts for the hue being in angular representation, i.e. if the angle goes beyond 360 degrees, it will start again at 0 degrees. The colorspace is finally converted back to `RGB` (default setting).

```
import imgaug.augmenters as iaa
aug = iaa.WithHueAndSaturation(
    iaa.WithChannels(0, iaa.Add((0, 50)))
)
```

Example. Create an augmenter that adds a random value sampled uniformly from the range `[-30, 10]` to the hue and multiplies the saturation by a random factor sampled uniformly from `[0.5, 1.5]`. It also modifies the contrast of the saturation channel. After these steps, the `HSV` image is converted back to `RGB`.

```
aug = iaa.WithHueAndSaturation([
    iaa.WithChannels(0, iaa.Add((-30, 10))),
    iaa.WithChannels(1, [
        iaa.Multiply((0.5, 1.5)),
        iaa.LinearContrast((0.75, 1.25))
    ])
])
```

9.7.7 MultiplyHueAndSaturation

Multipy hue and saturation by random values.

The augmenter first transforms images to HSV colorspace, then multiplies the pixel values in the H and S channels and afterwards converts back to RGB.

This augmenter is a wrapper around `WithHueAndSaturation`.

API link: *MultiplyHueAndSaturation()*

Example. Multiply hue and saturation by random values between `0.5` and `1.5` (independently per channel and the same value for all pixels within that channel). The hue will be automatically projected to an angular representation.

```
import imgaug.augmenters as iaa
aug = iaa.MultiplyHueAndSaturation((0.5, 1.5), per_channel=True)
```

Example. Multiply only the hue by random values between `0.5` and `1.5`.

```
aug = iaa.MultiplyHueAndSaturation(mul_hue=(0.5, 1.5))
```

Example. Multiply only the saturation by random values between `0.5` and `1.5`.

```
aug = iaa.MultiplyHueAndSaturation(mul_saturation=(0.5, 1.5))
```

9.7.8 MultiplyHue

Multiply the hue of images by random values.

The augmenter first transforms images to HSV colorspace, then multiplies the pixel values in the H channel and afterwards converts back to RGB.

This augmenter is a shortcut for `MultiplyHueAndSaturation(mul_hue=...)`.

API link: *MultiplyHue()*

Example. Multiply the hue channel of images using random values between `0.5` and `1.5`:

```
import imgaug.augmenters as iaa
aug = iaa.MultiplyHue((0.5, 1.5))
```

9.7.9 MultiplySaturation

Multiply the saturation of images by random values.

The augmenter first transforms images to HSV colorspace, then multiplies the pixel values in the H channel and afterwards converts back to RGB.

This augmenter is a shortcut for `MultiplyHueAndSaturation(mul_saturation=...)`.

API link: *MultiplySaturation()*

Example. Multiply the saturation channel of images using random values between `0.5` and `1.5`:

```
import imgaug.augmenters as iaa
aug = iaa.MultiplySaturation((0.5, 1.5))
```

9.7.10 RemoveSaturation

Decrease the saturation of images by varying degrees.

This creates images looking similar to *Grayscale*.

This augmenter is the same as `MultiplySaturation((0.0, 1.0))`.

API link: *RemoveSaturation()*

Example. Create an augmenter that decreases saturation by varying degrees:

```
import imgaug.augmenters as iaa
aug = iaa.RemoveSaturation()
```

Example. Create an augmenter that removes all saturation from input images. This is similar to *imgaug.augmenters.color.Grayscale*.

```
aug = iaa.RemoveSaturation(1.0)
```

Example. Create an augmenter that decreases saturation of images in `BGR` colorspace by varying degrees.

```
aug = iaa.RemoveSaturation(from_colorspace=iaa.CSPACE_BGR)
```

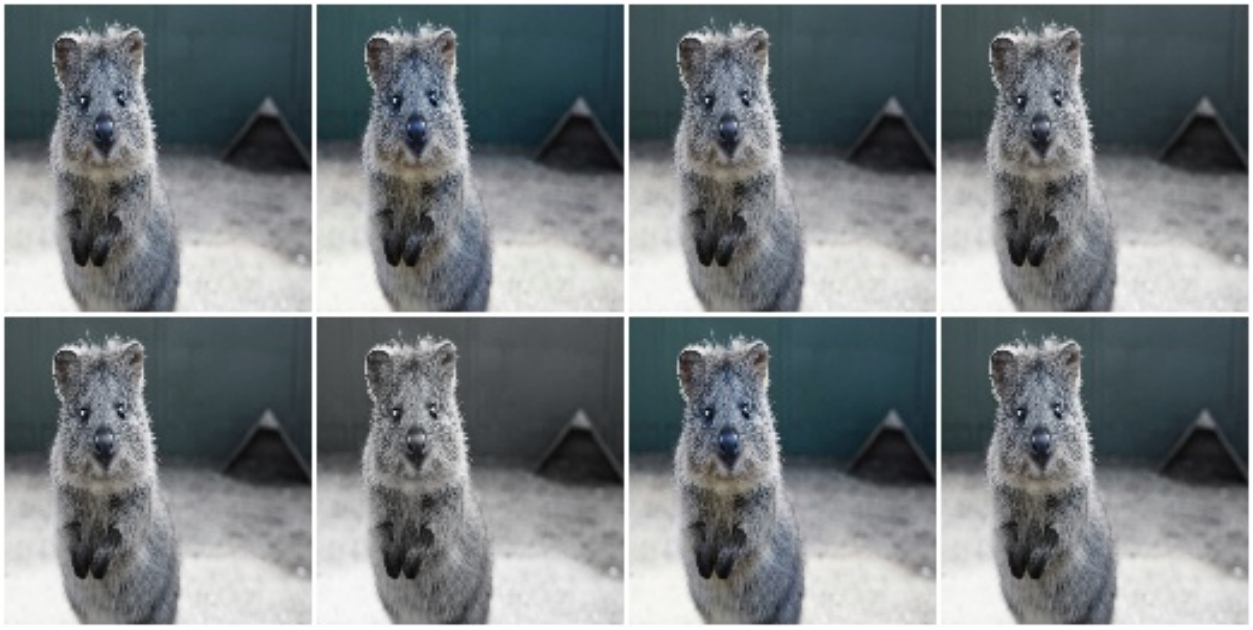

9.7.11 AddToHueAndSaturation

Increases or decreases hue and saturation by random values.

The augmenter first transforms images to HSV colorspace, then adds random values to the H and S channels and afterwards converts back to RGB.

This augmenter is faster than using `WithHueAndSaturation` in combination with `Add`.

API link: *AddToHueAndSaturation*

Example. Add random values between `-50` and `50` to the hue and saturation (independently per channel and the same value for all pixels within that channel):

```
import imgaug.augmenters as iaa
aug = iaa.AddToHueAndSaturation((-50, 50), per_channel=True)
```

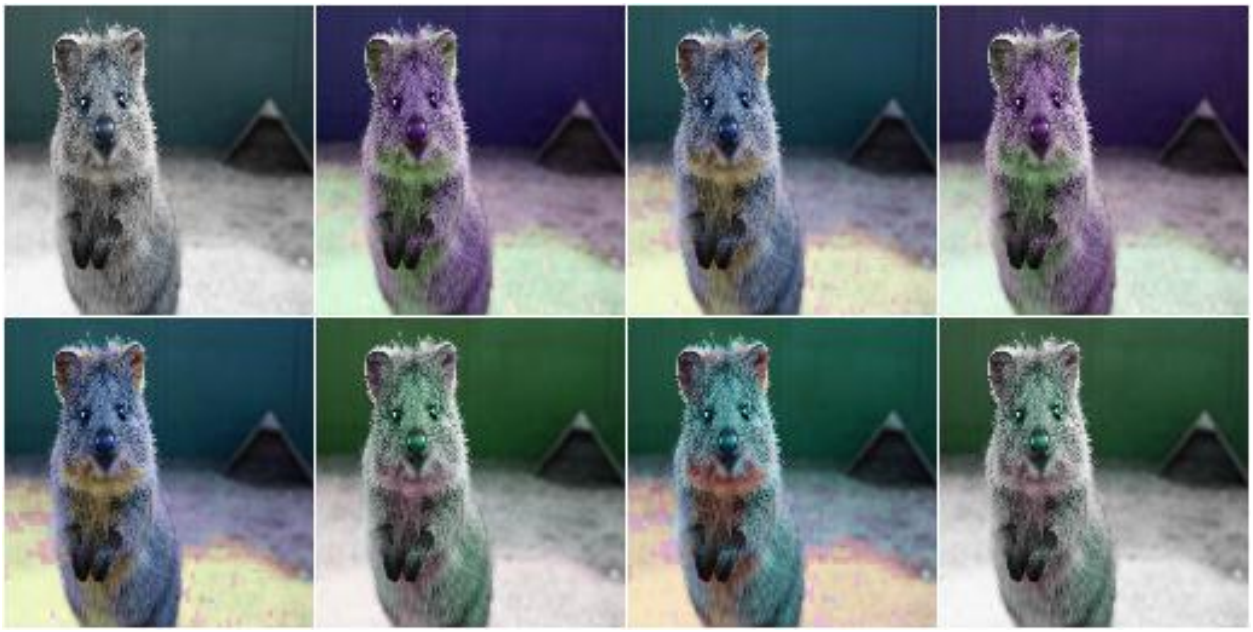

9.7.12 AddToHue

Add random values to the hue of images.

The augmenter first transforms images to HSV colorspace, then adds random values to the H channel and afterwards converts back to RGB.

If you want to change both the hue and the saturation, it is recommended to use `AddToHueAndSaturation` as otherwise the image will be converted twice to HSV and back to RGB.

This augmenter is a shortcut for `AddToHueAndSaturation(value_hue=...)`.

API link: *AddToHue()*

Example. Sample random values from the discrete uniform range `[-50..50]`, convert them to angular representation and add them to the hue, i.e. to the `H` channel in `HSV` colorspace:

```
import imgaug.augmenters as iaa
aug = iaa.AddToHue((-50, 50))
```

9.7.13 AddToSaturation

Add random values to the saturation of images.

The augmenter first transforms images to HSV colorspace, then adds random values to the S channel and afterwards converts back to RGB.

If you want to change both the hue and the saturation, it is recommended to use `AddToHueAndSaturation` as otherwise the image will be converted twice to HSV and back to RGB.

This augmenter is a shortcut for `AddToHueAndSaturation(value_saturation=...)`.

API link: *AddToSaturation()*

Example. Sample random values from the discrete uniform range `[-50..50]`, and add them to the saturation, i.e. to the `S` channel in `HSV` colorspace:

```
import imgaug.augmenters as iaa
aug = iaa.AddToSaturation((-50, 50))
```

9.7.14 ChangeColorspace

Augmenter to change the colorspace of images.

API link: *ChangeColorspace*

Example. The following example shows how to change the colorspace from RGB to HSV, then add 50-100 to the first channel, then convert back to RGB. This increases the hue value of each image.

```
import imgaug.augmenters as iaa
aug = iaa.Sequential([
    iaa.ChangeColorspace(from_colorspace="RGB", to_colorspace="HSV"),
    iaa.WithChannels(0, iaa.Add((50, 100))),
    iaa.ChangeColorspace(from_colorspace="HSV", to_colorspace="RGB")
])
```

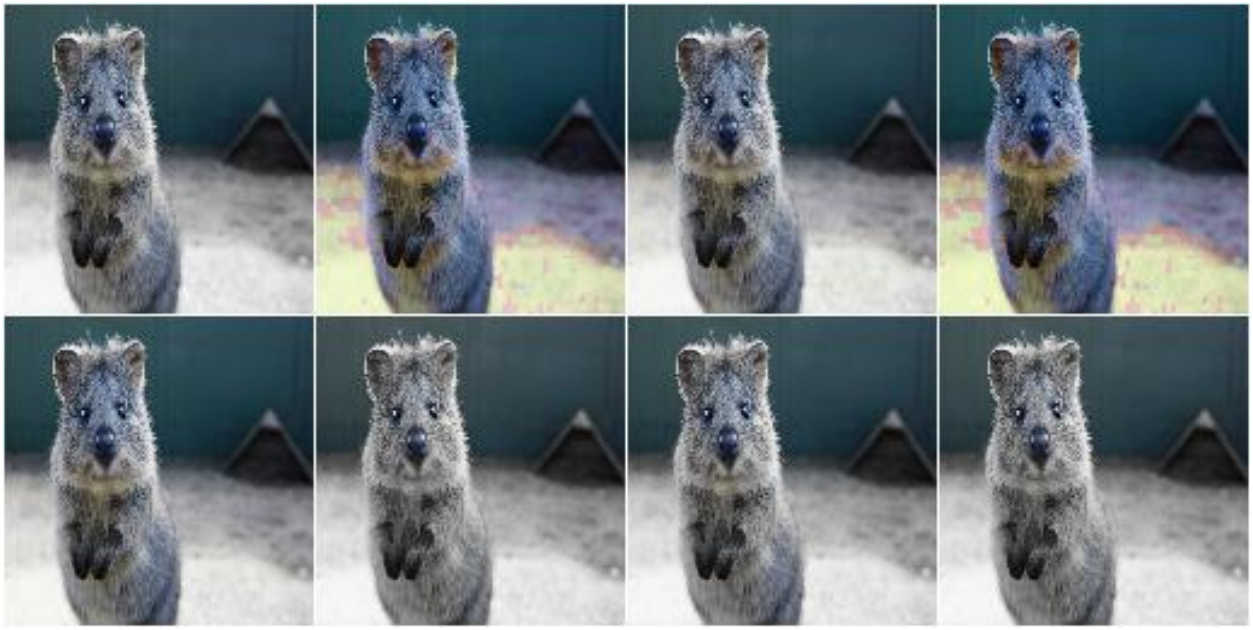

9.7.15 Grayscale

Augmenter to convert images to their grayscale versions.

API link: *Grayscale*

Example. Change images to grayscale and overlay them with the original image by varying strengths, effectively removing 0 to 100% of the color:

```
import imgaug.augmenters as iaa
aug = iaa.Grayscale(alpha=(0.0, 1.0))
```

Example. Visualization of increasing `alpha` from `0.0` to `1.0` in eight steps:

9.7.16 ChangeColorTemperature

Change the temperature to a provided Kelvin value.

Low Kelvin values around `1000` to `4000` will result in red, yellow or orange images. Kelvin values around `10000` to `40000` will result in progressively darker blue tones.

API link: *ChangeColorTemperature*

Example. Create an augmenter that changes the color temperature of images to a random value between `1100` and `10000` Kelvin:

```
import imgaug.augmenters as iaa
aug = iaa.ChangeColorTemperature((1100, 10000))
```

9.7.17 KMeansColorQuantization

Quantize colors using k-Means clustering.

This "collects" the colors from the input image, groups them into `k` clusters using k-Means clustering and replaces the colors in the input image using the cluster centroids.

This is slower than `UniformColorQuantization`, but adapts dynamically to the color range in the input image.

Note: This augmenter expects input images to be either grayscale or to have 3 or 4 channels and use colorspace *from_colorspace*. If images have 4 channels, it is assumed that the 4th channel is an alpha channel and it will not be quantized.

API link: *KMeansColorQuantization*

Example. Create an augmenter to apply k-Means color quantization to images using a random amount of colors, sampled uniformly from the interval `[2..16]`. It assumes the input image colorspace to be `RGB` and clusters colors randomly in `RGB` or `Lab` colorspace.

```
import imgaug.augmenters as iaa
aug = iaa.KMeansColorQuantization()
```

Example. Create an augmenter that quantizes images to (up to) eight colors:

```
aug = iaa.KMeansColorQuantization(n_colors=8)
```

Example. Create an augmenter that quantizes images to (up to) `n` colors, where `n` is randomly and uniformly sampled from the discrete interval `[4..32]`:

```
aug = iaa.KMeansColorQuantization(n_colors=(4, 16))
```

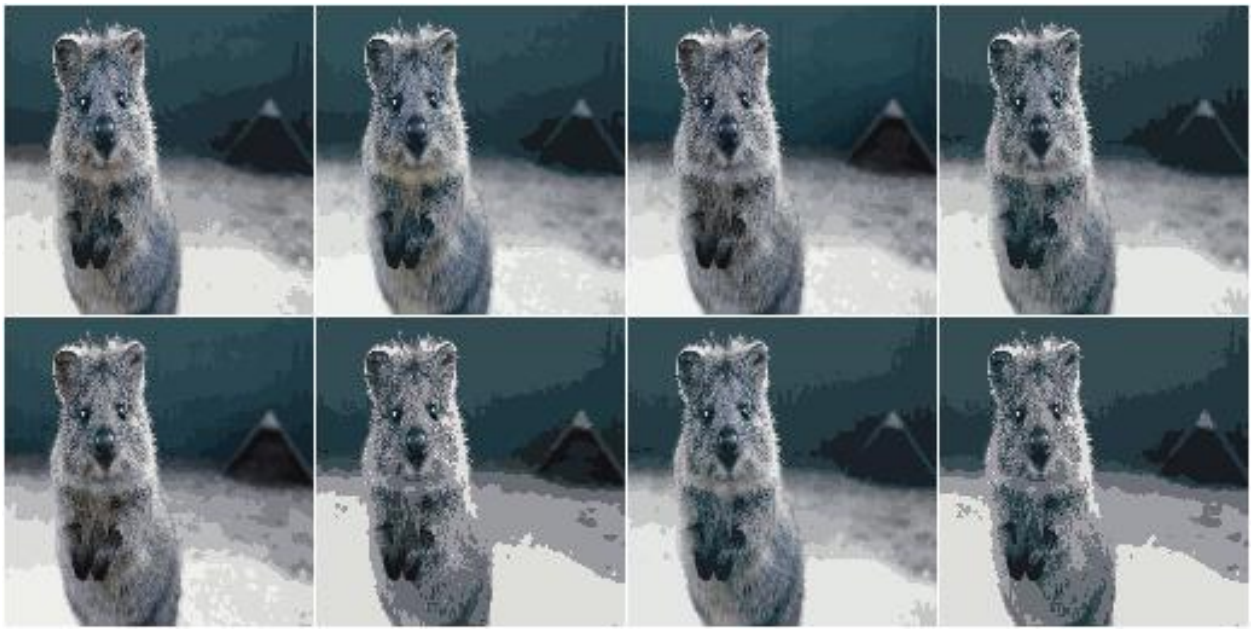

Example. Create an augmenter that quantizes input images that are in `BGR` colorspace. The quantization happens in `RGB` or `Lab` colorspace, into which the images are temporarily converted.

```
aug = iaa.KMeansColorQuantization(
    from_colorspace=iaa.ChangeColorspace.BGR)
```

Example. Create an augmenter that quantizes images by clustering colors randomly in either `RGB` or `HSV` colorspace. The assumed input colorspace of images is `RGB`.

```
aug = iaa.KMeansColorQuantization(
    to_colorspace=[iaa.ChangeColorspace.RGB, iaa.ChangeColorspace.HSV])
```

9.7.18 UniformColorQuantization

Quantize colors into N bins with regular distance.

For `uint8` images the equation is `floor(v/q)*q + q/2` with `q = 256/N`, where `v` is a pixel intensity value and `N` is the target number of colors after quantization.

This augmenter is faster than `KMeansColorQuantization`, but the set of possible output colors is constant (i.e. independent of the input images). It may produce unsatisfying outputs for input images that are made up of very similar colors.

Note: This augmenter expects input images to be either grayscale or to have 3 or 4 channels and use colorspace *from_colorspace*. If images have 4 channels, it is assumed that the 4th channel is an alpha channel and it will not be quantized.

API link: *UniformColorQuantization*

Example. Create an augmenter to apply uniform color quantization to images using a random amount of colors, sampled uniformly from the discrete interval `[2..16]`:

```
import imgaug.augmenters as iaa
aug = iaa.UniformColorQuantization()
```

Example. Create an augmenter that quantizes images to (up to) eight colors:

```
aug = iaa.UniformColorQuantization(n_colors=8)
```

Example. Create an augmenter that quantizes images to (up to) `n` colors, where `n` is randomly and uniformly sampled from the discrete interval `[4..32]`:

```
aug = iaa.UniformColorQuantization(n_colors=(4, 16))
```

Example. Create an augmenter that uniformly quantizes images in either `RGB` or `HSV` colorspace (randomly picked per image). The input colorspace of all images has to be `BGR`.

```
aug = iaa.UniformColorQuantization(
    from_colorspace=iaa.ChangeColorspace.BGR,
    to_colorspace=[iaa.ChangeColorspace.RGB, iaa.ChangeColorspace.HSV])
```

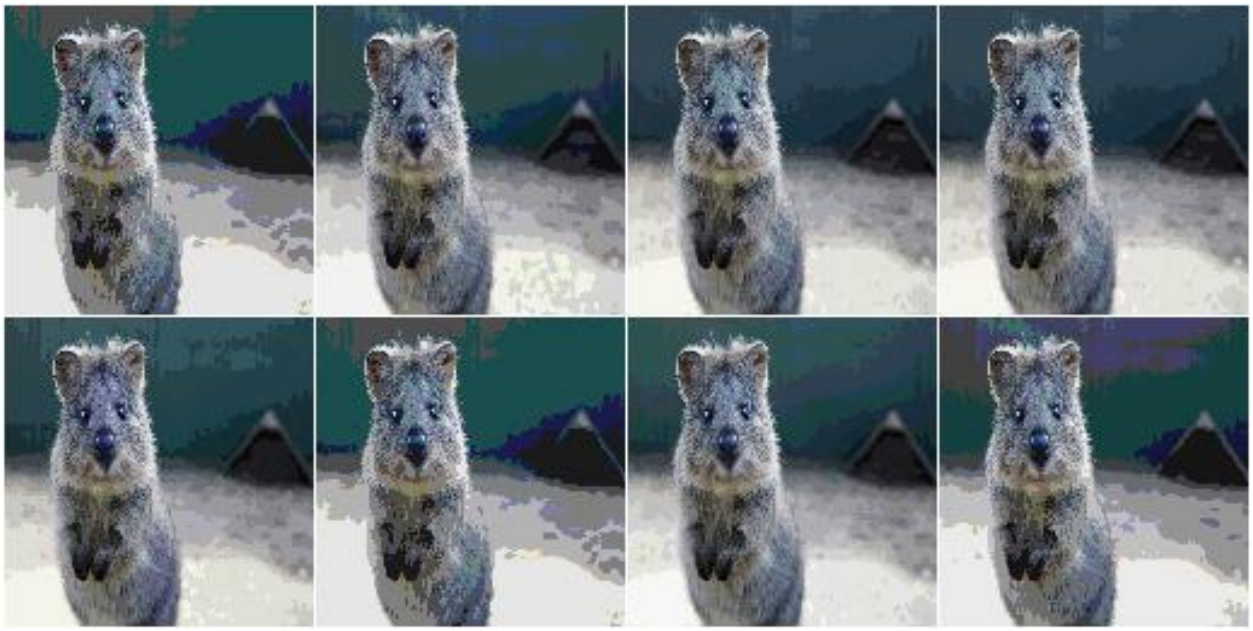

9.7.19 UniformColorQuantizationToNBits

Quantize images by setting `8-B` bits of each component to zero.

This augmenter sets the `8-B` highest frequency (rightmost) bits of each array component to zero. For `B` bits this is equivalent to changing each component's intensity value `v` to `v' = v & (2**(8-B) - 1)`, e.g. for `B=3` this results in `v' = c & ~(2**(3-1) - 1) = c & ~3 = c & ~0000 0011 = c & 1111 1100`.

This augmenter behaves for `B` similarly to `UniformColorQuantization(2**B)`, but quantizes each bin with interval `(a, b)` to `a` instead of to `a + (b-a)/2`.

This augmenter is comparable to `posterize()`.

Note: This augmenter expects input images to be either grayscale or to have 3 or 4 channels and use colorspace *from_colorspace*. If images have 4 channels, it is assumed that the 4th channel is an alpha channel and it will not be quantized.

API link: *UniformColorQuantizationToNBits*

Example. Create an augmenter to apply uniform color quantization to images using a random amount of bits to remove, sampled uniformly from the discrete interval `[1..8]`:

```
import imgaug.augmenters as iaa
aug = iaa.UniformColorQuantizationToNBits()
```

Example. Create an augmenter that quantizes images by removing `8-B` rightmost bits from each component, where `B` is uniformly sampled from the discrete interval `[2..8]`:

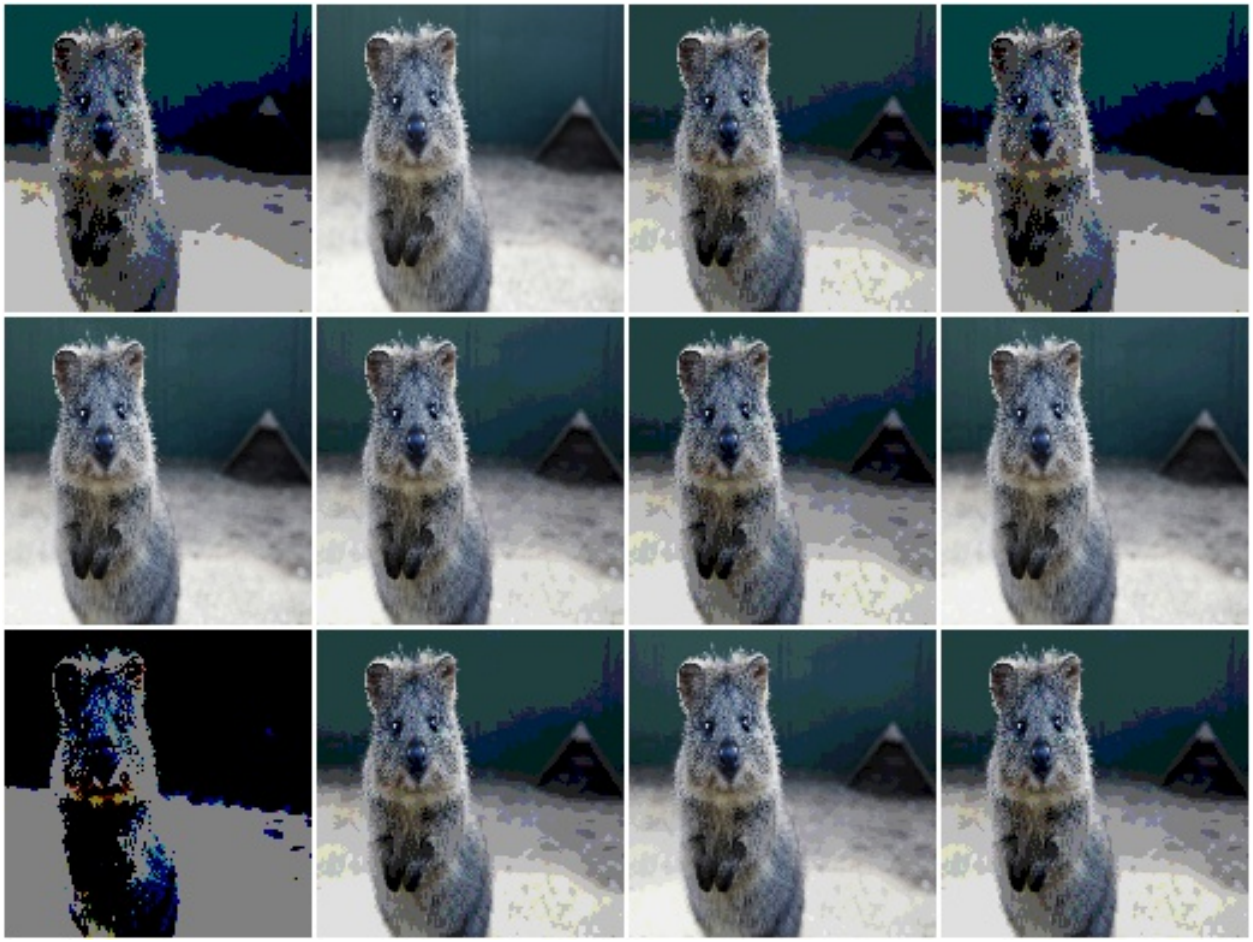

```
aug = iaa.UniformColorQuantizationToNBits(nb_bits=(2, 8))
```

Example. Create an augmenter that uniformly quantizes images in either `RGB` or `HSV` colorspace (randomly picked per image). The input colorspace of all images has to be `BGR`:

```
aug = iaa.UniformColorQuantizationToNBits(
    from_colorspace=iaa.CSPACE_BGR,
    to_colorspace=[iaa.CSPACE_RGB, iaa.CSPACE_HSV])
```

9.7.20 Posterize

Alias for *imgaug.augmenters.color.UniformColorQuantizationToNBits*.

API link: *Posterize*

9.8 augmenters.contrast

9.8.1 GammaContrast

Adjust image contrast by scaling pixel values to `255*((v/255)**gamma)`.

Values in the range `gamma=(0.5, 2.0)` seem to be sensible.

API link: *GammaContrast()*

Example. Modify the contrast of images according to `255*((v/255)**gamma)`, where `v` is a pixel value and `gamma` is sampled uniformly from the interval `[0.5, 2.0]` (once per image):

```
import imgaug.augmenters as iaa
aug = iaa.GammaContrast((0.5, 2.0))
```

Example. Same as in the previous example, but `gamma` is sampled once per image *and* channel:

```
aug = iaa.GammaContrast((0.5, 2.0), per_channel=True)
```

9.8.2 SigmoidContrast

Adjust image contrast to `255*1/(1+exp(gain*(cutoff-I_ij/255)))`.

Values in the range `gain=(5, 20)` and `cutoff=(0.25, 0.75)` seem to be sensible.

API link: *SigmoidContrast()*

Example. Modify the contrast of images according to `255*1/(1+exp(gain*(cutoff-v/255)))`, where `v` is a pixel value, `gain` is sampled uniformly from the interval `[3, 10]` (once per image) and `cutoff` is sampled uniformly from the interval `[0.4, 0.6]` (also once per image).

```
import imgaug.augmenters as iaa
aug = iaa.SigmoidContrast(gain=(3, 10), cutoff=(0.4, 0.6))
```

Example. Same as in the previous example, but `gain` and `cutoff` are each sampled once per image *and* channel:

```
aug = iaa.SigmoidContrast(
    gain=(3, 10), cutoff=(0.4, 0.6), per_channel=True)
```

9.8.3 LogContrast

Adjust image contrast by scaling pixels to `255*gain*log_2(1+v/255)`.

This augmenter is fairly similar to `imgaug.augmenters.arithmetic.Multiply`.

API link: *LogContrast()*

Example. Modify the contrast of images according to `255*gain*log_2(1+v/255)`, where `v` is a pixel value and `gain` is sampled uniformly from the interval `[0.6, 1.4]` (once per image):

```
import imgaug.augmenters as iaa
aug = iaa.LogContrast(gain=(0.6, 1.4))
```

Example. Same as in the previous example, but `gain` is sampled once per image *and* channel:

```
aug = iaa.LogContrast(gain=(0.6, 1.4), per_channel=True)
```

9.8.4 LinearContrast

Adjust contrast by scaling each pixel to `127 + alpha*(v-127)`.

API link: *LinearContrast()*

Example. Modify the contrast of images according to *127 + alpha*(v-127)*‘, where `v` is a pixel value and `alpha` is sampled uniformly from the interval `[0.4, 1.6]` (once per image):

```
import imgaug.augmenters as iaa
aug = iaa.LinearContrast((0.4, 1.6))
```

Example. Same as in the previous example, but `alpha` is sampled once per image *and* channel:

```
aug = iaa.LinearContrast((0.4, 1.6), per_channel=True)
```

9.8.5 AllChannelsCLAHE

Apply CLAHE to all channels of images in their original colorspaces.

CLAHE (Contrast Limited Adaptive Histogram Equalization) performs histogram equilization within image patches, i.e. over local neighbourhoods.

In contrast to `imgaug.augmenters.contrast.CLAHE`, this augmenter operates directly on all channels of the input images. It does not perform any colorspace transformations and does not focus on specific channels (e.g. `L` in `Lab` colorspace).

API link: *AllChannelsCLAHE*

Example. Create an augmenter that applies CLAHE to all channels of input images:

```
import imgaug.augmenters as iaa
aug = iaa.AllChannelsCLAHE()
```

Example. Same as in the previous example, but the *clip_limit* used by CLAHE is uniformly sampled per image from the interval `[1, 10]`. Some images will therefore have stronger contrast than others (i.e. higher clip limit values).

```
aug = iaa.AllChannelsCLAHE(clip_limit=(1, 10))
```

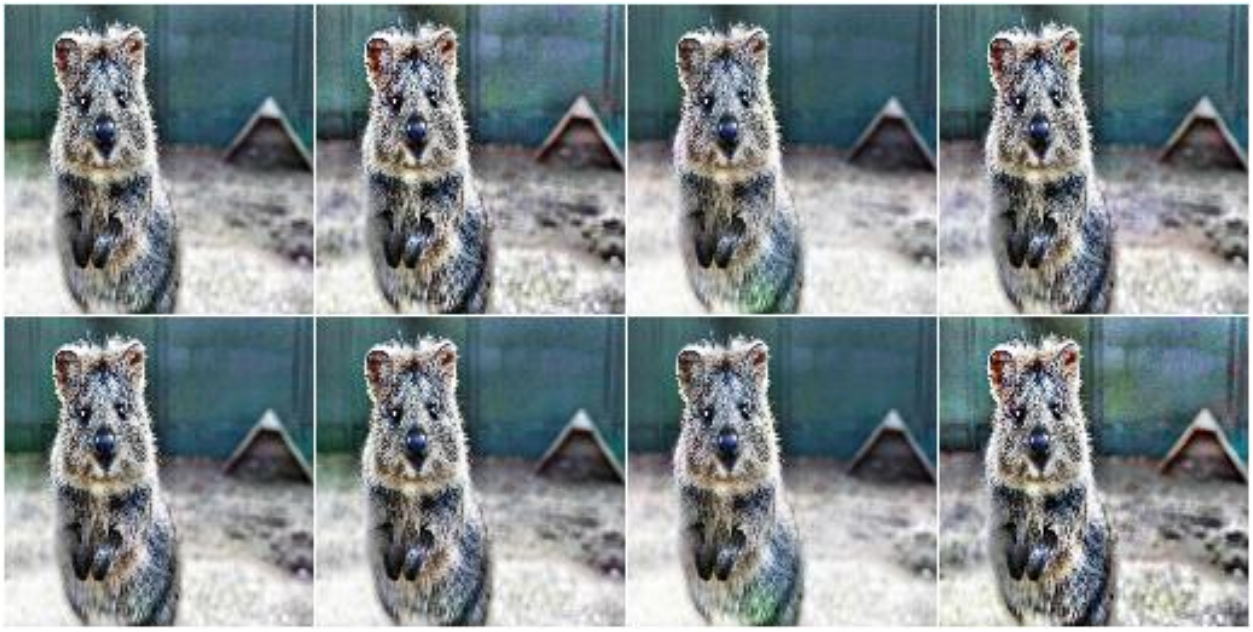

Example. Same as in the previous example, but the *clip_limit* is sampled per image *and* channel, leading to different levels of contrast for each channel:

```
aug = iaa.AllChannelsCLAHE(clip_limit=(1, 10), per_channel=True)
```

9.8.6 CLAHE

Apply CLAHE to L/V/L channels in HLS/HSV/Lab colorspaces.

This augmenter applies CLAHE (Contrast Limited Adaptive Histogram Equalization) to images, a form of histogram equalization that normalizes within local image patches. The augmenter transforms input images to a target colorspace (e.g. `Lab`), extracts an intensity-related channel from the converted images (e.g. `L` for `Lab`), applies CLAHE to the channel and then converts the resulting image back to the original colorspace.

Grayscale images (images without channel axis or with only one channel axis) are automatically handled, *from_colorspace* does not have to be adjusted for them. For images with four channels (e.g. `RGBA`), the fourth channel is ignored in the colorspace conversion (e.g. from an `RGBA` image, only the `RGB` part is converted, normalized, converted back and concatenated with the input `A` channel). Images with unusual channel numbers (2, 5 or more than 5) are normalized channel-by-channel (same behaviour as `AllChannelsCLAHE`, though a warning will be raised).

If you want to apply CLAHE to each channel of the original input image's colorspace (without any colorspace conversion), use `imgaug.augmenters.contrast.AllChannelsCLAHE` instead.

API link: *CLAHE*

Example. Create a standard CLAHE augmenter:

```
import imgaug.augmenters as iaa
aug = iaa.CLAHE()
```

Example. Create a CLAHE augmenter with a clip limit uniformly sampled from `[1..10]`, where `1` is rather low contrast and `10` is rather high contrast:

```
aug = iaa.CLAHE(clip_limit=(1, 10))
```

Example. Create a CLAHE augmenter with kernel sizes of `SxS`, where `S` is uniformly sampled from `[3..21]`. Sampling happens once per image.

```
aug = iaa.CLAHE(tile_grid_size_px=(3, 21))
```

Example. Create a CLAHE augmenter with kernel sizes of `SxS`, where `S` is sampled from `N(7, 2)`, but does not go below `3`:

```
import imgaug.parameters as iap
aug = iaa.CLAHE(
    tile_grid_size_px=iap.Discretize(iap.Normal(loc=7, scale=2)),
    tile_grid_size_px_min=3)
```

Example. Create a CLAHE augmenter with kernel sizes of `HxW`, where `H` is uniformly sampled from `[3..21]` and `W` is randomly picked from the list `[3, 5, 7]`:

```
aug = iaa.CLAHE(tile_grid_size_px=((3, 21), [3, 5, 7]))
```

Example. Create a CLAHE augmenter that converts images from BGR colorspace to HSV colorspace and then applies the local histogram equalization to the `V` channel of the images (before converting back to `BGR`). Alternatively, `Lab` (default) or `HLS` can be used as the target colorspace. Grayscale images (no channels / one channel) are never converted and are instead directly normalized (i.e. *from_colorspace* does not have to be changed for them).

```
aug = iaa.CLAHE(
    from_colorspace=iaa.CLAHE.BGR,
    to_colorspace=iaa.CLAHE.HSV)
```

9.8.7 AllChannelsHistogramEqualization

Apply Histogram Eq. to all channels of images in their original colorspaces.

In contrast to `imgaug.augmenters.contrast.HistogramEqualization`, this augmenter operates directly on all channels of the input images. It does not perform any colorspace transformations and does not focus on specific channels (e.g. `L` in `Lab` colorspace).

API link: *AllChannelsHistogramEqualization*

Example. Create an augmenter that applies histogram equalization to all channels of input images in the original colorspaces:

```
import imgaug.augmenters as iaa
aug = iaa.AllChannelsHistogramEqualization()
```

Example. Same as in the previous example, but alpha-blends the contrast-enhanced augmented images with the original input images using random blend strengths. This leads to random strengths of the contrast adjustment.

```
aug = iaa.Alpha((0.0, 1.0), iaa.AllChannelsHistogramEqualization())
```

9.8.8 HistogramEqualization

Apply Histogram Eq. to L/V/L channels of images in HLS/HSV/Lab colorspaces.

This augmenter is similar to `imgaug.augmenters.contrast.CLAHE`.

The augmenter transforms input images to a target colorspace (e.g. `Lab`), extracts an intensity-related channel from the converted images (e.g. `L` for `Lab`), applies Histogram Equalization to the channel and then converts the resulting image back to the original colorspace.

Grayscale images (images without channel axis or with only one channel axis) are automatically handled, *from_colorspace* does not have to be adjusted for them. For images with four channels (e.g. RGBA), the fourth channel is ignored in the colorspace conversion (e.g. from an `RGBA` image, only the `RGB` part is converted, normalized, converted back and concatenated with the input `A` channel). Images with unusual channel numbers (2, 5 or more than 5) are normalized channel-by-channel (same behaviour as `AllChannelsHistogramEqualization`, though a warning will be raised).

If you want to apply HistogramEqualization to each channel of the original input image's colorspace (without any colorspace conversion), use `imgaug.augmenters.contrast.AllChannelsHistogramEqualization` instead.

API link: *HistogramEqualization*

Example. Create an augmenter that converts images to `HLS`/`HSV`/`Lab` colorspaces, extracts intensity-related channels (i.e. `L`/`V`/`L`), applies histogram equalization to these channels and converts back to the input colorspace:

```
import imgaug.augmenters as iaa
aug = iaa.HistogramEqualization()
```

Example. Same as in the previous example, but alpha blends the result, leading to various strengths of contrast normalization:

```
aug = iaa.Alpha((0.0, 1.0), iaa.HistogramEqualization())
```

Example. Same as in the first example, but the colorspace of input images has to be `BGR` (instead of default `RGB`) and the histogram equalization is applied to the `V` channel in `HSV` colorspace:

```
aug = iaa.HistogramEqualization(
    from_colorspace=iaa.HistogramEqualization.BGR,
    to_colorspace=iaa.HistogramEqualization.HSV)
```

9.9 augmenters.convolutional

9.9.1 Convolve

Apply a Convolution to input images.

API link: *Convolve*

Example. Convolve each image with a 3x3 kernel:

```
import imgaug.augmenters as iaa
matrix = np.array([[0, -1, 0],
                   [-1, 4, -1],
                   [0, -1, 0]])
aug = iaa.Convolve(matrix=matrix)
```

Example. Convolve each image with a 3x3 kernel, which is chosen dynamically per image:

```
def gen_matrix(image, nb_channels, random_state):
      matrix_A = np.array([[0, -1, 0],
                           [-1, 4, -1],
                           [0, -1, 0]])
     matrix_B = np.array([[0, 0, 0],
                          [0, -4, 1],
                          [0, 2, 1]])
      if random_state.rand() < 0.5:
          return [matrix_A] * nb_channels
```

(continues on next page)

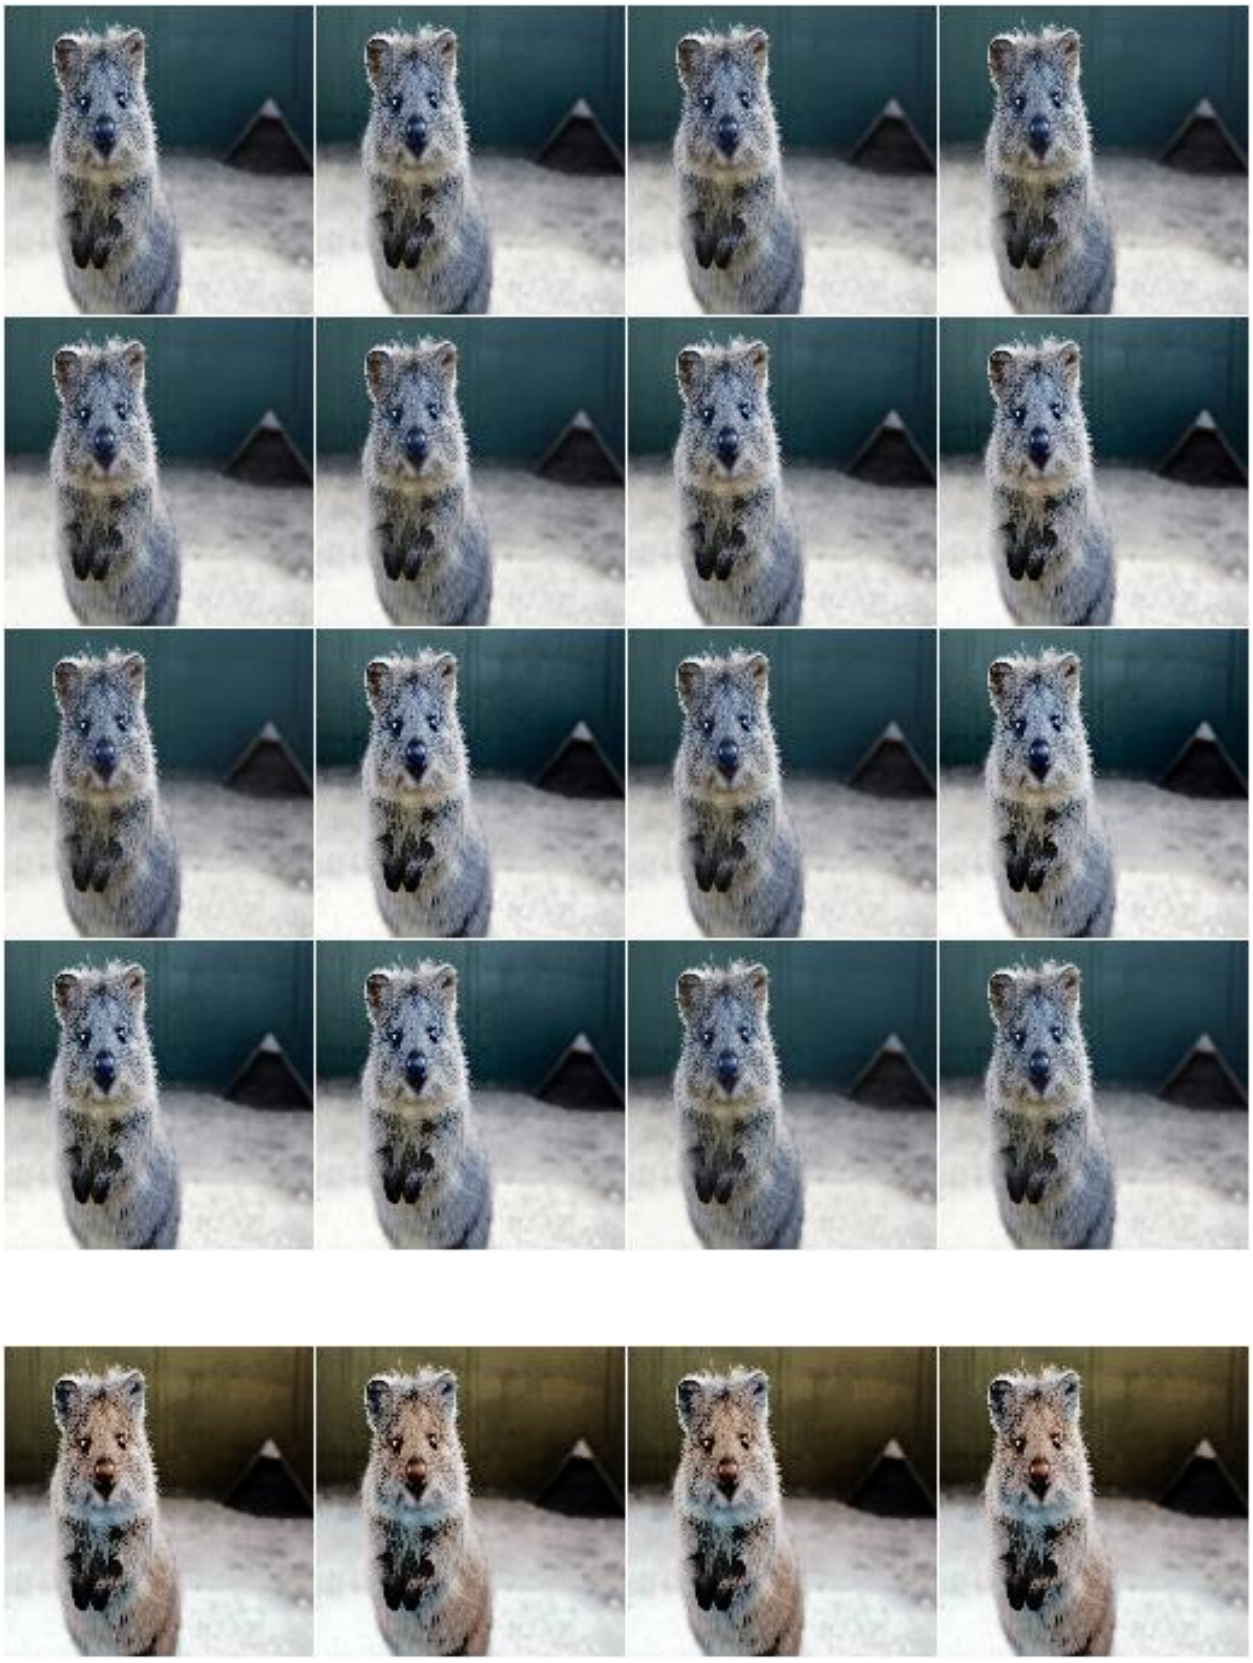

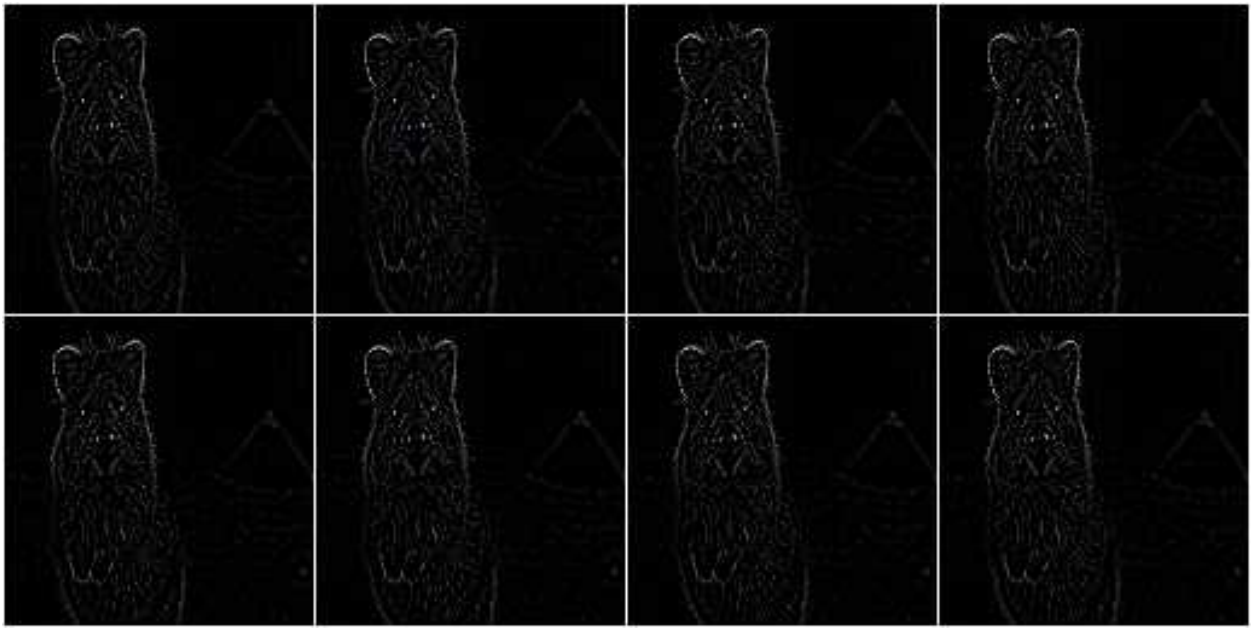

(continued from previous page)

```
        else:
            return [matrix_B] * nb_channels
aug = iaa.Convolve(matrix=gen_matrix)
```

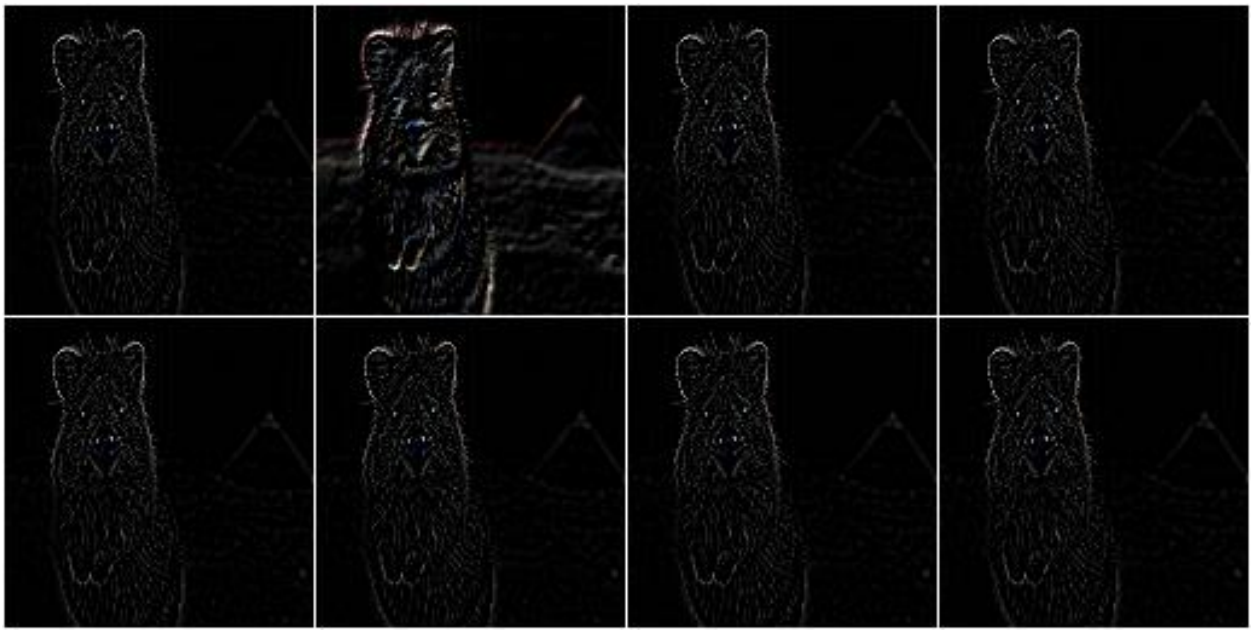

9.9.2 Sharpen

Augmenter that sharpens images and overlays the result with the original image.

API link: *Sharpen()*

Example. Sharpen an image, then overlay the results with the original using an alpha between 0.0 and 1.0:

```
import imgaug.augmenters as iaa
aug = iaa.Sharpen(alpha=(0.0, 1.0), lightness=(0.75, 2.0))
```

Example. Effects of keeping `lightness` fixed at `1.0` and then varying `alpha` between `0.0` and `1.0` in eight steps:

Example. Effects of keeping `alpha` fixed at `1.0` and then varying `lightness` between `0.75` and `1.5` in eight steps:

9.9.3 Emboss

Augmenter that embosses images and overlays the result with the original image.

API link: *Emboss()*

Example. Emboss an image, then overlay the results with the original using an alpha between `0.0` and `1.0`:

```
import imgaug.augmenters as iaa
aug = iaa.Emboss(alpha=(0.0, 1.0), strength=(0.5, 1.5))
```

Example. Effects of keeping `strength` fixed at `1.0` and then varying `alpha` between `0.0` and `1.0` in eight steps:

Example. Effects of keeping `alpha` fixed at `1.0` and then varying `strength` between `0.5` and `1.5` in eight steps:

9.9.4 EdgeDetect

Augmenter that detects all edges in images, marks them in a black and white image and then overlays the result with the original image.

API link: *EdgeDetect()*

Example. Detect edges in images, turning them into black and white images and then overlay these with the original images using random alphas between 0.0 and 1.0:

```
import imgaug.augmenters as iaa
aug = iaa.EdgeDetect(alpha=(0.0, 1.0))
```

Example. Effect of increasing `alpha` from `0.0` to `1.0` in eight steps:

9.9.5 DirectedEdgeDetect

Augmenter that detects edges that have certain directions and marks them in a black and white image and then overlays the result with the original image.

API link: *DirectedEdgeDetect()*

Example. Detect edges having random directions (0 to 360 degrees) in images, turning the images into black and white versions and then overlay these with the original images using random alphas between 0.0 and 1.0:

```
import imgaug.augmenters as iaa
aug = iaa.DirectedEdgeDetect(alpha=(0.0, 1.0), direction=(0.0, 1.0))
```

Example. Effect of fixing `direction` to 0.0 and then increasing `alpha` from `0.0` to `1.0` in eight steps:

Example. Effect of fixing `alpha` to `1.0` and then increasing `direction` from `0.0` to `1.0` (0 to 360 degrees) in eight steps:

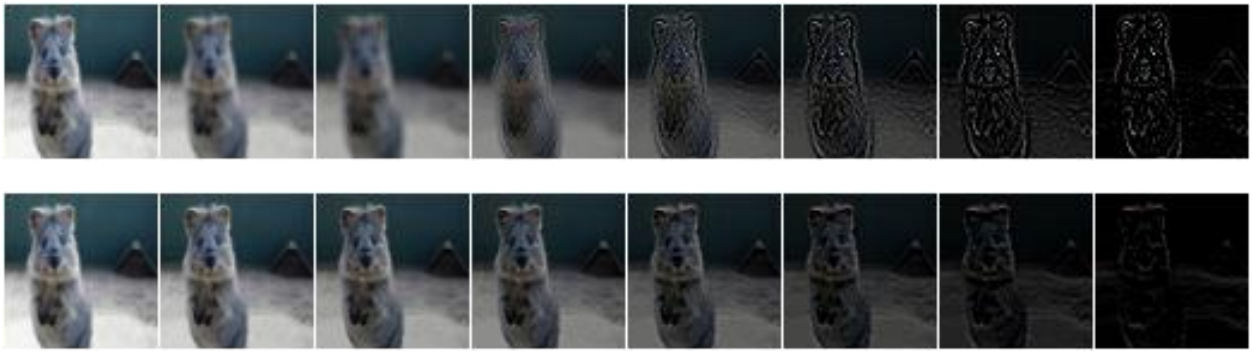

9.10 augmenters.debug

9.10.1 SaveDebugImageEveryNBatches

Visualize data in batches and save corresponding plots to a folder.

API link: *SaveDebugImageEveryNBatches()*

Example. Save a debug plot to a temporary folder every `100` batches. Set `folder_path` to a string, e.g. `/tmp/experiments/debug-images`, in order to save to that filepath instead of the temporary folder.

```
import imgaug.augmenters as iaa
import tempfile
with tempfile.TemporaryDirectory() as folder_path:
    seq = iaa.Sequential([
        iaa.Sequential([
            iaa.Fliplr(0.5),
            iaa.Crop(px=(0, 16))
        ], random_order=True),
        iaa.SaveDebugImageEveryNBatches(folder_path, 100)
    ])
```

9.11 augmenters.edges

9.11.1 Canny

Apply a canny edge detector to input images.

API link: *Canny*

Example. Create an augmenter that generates random blends between images and their canny edge representations:

```
import imgaug.augmenters as iaa
aug = iaa.Canny()
```

Images	SegMaps	Bounding Boxes
list of 4 arrays all shape (128, 128, 3) all dtype uint8 value range: 0 to 249	Channel 1 of 1 items for 4 images all arrays of shape (128, 128, 1) all on image shape (128, 128, 3) value range: 0 to 1 number of unique classes: 2	items for 4 images fewest items on image: 1 most items on image: 1 total items: 4 smallest area: 10466.9590 largest area: 11922.3184 unique labels: 1 most common labels: - None (4, 100.00%) coords out of image: 4 (50.00%) most extreme coord: (98.1, 144.1) all on image shape (128, 128, 3)

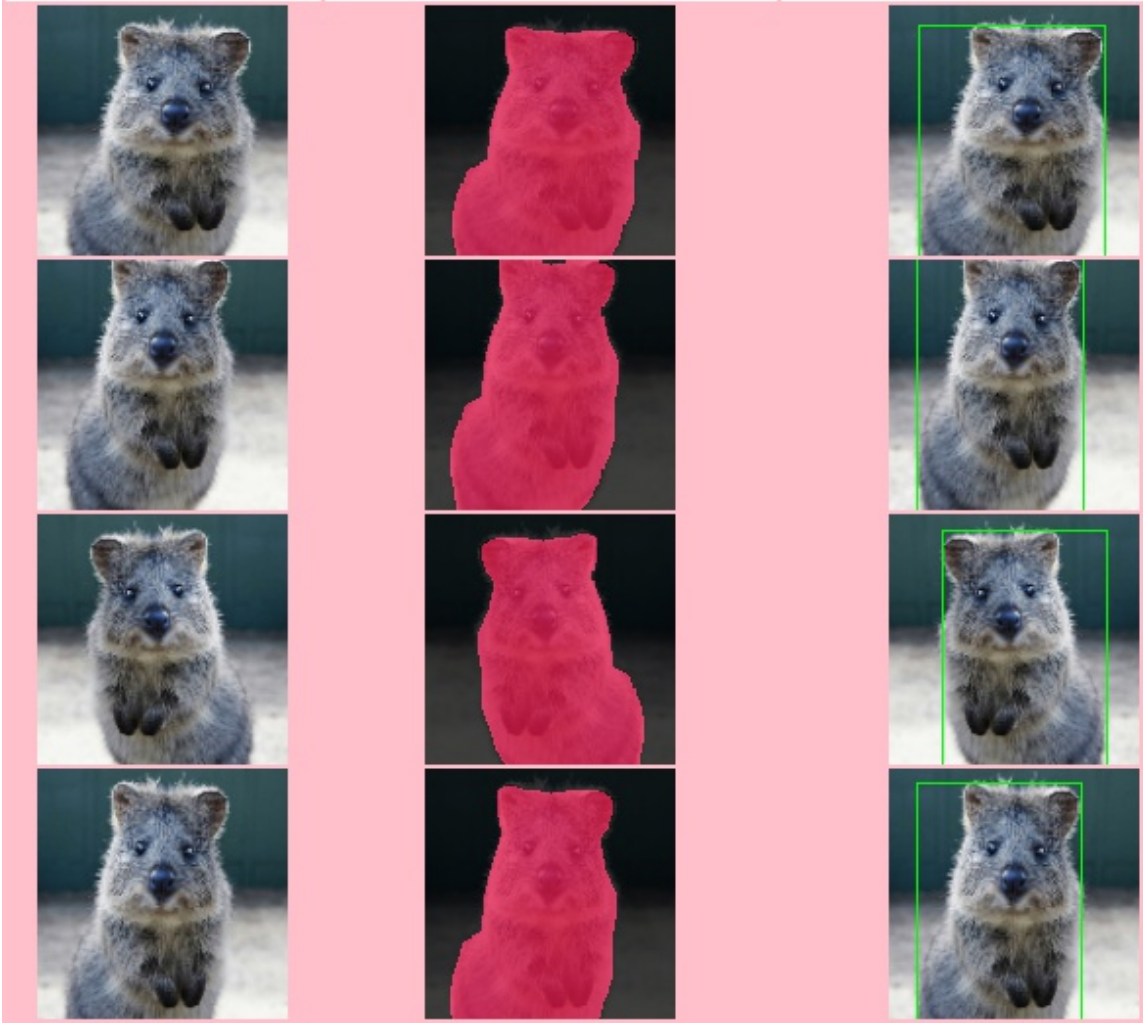

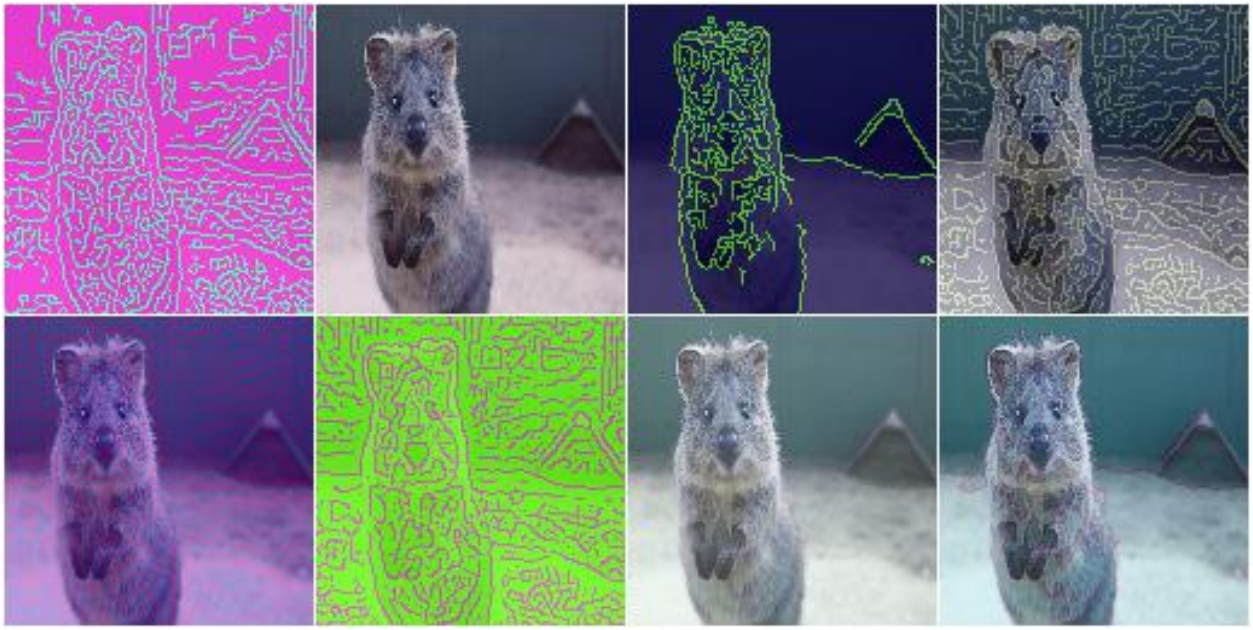

Example. Create a canny edge augmenter that generates edge images with a blending factor of max `50%`, i.e. the original (non-edge) image is always at least partially visible:

```
aug = iaa.Canny(alpha=(0.0, 0.5))
```

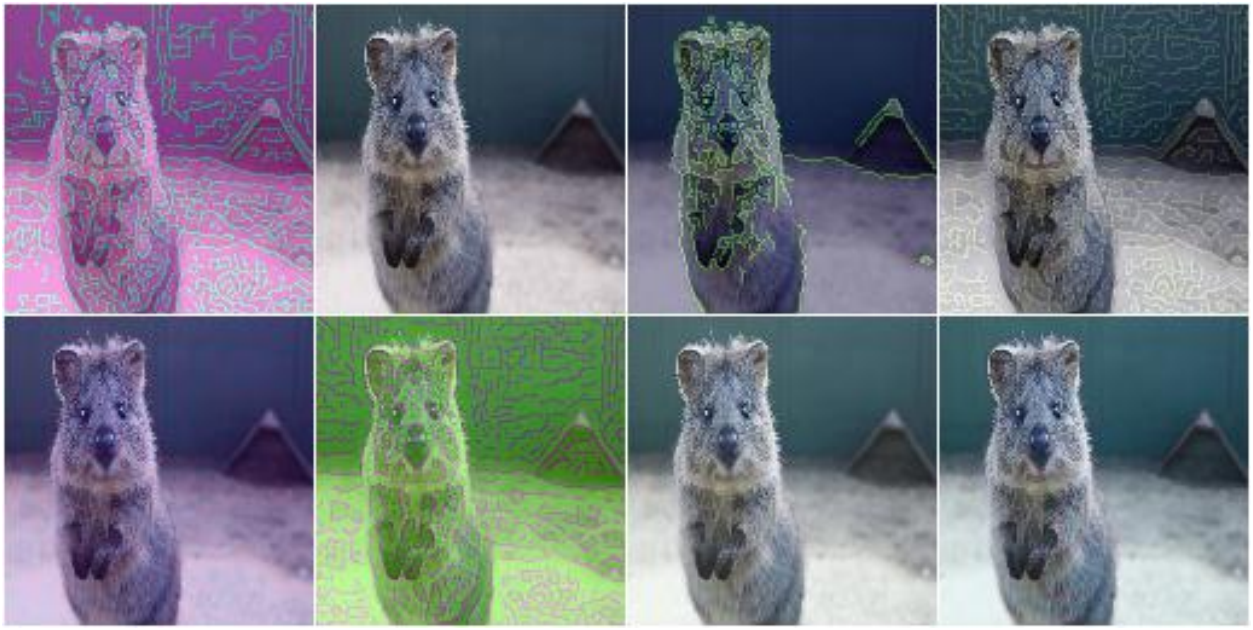

Example. Same as in the previous example, but the edge image always uses the color white for edges and black for the background:

```
aug = iaa.Canny(
    alpha=(0.0, 0.5),
    colorizer=iaa.RandomColorsBinaryImageColorizer(
        color_true=255,
```

(continues on next page)

(continued from previous page)

```
        color_false=0
    )
)
```

Example. Create a canny edge augmenter that initially preprocesses images using a sobel filter with kernel size of either `3x3` or `13x13` and alpha-blends with result using a strength of `50%` (both images equally visible) to `100%` (only edge image visible).

```
aug = iaa.Canny(alpha=(0.5, 1.0), sobel_kernel_size=[3, 7])
```

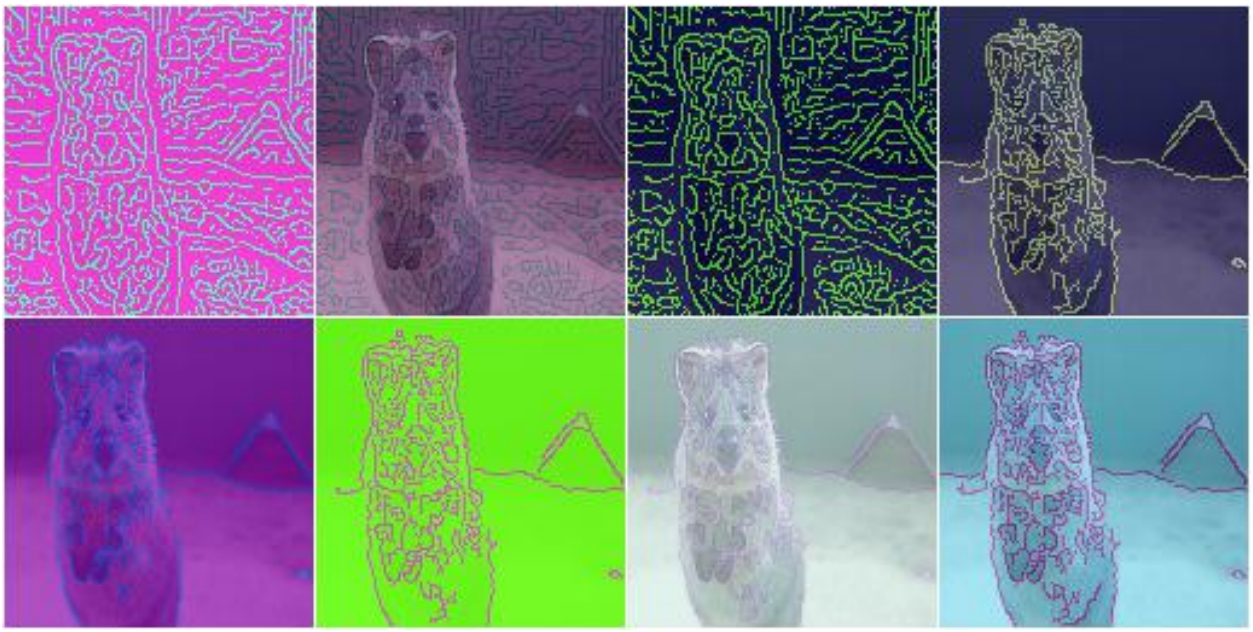

Example. Create an augmenter that blends a canny edge image with a median-blurred version of the input image. The median blur uses a fixed kernel size of `13x13` pixels.

```
aug = iaa.Alpha(
    (0.0, 1.0),
    iaa.Canny(alpha=1),
    iaa.MedianBlur(13)
)
```

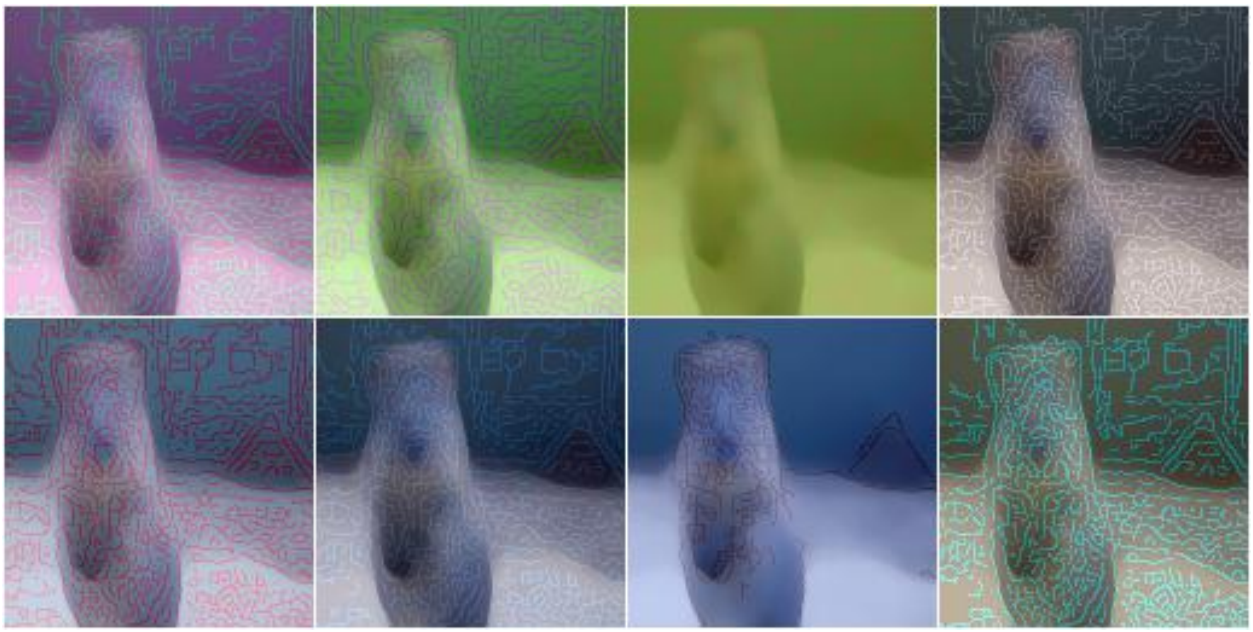

9.12 augmenters.flip

9.12.1 HorizontalFlip

Alias for `Fliplr`.

API link: *HorizontalFlip*

9.12.2 VericalFlip

Alias for `Flipud`.

API link: *VerticalFlip*

9.12.3 Fliplr

Flip/mirror input images horizontally.

Note: The default value for the probability is `1.0`, i.e. all images will be flipped.

API link: *Fliplr*

Example. Flip 50% of all images horizontally:

```
import imgaug.augmenters as iaa
aug = iaa.Fliplr(0.5)
```

9.12.4 Flipud

Flip/mirror input images vertically.

Note: The default value for the probability is `1.0`, i.e. all images will be flipped.

API link: *Flipud*

Example. Flip 50% of all images vertically:

```
aug = iaa.Flipud(0.5)
```

9.13 augmenters.geometric

9.13.1 Affine

Augmenter to apply affine transformations to images.

API link: *Affine*

Example. Scale images to a value of 50 to 150% of their original size:

```
import imgaug.augmenters as iaa
aug = iaa.Affine(scale=(0.5, 1.5))
```

Example. Scale images to a value of 50 to 150% of their original size, but do this independently per axis (i.e. sample two values per image):

```
aug = iaa.Affine(scale={"x": (0.5, 1.5), "y": (0.5, 1.5)})
```

Example. Translate images by -20 to +20% on x- and y-axis independently:

```
aug = iaa.Affine(translate_percent={"x": (-0.2, 0.2), "y": (-0.2, 0.2)})
```

Example. Translate images by -20 to 20 pixels on x- and y-axis independently:

```
aug = iaa.Affine(translate_px={"x": (-20, 20), "y": (-20, 20)})
```

Example. Rotate images by -45 to 45 degrees:

```
aug = iaa.Affine(rotate=(-45, 45))
```

Example. Shear images by -16 to 16 degrees:

```
aug = iaa.Affine(shear=(-16, 16))
```

Example. When applying affine transformations, new pixels are often generated, e.g. when translating to the left, pixels are generated on the right. Various modes exist to set how these pixels are ought to be filled. Below code shows an example that uses all modes, sampled randomly per image. If the mode is `constant` (fill all with one constant value), then a random brightness between 0 and 255 is used:

```
aug = iaa.Affine(translate_percent={"x": -0.20}, mode=ia.ALL, cval=(0, 255))
```

9.13.2 ScaleX

Apply affine scaling on the x-axis to input data.

This is a wrapper around *imgaug.augmenters.geometric.Affine*.

API link: *ScaleX*

Example. Create an augmenter that scales images along the width to sizes between `50%` and `150%`. This does not change the image shape (i.e. height and width), only the pixels within the image are remapped and potentially new ones are filled in.

```
import imgaug.augmenters as iaa
aug = iaa.ScaleX((0.5, 1.5))
```

9.13.3 ScaleY

Apply affine scaling on the y-axis to input data.

This is a wrapper around *imgaug.augmenters.geometric.Affine*.

API link: *ScaleY*

Example. Create an augmenter that scales images along the height to sizes between `50%` and `150%`. This does not change the image shape (i.e. height and width), only the pixels within the image are remapped and potentially new ones are filled in.

```
import imgaug.augmenters as iaa
aug = iaa.ScaleY((0.5, 1.5))
```

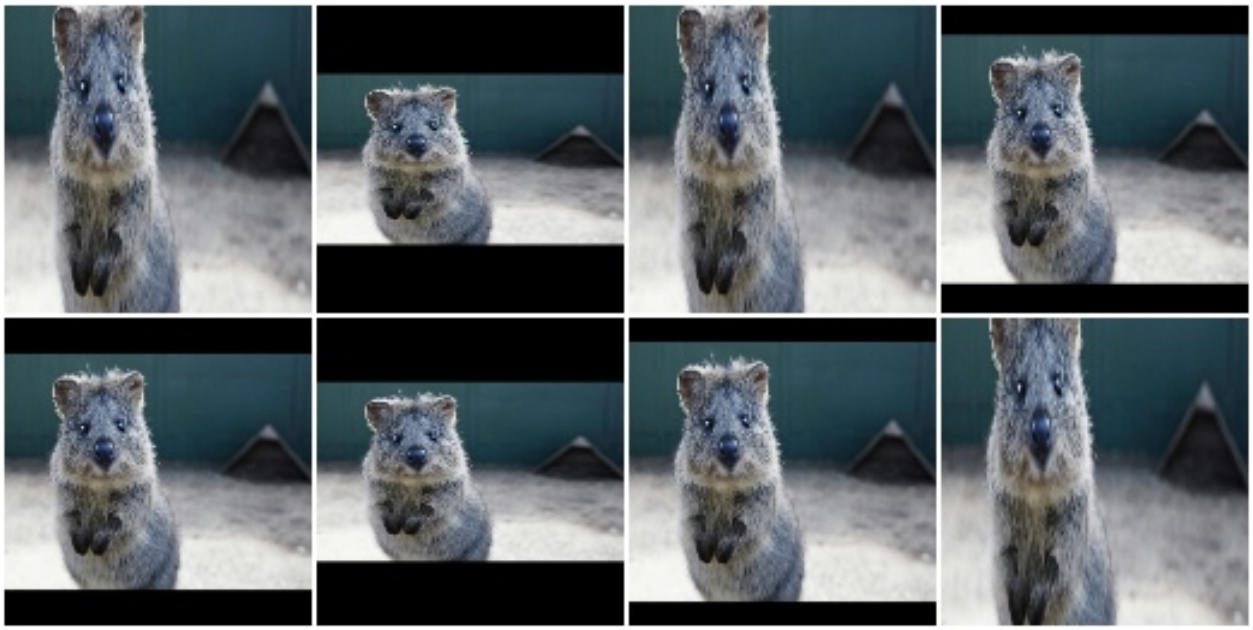

9.13.4 TranslateX

Apply affine translation on the x-axis to input data.

This is a wrapper around *imgaug.augmenters.geometric.Affine*.

API link: *TranslateX*

Example. Create an augmenter that translates images along the x-axis by `-20` to `20` pixels:

```
import imgaug.augmenters as iaa
aug = iaa.TranslateX(px=(-20, 20))
```

Example. Create an augmenter that translates images along the x-axis by `-10%` to `10%` (relative to the x-axis size):

```
aug = iaa.TranslateX(percent=(-0.1, 0.1))
```

9.13.5 TranslateY

Apply affine translation on the y-axis to input data.

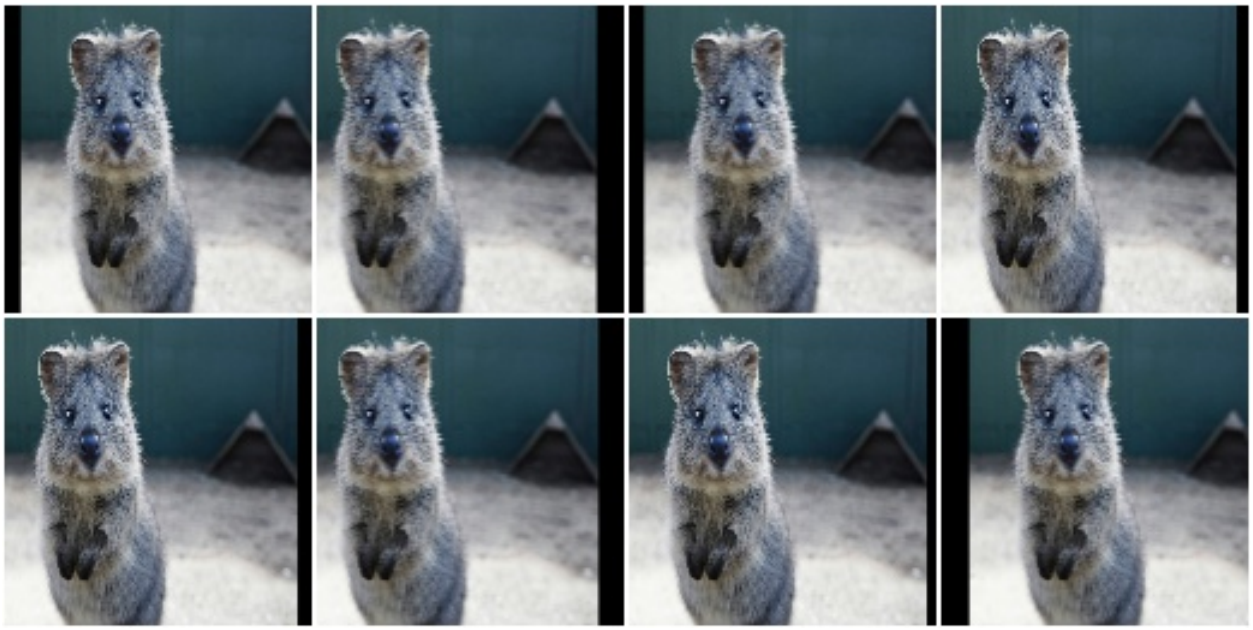

This is a wrapper around *imgaug.augmenters.geometric.Affine*.

API link: *TranslateY*

Example. Create an augmenter that translates images along the y-axis by `-20` to `20` pixels:

```
import imgaug.augmenters as iaa
aug = iaa.TranslateY(px=(-20, 20))
```

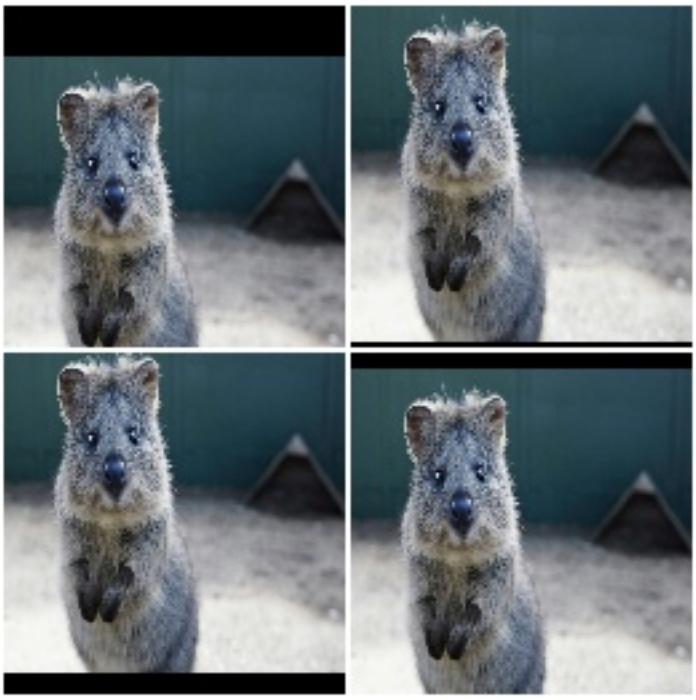

Example. Create an augmenter that translates images along the y-axis by `-10%` to `10%` (relative to the y-axis size):

```
aug = iaa.TranslateY(percent=(-0.1, 0.1))
```

9.13.6 Rotate

Apply affine rotation on the y-axis to input data.

This is a wrapper around *imgaug.augmenters.geometric.Affine*. It is the same as `Affine(rotate=<value>)`.

API link: *Rotate*

Example. Create an augmenter that rotates images by a random value between `-45` and `45` degress:

```
import imgaug.augmenters as iaa
aug = iaa.Rotate((-45, 45))
```

9.13.7 ShearX

Apply affine shear on the x-axis to input data.

This is a wrapper around `Affine`.

API link: *ShearX*

Example. Create an augmenter that shears images along the x-axis by random amounts between `-20` and `20` degrees:

```
import imgaug.augmenters as iaa
aug = iaa.ShearX((-20, 20))
```

9.13.8 ShearY

Apply affine shear on the y-axis to input data.

This is a wrapper around `Affine`.

API link: *ShearY*

Example. Create an augmenter that shears images along the y-axis by random amounts between `-20` and `20` degrees:

```
import imgaug.augmenters as iaa
aug = iaa.ShearY((-20, 20))
```

9.13.9 PiecewiseAffine

Apply affine transformations that differ between local neighbourhoods.

This augmenter places a regular grid of points on an image and randomly moves the neighbourhood of these point around via affine transformations. This leads to local distortions.

This is mostly a wrapper around scikit-image's `PiecewiseAffine`. See also `Affine` for a similar technique.

Note: This augmenter is very slow. See *Performance*. Try to use `ElasticTransformation` instead, which is at least 10x faster.

Note: For coordinate-based inputs (keypoints, bounding boxes, polygons, ...), this augmenter still has to perform an image-based augmentation, which will make it significantly slower for such inputs than other augmenters. See *Performance*.

API link: *PiecewiseAffine*

Example. Distort images locally by moving points around, each with a distance v (percent relative to image size), where v is sampled per point from `N(0, z)` `z` is sampled per image from the range 0.01 to 0.05:

```
import imgaug.augmenters as iaa
aug = iaa.PiecewiseAffine(scale=(0.01, 0.05))
```

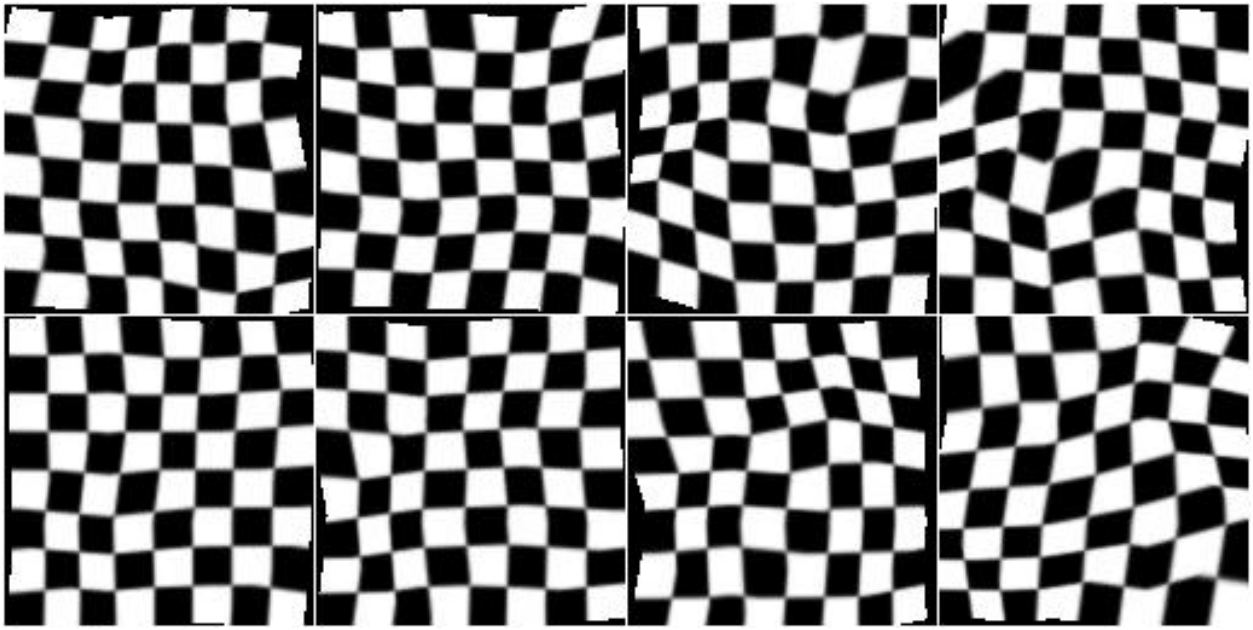

Example. Effect of increasing `scale` from `0.01` to `0.3` in eight steps:

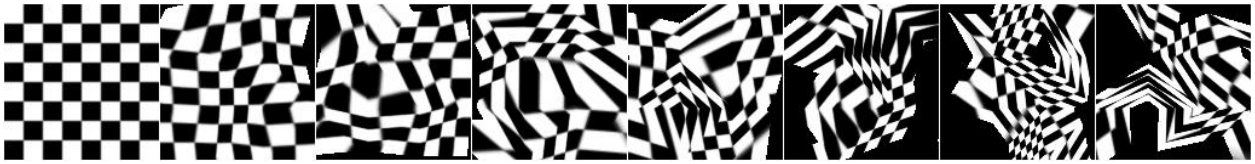

Example. PiecewiseAffine works by placing a regular grid of points on the image and moving them around. By default this grid consists of 4x4 points. The below image shows the effect of increasing that value from 2x2 to 16x16 in 8 steps:

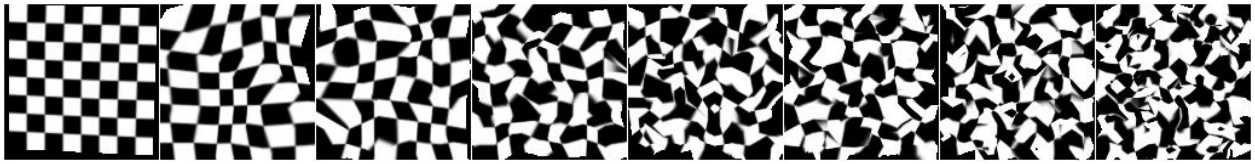

9.13.10 PerspectiveTransform

Apply random four point perspective transformations to images.

Each of the four points is placed on the image using a random distance from its respective corner. The distance is sampled from a normal distribution. As a result, most transformations don't change the image very much, while some "focus" on polygons far inside the image.

The results of this augmenter have some similarity with `Crop`.

API link: *`PerspectiveTransform`*

Example. Apply perspective transformations using a random scale between `0.01` and `0.15` per image, where the scale is roughly a measure of how far the perspective transformation's corner points may be distanced from the image's corner points:

```
import imgaug.augmenters as iaa
aug = iaa.PerspectiveTransform(scale=(0.01, 0.15))
```

Example. Same as in the previous example, but images are not resized back to the input image size after augmentation. This will lead to smaller output images.

```
aug = iaa.PerspectiveTransform(scale=(0.01, 0.15), keep_size=False)
```

9.13.11 ElasticTransformation

Transform images by moving pixels locally around using displacement fields.

The augmenter has the parameters `alpha` and `sigma`. `alpha` controls the strength of the displacement: higher values mean that pixels are moved further. `sigma` controls the smoothness of the displacement: higher values lead to smoother patterns – as if the image was below water – while low values will cause indivdual pixels to be moved very differently from their neighbours, leading to noisy and pixelated images.

A relation of 10:1 seems to be good for `alpha` and `sigma`, e.g. `alpha=10` and `sigma=1` or `alpha=50`, `sigma=5`. For `128x128` a setting of `alpha=(0, 70.0)`, `sigma=(4.0, 6.0)` may be a good choice and will lead to a water-like effect.

For a detailed explanation, see

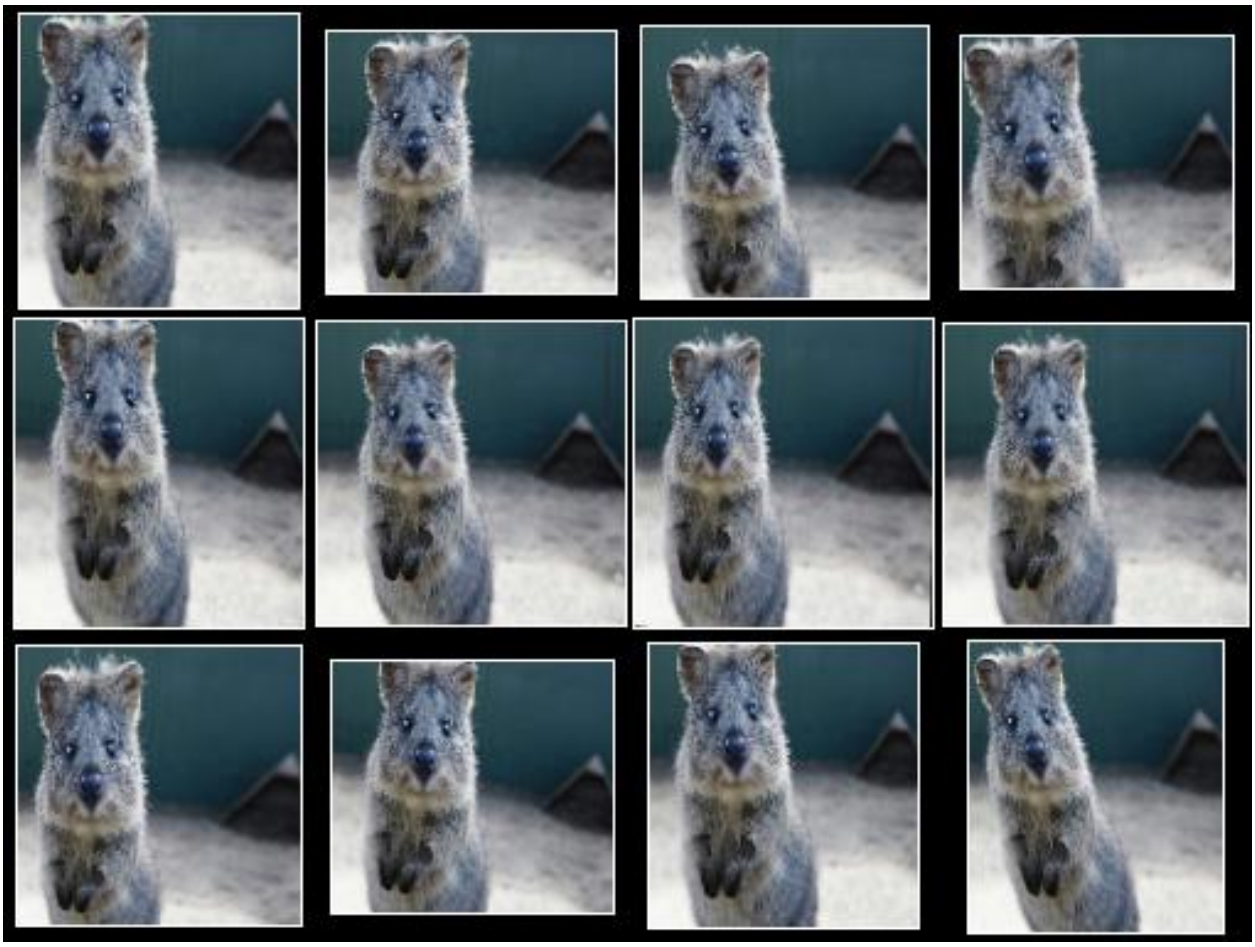

Fig. 1: `PerspectiveTransform` with `keep_size` set to `False`. Note that the individual images are here padded after augmentation in order to align them in a grid (i.e. purely for visualization purposes).

```
Simard, Steinkraus and Platt
Best Practices for Convolutional Neural Networks applied to Visual
Document Analysis
in Proc. of the International Conference on Document Analysis and
Recognition, 2003
```

Note: For coordinate-based inputs (keypoints, bounding boxes, polygons, ...), this augmenter still has to perform an image-based augmentation, which will make it significantly slower for such inputs than other augmenters. See *Performance*.

API link: *ElasticTransformation*

Example. Distort images locally by moving individual pixels around following a distortions field with strength 0.25. The strength of the movement is sampled per pixel from the range 0 to 5.0:

```
import imgaug.augmenters as iaa
aug = iaa.ElasticTransformation(alpha=(0, 5.0), sigma=0.25)
```

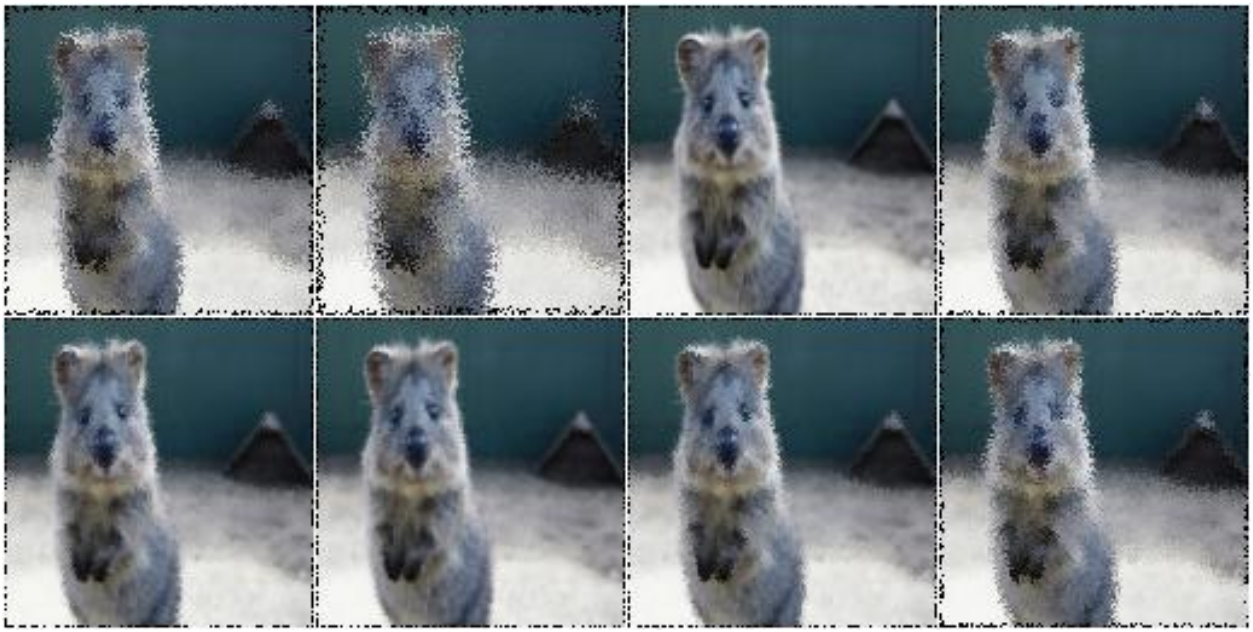

Example. Effect of keeping sigma fixed at `0.25` and increasing alpha from `0` to `5.0` in eight steps:

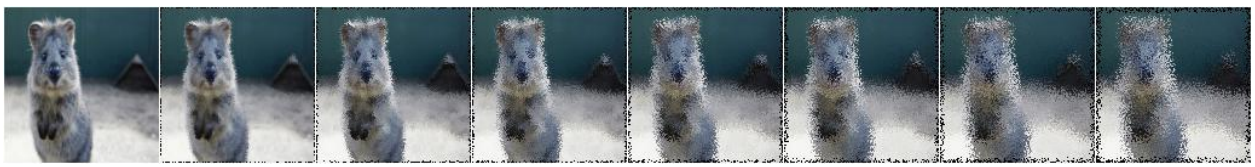

Example. Effect of keeping alpha fixed at `2.5` and increasing sigma from `0.01` to `1.0` in eight steps:

9.13.12 Rot90

Rotate images clockwise by multiples of 90 degrees.

This could also be achieved using `Affine`, but `Rot90` is significantly more efficient.

API link: *Rot90*

Fig. 2: The below examples use this input image, which slightly deviates from the examples for other augmenters (i.e. it is not square).

Example. Rotate all images by 90 degrees. Resize these images afterwards to keep the size that they had before augmentation. This may cause the images to look distorted.

```
import imgaug.augmenters as iaa
aug = iaa.Rot90(1)
```

Example. Rotate all images by 90 or 270 degrees. Resize these images afterwards to keep the size that they had before augmentation. This may cause the images to look distorted.

```
aug = iaa.Rot90([1, 3])
```

Example. Rotate all images by 90, 180 or 270 degrees. Resize these images afterwards to keep the size that they had before augmentation. This may cause the images to look distorted.

```
aug = iaa.Rot90((1, 3))
```

Example. Rotate all images by 90, 180 or 270 degrees. Does not resize to the original image size afterwards, i.e. each image's size may change.

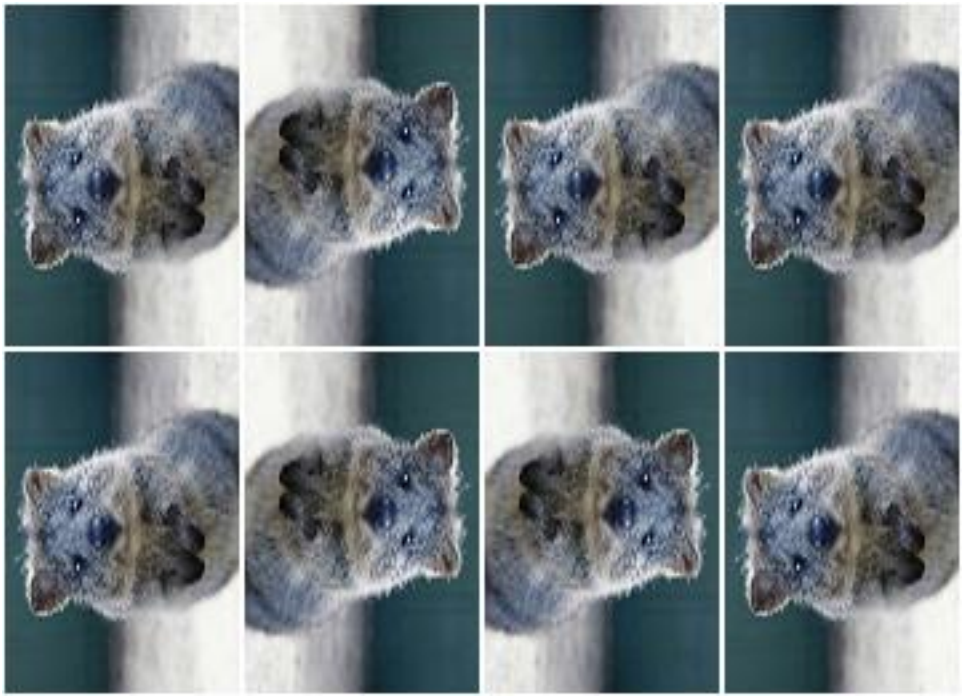

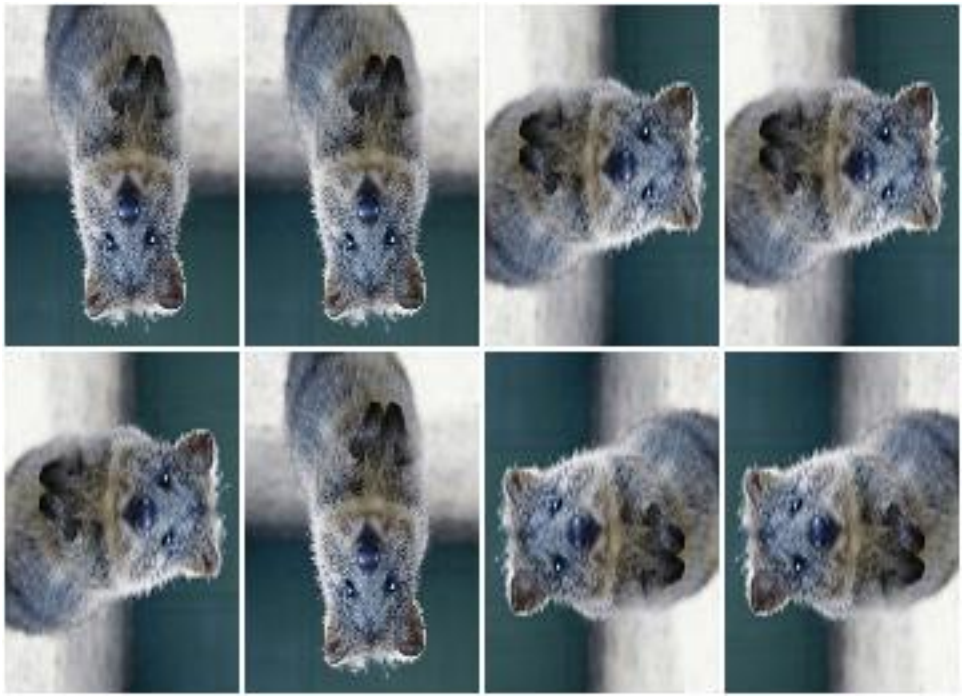

```
aug = iaa.Rot90((1, 3), keep_size=False)
```

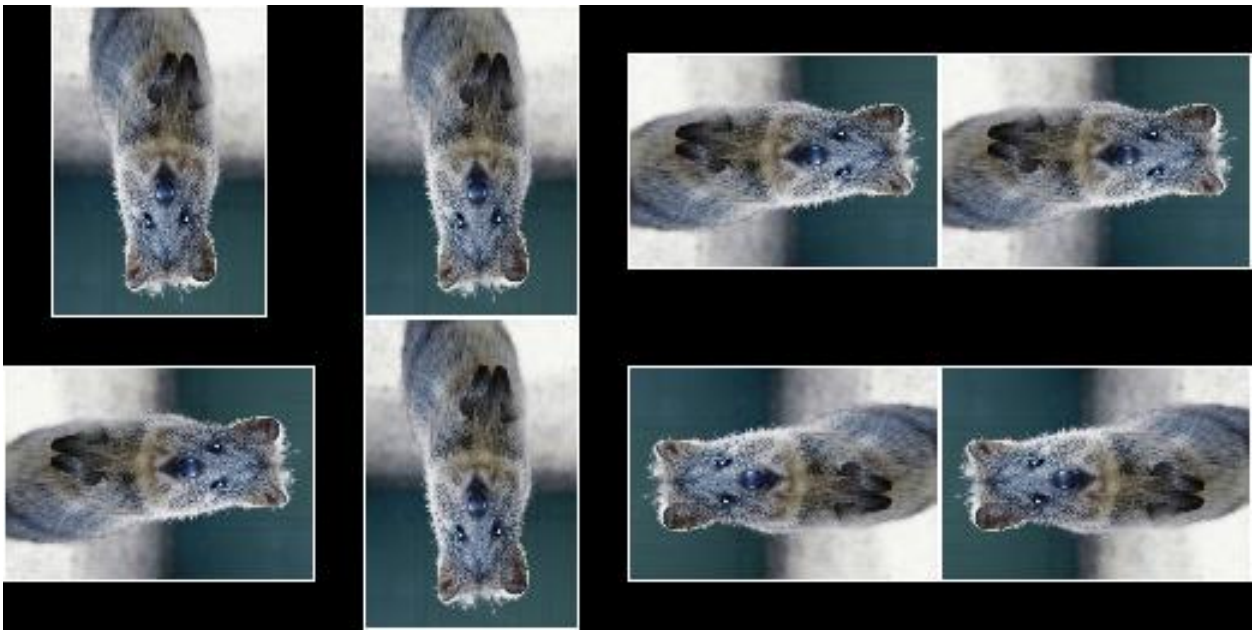

Fig. 3: `Rot90` with `keep_size` set to `False`. Note that the individual images are here padded after augmentation in order to align them in a grid (i.e. purely for visualization purposes).

9.13.13 WithPolarWarping

Augmenter that applies other augmenters in a polar-transformed space.

This augmenter first transforms an image into a polar representation, then applies its child augmenter, then transforms back to cartesian space. The polar representation is still in the image's input dtype (i.e. `uint8` stays `uint8`) and can be visualized. It can be thought of as an "unrolled" version of the image, where previously circular lines appear straight. Hence, applying child augmenters in that space can lead to circular effects. E.g. replacing rectangular pixel areas in the polar representation with black pixels will lead to curved black areas in the cartesian result.

This augmenter can create new pixels in the image. It will fill these with black pixels. For segmentation maps it will fill with class id `0`. For heatmaps it will fill with `0.0`.

This augmenter is limited to arrays with a height and/or width of `32767` or less.

Warning: When augmenting coordinates in polar representation, it is possible that these are shifted outside of the polar image, but are inside the image plane after transforming back to cartesian representation, usually on newly created pixels (i.e. black backgrounds). These coordinates are currently not removed. It is recommended to not use very strong child transformations when also augmenting coordinate-based augmentables.

Warning: For bounding boxes, this augmenter suffers from the same problem as affine rotations applied to bounding boxes, i.e. the resulting bounding boxes can have unintuitive (seemingly wrong) appearance. This is due to coordinates being "rotated" that are inside the bounding box, but do not fall on the object and actually are background. It is recommended to use this augmenter with caution when augmenting bounding boxes.

Warning: For polygons, this augmenter should not be combined with augmenters that perform automatic polygon recovery for invalid polygons, as the polygons will frequently appear broken in polar representation and their "fixed" version will be very broken in cartesian representation. Augmenters that perform such polygon recovery are currently `PerspectiveTransform`, `PiecewiseAffine` and `ElasticTransformation`.

API link: *WithPolarWarping*

Example. Apply cropping and padding in polar representation, then warp back to cartesian representation:

```
import imgaug.augmenters as iaa
aug = iaa.WithPolarWarping(iaa.CropAndPad(percent=(-0.1, 0.1)))
```

Example. Apply affine translations in polar representation:

```
aug = iaa.WithPolarWarping(
    iaa.Affine(
        translate_percent={"x": (-0.1, 0.1), "y": (-0.1, 0.1)}
    )
)
```

Example. Apply average pooling in polar representation. This leads to circular bins:

```
aug = iaa.WithPolarWarping(iaa.AveragePooling((2, 8)))
```

9.13.14 Jigsaw

Move cells within images similar to jigsaw patterns.

Note: This augmenter will by default pad images until their height is a multiple of *nb_rows*. Analogous for *nb_cols*.

Note: This augmenter will resize heatmaps and segmentation maps to the image size, then apply similar padding as for the corresponding images and resize back to the original map size. That also means that images may change in shape (due to padding), but heatmaps/segmaps will not change. For heatmaps/segmaps, this deviates from pad augmenters that will change images and heatmaps/segmaps in corresponding ways and then keep the heatmaps/segmaps at the new size.

Warning: This augmenter currently only supports augmentation of images, heatmaps, segmentation maps and keypoints. Other augmentables, i.e. bounding boxes, polygons and line strings, will result in errors.

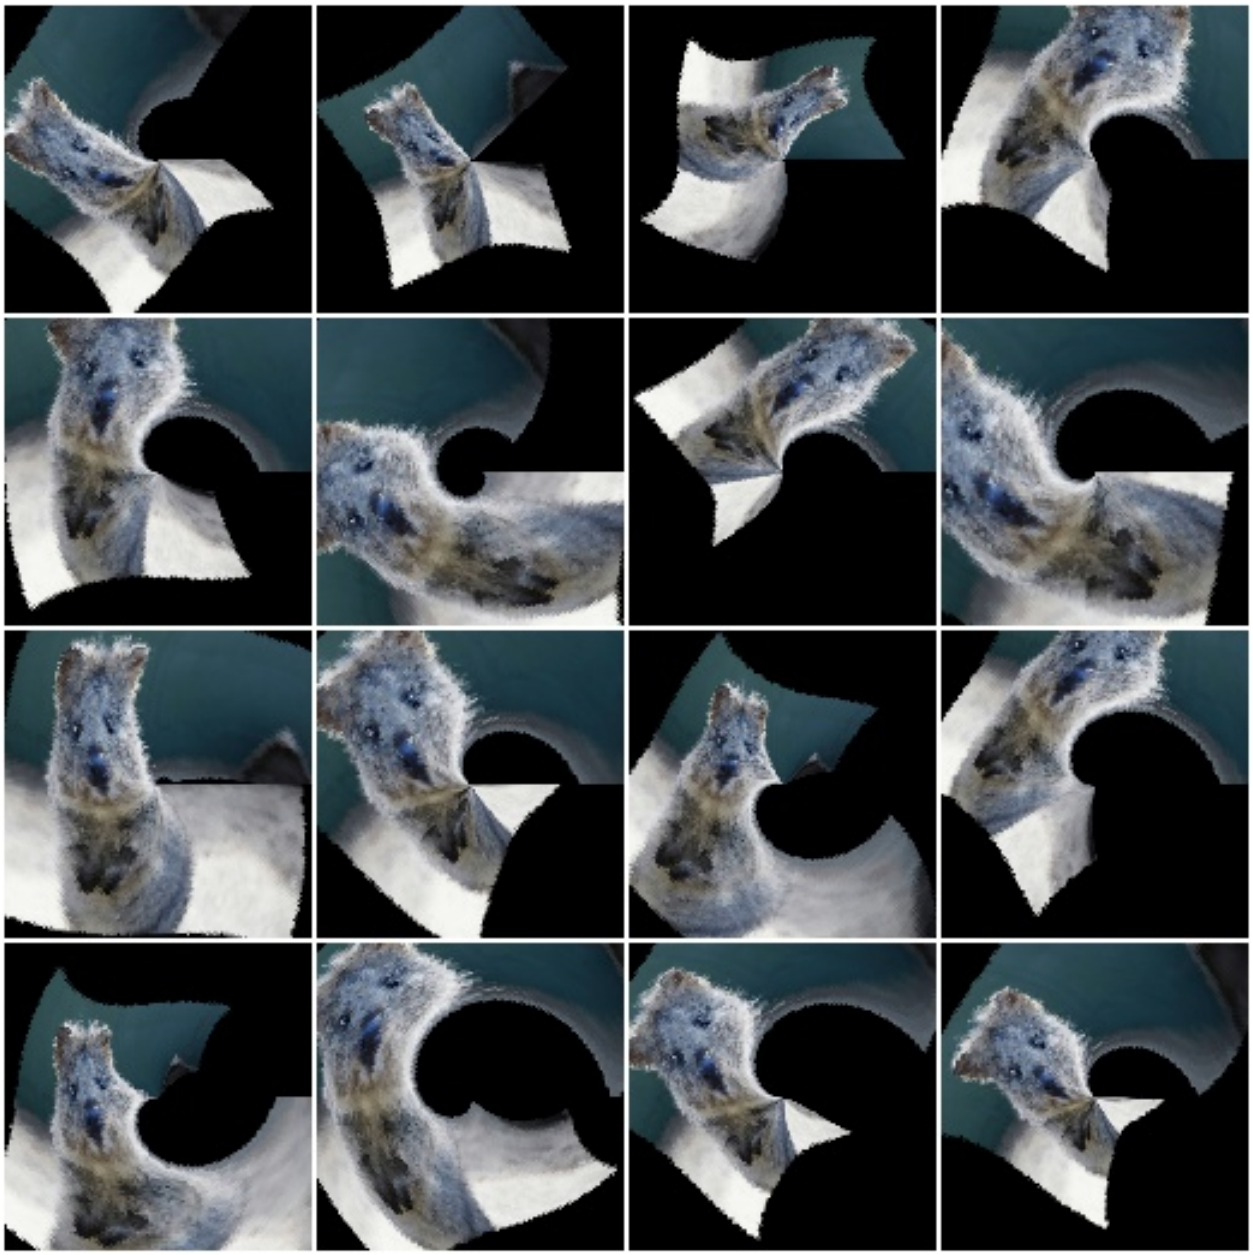

API link: *Jigsaw*

Example. Create a jigsaw augmenter that splits images into `10x10` cells and shifts them around by `0` to `2` steps (default setting):

```
import imgaug.augmenters as iaa
aug = iaa.Jigsaw(nb_rows=10, nb_cols=10)
```

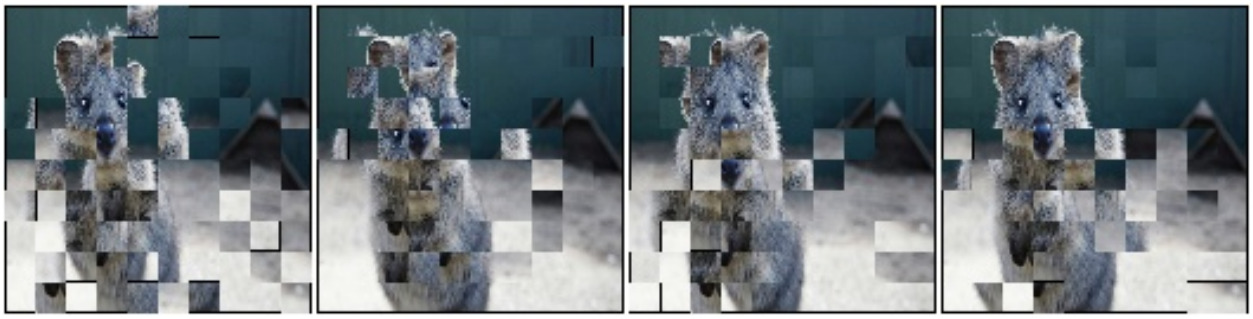

Example. Create a jigsaw augmenter that splits each image into `1` to `4` cells along each axis:

```
aug = iaa.Jigsaw(nb_rows=(1, 4), nb_cols=(1, 4))
```

Example. Create a jigsaw augmenter that moves the cells in each image by a random amount between `1` and `5` times (decided per image). Some images will be barely changed, some will be fairly distorted.

```
aug = iaa.Jigsaw(nb_rows=10, nb_cols=10, max_steps=(1, 5))
```

9.14 augmenters.imgcorruptlike

9.14.1 GaussianNoise

Wrapper around `gaussian_noise()`.

Note: This augmenter only affects images. Other data is not changed.

API link: *GaussianNoise*

The image below visualizes severities `1` to `5` – one severity per row:

Example. Create an augmenter around `gaussian_noise()`. Apply it to images using e.g. `aug(images=[image1, image2, ...])`:

```
import imgaug.augmenters as iaa
aug = iaa.imgcorruptlike.GaussianNoise(severity=2)
```

9.14.2 ShotNoise

Wrapper around `shot_noise()`.

Note: This augmenter only affects images. Other data is not changed.

API link: *ShotNoise*

The image below visualizes severities `1` to `5` – one severity per row:

Example. Create an augmenter around `shot_noise()`. Apply it to images using e.g. `aug(images=[image1, image2, ...])`:

```
import imgaug.augmenters as iaa
aug = iaa.imgcorruptlike.ShotNoise(severity=2)
```

9.14.3 ImpulseNoise

Wrapper around `impulse_noise()`.

Note: This augmenter only affects images. Other data is not changed.

API link: *ImpulseNoise*

The image below visualizes severities `1` to `5` – one severity per row:

Example. Create an augmenter around `impulse_noise()`. Apply it to images using e.g. `aug(images=[image1, image2, ...])`:

```
import imgaug.augmenters as iaa
aug = iaa.imgcorruptlike.ImpulseNoise(severity=2)
```

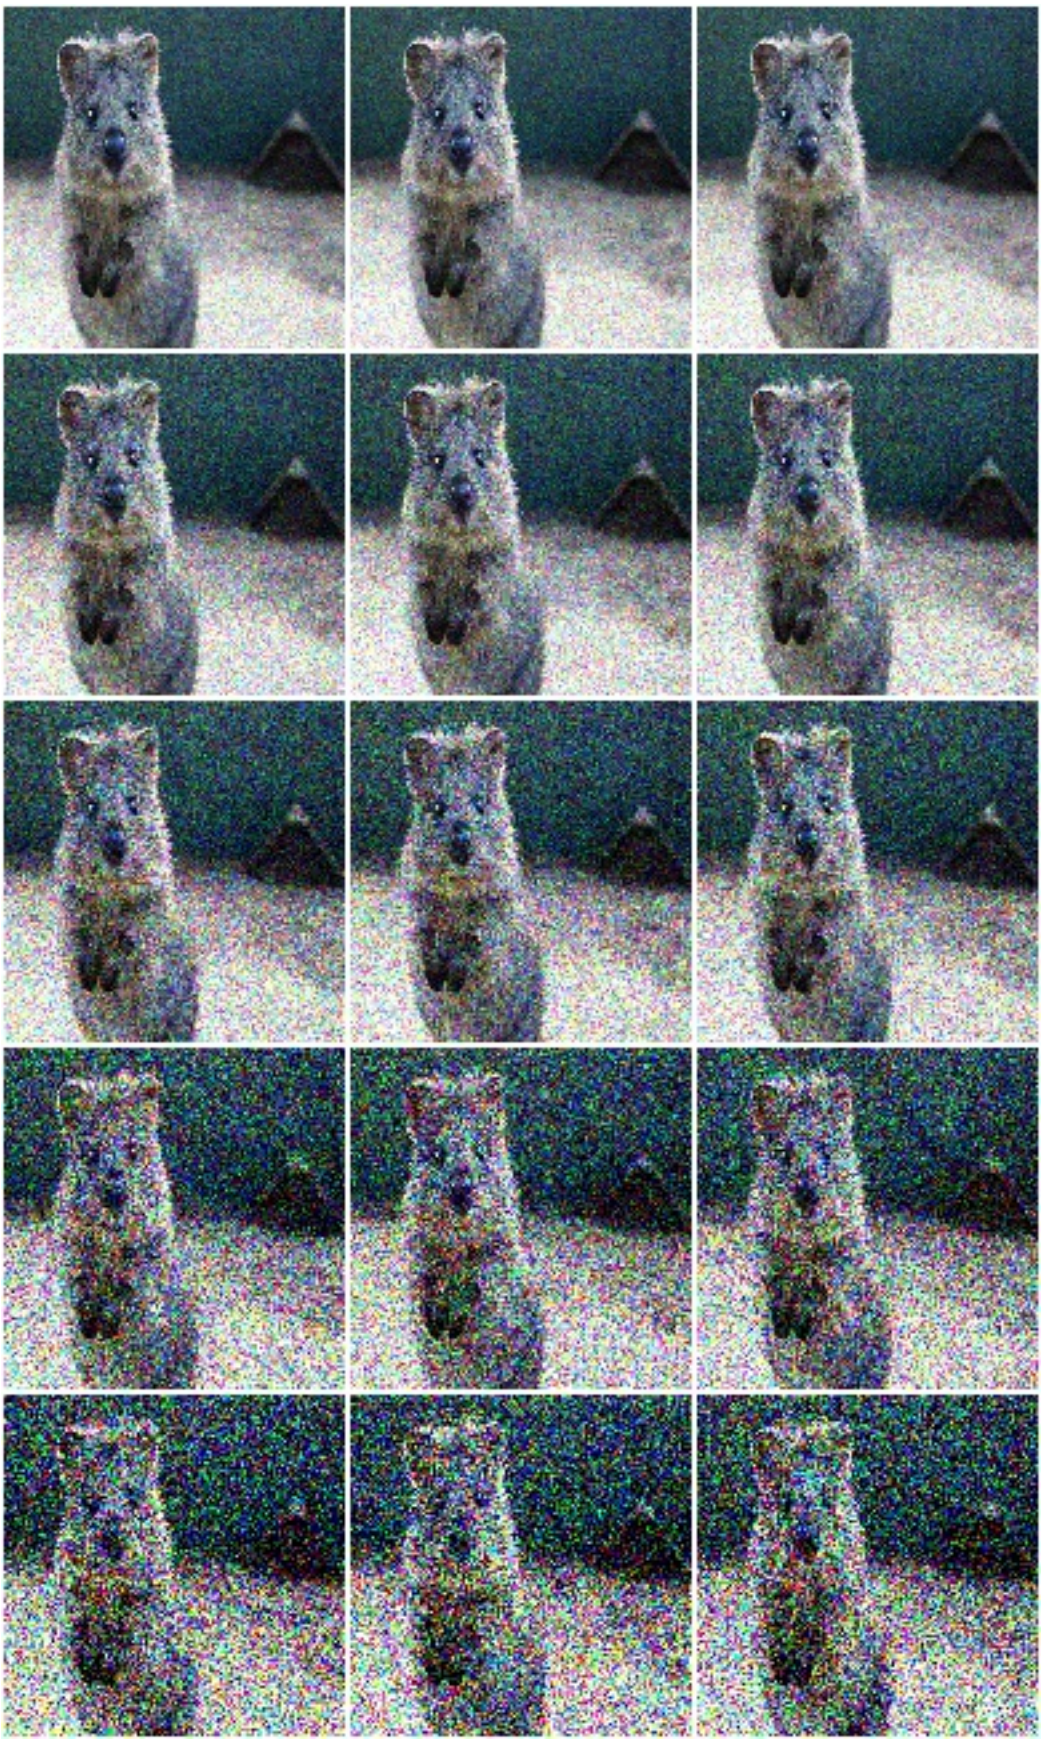

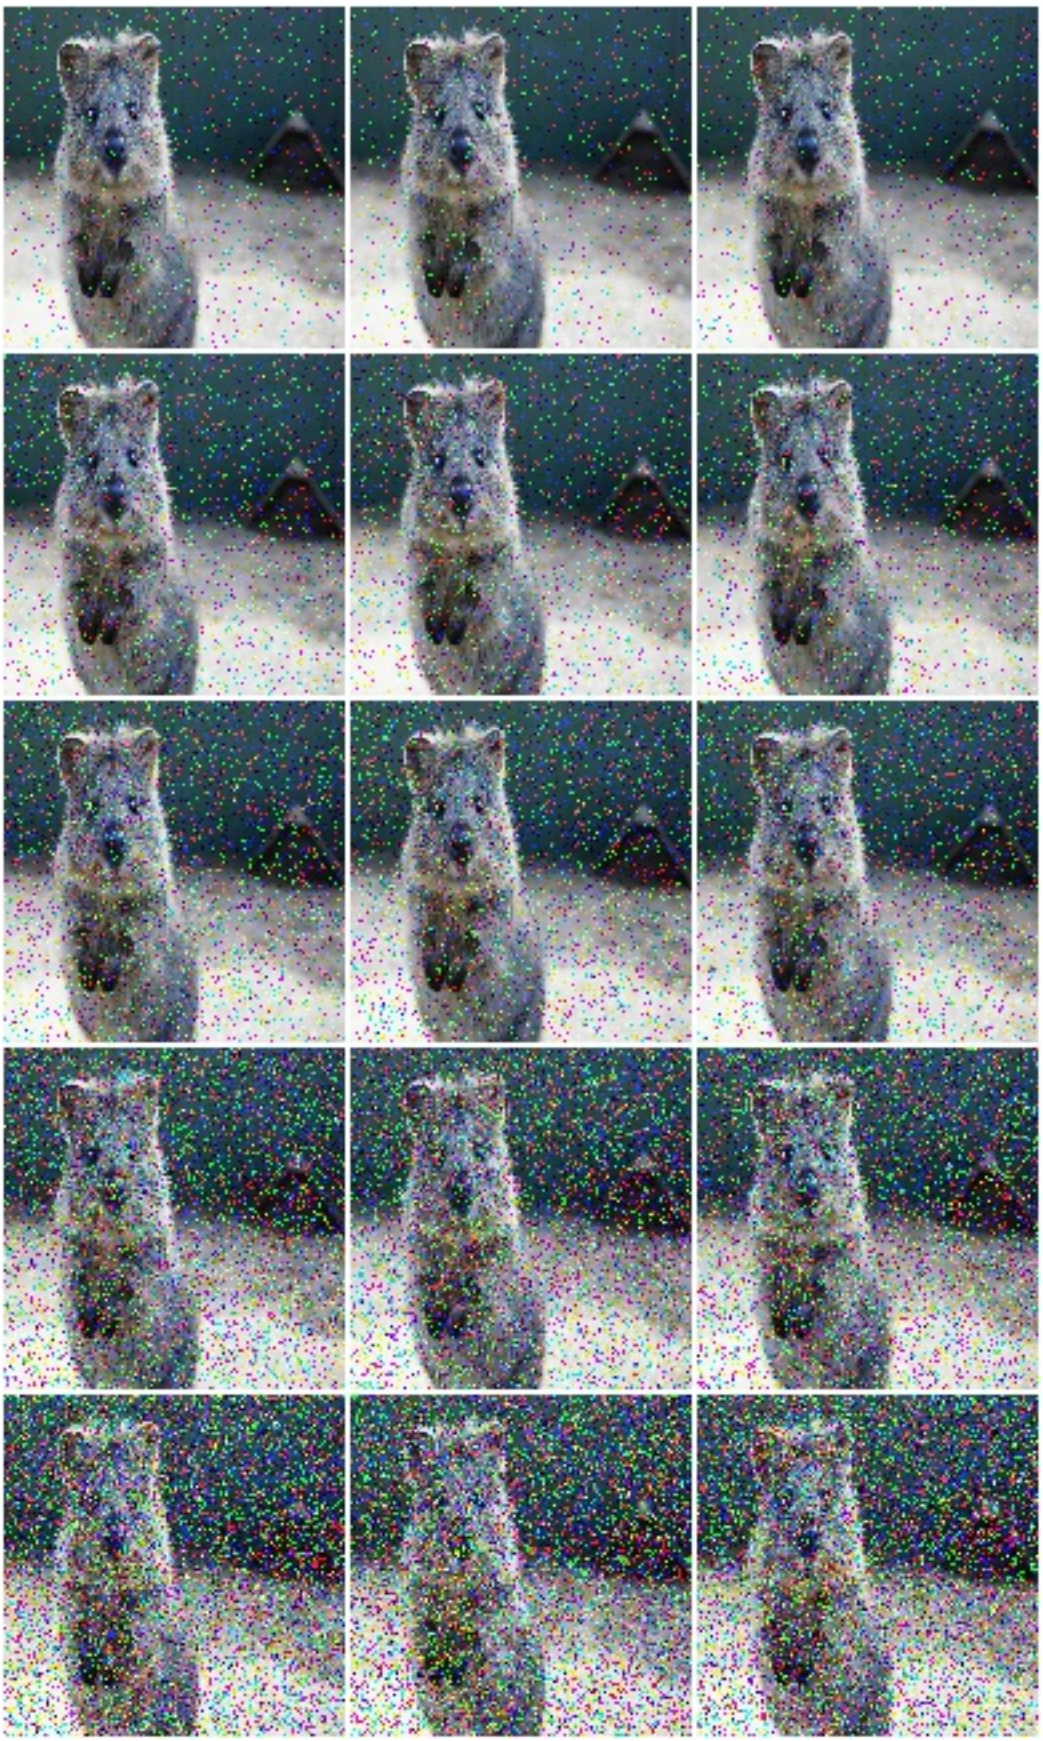

9.14.4 SpeckleNoise

Wrapper around `speckle_noise()`.

Note: This augmenter only affects images. Other data is not changed.

API link: *SpeckleNoise*

The image below visualizes severities `1` to `5` – one severity per row:

Example. Create an augmenter around `speckle_noise()`. Apply it to images using e.g. `aug(images=[image1, image2, ...])`:

```
import imgaug.augmenters as iaa
aug = iaa.imgcorruptlike.SpeckleNoise(severity=2)
```

9.14.5 GaussianBlur

Wrapper around `gaussian_blur()`.

Note: This augmenter only affects images. Other data is not changed.

API link: *GaussianBlur*

The image below visualizes severities `1` to `5` – one severity per row:

Example. Create an augmenter around `gaussian_blur()`. Apply it to images using e.g. `aug(images=[image1, image2, ...])`:

```
import imgaug.augmenters as iaa
aug = iaa.imgcorruptlike.GaussianBlur(severity=2)
```

9.14.6 GlassBlur

Wrapper around `glass_blur()`.

Note: This augmenter only affects images. Other data is not changed.

API link: *GlassBlur*

The image below visualizes severities `1` to `5` – one severity per row:

Example. Create an augmenter around `glass_blur()`. Apply it to images using e.g. `aug(images=[image1, image2, ...])`:

```
import imgaug.augmenters as iaa
aug = iaa.imgcorruptlike.GlassBlur(severity=2)
```

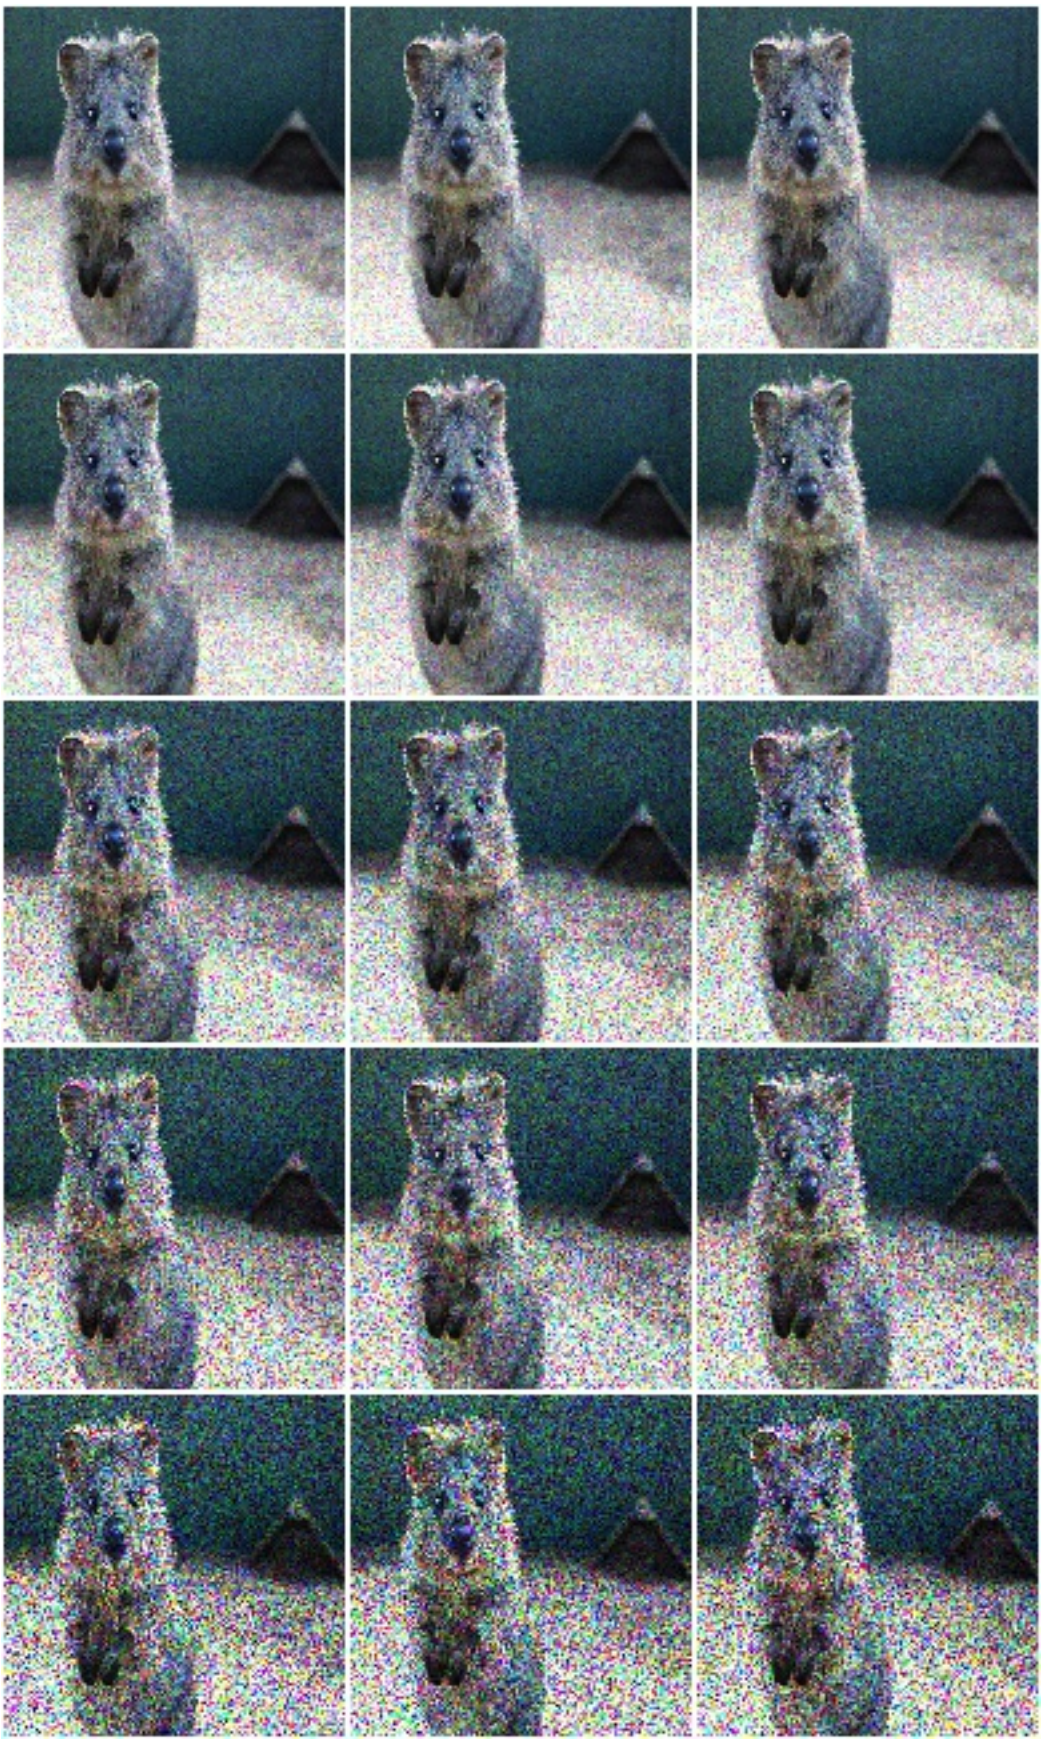

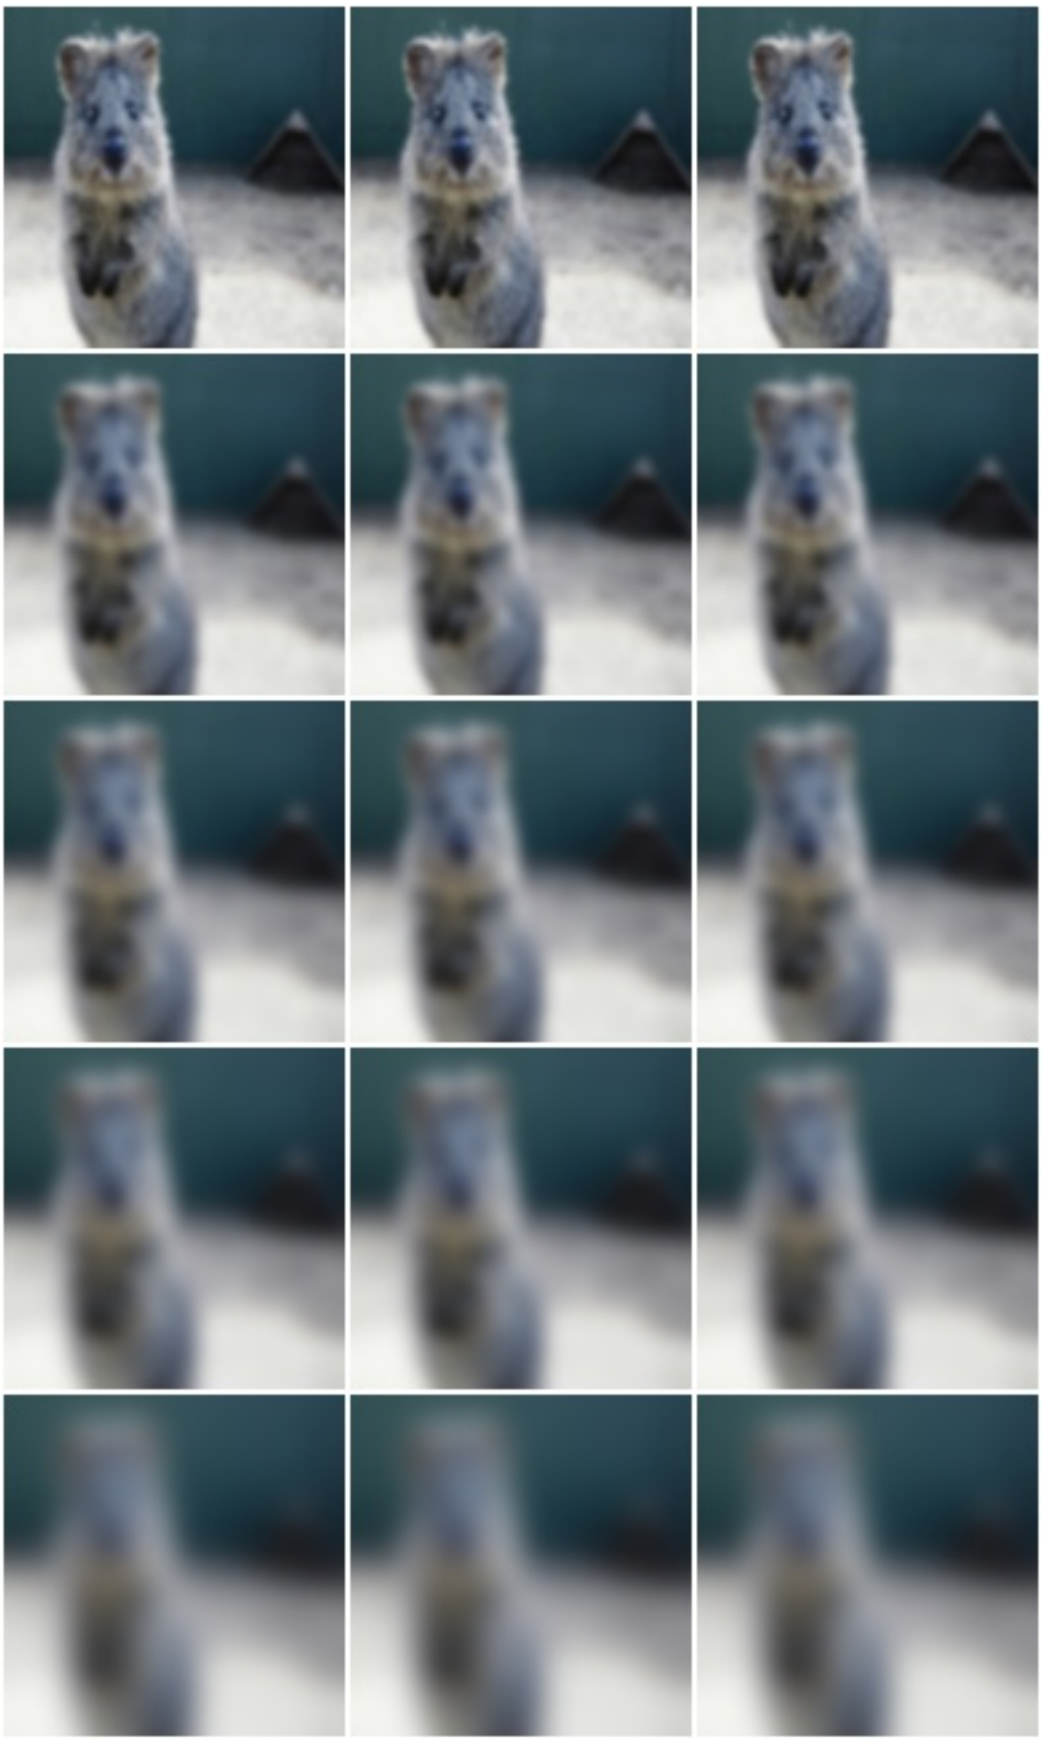

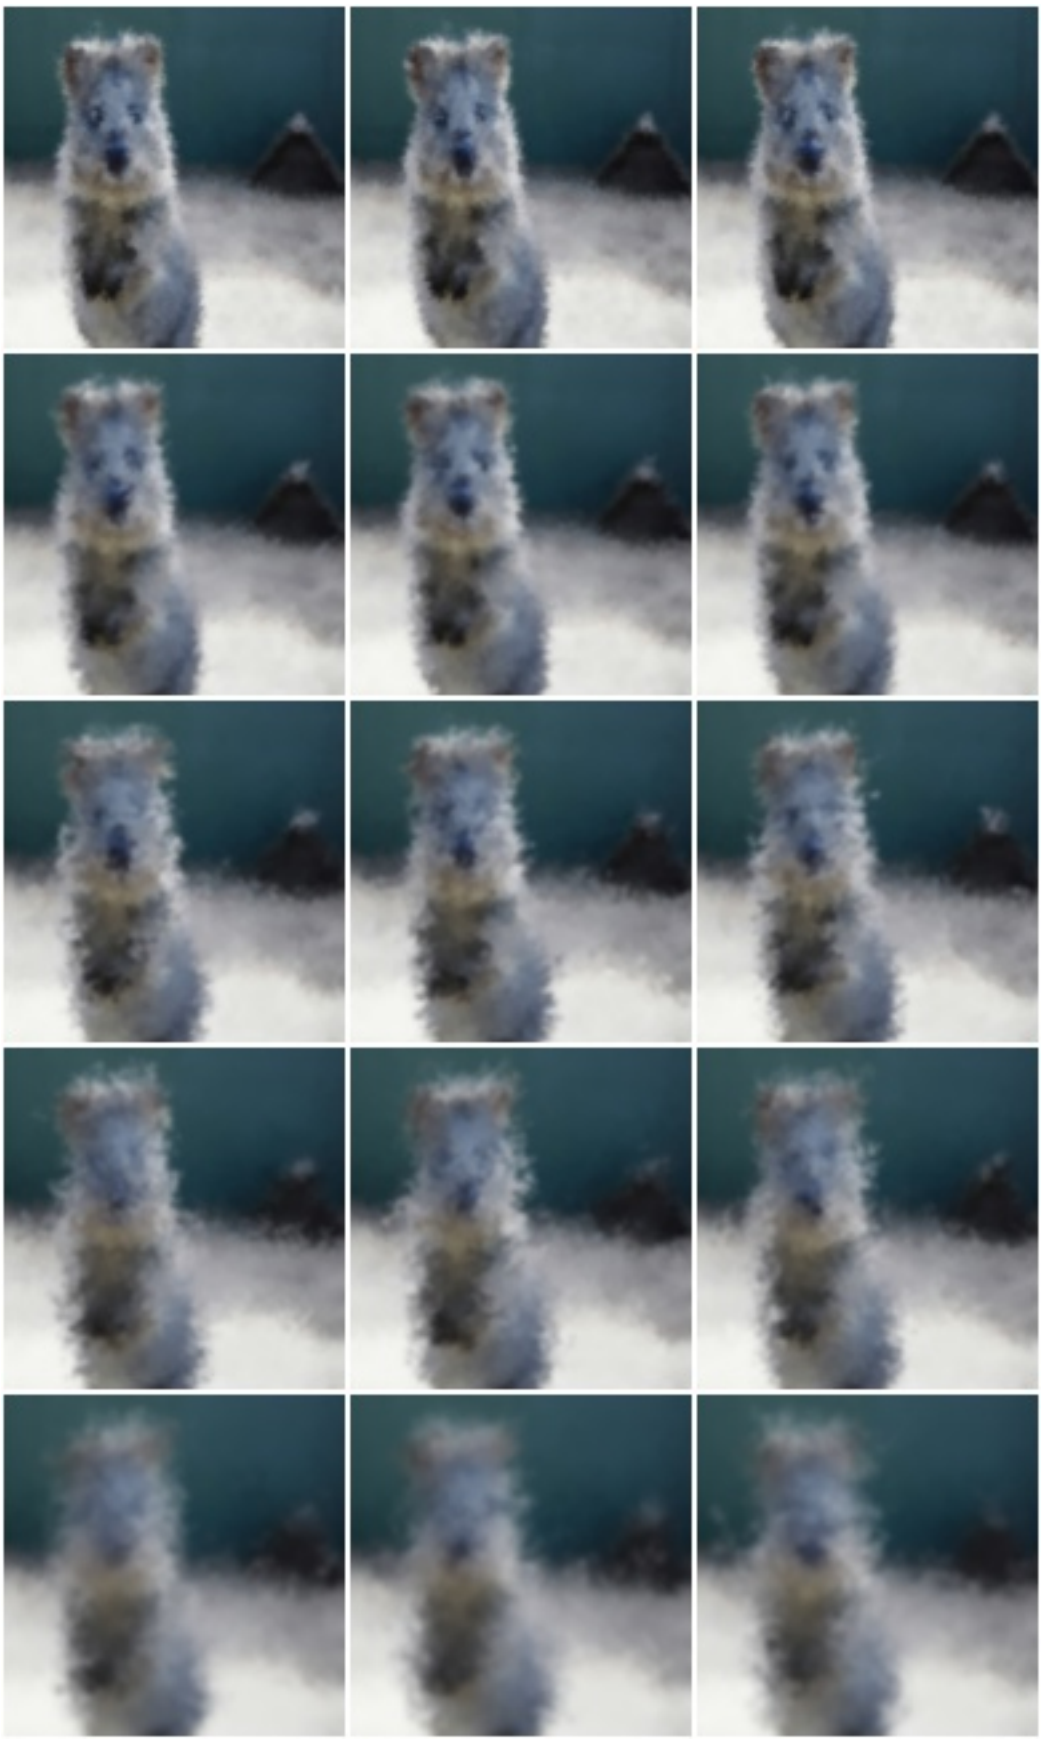

9.14.7 DefocusBlur

Wrapper around `defocus_blur()`.

Note: This augmenter only affects images. Other data is not changed.

API link: *DefocusBlur*

The image below visualizes severities `1` to `5` – one severity per row:

Example. Create an augmenter around `defocus_blur()`. Apply it to images using e.g. `aug(images=[image1, image2, ...])`:

```
import imgaug.augmenters as iaa
aug = iaa.imgcorruptlike.DefocusBlur(severity=2)
```

9.14.8 MotionBlur

Wrapper around `motion_blur()`.

Note: This augmenter only affects images. Other data is not changed.

API link: *MotionBlur*

The image below visualizes severities `1` to `5` – one severity per row:

Example. Create an augmenter around `motion_blur()`. Apply it to images using e.g. `aug(images=[image1, image2, ...])`:

```
import imgaug.augmenters as iaa
aug = iaa.imgcorruptlike.MotionBlur(severity=2)
```

9.14.9 ZoomBlur

Wrapper around `zoom_blur()`.

Note: This augmenter only affects images. Other data is not changed.

API link: *ZoomBlur*

The image below visualizes severities `1` to `5` – one severity per row:

Example. Create an augmenter around `zoom_blur()`. Apply it to images using e.g. `aug(images=[image1, image2, ...])`:

```
import imgaug.augmenters as iaa
aug = iaa.imgcorruptlike.ZoomBlur(severity=2)
```

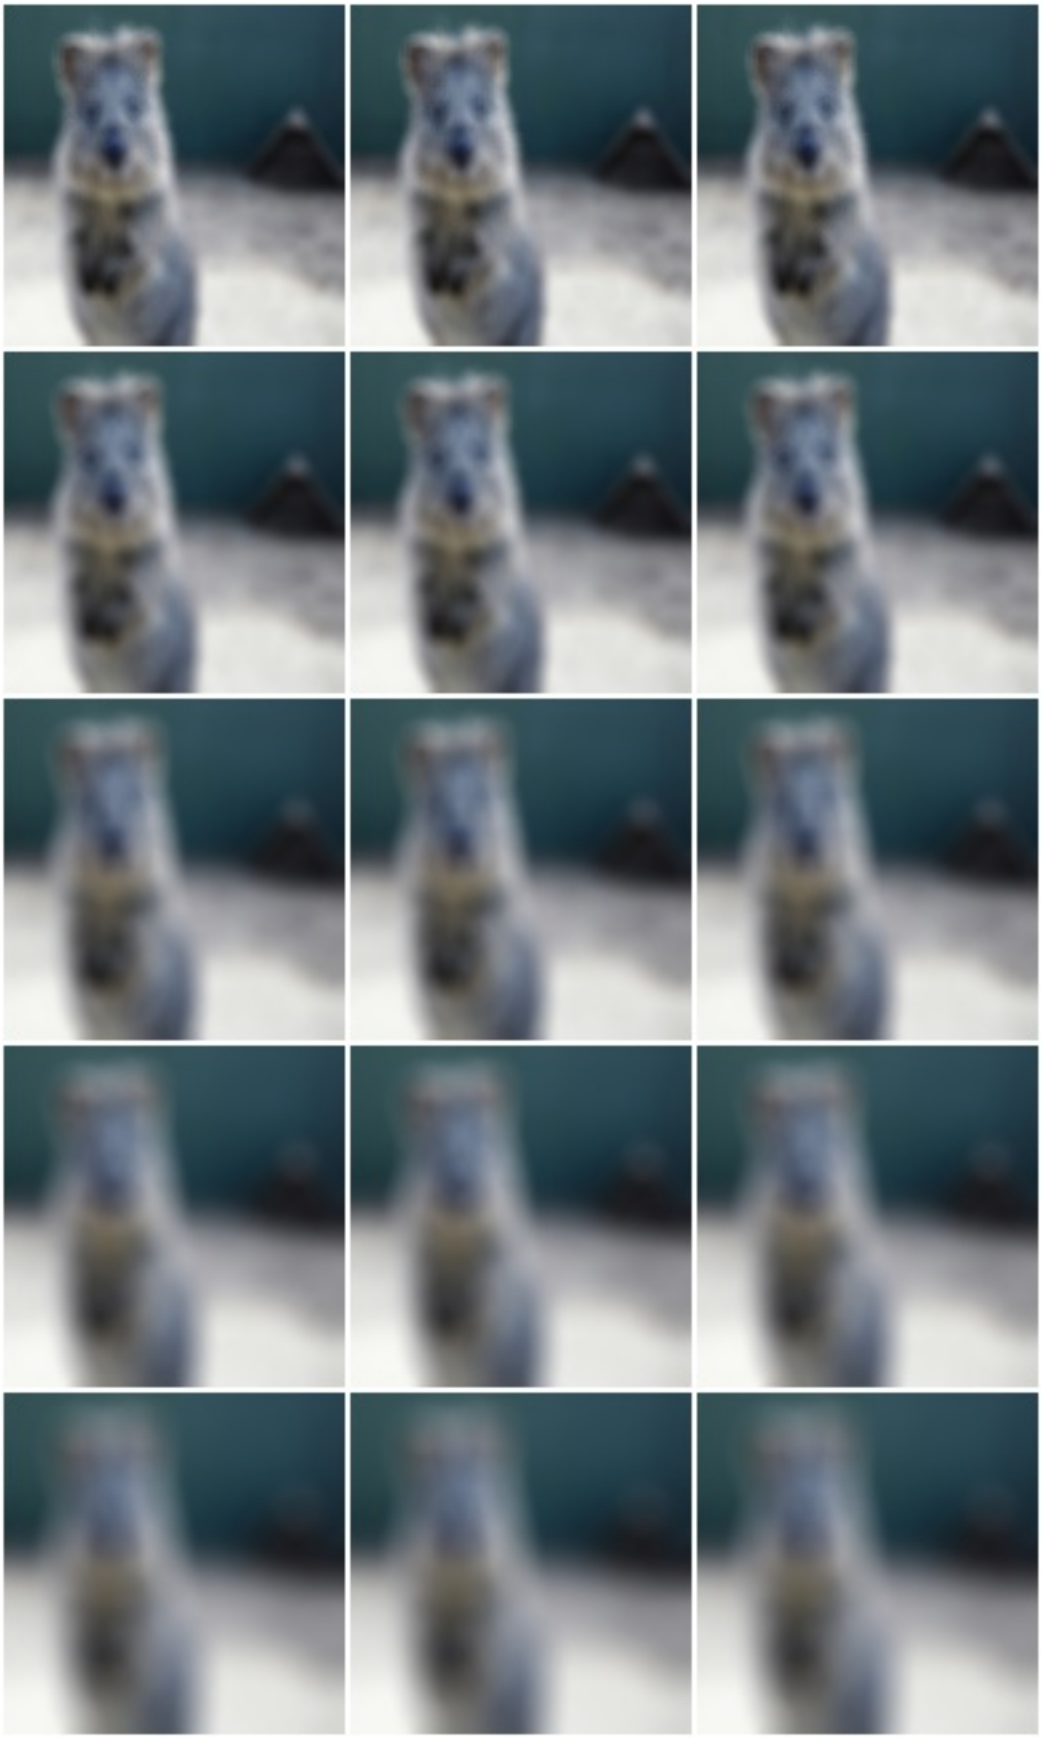

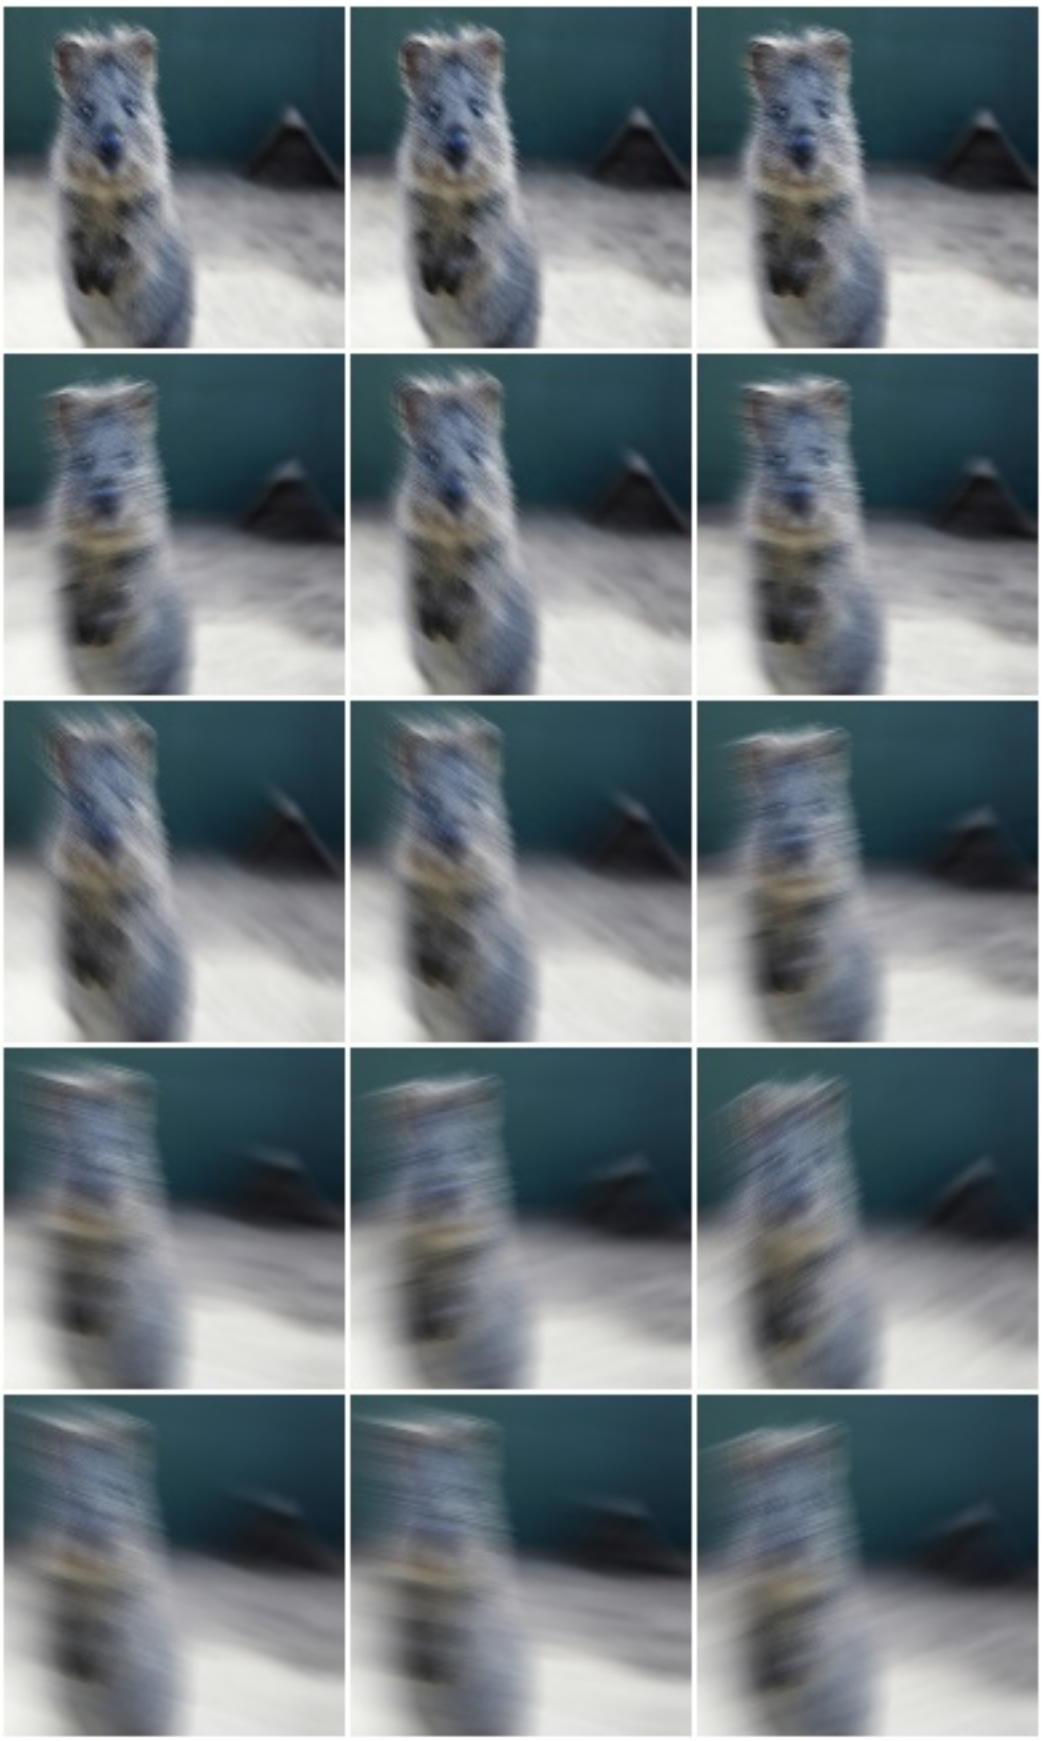

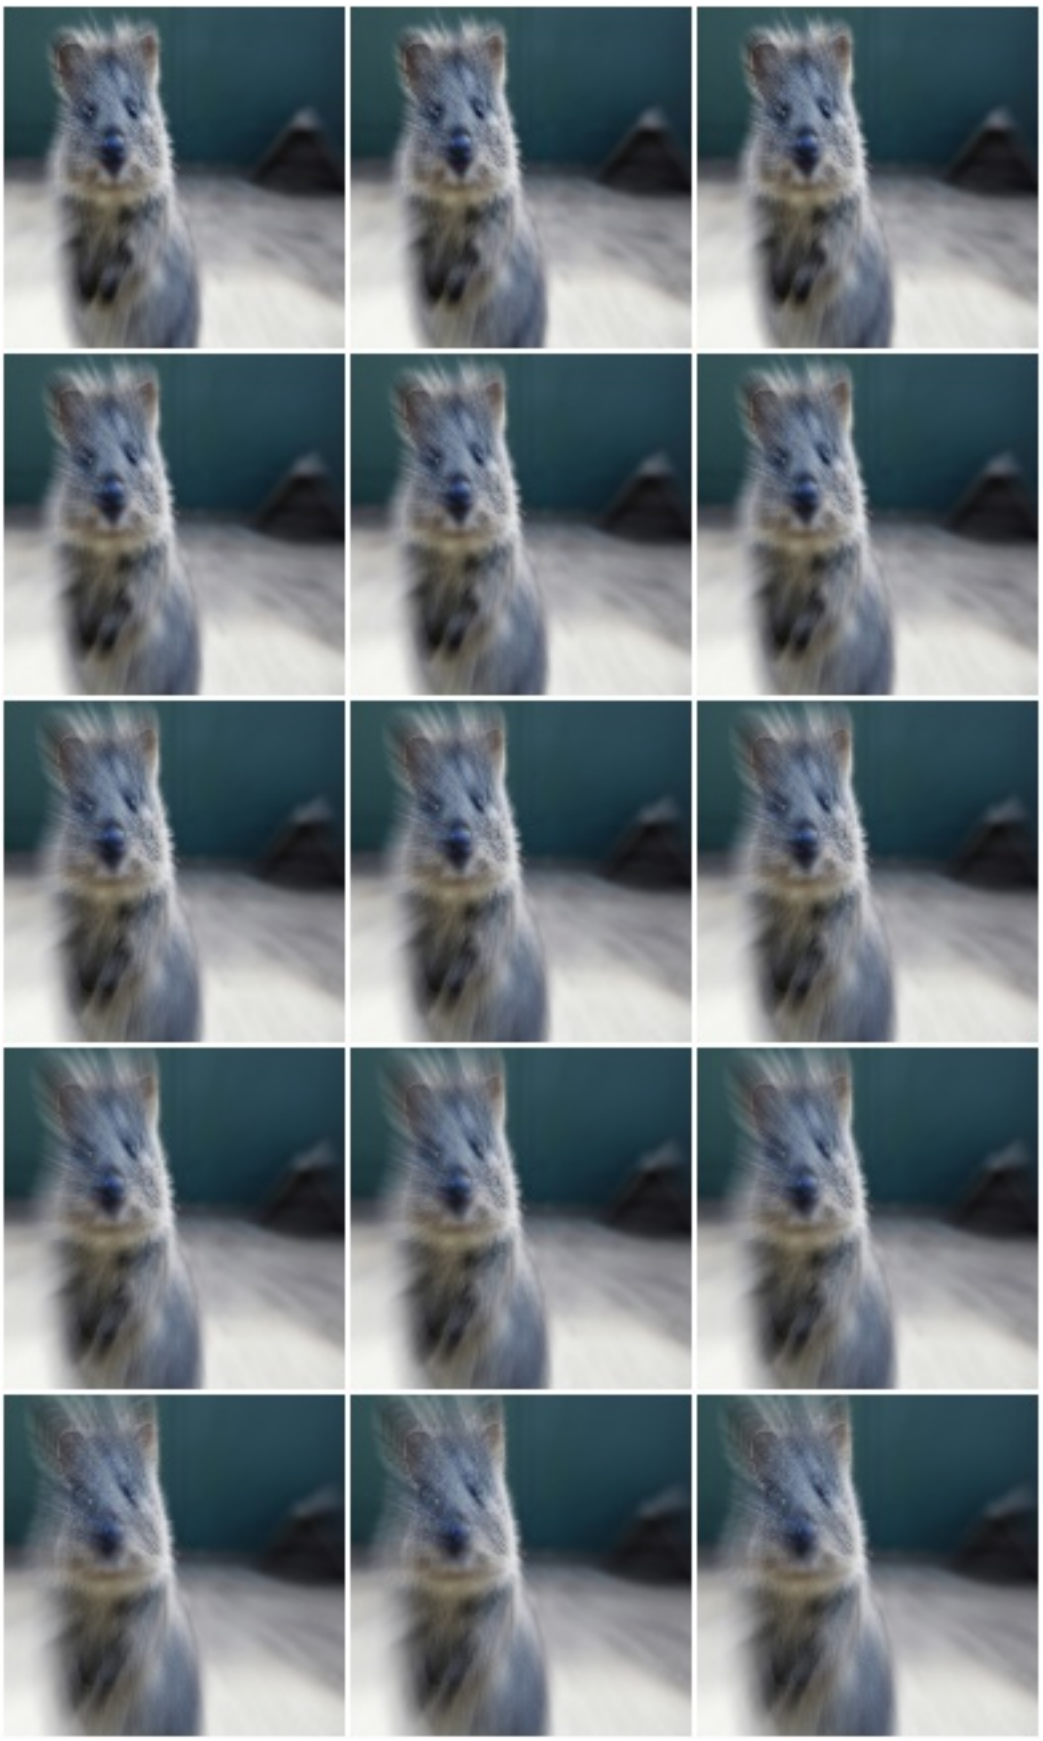

9.14.10 Fog

Wrapper around `fog()`.

Note: This augmenter only affects images. Other data is not changed.

API link: *Fog*

The image below visualizes severities `1` to `5` – one severity per row:

Example. Create an augmenter around `fog()`. Apply it to images using e.g. `aug(images=[image1, image2, ...])`:

```
import imgaug.augmenters as iaa
aug = iaa.imgcorruptlike.Fog(severity=2)
```

9.14.11 Frost

Wrapper around `frost()`.

Note: This augmenter only affects images. Other data is not changed.

API link: *Frost*

The image below visualizes severities `1` to `5` – one severity per row:

Example. Create an augmenter around `frost()`. Apply it to images using e.g. `aug(images=[image1, image2, ...])`:

```
import imgaug.augmenters as iaa
aug = iaa.imgcorruptlike.Frost(severity=2)
```

9.14.12 Snow

Wrapper around `snow()`.

Note: This augmenter only affects images. Other data is not changed.

API link: *Snow*

The image below visualizes severities `1` to `5` – one severity per row:

Example. Create an augmenter around `snow()`. Apply it to images using e.g. `aug(images=[image1, image2, ...])`:

```
import imgaug.augmenters as iaa
aug = iaa.imgcorruptlike.Snow(severity=2)
```

9.14.13 Spatter

Wrapper around `spatter()`.

Note: This augmenter only affects images. Other data is not changed.

API link: *Spatter*

The image below visualizes severities `1` to `5` – one severity per row:

Example. Create an augmenter around `spatter()`. Apply it to images using e.g. `aug(images=[image1, image2, ...])`:

```
import imgaug.augmenters as iaa
aug = iaa.imgcorruptlike.Spatter(severity=2)
```

9.14.14 Contrast

Wrapper around `contrast()`.

Note: This augmenter only affects images. Other data is not changed.

API link: *Contrast*

The image below visualizes severities `1` to `5` – one severity per row:

Example. Create an augmenter around `contrast()`. Apply it to images using e.g. `aug(images=[image1, image2, ...])`:

```
import imgaug.augmenters as iaa
aug = iaa.imgcorruptlike.Contrast(severity=2)
```

9.14.15 Brightness

Wrapper around `brightness()`.

Note: This augmenter only affects images. Other data is not changed.

API link: *Brightness*

The image below visualizes severities `1` to `5` – one severity per row:

Example. Create an augmenter around `brightness()`. Apply it to images using e.g. `aug(images=[image1, image2, ...])`:

```
import imgaug.augmenters as iaa
aug = iaa.imgcorruptlike.Brightness(severity=2)
```

9.14.16 Saturate

Wrapper around `saturate()`.

Note: This augmenter only affects images. Other data is not changed.

API link: *Saturate*

The image below visualizes severities `1` to `5` – one severity per row:

Example. Create an augmenter around `saturate()`. Apply it to images using e.g. `aug(images=[image1, image2, ...])`:

```
import imgaug.augmenters as iaa
aug = iaa.imgcorruptlike.Saturate(severity=2)
```

9.14.17 JpegCompression

Wrapper around `jpeg_compression()`.

Note: This augmenter only affects images. Other data is not changed.

API link: *JpegCompression*

The image below visualizes severities `1` to `5` – one severity per row:

Example. Create an augmenter around `jpeg_compression()`. Apply it to images using e.g. `aug(images=[image1, image2, ...])`:

```
import imgaug.augmenters as iaa
aug = iaa.imgcorruptlike.JpegCompression(severity=2)
```

9.14.18 Pixelate

Wrapper around `jpeg_compression()`.

Note: This augmenter only affects images. Other data is not changed.

Wrapper around `pixelate()`.

Note: This augmenter only affects images. Other data is not changed.

API link: *Pixelate*

The image below visualizes severities `1` to `5` – one severity per row:

Example. Create an augmenter around `pixelate()`. Apply it to images using e.g. `aug(images=[image1, image2, ...])`:

```
import imgaug.augmenters as iaa
aug = iaa.imgcorruptlike.Pixelate(severity=2)
```

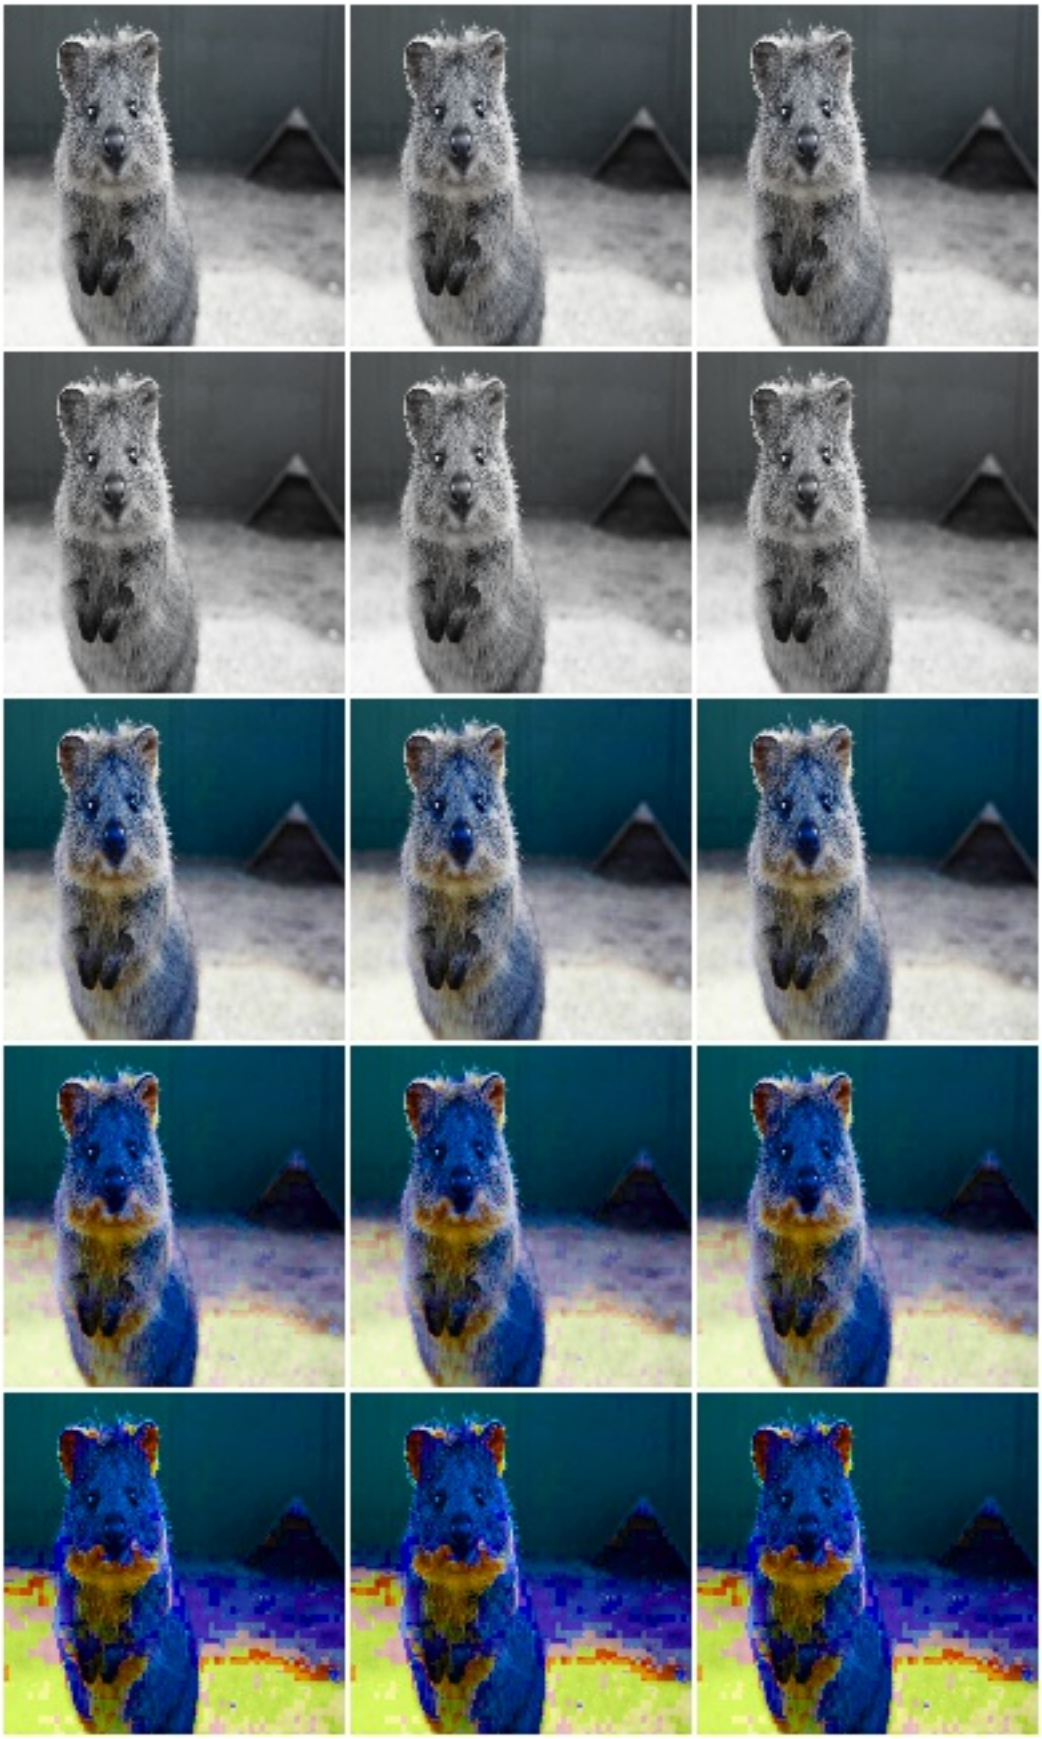

9.14.19 ElasticTransform

Wrapper around `elastic_transform()`.

Note: This augmenter only affects images. Other data is not changed.

API link: *ElasticTransform*

The image below visualizes severities `1` to `5` – one severity per row:

Example. Create an augmenter around `elastic_transform()`. Apply it to images using e.g. `aug(images=[image1, image2, ...])`:

```
import imgaug.augmenters as iaa
aug = iaa.imgcorruptlike.ElasticTransform(severity=2)
```

9.15 augmenters.pillike

9.15.1 Solarize

Augmenter with identical outputs to PIL's `solarize()` function.

This augmenter inverts all pixel values above a threshold.

The outputs are identical to PIL's `solarize()`.

API link: *Solarize()*

Example. Invert the colors in `50` percent of all images for pixels with a value between `32` and `128` or more. The threshold is sampled once per image. The thresholding operation happens per channel.

```
import imgaug.augmenters as iaa
aug = iaa.Solarize(0.5, threshold=(32, 128))
```

9.15.2 Posterize

Augmenter with identical outputs to PIL's `posterize()` function.

This augmenter quantizes each array component to `N` bits.

This class is currently an alias for *Posterize*, which again is an alias for *UniformColorQuantizationToNBits*, i.e. all three classes are right now guarantueed to have the same outputs as PIL's function.

API link: *Posterize()*

9.15.3 Equalize

Equalize the image histogram.

This augmenter has identical outputs to `equalize()`.

API link: *Equalize()*

Example. Equalize the histograms of all input images:

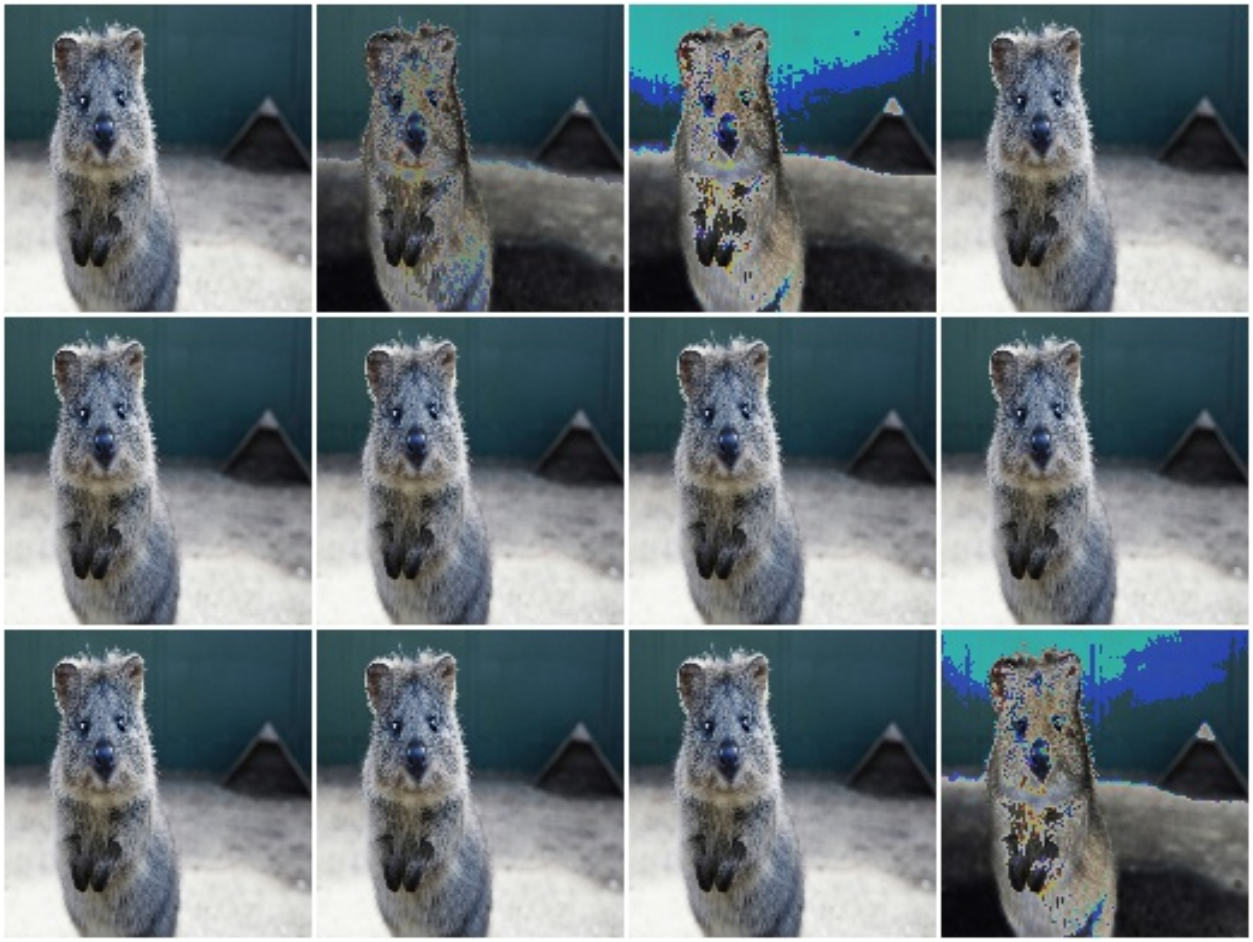

```
import imgaug.augmenters as iaa
aug = iaa.pillike.Equalize()
```

9.15.4 Autocontrast

Adjust contrast by cutting off `p%` of lowest/highest histogram values.

This augmenter has identical outputs to `autocontrast()`.

See *`autocontrast()`* for more details.

API link: *`Autocontrast()`*

Example. Modify the contrast of images by cutting off the `0` to `20%` lowest and highest values from the histogram, then stretching it to full length:

```
import imgaug.augmenters as iaa
aug = iaa.pillike.Autocontrast()
```

Example. Modify the contrast of images by cutting off the `10` to `20%` lowest and highest values from the histogram, then stretching it to full length. The cutoff value is sampled per *channel* instead of per *image*.

```
aug = iaa.pillike.Autocontrast((10, 20), per_channel=True)
```

9.15.5 EnhanceColor

Convert images to grayscale.

This augmenter has identical outputs to `Color`.

API link: *`EnhanceColor()`*

Example. Create an augmenter to remove a random fraction of color from input images:

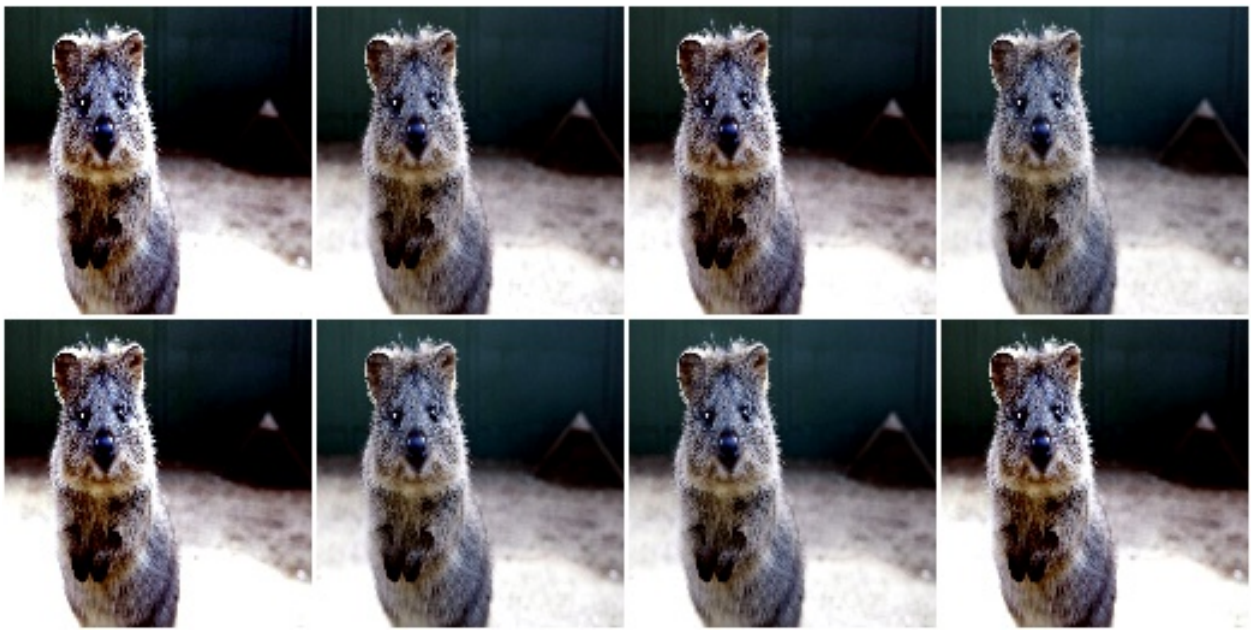

```
import imgaug.augmenters as iaa
aug = iaa.pillike.EnhanceColor()
```

9.15.6 EnhanceContrast

Change the contrast of images.

This augmenter has identical outputs to `Contrast`.

API link: *EnhanceContrast()*

Example. Create an augmenter that worsens the contrast of an image by a random factor:

```
import imgaug.augmenters as iaa
aug = iaa.pillike.EnhanceContrast()
```

9.15.7 EnhanceBrightness

Change the brightness of images.

This augmenter has identical outputs to `Brightness`.

API link: *EnhanceBrightness()*

Example. Create an augmenter that worsens the brightness of an image by a random factor:

```
import imgaug.augmenters as iaa
aug = iaa.pillike.EnhanceBrightness()
```

9.15.8 EnhanceSharpness

Change the sharpness of images.

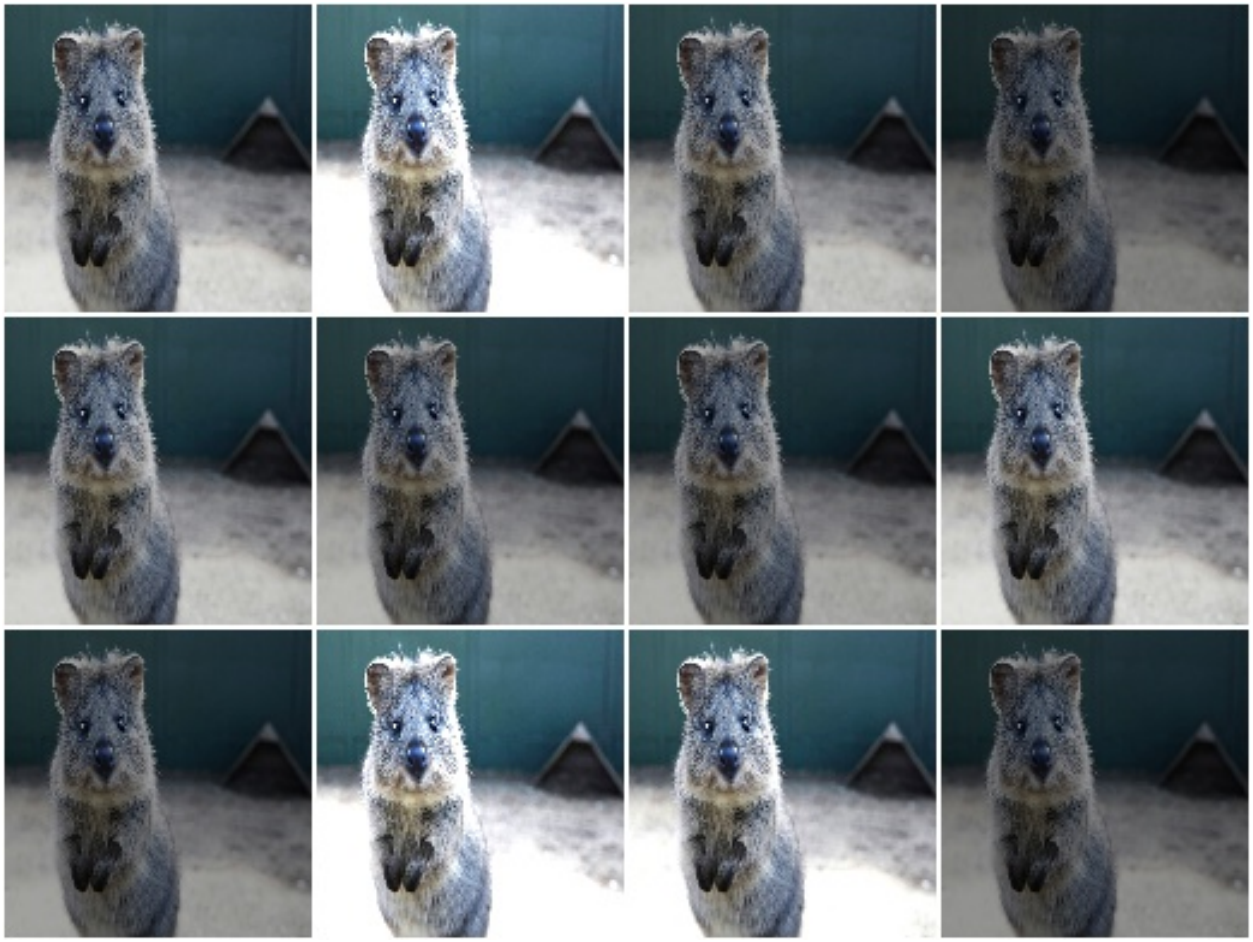

This augmenter has identical outputs to `Sharpness`.

API link: *EnhanceSharpness()*

Example. Create an augmenter that randomly decreases or increases the sharpness of an image:

```
import imgaug.augmenters as iaa
aug = iaa.pillike.EnhanceSharpness()
```

9.15.9 FilterBlur

Apply a blur filter kernel to images.

This augmenter has identical outputs to calling `filter()` with kernel `PIL.ImageFilter.BLUR`.

API link: *FilterBlur()*

Example. Create an augmenter that applies a blur filter kernel to images:

```
import imgaug.augmenters as iaa
aug = iaa.pillike.FilterBlur()
```

9.15.10 FilterSmooth

Apply a smoothening filter kernel to images.

This augmenter has identical outputs to calling `filter()` with kernel `PIL.ImageFilter.SMOOTH`.

API link: *FilterSmooth()*

Example. Create an augmenter that applies a smoothening filter kernel to images:

```
import imgaug.augmenters as iaa
aug = iaa.pillike.FilterSmooth()
```

9.15.11 FilterSmoothMore

Apply a strong smoothening filter kernel to images.

This augmenter has identical outputs to calling `filter()` with kernel `PIL.ImageFilter.BLUR`.

API link: *FilterSmoothMore()*

Example. Create an augmenter that applies a strong smoothening filter kernel to images:

```
import imgaug.augmenters as iaa
aug = iaa.pillike.FilterSmoothMore()
```

9.15.12 FilterEdgeEnhance

Apply an edge enhance filter kernel to images.

This augmenter has identical outputs to calling `filter()` with kernel `PIL.ImageFilter.EDGE_ENHANCE`.

API link: *FilterEdgeEnhance()*

Example. Create an augmenter that applies a edge enhancement filter kernel to images:

```
import imgaug.augmenters as iaa
aug = iaa.pillike.FilterEdgeEnhance()
```

9.15.13 FilterEdgeEnhanceMore

Apply a strong edge enhancement filter kernel to images.

This augmenter has identical outputs to calling `filter()` with kernel `PIL.ImageFilter.EDGE_ENHANCE_MORE`.

API link: *FilterEdgeEnhanceMore()*

Example. Create an augmenter that applies a strong edge enhancement filter kernel to images:

```
import imgaug.augmenters as iaa
aug = iaa.pillike.FilterEdgeEnhanceMore()
```

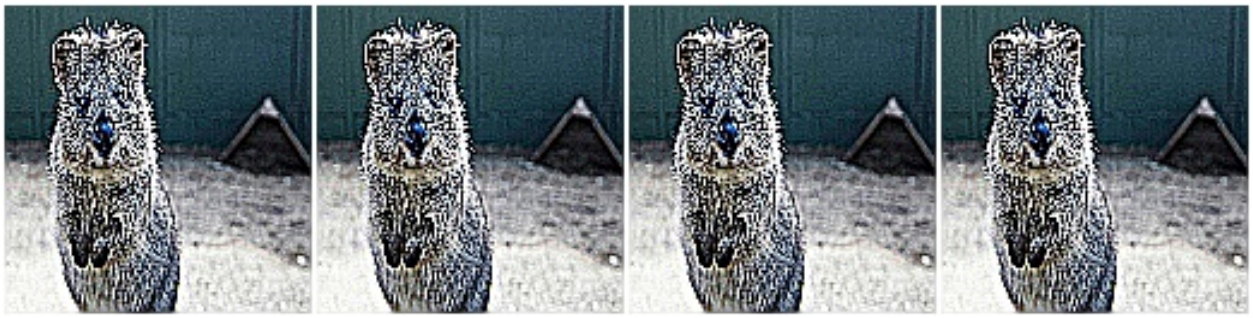

9.15.14 FilterFindEdges

Apply a edge detection kernel to images.

This augmenter has identical outputs to calling `filter()` with kernel `PIL.ImageFilter.FIND_EDGES`.

API link: *FilterFindEdges()*

Example. Create an augmenter that applies an edge detection filter kernel to images:

```
import imgaug.augmenters as iaa
aug = iaa.pillike.FilterFindEdges()
```

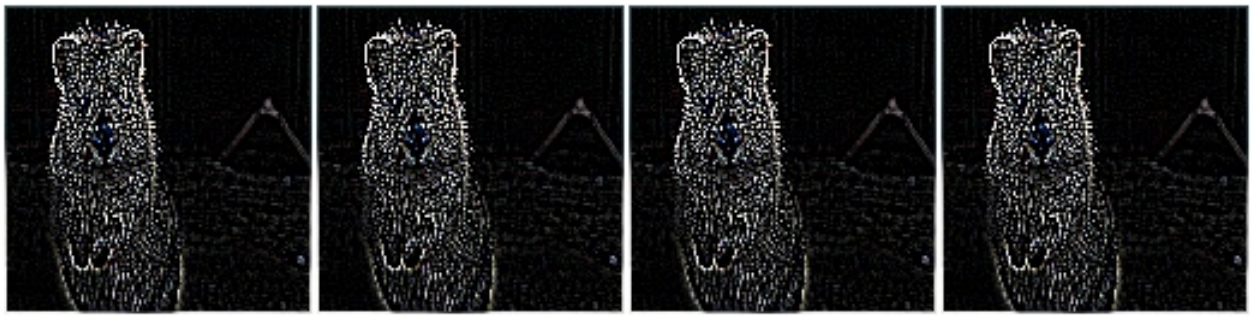

9.15.15 FilterContour

Apply a contour detection filter kernel to images.

This augmenter has identical outputs to calling `filter()` with kernel `PIL.ImageFilter.CONTOUR`.

API link: *FilterContour()*

Example. Create an augmenter that applies a contour detection filter kernel to images:

```
import imgaug.augmenters as iaa
aug = iaa.pillike.FilterContour()
```

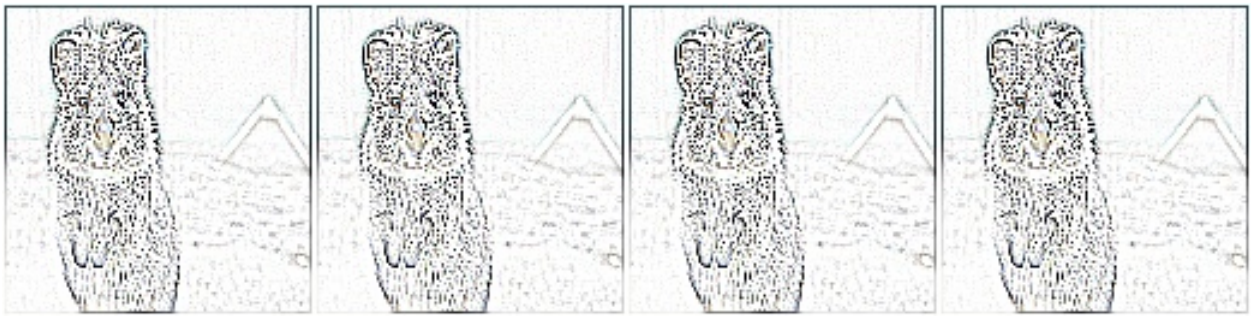

9.15.16 FilterEmboss

Apply an emboss filter kernel to images.

This augmenter has identical outputs to calling `filter()` with kernel `PIL.ImageFilter.EMBOSS`.

API link: *FilterEmboss()*

Example. Create an augmenter that applies an emboss filter kernel to images:

```
import imgaug.augmenters as iaa
aug = iaa.pillike.FilterEmboss()
```

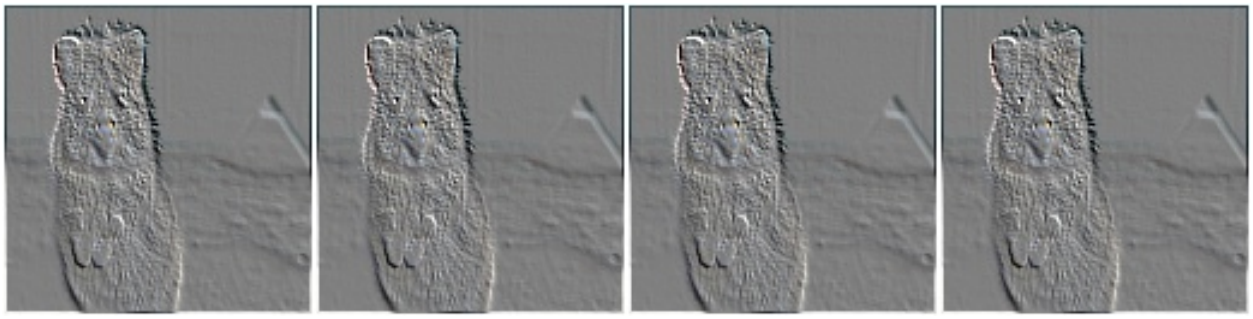

9.15.17 FilterSharpen

Apply a sharpening filter kernel to images.

This augmenter has identical outputs to calling `filter()` with kernel `PIL.ImageFilter.SHARPEN`.

API link: *FilterSharpen()*

Example. Create an augmenter that applies a sharpening filter kernel to images:

```
import imgaug.augmenters as iaa
aug = iaa.pillike.FilterSharpen()
```

9.15.18 FilterDetail

Apply a detail enhancement filter kernel to images.

This augmenter has identical outputs to calling `filter()` with kernel `PIL.ImageFilter.DETAIL`.

API link: *FilterDetail()*

Example. Create an augmenter that applies a detail enhancement filter kernel to images:

```
import imgaug.augmenters as iaa
aug = iaa.pillike.FilterDetail()
```

9.15.19 Affine

Apply PIL-like affine transformations to images.

This augmenter has identical outputs to `transform()` with parameter `method=PIL.Image.AFFINE`.

Note: This augmenter can currently only transform image-data. Batches containing heatmaps, segmentation maps and coordinate-based augmentables will be rejected with an error. Use *`Affine`* if you have to transform such inputs.

Note: This augmenter uses the image center as the transformation center. This has to be explicitly enforced in PIL using corresponding translation matrices. Without such translation, PIL uses the image top left corner as the transformation center. To mirror that behaviour, use `center=(0.0, 0.0)`.

API link: *`Affine()`*

Example. Create an augmenter that applies affine scaling (zoom in/out) to images. Along the x-axis they are scaled to 80-120% of their size, along the y-axis to 50-150% (both values randomly and uniformly chosen per image).

```
import imgaug.augmenters as iaa
aug = iaa.pillike.Affine(scale={"x": (0.8, 1.2), "y": (0.5, 1.5)})
```

Example. Create an augmenter that translates images along the y-axis by either `-10px` or `10px`. Newly created pixels are always filled with the value `128` (along all channels).

```
aug = iaa.pillike.Affine(translate_px={"x": 0, "y": [-10, 10]},
                         fillcolor=128)
```

Example. Rotate an image by `-20` to `20` degress and fill up all newly created pixels with a random RGB color:

```
aug = iaa.pillike.Affine(rotate=(-20, 20), fillcolor=(0, 256))
```

See the similar augmenter *Affine* for more examples.

9.16 augmenters.pooling

9.16.1 AveragePooling

Apply average pooling to images.

This augmenter pools images with kernel sizes `H x W` by averaging the pixel values within these windows. For e.g. `2 x 2` this halves the image size. Optionally, the augmenter will automatically re-upscale the image to the input size (by default this is activated).

This augmenter does not affect heatmaps, segmentation maps or coordinates-based augmentables (e.g. keypoints, bounding boxes, ...).

Note that this augmenter is very similar to `AverageBlur`. `AverageBlur` applies averaging within windows of given kernel size *without* striding, while `AveragePooling` applies striding corresponding to the kernel size, with optional upscaling afterwards. The upscaling is configured to create "pixelated"/"blocky" images by default.

API link: *AveragePooling*

Example. Create an augmenter that always pools with a kernel size of `2 x 2`:

```
import imgaug.augmenters as iaa
aug = iaa.AveragePooling(2)
```

Example. Create an augmenter that always pools with a kernel size of `2 x 2` and does *not* resize back to the input image size, i.e. the resulting images have half the resolution:

```
aug = iaa.AveragePooling(2, keep_size=False)
```

Example. Create an augmenter that always pools either with a kernel size of `2 x 2` or `8 x 8`:

```
aug = iaa.AveragePooling([2, 8])
```

Example. Create an augmenter that always pools with a kernel size of `1 x 1` (does nothing) to `7 x 7`. The kernel sizes are always symmetric.

```
aug = iaa.AveragePooling((1, 7))
```

Example. Create an augmenter that always pools with a kernel size of `H x W` where `H` and `W` are both sampled independently from the range `[1..7]`. E.g. resulting kernel sizes could be `3 x 7` or `5 x 1`.

```
aug = iaa.AveragePooling(((1, 7), (1, 7)))
```

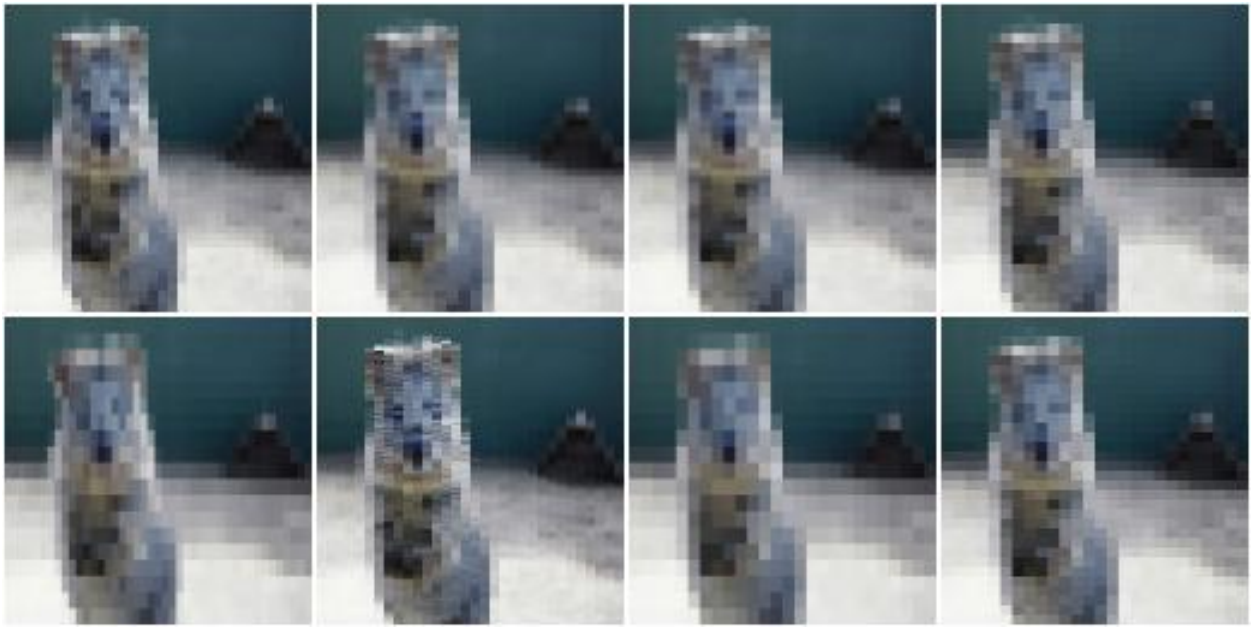

9.16.2 MaxPooling

Apply max pooling to images.

This augmenter pools images with kernel sizes `H x W` by taking the maximum pixel value over windows. For e.g. `2 x 2` this halves the image size. Optionally, the augmenter will automatically re-upscale the image to the input size (by default this is activated).

The maximum within each pixel window is always taken channelwise.

This augmenter does not affect heatmaps, segmentation maps or coordinates-based augmentables (e.g. keypoints, bounding boxes, ...).

API link: *MaxPooling*

Example. Create an augmenter that always pools with a kernel size of `2 x 2`:

```
import imgaug.augmenters as iaa
aug = iaa.MaxPooling(2)
```

Example. Create an augmenter that always pools with a kernel size of `2 x 2` and does *not* resize back to the input image size, i.e. the resulting images have half the resolution:

```
aug = iaa.MaxPooling(2, keep_size=False)
```

Example. Create an augmenter that always pools either with a kernel size of `2 x 2` or `8 x 8`:

```
aug = iaa.MaxPooling([2, 8])
```

Example. Create an augmenter that always pools with a kernel size of `1 x 1` (does nothing) to `7 x 7`. The kernel sizes are always symmetric.

```
aug = iaa.MaxPooling((1, 7))
```

Example. Create an augmenter that always pools with a kernel size of `H x W` where `H` and `W` are both sampled independently from the range `[1..7]`. E.g. resulting kernel sizes could be `3 x 7` or `5 x 1`.

```
aug = iaa.MaxPooling(((1, 7), (1, 7)))
```

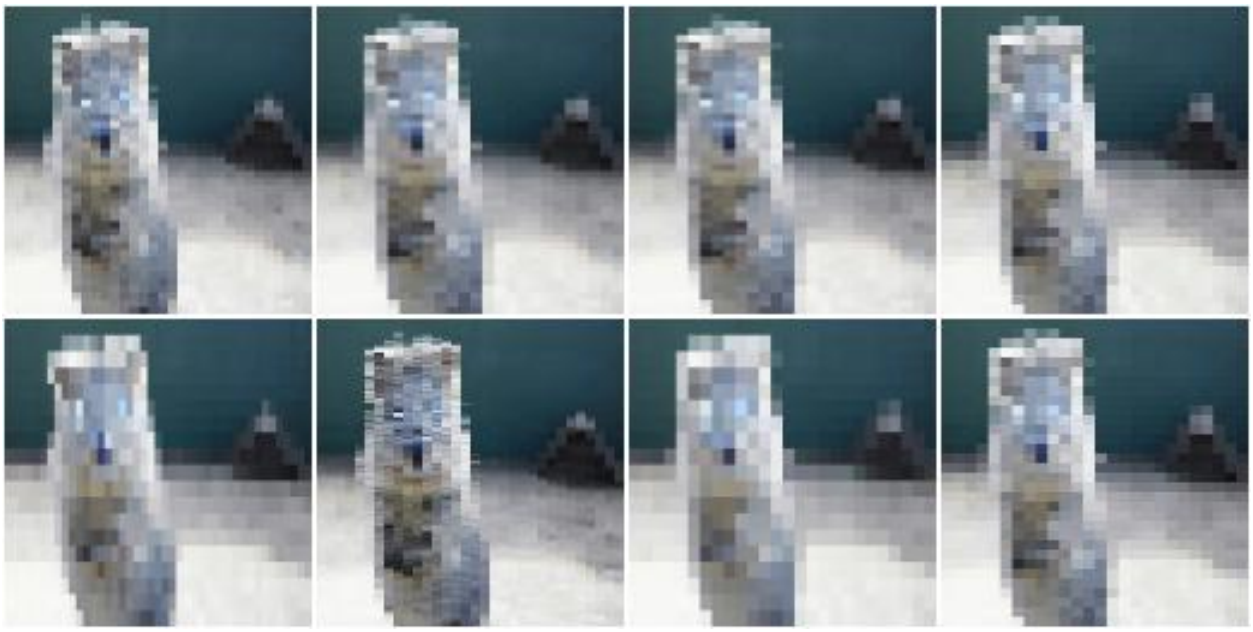

9.16.3 MinPooling

Apply minimum pooling to images.

This augmenter pools images with kernel sizes `H x W` by taking the minimum pixel value over windows. For e.g. `2 x 2` this halves the image size. Optionally, the augmenter will automatically re-upscale the image to the input size (by default this is activated).

The minimum within each pixel window is always taken channelwise.

This augmenter does not affect heatmaps, segmentation maps or coordinates-based augmentables (e.g. keypoints, bounding boxes, ...).

API link: *MinPooling*

Example. Create an augmenter that always pools with a kernel size of `2 x 2`:

```
import imgaug.augmenters as iaa
aug = iaa.MinPooling(2)
```

Example. Create an augmenter that always pools with a kernel size of `2 x 2` and does *not* resize back to the input image size, i.e. the resulting images have half the resolution.

```
aug = iaa.MinPooling(2, keep_size=False)
```

Example. Create an augmenter that always pools either with a kernel size of `2 x 2` or `8 x 8`:

```
aug = iaa.MinPooling([2, 8])
```

Example. Create an augmenter that always pools with a kernel size of `1 x 1` (does nothing) to `7 x 7`. The kernel sizes are always symmetric.

```
aug = iaa.MinPooling((1, 7))
```

Example. Create an augmenter that always pools with a kernel size of `H x W` where `H` and `W` are both sampled independently from the range `[1..7]`. E.g. resulting kernel sizes could be `3 x 7` or `5 x 1`.

```
aug = iaa.MinPooling(((1, 7), (1, 7)))
```

9.16.4 MedianPooling

Apply median pooling to images.

This augmenter pools images with kernel sizes `H x W` by taking the median pixel value over windows. For e.g. `2 x 2` this halves the image size. Optionally, the augmenter will automatically re-upscale the image to the input size (by default this is activated).

The median within each pixel window is always taken channelwise.

This augmenter does not affect heatmaps, segmentation maps or coordinates-based augmentables (e.g. keypoints, bounding boxes, ...).

API link: *MedianPooling*

Example. Create an augmenter that always pools with a kernel size of `2 x 2`:

```
import imgaug.augmenters as iaa
aug = iaa.MedianPooling(2)
```

Example. Create an augmenter that always pools with a kernel size of `2 x 2` and does *not* resize back to the input image size, i.e. the resulting images have half the resolution:

```
aug = iaa.MedianPooling(2, keep_size=False)
```

Example. Create an augmenter that always pools either with a kernel size of `2 x 2` or `8 x 8`:

```
aug = iaa.MedianPooling([2, 8])
```

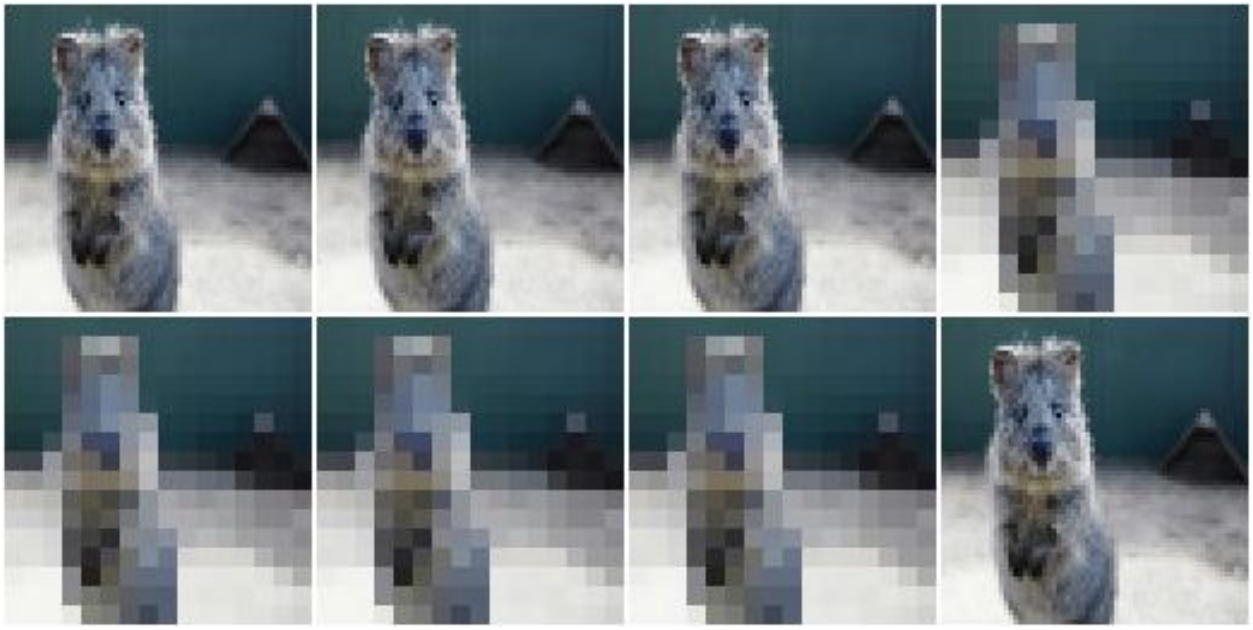

Example. Create an augmenter that always pools with a kernel size of `1 x 1` (does nothing) to `7 x 7`. The kernel sizes are always symmetric.

```
aug = iaa.MedianPooling((1, 7))
```

Example. Create an augmenter that always pools with a kernel size of `H x W` where `H` and `W` are both sampled independently from the range `[1..7]`. E.g. resulting kernel sizes could be `3 x 7` or `5 x 1`.

```
aug = iaa.MedianPooling(((1, 7), (1, 7)))
```

9.17 augmenters.segmentation

9.17.1 Superpixels

Completely or partially transform images to their superpixel representation.

Note: This augmenter is fairly slow. See *Performance*.

API link: *Superpixels*

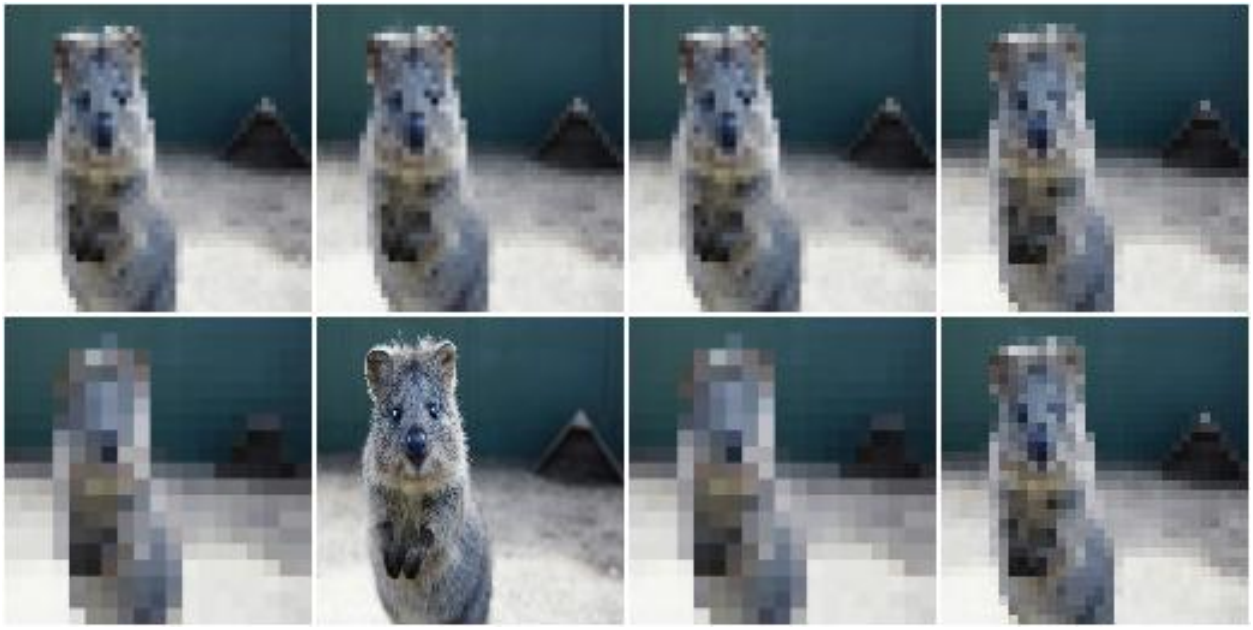

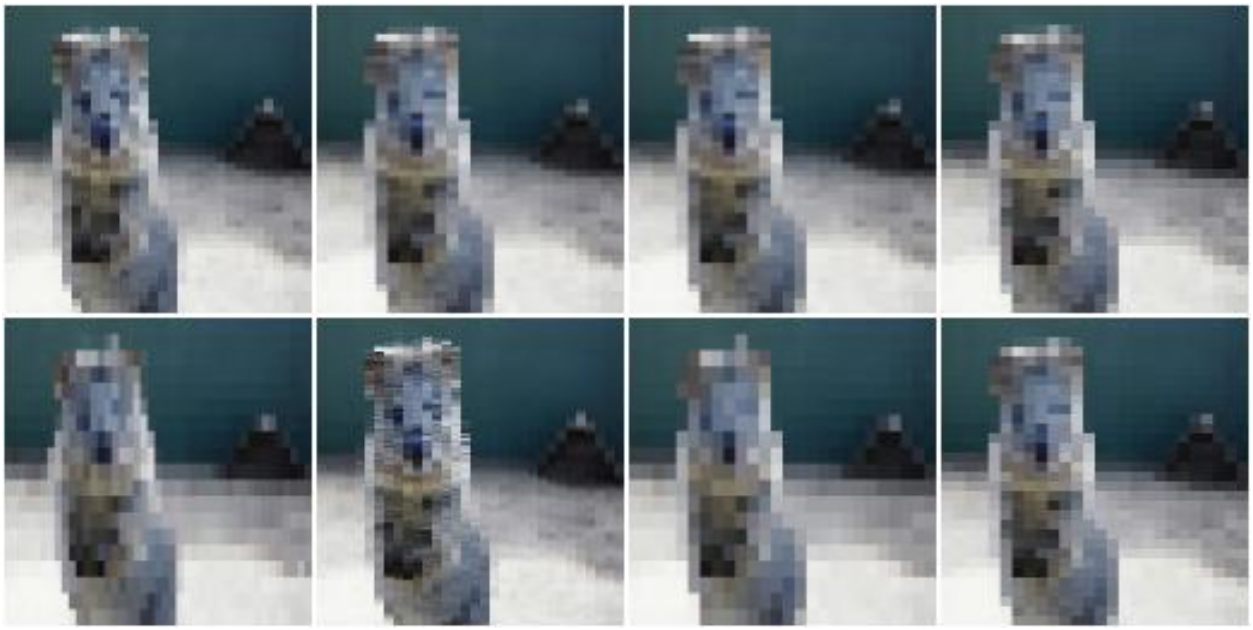

Example. Generate about `64` superpixels per image. Replace each one with a probability of `50%` by its average pixel color.

```
import imgaug.augmenters as iaa
aug = iaa.Superpixels(p_replace=0.5, n_segments=64)
```

Example. Generate `16` to `128` superpixels per image. Replace each superpixel with a probability between `10` and `100%` (sampled once per image) by its average pixel color.

```
aug = iaa.Superpixels(p_replace=(0.1, 1.0), n_segments=(16, 128))
```

Example. Effect of setting `n_segments` to a fixed value of `64` and then increasing `p_replace` from `0.0` and `1.0`:

Example. Effect of setting `p_replace` to a fixed value of `1.0` and then increasing `n_segments` from `1*16` to `9*16=144`:

9.17.2 Voronoi

Average colors of an image within Voronoi cells.

This augmenter performs the following steps:

1. Query *points_sampler* to sample random coordinates of cell centers. On the image.
2. Estimate for each pixel to which voronoi cell (i.e. segment) it belongs. Each pixel belongs to the cell with the closest center coordinate (euclidean distance).
3. Compute for each cell the average color of the pixels within it.
4. Replace the pixels of *p_replace* percent of all cells by their average color. Do not change the pixels of `(1 - p_replace)` percent of all cells. (The percentages are average values over many images. Some images may get more/less cells replaced by their average color.)

API link: *Voronoi*

Example. Create an augmenter that places a `20x40` (`HxW`) grid of cells on the image and replaces all pixels within each cell by the cell's average color. The process is performed at an image size not exceeding 128px on any side. If necessary, the downscaling is performed using linear interpolation.

```
import imgaug.augmenters as iaa
points_sampler = iaa.RegularGridPointsSampler(n_cols=20, n_rows=40)
aug = iaa.Voronoi(points_sampler)
```

Example. Create a voronoi augmenter that generates a grid of cells dynamically adapted to the image size. Larger images get more cells. On the x-axis, the distance between two cells is `w * W` pixels, where `W` is the width of the image and `w` is always `0.1`. On the y-axis, the distance between two cells is `h * H` pixels, where `H` is the height of the image and `h` is sampled uniformly from the interval `[0.05, 0.2]`. To make the voronoi pattern less regular, about `20` percent of the cell coordinates are randomly dropped (i.e. the remaining cells grow in size). In contrast to the first example, the image is not resized (if it was, the sampling would happen *after* the resizing, which would affect `W` and `H`). Not all voronoi cells are replaced by their average color, only around `90` percent of them. The remaining `10` percent's pixels remain unchanged.

```
points_sampler = iaa.DropoutPointsSampler(
    iaa.RelativeRegularGridPointsSampler(
        n_cols_frac=(0.05, 0.2),
        n_rows_frac=0.1),
```

(continues on next page)

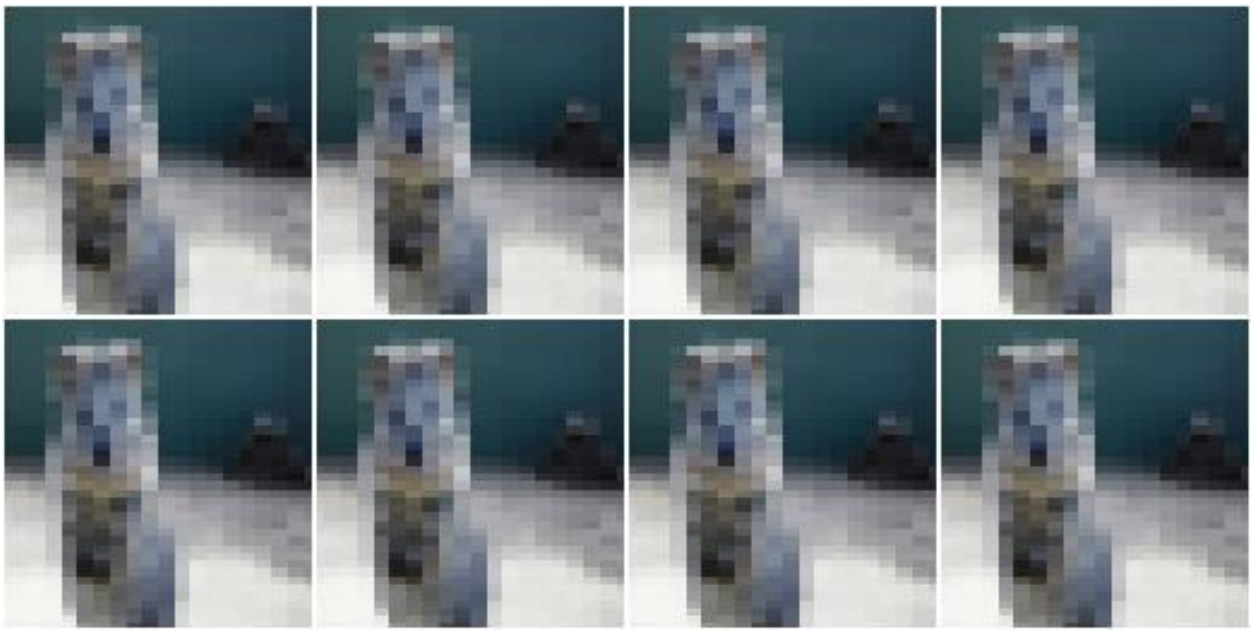

(continued from previous page)

```
    0.2)
aug = iaa.Voronoi(points_sampler, p_replace=0.9, max_size=None)
```

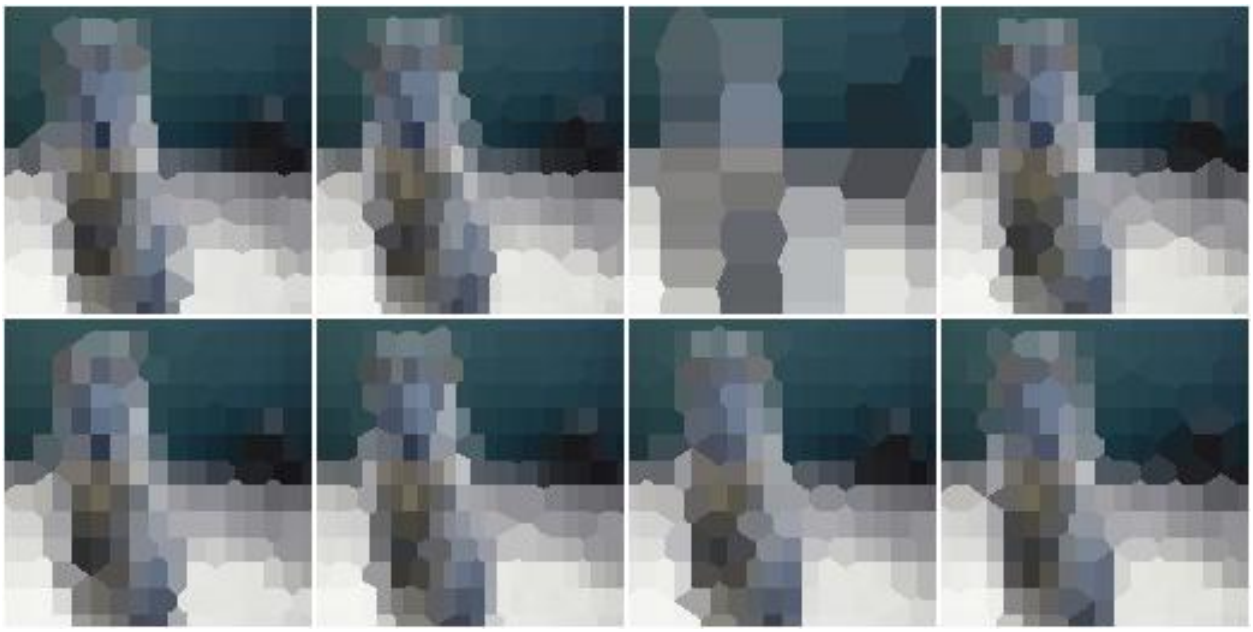

9.17.3 UniformVoronoi

Uniformly sample Voronoi cells on images and average colors within them.

This augmenter is a shortcut for the combination of `Voronoi` with `UniformPointsSampler`. Hence, it generates a fixed amount of `N` random coordinates of voronoi cells on each image. The cell coordinates are sampled uniformly using the image height and width as maxima.

API link: *UniformVoronoi*

Example. Sample for each image uniformly the number of voronoi cells `N` from the interval `[100, 500]`. Then generates `N` coordinates by sampling uniformly the x-coordinates from `[0, W]` and the y-coordinates from `[0, H]`, where `H` is the image height and `W` the image width. Then uses these coordinates to group the image pixels into voronoi cells and averages the colors within them. The process is performed at an image size not exceeding 128px on any side. If necessary, the downscaling is performed using linear interpolation.

```
import imgaug.augmenters as iaa
aug = iaa.UniformVoronoi((100, 500))
```

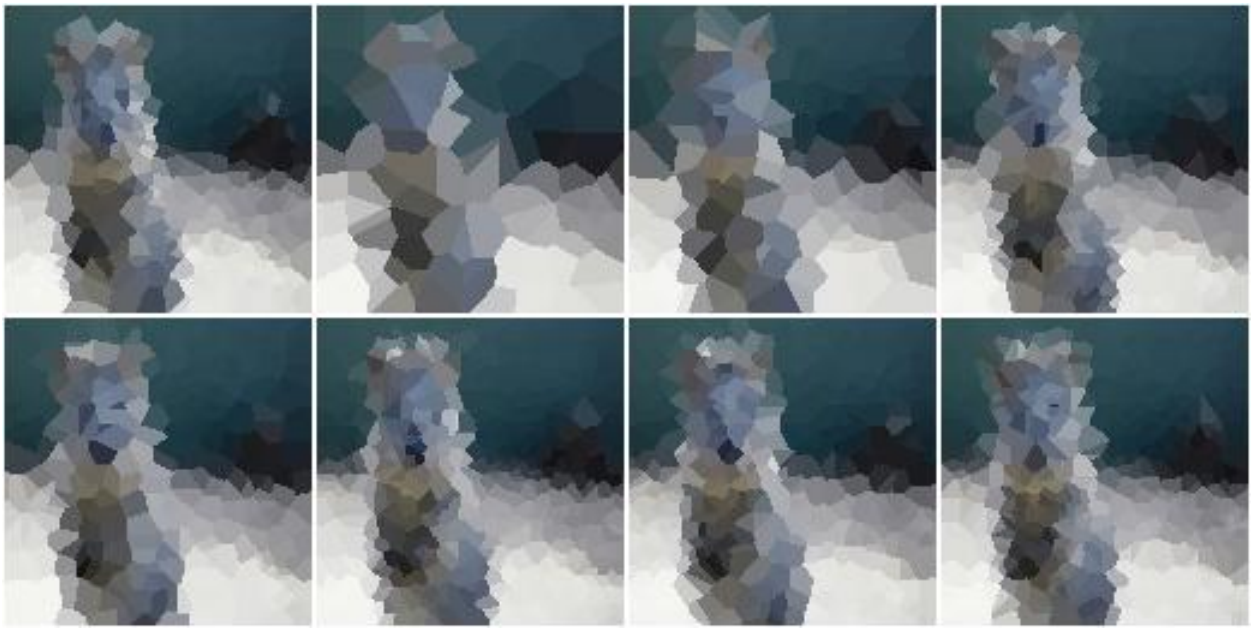

Example. Same as above, but always samples `N=250` cells, replaces only `90` percent of them with their average color (the pixels of the remaining `10` percent are not changed) and performs the transformation at the original image size.

```
aug = iaa.UniformVoronoi(250, p_replace=0.9, max_size=None)
```

9.17.4 RegularGridVoronoi

Sample Voronoi cells from regular grids and color-average them.

This augmenter is a shortcut for the combination of `Voronoi`, `RegularGridPointsSampler` and `DropoutPointsSampler`. Hence, it generates a regular grid with `R` rows and `C` columns of coordinates on each image. Then, it drops `p` percent of the `R*C` coordinates to randomize the grid. Each image pixel then belongs to the voronoi cell with the closest coordinate.

API link: *RegularGridVoronoi*

Example. Place a regular grid of `10x20` (`height x width`) coordinates on each image. Randomly drop on average `20` percent of these points to create a less regular pattern. Then use the remaining coordinates to group the image pixels into voronoi cells and average the colors within them. The process is performed at an image size not exceeding 128px on any side. If necessary, the downscaling is performed using linear interpolation.

```
import imgaug.augmenters as iaa
aug = iaa.RegularGridVoronoi(10, 20)
```

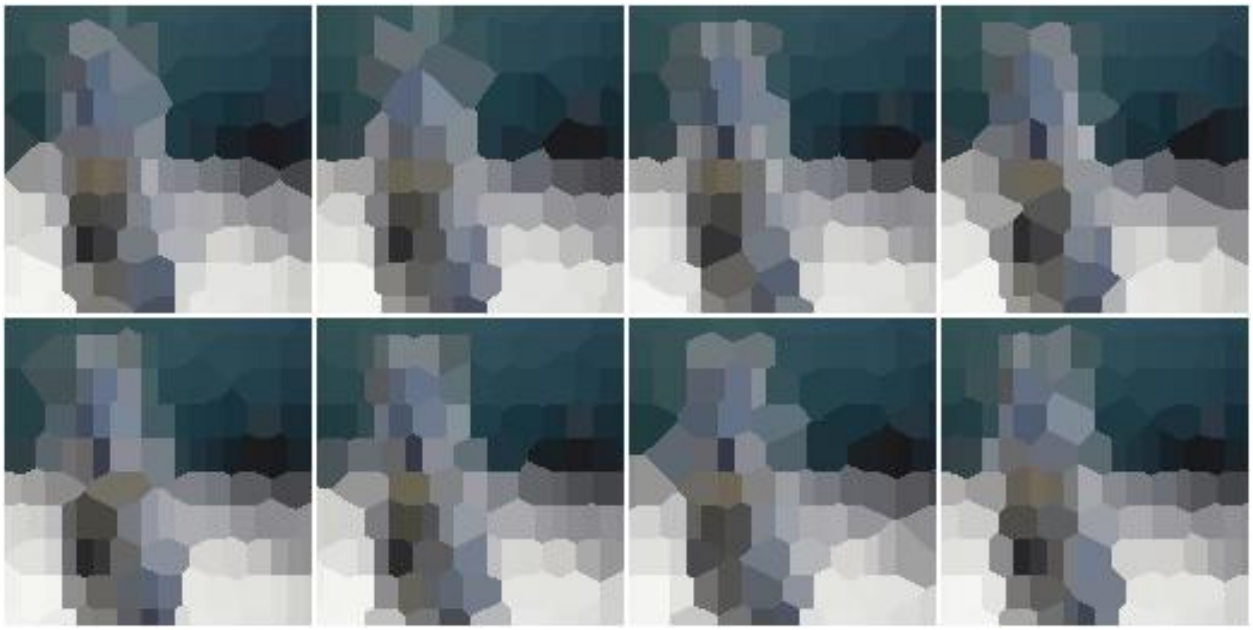

Example. Same as above, generates a grid with randomly `10` to `30` rows, drops none of the generated points, replaces only `90` percent of the voronoi cells with their average color (the pixels of the remaining `10` percent are not changed) and performs the transformation at the original image size.

```
aug = iaa.RegularGridVoronoi(
    (10, 30), 20, p_drop_points=0.0, p_replace=0.9, max_size=None)
```

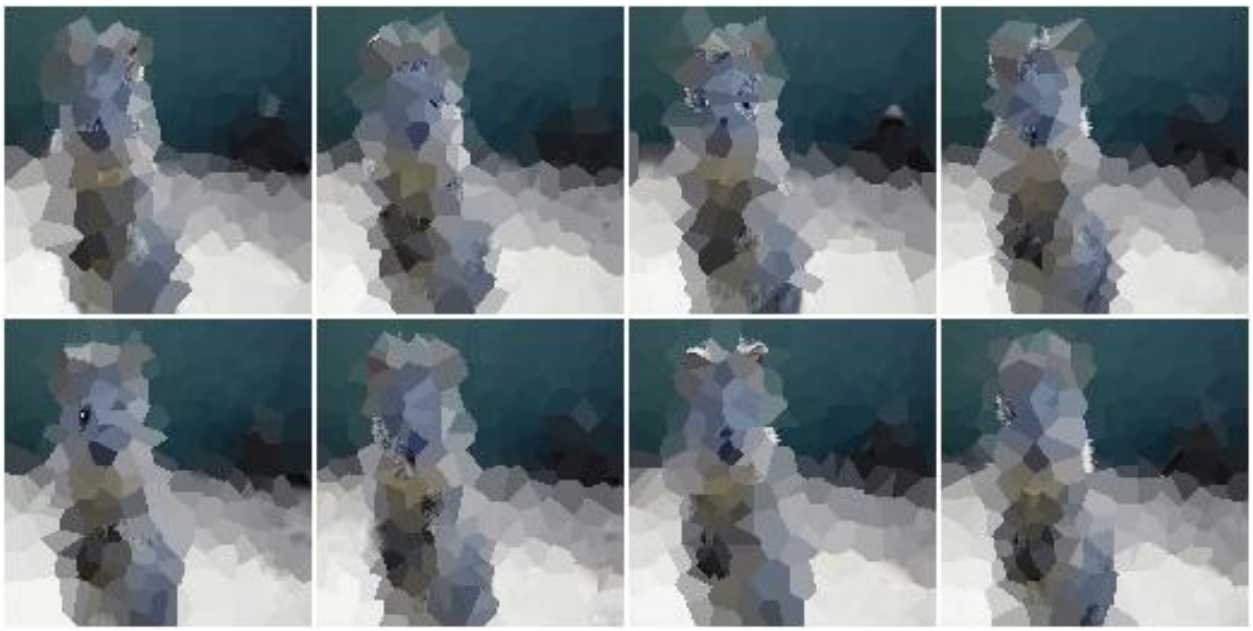

9.17.5 RelativeRegularGridVoronoi

Sample Voronoi cells from image-dependent grids and color-average them.

This augmenter is a shortcut for the combination of `Voronoi`, `RegularGridPointsSampler` and `DropoutPointsSampler`. Hence, it generates a regular grid with `R` rows and `C` columns of coordinates on each image. Then, it drops `p` percent of the `R*C` coordinates to randomize the grid. Each image pixel then belongs to the voronoi cell with the closest coordinate.

Note: In contrast to the other Voronoi augmenters, this one uses `None` as the default value for *max_size*, i.e. the color averaging is always performed at full resolution. This enables the augmenter to make most use of the added points for larger images. It does however slow down the augmentation process.

API link: *RelativeRegularGridVoronoi*

Example. Place a regular grid of `R x C` coordinates on each image, where `R` is the number of rows and computed as `R=0.1*H` with `H` being the height of the input image. `C` is the number of columns and analogously estimated from the image width `W` as `C=0.25*W`. Larger images will lead to larger `R` and `C` values. On average, `20` percent of these grid coordinates are randomly dropped to create a less regular pattern. Then, the remaining coordinates are used to group the image pixels into voronoi cells and the colors within them are averaged.

```
import imgaug.augmenters as iaa
aug = iaa.RelativeRegularGridVoronoi(0.1, 0.25)
```

Example. Same as above, generates a grid with randomly `R=r*H` rows, where `r` is sampled uniformly from the interval `[0.03, 0.1]` and `C=0.1*W` rows. No points are dropped. The augmenter replaces only `90` percent of the

voronoi cells with their average color (the pixels of the remaining `10` percent are not changed). Images larger than `512px` are temporarily downscaled (*before* sampling the grid points) so that no side exceeds `512px`. This improves performance, but degrades the quality of the resulting image.

```
aug = iaa.RelativeRegularGridVoronoi(
    (0.03, 0.1), 0.1, p_drop_points=0.0, p_replace=0.9, max_size=512)
```

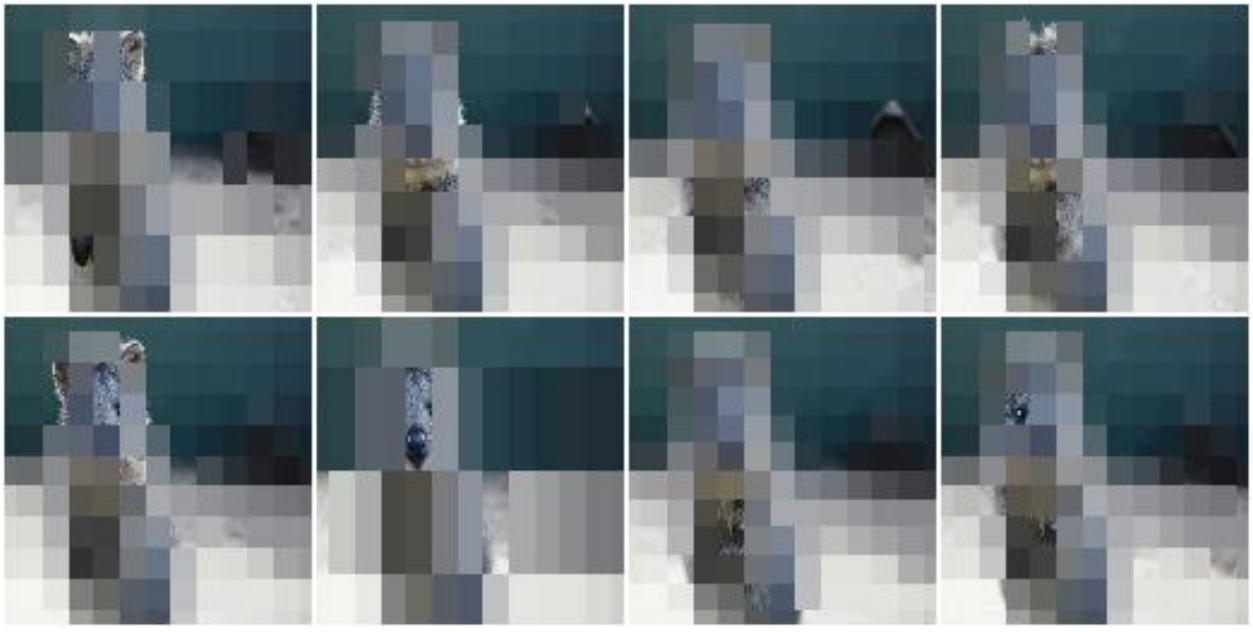

9.18 augmenters.size

9.18.1 Resize

Augmenter that resizes images to specified heights and widths.

API link: *Resize*

Example. Resize each image to height=32 and width=64:

```
import imgaug.augmenters as iaa
aug = iaa.Resize({"height": 32, "width": 64})
```

Example. Resize each image to height=32 and keep the aspect ratio for width the same:

```
aug = iaa.Resize({"height": 32, "width": "keep-aspect-ratio"})
```

Example. Resize each image to something between 50 and 100% of its original size:

```
aug = iaa.Resize((0.5, 1.0))
```

Example. Resize each image's height to 50-75% of its original size and width to either 16px or 32px or 64px:

```
aug = iaa.Resize({"height": (0.5, 0.75), "width": [16, 32, 64]})
```

9.18.2 CropAndPad

Crop/pad images by pixel amounts or fractions of image sizes.

Cropping removes pixels at the sides (i.e. extracts a subimage from a given full image). Padding adds pixels to the sides (e.g. black pixels).

Note: This augmenter automatically resizes images back to their original size after it has augmented them. To deactivate this, add the parameter `keep_size=False`.

API link: *CropAndPad*

Example. Crop or pad each side by up to 10 percent relative to its original size (negative values result in cropping, positive in padding):

```
import imgaug.augmenters as iaa
aug = iaa.CropAndPad(percent=(-0.25, 0.25))
```

Example. Pad each side by 0 to 20 percent. This adds new pixels to the sides. These pixels will either be filled with a constant value (mode=constant) or filled with the value on the closest edge (mode=edge). If a constant value is used, it will be a random value between 0 and 128 (sampled per image).

```
aug = iaa.CropAndPad(
    percent=(0, 0.2),
    pad_mode=["constant", "edge"],
    pad_cval=(0, 128)
)
```

Example. Pad the top side of each image by 0 to 30 pixels, the right side by 0-10px, bottom side by 0-30px and left side by 0-10px. Use any of the available modes to fill new pixels and if the mode is *constant* then use a constant value between 0 and 128.

```
aug = iaa.CropAndPad(
    px=((0, 30), (0, 10), (0, 30), (0, 10)),
    pad_mode=ia.ALL,
    pad_cval=(0, 128)
)
```

Example. Crop/pad each side by up to 10px. The value will be sampled once per image and used for all sides (i.e. all sides gain/lose the same number of rows/colums).

```
aug = iaa.CropAndPad(
    px=(-10, 10),
    sample_independently=False
)
```

9.18.3 Pad

Pad images, i.e. adds columns/rows of pixels to them.

This is a shortcut for `CropAndPad`. It only accepts positive pixel/percent values.

API link: *Pad*

9.18.4 Crop

Crop images, i.e. remove columns/rows of pixels at the sides of images.

This is a shortcut for `CropAndPad`. It only accepts positive pixel/percent values and transfers them as negative values to `CropAndPad`.

API link: *Crop*

9.18.5 PadToFixedSize

Pad images to minimum width/height.

If images are already at the minimum width/height or are larger, they will not be padded. Note that this also means that images will not be cropped if they exceed the required width/height.

The augmenter randomly decides per image how to distribute the required padding amounts over the image axis. E.g. if 2px have to be padded on the left or right to reach the required width, the augmenter will sometimes add 2px to the left and 0px to the right, sometimes add 2px to the right and 0px to the left and sometimes add 1px to both sides. Set *position* to `center` to prevent that.

API link: *PadToFixedSize*

Example. For image sides smaller than `100` pixels, pad to `100` pixels. Do nothing for the other edges. The padding is randomly (uniformly) distributed over the sides, so that e.g. sometimes most of the required padding is applied to the left, sometimes to the right (analogous top/bottom). The input image here has a size of `80x80`.

```
import imgaug.augmenters as iaa
aug = iaa.PadToFixedSize(width=100, height=100)
```

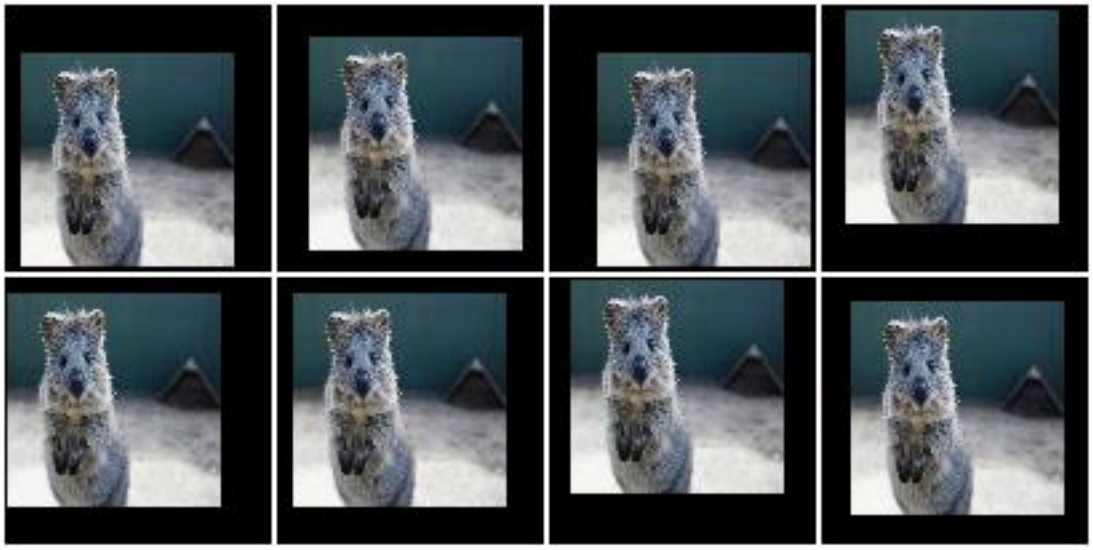

Example. For image sides smaller than `100` pixels, pad to `100` pixels. Do nothing for the other image sides. The padding is always equally distributed over the left/right and top/bottom sides. The input image here has a size of `80x80`.

```
aug = iaa.PadToFixedSize(width=100, height=100, position="center")
```

Example. For image sides smaller than `100` pixels, pad to `100` pixels and use any possible padding mode for that. Do nothing for the other image sides. The padding is always equally distributed over the left/right and top/bottom sides. The input image here has a size of `80x80`.

```
aug = iaa.PadToFixedSize(width=100, height=100, pad_mode=ia.ALL)
```

Example. Pad images smaller than `100x100` until they reach `100x100`. Analogously, crop images larger than `100x100` until they reach `100x100`. The output images therefore have a fixed size of `100x100`. The input image here has a size of `80x120`, so that the top/bottom sides have to be cropped and the left/right sides have to be padded. Note that the original image was resized to `80x120`, leading to a bit of an distorted appearance.

```
aug = iaa.Sequential([
    iaa.PadToFixedSize(width=100, height=100),
    iaa.CropToFixedSize(width=100, height=100)
])
```

9.18.6 CropToFixedSize

Crop images down to a fixed maximum width/height.

If images are already at the maximum width/height or are smaller, they will not be cropped. Note that this also means that images will not be padded if they are below the required width/height.

The augmenter randomly decides per image how to distribute the required cropping amounts over the image axis. E.g. if 2px have to be cropped on the left or right to reach the required width, the augmenter will sometimes remove 2px from the left and 0px from the right, sometimes remove 2px from the right and 0px from the left and sometimes remove 1px from both sides. Set *position* to `center` to prevent that.

API link: *CropToFixedSize*

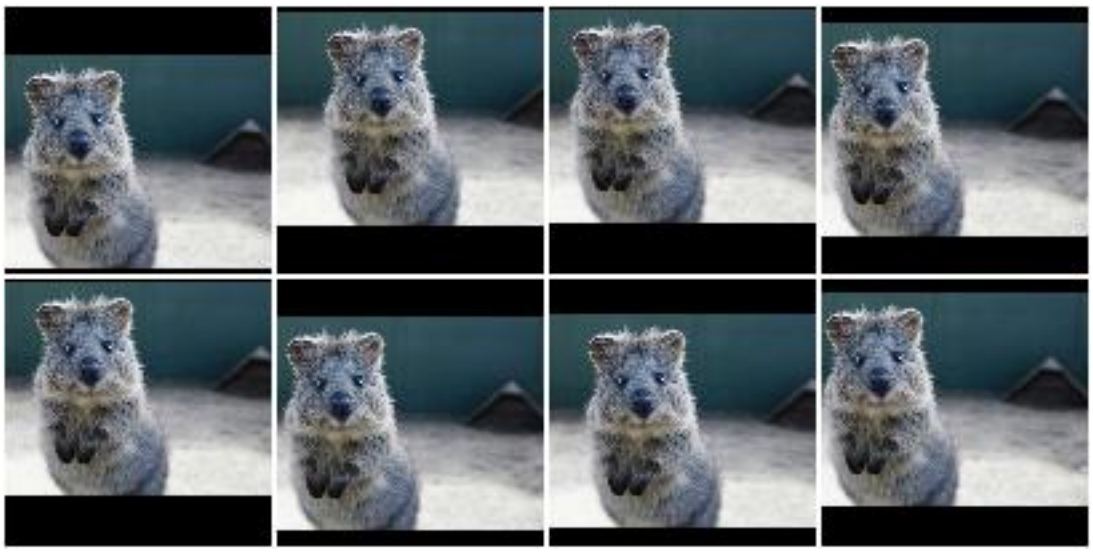

Example. For image sides larger than `100` pixels, crop to `100` pixels. Do nothing for the other sides. The cropping amounts are randomly (and uniformly) distributed over the sides of the image. The input image here has a size of `120x120`.

```
import imgaug.augmenters as iaa
aug = iaa.CropToFixedSize(width=100, height=100)
```

Example. For sides larger than `100` pixels, crop to `100` pixels. Do nothing for the other sides. The cropping amounts are always equally distributed over the left/right sides of the image (and analogously for top/bottom). The input image here has a size of `120x120`.

```
aug = iaa.CropToFixedSize(width=100, height=100, position="center")
```

Example. Pad images smaller than `100x100` until they reach `100x100`. Analogously, crop images larger than `100x100` until they reach `100x100`. The output images therefore have a fixed size of `100x100`. The input image here has a size of `80x120`, so that the top/bottom sides have to be cropped and the left/right sides have to be padded. Note that the original image was resized to `80x120`, leading to a bit of an distorted appearance.

```
aug = iaa.Sequential([
    iaa.PadToFixedSize(width=100, height=100),
    iaa.CropToFixedSize(width=100, height=100)
])
```

9.18.7 PadToMultiplesOf

Pad images until their height/width is a multiple of a value.

API link: *PadToMultiplesOf*

Example. Create an augmenter that pads images to multiples of `10` along the y-axis (i.e. 10, 20, 30, ...) and to multiples of `6` along the x-axis (i.e. 6, 12, 18, ...). The rows to be padded will be spread *randomly* over the top and bottom sides (analogous for the left/right sides).

```
import imgaug.augmenters as iaa
aug = iaa.PadToMultiplesOf(height_multiple=10, width_multiple=6)
```

9.18.8 CropToMultiplesOf

Crop images down until their height/width is a multiple of a value.

Note: For a given axis size `A` and multiple `M`, if `A` is in the interval `[0 .. M]`, the axis will not be changed. As a result, this augmenter can still produce axis sizes that are not multiples of the given values.

API link: *CropToMultiplesOf*

Example. Create an augmenter that crops images to multiples of `10` along the y-axis (i.e. 10, 20, 30, ...) and to multiples of `6` along the x-axis (i.e. 6, 12, 18, ...). The rows to be cropped will be spread *randomly* over the top and bottom sides (analogous for the left/right sides).

```
import imgaug.augmenters as iaa
aug = iaa.CropToMultiplesOf(height_multiple=10, width_multiple=6)
```

9.18.9 CropToPowersOf

Crop images until their height/width is a power of a base.

This augmenter removes pixels from an axis with size `S` leading to the new size `S'` until `S' = B^E` is fulfilled, where `B` is a provided base (e.g. `2`) and `E` is an exponent from the discrete interval `[1 .. inf)`.

Note: This augmenter does nothing for axes with size less than `B^1 = B`. If you have images with `S < B^1`, it is recommended to combine this augmenter with a padding augmenter that pads each axis up to `B`.

API link: *CropToPowersOf*

Example. Create an augmenter that crops each image down to powers of `3` along the y-axis (i.e. 3, 9, 27, ...) and powers of `2` along the x-axis (i.e. 2, 4, 8, 16, ...). The rows to be cropped will be spread *randomly* over the top and bottom sides (analogous for the left/right sides).

```
import imgaug.augmenters as iaa
aug = iaa.CropToPowersOf(height_base=3, width_base=2)
```

9.18.10 PadToPowersOf

Pad images until their height/width is a power of a base.

This augmenter adds pixels to an axis with size `S` leading to the new size `S'` until `S' = B^E` is fulfilled, where `B` is a provided base (e.g. `2`) and `E` is an exponent from the discrete interval `[1 .. inf)`.

API link: *PadToPowersOf*

Example. Create an augmenter that pads each image to powers of `3` along the y-axis (i.e. 3, 9, 27, ...) and powers of `2` along the x-axis (i.e. 2, 4, 8, 16, ...). The rows to be padded will be spread *randomly* over the top and bottom sides (analogous for the left/right sides).

```
import imgaug.augmenters as iaa
aug = iaa.PadToPowersOf(height_base=3, width_base=2)
```

9.18.11 CropToAspectRatio

Crop images until their width/height matches an aspect ratio.

This augmenter removes either rows or columns until the image reaches the desired aspect ratio given in `width / height`. The cropping operation is stopped once the desired aspect ratio is reached or the image side to crop reaches a size of `1`. If any side of the image starts with a size of `0`, the image will not be changed.

API link: *CropToAspectRatio*

Example. Create an augmenter that crops each image until its aspect ratio is as close as possible to `2.0` (i.e. two times as many pixels along the x-axis than the y-axis). The rows to be cropped will be spread *randomly* over the top and bottom sides (analogous for the left/right sides).

```
import imgaug.augmenters as iaa
aug = iaa.CropToAspectRatio(2.0)
```

9.18.12 PadToAspectRatio

Pad images until their width/height matches an aspect ratio.

This augmenter adds either rows or columns until the image reaches the desired aspect ratio given in `width / height`.

API link: *PadToAspectRatio*

Example. Create an augmenter that pads each image until its aspect ratio is as close as possible to `2.0` (i.e. two times as many pixels along the x-axis than the y-axis). The rows to be padded will be spread *randomly* over the top and bottom sides (analogous for the left/right sides).

```
import imgaug.augmenters as iaa
aug = iaa.PadToAspectRatio(2.0)
```

9.18.13 CropToSquare

Crop images until their width and height are identical.

This is identical to *imgaug.augmenters.size.CropToAspectRatio* with `aspect_ratio=1.0`.

Images with axis sizes of `0` will not be altered.

API link: *CropToSquare*

Example. Create an augmenter that crops each image until its square, i.e. height and width match. The rows to be cropped will be spread *randomly* over the top and bottom sides (analogous for the left/right sides).

```
import imgaug.augmenters as iaa
aug = iaa.CropToSquare()
```

9.18.14 PadToSquare

Pad images until their height and width are identical.

This augmenter is identical to *imgaug.augmenters.size.PadToAspectRatio* with `aspect_ratio=1.0`.

API link: *PadToSquare*

Example. Create an augmenter that pads each image until its square, i.e. height and width match. The rows to be padded will be spread *randomly* over the top and bottom sides (analogous for the left/right sides).

```
import imgaug.augmenters as iaa
aug = iaa.PadToSquare()
```

9.18.15 CenterPadToFixedSize

Pad images equally on all sides up to given minimum heights/widths.

This is an alias for *imgaug.augmenters.size.PadToFixedSize* with `position="center"`. It spreads the pad amounts equally over all image sides, while *imgaug.augmenters.size.PadToFixedSize* by defaults spreads them randomly.

API link: *CenterPadToFixedSize*

Example. Create an augmenter that pads images up to `20x30`, with the padded rows added *equally* on the top and bottom (analogous for the padded columns).

```
import imgaug.augmenters as iaa
aug = iaa.CenterPadToFixedSize(height=20, width=30)
```

9.18.16 CenterCropToFixedSize

Take a crop from the center of each image.

This is an alias for *imgaug.augmenters.size.CropToFixedSize* with `position="center"`.

Note: If images already have a width and/or height below the provided width and/or height then this augmenter will do nothing for the respective axis. Hence, resulting images can be smaller than the provided axis sizes.

API link: *CenterCropToFixedSize*

Example. Create an augmenter that takes `20x10` sized crops from the center of images:

```
import imgaug.augmenters as iaa
crop = iaa.CenterCropToFixedSize(height=20, width=10)
```

9.18.17 CenterCropToMultiplesOf

Crop images equally on all sides until H/W are multiples of given values.

This is the same as *imgaug.augmenters.size.CropToMultiplesOf*, but uses `position="center"` by default, which spreads the crop amounts equally over all image sides, while *imgaug.augmenters.size.CropToMultiplesOf* by default spreads them randomly.

API link: *CenterCropToMultiplesOf*

Example. Create an augmenter that crops images to multiples of `10` along the y-axis (i.e. 10, 20, 30, ...) and to multiples of `6` along the x-axis (i.e. 6, 12, 18, ...). The rows to be cropped will be spread *equally* over the top and bottom sides (analogous for the left/right sides).

```
import imgaug.augmenters as iaa
aug = iaa.CenterCropToMultiplesOf(height_multiple=10, width_multiple=6)
```

9.18.18 CenterPadToMultiplesOf

Pad images equally on all sides until H/W are multiples of given values.

This is the same as *imgaug.augmenters.size.PadToMultiplesOf*, but uses `position="center"` by default, which spreads the pad amounts equally over all image sides, while *imgaug.augmenters.size.PadToMultiplesOf* by default spreads them randomly.

API link: *CenterPadToMultiplesOf*

Example. Create an augmenter that pads images to multiples of `10` along the y-axis (i.e. 10, 20, 30, ...) and to multiples of `6` along the x-axis (i.e. 6, 12, 18, ...). The rows to be padded will be spread *equally* over the top and bottom sides (analogous for the left/right sides).

```
import imgaug.augmenters as iaa
aug = iaa.CenterPadToMultiplesOf(height_multiple=10, width_multiple=6)
```

9.18.19 CenterCropToPowersOf

Crop images equally on all sides until H/W is a power of a base.

This is the same as *imgaug.augmenters.size.CropToPowersOf*, but uses `position="center"` by default, which spreads the crop amounts equally over all image sides, while *imgaug.augmenters.size.CropToPowersOf* by default spreads them randomly.

API link: *CenterCropToPowersOf*

Example. Create an augmenter that crops each image down to powers of `3` along the y-axis (i.e. 3, 9, 27, ...) and powers of `2` along the x-axis (i.e. 2, 4, 8, 16, ...). The rows to be cropped will be spread *equally* over the top and bottom sides (analogous for the left/right sides).

```
import imgaug.augmenters as iaa
aug = iaa.CropToPowersOf(height_base=3, width_base=2)
```

9.18.20 CenterPadToPowersOf

Pad images equally on all sides until H/W is a power of a base.

This is the same as *imgaug.augmenters.size.PadToPowersOf*, but uses `position="center"` by default, which spreads the pad amounts equally over all image sides, while *imgaug.augmenters.size.PadToPowersOf* by default spreads them randomly.

API link: *CenterPadToPowersOf*

Example. Create an augmenter that pads each image to powers of `3` along the y-axis (i.e. 3, 9, 27, ...) and powers of `2` along the x-axis (i.e. 2, 4, 8, 16, ...). The rows to be padded will be spread *equally* over the top and bottom sides (analogous for the left/right sides).

```
import imgaug.augmenters as iaa
aug = iaa.CenterPadToPowersOf(height_base=3, width_base=2)
```

9.18.21 CenterCropToAspectRatio

Crop images equally on all sides until they reach an aspect ratio.

This is the same as *imgaug.augmenters.size.CropToAspectRatio*, but uses `position="center"` by default, which spreads the crop amounts equally over all image sides, while *imgaug.augmenters.size.CropToAspectRatio* by default spreads them randomly.

API link: *CenterCropToAspectRatio*

Example. Create an augmenter that crops each image until its aspect ratio is as close as possible to `2.0` (i.e. two times as many pixels along the x-axis than the y-axis). The rows to be cropped will be spread *equally* over the top and bottom sides (analogous for the left/right sides).

```
import imgaug.augmenters as iaa
aug = iaa.CenterCropToAspectRatio(2.0)
```

9.18.22 CenterPadToAspectRatio

Pad images equally on all sides until H/W matches an aspect ratio.

This is the same as *imgaug.augmenters.size.PadToAspectRatio*, but uses `position="center"` by default, which spreads the pad amounts equally over all image sides, while *imgaug.augmenters.size.PadToAspectRatio* by default spreads them randomly.

API link: *CenterPadToAspectRatio*

Example. Create am augmenter that pads each image until its aspect ratio is as close as possible to `2.0` (i.e. two times as many pixels along the x-axis than the y-axis). The rows to be padded will be spread *equally* over the top and bottom sides (analogous for the left/right sides).

```
import imgaug.augmenters as iaa
aug = iaa.PadToAspectRatio(2.0)
```

9.18.23 CenterCropToSquare

Crop images equally on all sides until their height/width are identical.

In contrast to *imgaug.augmenters.size.CropToSquare*, this augmenter always tries to spread the columns/rows to remove equally over both sides of the respective axis to be cropped. *imgaug.augmenters.size.CropToAspectRatio* by default spreads the croppings randomly.

This augmenter is identical to *imgaug.augmenters.size.CropToSquare* with `position="center"`, and thereby the same as *imgaug.augmenters.size.CropToAspectRatio* with `aspect_ratio=1.0, position="center"`.

Images with axis sizes of `0` will not be altered.

API link: *CenterCropToSquare*

Example. Create an augmenter that crops each image until its square, i.e. height and width match. The rows to be cropped will be spread *equally* over the top and bottom sides (analogous for the left/right sides).

```
import imgaug.augmenters as iaa
aug = iaa.CenterCropToSquare()
```

9.18.24 CenterPadToSquare

Pad images equally on all sides until their height & width are identical.

This is the same as *imgaug.augmenters.size.PadToSquare*, but uses `position="center"` by default, which spreads the pad amounts equally over all image sides, while *imgaug.augmenters.size.PadToSquare* by default spreads them randomly. This augmenter is thus also identical to *imgaug.augmenters.size.PadToAspectRatio* with `aspect_ratio=1.0, position="center"`.

API link: *CenterPadToSquare*

Example. Create an augmenter that pads each image until its square, i.e. height and width match. The rows to be padded will be spread *equally* over the top and bottom sides (analogous for the left/right sides).

```
import imgaug.augmenters as iaa
aug = iaa.CenterPadToSquare()
```

9.18.25 KeepSizeByResize

Resize images back to their input sizes after applying child augmenters.

Combining this with e.g. a cropping augmenter as the child will lead to images being resized back to the input size after the crop operation was applied. Some augmenters have a `keep_size` argument that achieves the same goal (if set to `True`), though this augmenter offers control over the interpolation mode and which augmentables to resize (images, heatmaps, segmentation maps).

API link: *KeepSizeByResize*

Example. Apply random cropping to input images, then resize them back to their original input sizes. The resizing is done using this augmenter instead of the corresponding internal resizing operation in `Crop`.

```
import imgaug.augmenters as iaa
aug = iaa.KeepSizeByResize(
    iaa.Crop((20, 40), keep_size=False)
)
```

Example. Same as in the previous example, but images are now always resized using nearest neighbour interpolation.

```
aug = iaa.KeepSizeByResize(
    iaa.Crop((20, 40), keep_size=False),
    interpolation="nearest"
)
```

Example. Similar to the previous example, but images are now sometimes resized using linear interpolation and sometimes using nearest neighbour interpolation. Heatmaps are resized using the same interpolation as was used for the corresponding image. Segmentation maps are not resized and will therefore remain at their size after cropping.

```
aug = iaa.KeepSizeByResize(
    iaa.Crop((20, 40), keep_size=False),
    interpolation=["nearest", "cubic"],
    interpolation_heatmaps=iaa.KeepSizeByResize.SAME_AS_IMAGES,
    interpolation_segmaps=iaa.KeepSizeByResize.NO_RESIZE
)
```

9.19 augmenters.weather

Note: All examples below use the following input image:

9.19.1 FastSnowyLandscape

Convert non-snowy landscapes to snowy ones.

This augmenter expects to get an image that roughly shows a landscape.

API link: *FastSnowyLandscape*

Example. Search for all pixels in the image with a lightness value in HLS colorspace of less than `140` and increase their lightness by a factor of `2.5`.

```
import imgaug.augmenters as iaa
aug = iaa.FastSnowyLandscape(
    lightness_threshold=140,
    lightness_multiplier=2.5
)
```

Example. Search for all pixels in the image with a lightness value in HLS colorspace of less than `128` or less than `200` (one of these values is picked per image) and multiply their lightness by a factor of `x` with `x` being sampled from `uniform(1.5, 3.5)` (once per image).

```
aug = iaa.FastSnowyLandscape(
    lightness_threshold=[128, 200],
    lightness_multiplier=(1.5, 3.5)
)
```

Example. Similar to the previous example, but the lightness threshold is sampled from `uniform(100, 255)` (per image) and the multiplier from `uniform(1.0, 4.0)` (per image). This seems to produce good and varied results.

```
aug = iaa.FastSnowyLandscape(
    lightness_threshold=(100, 255),
    lightness_multiplier=(1.0, 4.0)
)
```

9.19.2 Clouds

Add clouds to images.

This is a wrapper around `CloudLayer`. It executes 1 to 2 layers per image, leading to varying densities and frequency patterns of clouds.

This augmenter seems to be fairly robust w.r.t. the image size. Tested with `96x128`, `192x256` and `960x1280`.

API link: *`Clouds()`*

Example. Create an augmenter that adds clouds to images:

```
import imgaug.augmenters as iaa
aug = iaa.Clouds()
```

9.19.3 Fog

Add fog to images.

This is a wrapper around `CloudLayer`. It executes a single layer per image with a configuration leading to fairly dense clouds with low-frequency patterns.

This augmenter seems to be fairly robust w.r.t. the image size. Tested with `96x128`, `192x256` and `960x1280`.

API link: *Fog()*

Example. Create an augmenter that adds fog to images:

```
import imgaug.augmenters as iaa
aug = iaa.Fog()
```

9.19.4 CloudLayer

Add a single layer of clouds to an image.

API link: *CloudLayer*

9.19.5 Snowflakes

Add falling snowflakes to images.

This is a wrapper around `SnowflakesLayer`. It executes 1 to 3 layers per image.

API link: *Snowflakes()*

Example. Add snowflakes to small images (around `96x128`):

```
import imgaug.augmenters as iaa
aug = iaa.Snowflakes(flake_size=(0.1, 0.4), speed=(0.01, 0.05))
```

Example. Add snowflakes to medium-sized images (around `192x256`):

```
aug = iaa.Snowflakes(flake_size=(0.2, 0.7), speed=(0.007, 0.03))
```

Example. Add snowflakes to large images (around `960x1280`):

```
aug = iaa.Snowflakes(flake_size=(0.7, 0.95), speed=(0.001, 0.03))
```

9.19.6 SnowflakesLayer

Add a single layer of falling snowflakes to images.

API link: *SnowflakesLayer*

9.19.7 Rain

Add falling snowflakes to images.

This is a wrapper around *RainLayer*. It executes 1 to 3 layers per image.

Note: This augmenter currently seems to work best for medium-sized images around `192x256`. For smaller images, you may want to increase the *speed* value to e.g. `(0.1, 0.3)`, otherwise the drops tend to look like snowflakes. For larger images, you may want to increase the *drop_size* to e.g. `(0.10, 0.20)`.

API link: *Rain*

Example. Add rain to small images (around `96x128`):

```
import imgaug.augmenters as iaa
aug = iaa.Rain(speed=(0.1, 0.3))
```

Example. Add rain to medium sized images (around `192x256`):

```
aug = iaa.Rain()
```

Example. Add rain to large images (around `960x1280`):

```
aug = iaa.Rain(drop_size=(0.10, 0.20))
```

9.19.8 RainLayer

Add a single layer of falling raindrops to images.

API link: *RainLayer*

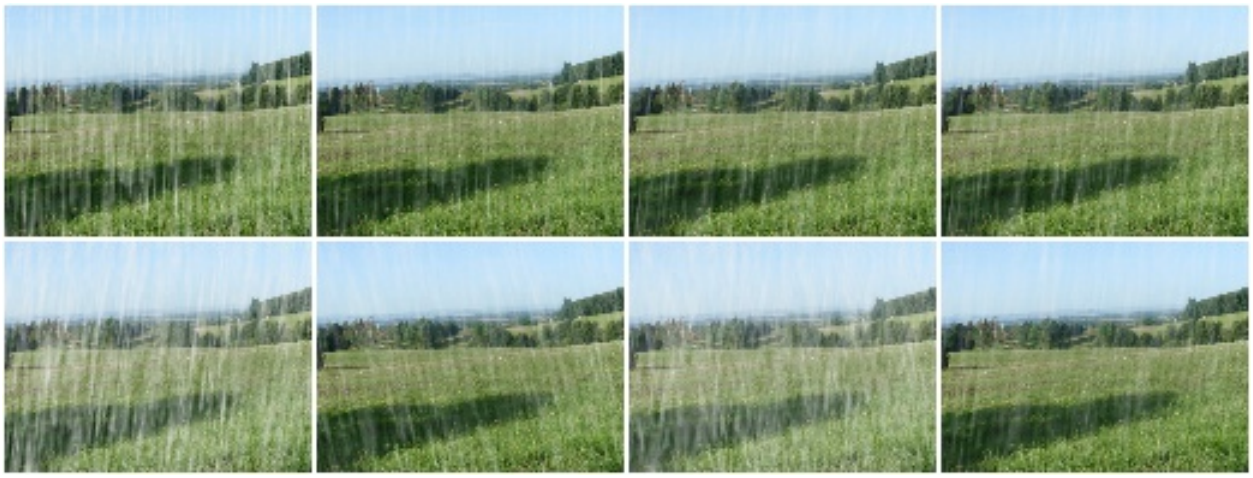

CHAPTER 10

Performance

Below are performance measurements of each augmenter for image augmentation (`augment_images()`), heatmap augmentation (`augment_heatmaps()`) and keypoint/landmark augmentation (`augment_keypoints()`). (Last updated for `0.3.0`)

System: The numbers were computed based on a haswell-generation i7 3.2Ghz CPU with DDR3 memory. That is a rather dated system by today's standards. A modern, high-end system should achieve higher bandwidths.

All experiments were conducted using python 3.7 and numpy 1.17.0. Note that the precise python/numpy version can have significant impact on your performance.

Experiments Settings: All augmenters were run with reasonable parameter choices that should reflect expected real-world usage, while avoiding too simple parameter values that would lead to inflated scores. Some parameter choices are listed below, the remaining ones can be looked up in `measure_performance.py`. Kernel sizes were all set to `3x3`, unless otherwise mentioned. The inputs focused on a small and large image-size setting, using `64x64x3` and `224x224x3` as the respective sizes. The base image was taken from skimage.data.astronaut, which should be a representative real-world image. Batch sizes of `1` and `128` were tested. Each augmenter was run at least `40` times on the generated input and the average of the measured runtimes was computed to derive bandwidth in mbit per second and th raw number of augmented items (e.g. images) per second.

10.1 Results Overview

From the results, the following points can be derived.

Inputs:

- Use large batch sizes whenever possible. Many augmenters are significantly faster with these.
- Large image sizes lead to higher throughput based on mbit/sec. Smaller images lead to lower throughput, but significantly more items/sec (roughly 4-10x more). Use small images whenever possible.
- For keypoint-based and heatmap-based augmentation, try to increase the number of items per augmented instance. E.g. `augment_keypoints()` accepts a list of `KeypointsOnImage` instances, with each such instance representing the keypoints on an image. Try to place for each image all keypoints in the respective

`KeypointsOnImage` instance instead of splitting them into multiple such instances (which would be more work anyways). The same is true for bounding boxes, heatmaps and segmentation maps.

- Keypoint- and heatmap-based inputs are only affected by augmenters that change the geometry of the image (e.g. `Crop` or `Affine`). Other augmenters are essentially free to execute as they do not perform any changes.
- Keypoint-based augmentation is very fast for almost all augmenters, reaching several 100k keypoints per second. Slower augmenters are `ElasticTransformation` and `PiecewiseAffine`, as these currently have to fall back to image-based algorithms.

Parameter choices:

- When possible, nearest neighbour interpolation or linear interpolation should be used as these are significantly faster than other options. Most augmenters that use interpolation offer either an `order` parameter (0=nearest neighbour, 1=linear) or an `interpolation` parameter ("nearest", "linear").
- Using `keep_size=True` is the default setting in all augmenters that change image sizes. It is convenient, as it ensures that image sizes are not altered by the augmentation. It does however incur a significant performance penalty, often more than halving the bandwidth. Try `keep_size=False` when possible. You can still resize images manually after augmentation or by using `KeepSizeByResize(Sequential(<augmenters>))`.
- When augmenters offer modes to fill newly created pixels in user-defined ways (e.g. `pad_mode=constant` in `Pad` to fill up all padded pixels with a specified constant color), using `edge` instead of `constant` will usually not incur a significant performance penalty.

Specific Augmenter suggestions:

- For augmenters where an elementwise sibling exists (e.g. `Multiply` and `MultiplyElementwise`), the elementwise augmenter is usually significantly slower than the non-elementwise one.
- If blurring is required, `AverageBlur` is the fastest choice, followed by `GaussianBlur`.
- Augmenters that operate on coarser images (e.g. `CoarseDropout` vs `Dropout`) can be significantly faster than their non-coarse siblings.
- Contrast normalizing augmenters are all comparable in performance, except for histogram-based ones, which are significantly slower.
- `PiecewiseAffine` is a very slow augmenter and should usually be replaced by ElasticTransformation, which achieves similar outputs and is quite a bit faster.
- `Superpixels` is a fairly slow augmenter and should usually be wrapped in e.g. `Sometimes` to not apply it very often and reduce its performance impact.
- Weather augmenters other than `FastSnowyLandscape` are rather slow and should only be used when sensible.

10.2 Images

Numbers below are for small images (`64x64x3`) and large images (`224x224x3`). `B=1` denotes a batch size of `1`, `B=128` one of `128`.

In mbit/sec:

	64x64x3, uint8		224x224x3
Augmenter	B=1	B=128	B=1
Sequential (2xIdentity)	1114.3	24601.3	9815.7

Table 1 – continued from previous page

Sequential (2xIdentity, random_order)	903.3	25450.9	8697.9
SomeOf (1-3, 3xIdentity)	226.2	3800.6	2114.2
SomeOf (1-3, 3xIdentity, random_order)	220.9	3717.7	2037.7
OneOf (3xIdentity)	533.9	7941.3	4486.2
Sometimes (Identity)	367.3	12894.4	3763.7
WithChannels ([1,2], Identity)	541.2	4932.7	3668.9
Identity	1364.5	26741.2	11791.9
Noop	1341.4	27020.1	11611.5
Lambda (return input)	1262.1	24919.0	10837.8
AssertLambda (return True)	1244.8	26346.3	10864.2
AssertShape (None, H, W, None)	1068.8	14590.1	9860.2
ChannelShuffle (0.5)	418.4	3159.3	3285.1
Add	137.7	595.5	1972.4
AddElementwise	201.3	307.5	909.8
AdditiveGaussianNoise	167.5	218.5	695.9
AdditiveLaplaceNoise	147.2	185.2	419.3
AdditivePoissonNoise	122.5	151.0	294.5
Multiply	240.3	770.4	2156.3
MultiplyElementwise	188.0	276.9	876.7
Cutout (1 iter, constant fill)	151.0	1288.6	1723.9
Dropout (1-5%)	225.9	353.6	995.3
CoarseDropout (1-5%, size=1-10%)	133.4	172.6	1039.0
Dropout2d (10%)	324.1	5696.8	3569.2
TotalDropout (10%)	450.4	19944.9	4948.5
ReplaceElementwise	129.9	161.9	676.8
ImpulseNoise	112.4	135.2	469.9
SaltAndPepper	118.9	141.1	643.7
CoarseSaltAndPepper	86.9	98.6	666.4
Salt	99.8	114.2	590.8
CoarseSalt	78.6	86.5	606.6
Pepper	97.9	105.7	589.5
CoarsePepper	78.7	86.1	605.5
Invert (10%)	266.5	5468.8	2992.8
JpegCompression (50-99%)	81.9	103.7	420.6
Cartoon	6.0	5.8	7.1
BlendAlpha (Identity)	128.7	402.0	810.2
BlendAlphaElementwise (Identity)	130.7	207.0	450.3
BlendAlphaSimplexNoise (Identity)	24.7	28.3	175.4
BlendAlphaFrequencyNoise (Identity)	33.2	36.5	210.8
BlendAlphaSomeColors (Identity)	64.2	161.0	330.7
BlendAlphaHorizontalL.Grad. (Identity)	87.5	238.1	416.3
BlendAlphaVerticalL.Grad. (Identity)	87.9	231.9	407.5
BlendAlphaRegularGrid (Identity)	85.0	230.2	398.7
BlendAlphaCheckerboard (Identity)	86.0	200.2	399.0
GaussianBlur (sigma=(1,5))	191.7	532.7	1528.9
AverageBlur	245.4	1474.4	2021.4
MedianBlur	129.8	257.4	267.7
BilateralBlur	101.3	269.3	281.9
MotionBlur	56.2	64.1	541.7
MeanShiftBlur	1.4	1.4	1.3

Table 1 – continued from previous page

RandAugment (n=2, m=(6,12))	24.0	128.1	222.4
WithColorspace (HSV, Identity)	291.6	974.1	1691.5
WithBrightnessChannels (Identity)	43.5	736.9	1097.9
MultiplyAndAddToBrightness	71.4	251.4	665.1
MultiplyBrightness	113.4	396.8	850.6
AddToBrightness	109.3	347.8	841.2
WithHueAndSaturation	168.1	334.5	687.6
MultiplyHueAndSaturation	82.5	152.1	440.5
MultiplyHue	74.3	150.1	438.3
MultiplySaturation	57.6	150.4	442.1
RemoveSaturation	70.9	150.6	433.1
AddToHueAndSaturation	131.1	443.1	828.5
AddToHue	150.2	455.8	858.3
AddToSaturation	139.9	460.5	865.3
ChangeColorspace (HSV)	257.9	923.8	2258.6
Grayscale	143.1	308.5	632.4
KMeansColorQuantization (2-16 colors)	30.3	37.5	183.5
UniformColorQuantization (2-16 colors)	127.9	354.5	1512.0
UniformColorQuant.NBits (1-7 bits)	142.7	357.3	1508.6
Posterize (1-7 bits)	136.4	356.7	1506.4
GammaContrast	169.9	304.3	1832.4
SigmoidContrast	153.9	234.2	1551.7
LogContrast	183.9	303.1	1819.4
LinearContrast	214.2	391.9	2048.0
AllChannelsHistogramEqualization	519.5	1559.9	1858.7
HistogramEqualization	268.5	892.9	1470.4
AllChannelsCLAHE	112.9	326.2	878.5
CLAHE	112.7	442.1	824.7
Convolve (3x3)	430.6	1442.0	2833.3
Sharpen	176.7	261.5	1282.2
Emboss	176.9	262.5	1334.3
EdgeDetect	234.2	392.3	1696.6
DirectedEdgeDetect	90.3	107.5	827.8
Canny	54.6	103.9	280.3
Fliplr (p=100%)	446.1	2507.3	3359.3
Flipud (p=100%)	564.7	9721.4	5475.0
Affine (order=0, constant)	75.8	255.3	856.4
Affine (order=1, constant)	75.5	236.2	773.5
Affine (order=3, constant)	69.3	194.6	473.5
Affine (order=1, edge)	68.1	229.4	744.3
Affine (order=1, constant, skimage)	39.2	73.6	180.0
PiecewiseAffine (4x4, order=1, constant)	3.9	4.3	25.6
PiecewiseAffine (4x4, order=0, constant)	4.3	4.5	30.6
PiecewiseAffine (4x4, order=1, edge)	4.1	4.4	25.4
PiecewiseAffine (8x8, order=1, constant)	0.9	0.9	8.4
PerspectiveTransform	96.5	473.9	885.5
PerspectiveTransform (keep_size)	92.5	313.1	688.3
ElasticTransformation (order=0, constant)	82.5	158.3	543.1
ElasticTransformation (order=1, constant)	80.6	149.2	499.7
ElasticTransformation (order=1, nearest)	80.5	150.1	494.4

Table 1 – continued from previous page

ElasticTransformation (order=1, reflect)	81.5	149.1	500.4
Rot90	273.5	3981.3	3416.2
Rot90 (keep_size)	265.8	2193.9	1983.7
WithPolarWarping (Identity)	259.1	639.9	948.2
Jigsaw (rows/cols=(3,8), 1 step)	62.5	107.0	728.7
AveragePooling	97.1	175.8	434.7
AveragePooling (keep_size)	91.1	148.5	392.8
MaxPooling	100.7	187.6	460.6
MaxPooling (keep_size)	92.7	158.6	431.4
MinPooling	102.7	187.6	467.5
MinPooling (keep_size)	95.1	157.5	426.1
MedianPooling	94.3	171.6	557.5
MedianPooling (keep_size)	90.1	146.4	513.8
imgcorruptlike.GaussianNoise((1,5))	70.9	130.5	180.7
imgcorruptlike.ShotNoise((1,5))	52.0	68.0	83.9
imgcorruptlike.ImpulseNoise((1,5))	62.6	89.1	143.0
imgcorruptlike.SpeckleNoise((1,5))	81.5	130.8	175.4
imgcorruptlike.GaussianBlur((1,5))	72.0	114.9	195.8
imgcorruptlike.GlassBlur((1,5))	2.0	2.3	2.1
imgcorruptlike.DefocusBlur((1,5))	56.8	82.6	134.5
imgcorruptlike.MotionBlur((1,5))	19.2	22.6	62.5
imgcorruptlike.ZoomBlur((1,5))	8.0	9.1	11.6
imgcorruptlike.Fog((1,5))	33.1	50.9	171.0
imgcorruptlike.Frost((1,5))	10.0	11.2	113.2
imgcorruptlike.Snow((1,5))	26.1	29.3	68.9
imgcorruptlike.Spatter((1,5))	54.2	69.0	135.7
imgcorruptlike.Contrast((1,5))	113.7	206.4	364.7
imgcorruptlike.Brightness((1,5))	38.1	46.2	48.1
imgcorruptlike.Saturate((1,5))	34.7	46.0	48.4
imgcorruptlike.JpegCompression((1,5))	82.7	165.2	473.9
imgcorruptlike.Pixelate((1,5))	141.5	321.1	1013.5
imgcorruptlike.ElasticTransform((1,5))	36.1	44.1	56.2
pillike.Solarize (p=1.0)	183.2	843.5	1801.5
pillike.Posterize (1-7 bits)	120.7	360.9	1449.0
pillike.Equalize	163.9	288.2	1349.4
pillike.Autocontrast	69.5	98.6	748.8
pillike.EnhanceColor	190.3	587.5	937.4
pillike.EnhanceContrast	164.2	370.0	842.4
pillike.EnhanceBrightness	212.9	630.2	1017.1
pillike.EnhanceSharpness	178.8	422.3	590.4
pillike.FilterBlur	233.6	375.4	459.2
pillike.FilterSmooth	327.7	588.8	911.2
pillike.FilterSmoothMore	250.2	374.5	465.4
pillike.FilterEdgeEnhance	288.4	530.9	817.9
pillike.FilterEdgeEnhanceMore	293.1	523.0	791.9
pillike.FilterFindEdges	297.9	530.3	756.3
pillike.FilterContour	285.2	523.1	746.7
pillike.FilterEmboss	282.2	586.0	910.7
pillike.FilterSharpen	256.6	579.5	868.7
pillike.FilterDetail	304.5	586.4	880.1

Table 1 – continued from previous page

pillike.Affine	66.3	302.8	709.5
Superpixels (max_size=64, cubic)	9.4	10.4	118.4
Superpixels (max_size=64, linear)	9.9	10.4	118.7
Superpixels (max_size=128, linear)	8.0	10.6	49.5
Superpixels (max_size=224, linear)	7.6	10.6	19.5
UniformVoronoi (250-1000k points, linear)	2.7	3.6	12.1
RegularGridVoronoi (16-31 rows/cols)	3.5	3.6	12.0
RelativeRegularGridVoronoi (7%-14% rows/cols)	3.7	3.7	4.0
Resize (nearest)	186.3	735.5	1988.3
Resize (linear)	176.0	629.9	1537.8
Resize (cubic)	177.0	559.4	1187.7
CropAndPad	118.9	700.3	1422.4
CropAndPad (edge)	118.9	705.6	1449.5
CropAndPad (keep_size)	104.7	376.3	1018.1
Crop	153.0	1293.5	1974.8
Crop (keep_size)	130.8	501.6	1275.2
Pad	122.2	678.8	1384.0
Pad (edge)	118.7	683.5	1390.6
Pad (keep_size)	101.6	371.7	954.2
PadToFixedSize	130.8	916.5	1653.7
CropToFixedSize	228.9	3102.1	2756.7
KeepSizeByResize (CropToFixedSize(nearest))	139.8	880.7	1471.7
KeepSizeByResize (CropToFixedSize(linear))	134.2	761.3	1230.6
KeepSizeByResize (CropToFixedSize(cubic))	133.0	660.3	1002.8
FastSnowyLandscape	116.8	243.5	483.0
Clouds	16.9	20.3	61.7
Fog	31.3	33.9	98.3
CloudLayer	30.7	33.0	99.1
Snowflakes	14.2	15.1	91.3
SnowflakesLayer	28.5	30.3	173.5
Rain	11.1	11.6	58.5
RainLayer	22.0	23.5	110.4

In images/sec:

	64x64x3, uint8		224x224x
Augmenter	B=1	B=128	B=1
Sequential (2xIdentity)	11885.5	262413.9	8547.0
Sequential (2xIdentity, random_order)	9635.6	271476.0	7573.7
SomeOf (1-3, 3xIdentity)	2412.6	40539.6	1840.9
SomeOf (1-3, 3xIdentity, random_order)	2356.3	39655.4	1774.3
OneOf (3xIdentity)	5695.0	84707.6	3906.3
Sometimes (Identity)	3917.9	137540.7	3277.2
WithChannels ([1,2], Identity)	5772.9	52615.6	3194.7
Identity	14554.8	285239.9	10267.7
Noop	14308.5	288214.7	10110.7
Lambda (return input)	13462.3	265802.2	9437.0
AssertLambda (return True)	13277.6	281027.4	9460.0
AssertShape (None, H, W, None)	11400.9	155628.0	8585.7
ChannelShuffle (0.5)	4462.7	33699.4	2860.5

Table 2 – continued from previous page

Add	1468.7	6351.5	1717.5
AddElementwise	2147.0	3279.9	792.2
AdditiveGaussianNoise	1787.2	2330.6	605.9
AdditiveLaplaceNoise	1570.5	1975.4	365.1
AdditivePoissonNoise	1306.5	1610.3	256.5
Multiply	2563.5	8217.8	1877.6
MultiplyElementwise	2005.8	2953.5	763.4
Cutout (1 iter, constant fill)	1611.1	13745.5	1501.1
Dropout (1-5%)	2409.4	3771.3	866.7
CoarseDropout (1-5%, size=1-10%)	1423.1	1840.8	904.7
Dropout2d (10%)	3457.4	60765.5	3107.9
TotalDropout (10%)	4804.7	212745.7	4308.9
ReplaceElementwise	1385.7	1727.3	589.3
ImpulseNoise	1199.5	1442.0	409.2
SaltAndPepper	1267.9	1504.9	560.5
CoarseSaltAndPepper	926.8	1051.4	580.3
Salt	1064.4	1218.0	514.5
CoarseSalt	838.3	923.2	528.2
Pepper	1043.9	1127.7	513.3
CoarsePepper	839.9	918.5	527.3
Invert (10%)	2842.5	58334.2	2606.0
JpegCompression (50-99%)	873.3	1106.0	366.2
Cartoon	64.0	61.6	6.2
BlendAlpha (Identity)	1373.3	4288.0	705.5
BlendAlphaElementwise (Identity)	1393.6	2207.6	392.1
BlendAlphaSimplexNoise (Identity)	263.7	302.2	152.8
BlendAlphaFrequencyNoise (Identity)	354.2	389.8	183.6
BlendAlphaSomeColors (Identity)	684.8	1717.7	288.0
BlendAlphaHorizontalL.Grad. (Identity)	933.1	2539.4	362.5
BlendAlphaVerticalL.Grad. (Identity)	937.3	2473.5	354.8
BlendAlphaRegularGrid (Identity)	906.5	2455.6	347.2
BlendAlphaCheckerboard (Identity)	917.6	2135.5	347.4
GaussianBlur (sigma=(1,5))	2045.3	5681.7	1331.3
AverageBlur	2617.3	15727.3	1760.1
MedianBlur	1384.8	2745.8	233.1
BilateralBlur	1080.3	2872.0	245.5
MotionBlur	599.6	683.8	471.6
MeanShiftBlur	15.3	14.5	1.1
RandAugment (n=2, m=(6,12))	255.7	1366.3	193.7
WithColorspace (HSV, Identity)	3110.5	10389.9	1472.9
WithBrightnessChannels (Identity)	464.3	7860.5	956.0
MultiplyAndAddToBrightness	761.1	2682.0	579.1
MultiplyBrightness	1209.1	4232.3	740.7
AddToBrightness	1165.7	3710.0	732.5
WithHueAndSaturation	1793.2	3568.3	598.7
MultiplyHueAndSaturation	880.0	1622.2	383.5
MultiplyHue	792.7	1601.0	381.6
MultiplySaturation	614.6	1604.2	384.9
RemoveSaturation	756.1	1606.7	377.1
AddToHueAndSaturation	1398.0	4726.3	721.5

Table 2 – continued from previous page

AddToHue	1602.0	4862.4	747.3
AddToSaturation	1491.8	4912.4	753.4
ChangeColorspace (HSV)	2750.6	9853.8	1966.7
Grayscale	1526.8	3290.8	550.6
KMeansColorQuantization (2-16 colors)	323.3	400.4	159.7
UniformColorQuantization (2-16 colors)	1364.2	3781.6	1316.6
UniformColorQuant.NBits (1-7 bits)	1521.9	3811.1	1313.6
Posterize (1-7 bits)	1455.0	3805.0	1311.7
GammaContrast	1812.4	3245.3	1595.6
SigmoidContrast	1641.5	2498.6	1351.1
LogContrast	1962.1	3233.2	1584.2
LinearContrast	2285.2	4180.7	1783.3
AllChannelsHistogramEqualization	5540.9	16639.1	1618.4
HistogramEqualization	2863.9	9524.3	1280.3
AllChannelsCLAHE	1204.0	3480.0	765.0
CLAHE	1202.2	4716.2	718.1
Convolve (3x3)	4593.4	15381.1	2467.1
Sharpen	1885.1	2789.7	1116.5
Emboss	1887.4	2799.9	1161.9
EdgeDetect	2497.6	4184.9	1477.3
DirectedEdgeDetect	963.3	1146.6	720.8
Canny	582.2	1108.6	244.0
Fliplr (p=100%)	4758.0	26744.7	2925.1
Flipud (p=100%)	6023.3	103695.3	4767.3
Affine (order=0, constant)	808.8	2723.2	745.7
Affine (order=1, constant)	805.8	2519.9	673.5
Affine (order=3, constant)	739.2	2075.6	412.3
Affine (order=1, edge)	726.5	2447.5	648.1
Affine (order=1, constant, skimage)	417.8	785.3	156.7
PiecewiseAffine (4x4, order=1, constant)	41.9	46.4	22.3
PiecewiseAffine (4x4, order=0, constant)	45.4	47.9	26.7
PiecewiseAffine (4x4, order=1, edge)	43.6	46.4	22.1
PiecewiseAffine (8x8, order=1, constant)	9.6	10.0	7.3
PerspectiveTransform	1029.6	5054.7	771.1
PerspectiveTransform (keep_size)	986.5	3340.2	599.4
ElasticTransformation (order=0, constant)	880.3	1688.1	472.9
ElasticTransformation (order=1, constant)	859.3	1591.7	435.1
ElasticTransformation (order=1, nearest)	858.2	1601.2	430.5
ElasticTransformation (order=1, reflect)	868.9	1590.0	435.7
Rot90	2917.1	42467.5	2974.6
Rot90 (keep_size)	2835.5	23402.1	1727.3
WithPolarWarping (Identity)	2764.2	6825.5	825.7
Jigsaw (rows/cols=(3,8), 1 step)	666.4	1141.2	634.5
AveragePooling	1035.7	1875.7	378.5
AveragePooling (keep_size)	971.3	1584.3	342.0
MaxPooling	1074.4	2000.8	401.1
MaxPooling (keep_size)	988.8	1691.9	375.7
MinPooling	1095.3	2000.8	407.1
MinPooling (keep_size)	1014.9	1679.8	371.0
MedianPooling	1006.0	1830.6	485.5

Table 2 – continued from previous page

MedianPooling (keep_size)	961.1	1561.4	447.4
imgcorruptlike.GaussianNoise((1,5))	756.7	1391.8	157.4
imgcorruptlike.ShotNoise((1,5))	554.9	725.7	73.0
imgcorruptlike.ImpulseNoise((1,5))	667.6	950.0	124.5
imgcorruptlike.SpeckleNoise((1,5))	869.7	1395.3	152.7
imgcorruptlike.GaussianBlur((1,5))	768.0	1225.9	170.5
imgcorruptlike.GlassBlur((1,5))	21.7	25.0	1.8
imgcorruptlike.DefocusBlur((1,5))	606.1	881.0	117.1
imgcorruptlike.MotionBlur((1,5))	204.5	241.3	54.5
imgcorruptlike.ZoomBlur((1,5))	85.7	97.2	10.1
imgcorruptlike.Fog((1,5))	352.7	543.3	148.9
imgcorruptlike.Frost((1,5))	107.0	120.0	98.6
imgcorruptlike.Snow((1,5))	278.7	312.6	60.0
imgcorruptlike.Spatter((1,5))	578.1	735.7	118.1
imgcorruptlike.Contrast((1,5))	1212.6	2201.8	317.5
imgcorruptlike.Brightness((1,5))	406.3	493.3	41.9
imgcorruptlike.Saturate((1,5))	369.8	490.8	42.1
imgcorruptlike.JpegCompression((1,5))	882.6	1761.9	412.6
imgcorruptlike.Pixelate((1,5))	1509.0	3425.1	882.5
imgcorruptlike.ElasticTransform((1,5))	384.6	470.0	48.9
pillike.Solarize (p=1.0)	1954.2	8997.3	1568.6
pillike.Posterize (1-7 bits)	1288.0	3849.2	1261.7
pillike.Equalize	1748.5	3074.6	1175.0
pillike.Autocontrast	741.4	1052.1	652.0
pillike.EnhanceColor	2029.5	6266.6	816.2
pillike.EnhanceContrast	1751.2	3946.7	733.6
pillike.EnhanceBrightness	2271.3	6722.4	885.6
pillike.EnhanceSharpness	1907.0	4504.8	514.1
pillike.FilterBlur	2491.3	4004.4	399.8
pillike.FilterSmooth	3495.0	6280.4	793.5
pillike.FilterSmoothMore	2669.2	3995.1	405.2
pillike.FilterEdgeEnhance	3076.8	5662.9	712.2
pillike.FilterEdgeEnhanceMore	3126.6	5579.1	689.6
pillike.FilterFindEdges	3177.4	5656.8	658.6
pillike.FilterContour	3042.2	5580.0	650.2
pillike.FilterEmboss	3010.1	6251.0	793.0
pillike.FilterSharpen	2737.3	6181.6	756.4
pillike.FilterDetail	3248.1	6255.5	766.3
pillike.Affine	707.5	3229.4	617.8
Superpixels (max_size=64, cubic)	100.2	111.0	103.1
Superpixels (max_size=64, linear)	106.0	111.4	103.4
Superpixels (max_size=128, linear)	84.9	112.9	43.1
Superpixels (max_size=224, linear)	81.6	113.1	17.0
UniformVoronoi (250-1000k points, linear)	28.4	38.8	10.5
RegularGridVoronoi (16-31 rows/cols)	37.7	38.7	10.4
RelativeRegularGridVoronoi (7%-14% rows/cols)	39.2	39.0	3.4
Resize (nearest)	1987.7	7844.9	1731.3
Resize (linear)	1877.8	6718.6	1339.0
Resize (cubic)	1887.5	5966.6	1034.2
CropAndPad	1268.3	7470.3	1238.6

Table 2 – continued from previous page

CropAndPad (edge)	1268.3	7526.5	1262.1
CropAndPad (keep_size)	1117.2	4013.7	886.5
Crop	1632.5	13797.2	1719.6
Crop (keep_size)	1395.5	5350.3	1110.3
Pad	1303.7	7240.5	1205.1
Pad (edge)	1266.5	7290.4	1210.8
Pad (keep_size)	1084.1	3964.9	830.9
PadToFixedSize	1394.9	9776.5	1439.9
CropToFixedSize	2441.5	33089.5	2400.4
KeepSizeByResize (CropToFixedSize(nearest))	1491.3	9393.7	1281.5
KeepSizeByResize (CropToFixedSize(linear))	1431.7	8121.0	1071.5
KeepSizeByResize (CropToFixedSize(cubic))	1418.5	7043.0	873.2
FastSnowyLandscape	1246.1	2597.6	420.6
Clouds	180.6	216.2	53.7
Fog	334.1	361.5	85.6
CloudLayer	327.1	352.4	86.3
Snowflakes	151.5	161.2	79.5
SnowflakesLayer	304.2	323.6	151.1
Rain	118.4	123.9	50.9
RainLayer	234.6	251.1	96.1

10.3 Heatmaps

Numbers below are for heatmaps on large images, i.e. `224x224x3`. Smaller images were skipped for brevity. The heatmaps themselves can be small (`64x64xN`) or large (`224x224xN`), with `N` denoting the number of heatmaps per `HeatmapsOnImage` instance (i.e. the number of channels in the heatmaps array), for which below `1` and `5` are used. `B=1` denotes a batch size of `1` , `B=128` one of `128`.

mbit/sec for 64x64x5 or 224x224x5 heatmaps on 224x224x3 images:

	64x64x5 on 224x224x3		224x224x5 o
Augmenter	B=1	B=128	B=1
Sequential (2xIdentity)	1184.9	5580.5	10238.2
Sequential (2xIdentity, random_order)	1104.8	5543.8	9429.3
SomeOf (1-3, 3xIdentity)	720.2	5165.0	6854.4
SomeOf (1-3, 3xIdentity, random_order)	706.3	5160.3	6742.9
OneOf (3xIdentity)	1025.9	5388.7	9095.4
Sometimes (Identity)	831.1	5479.9	7836.0
WithChannels ([1,2], Identity)	901.2	2622.7	6464.3
Identity	1329.3	5606.5	10585.6
Noop	1351.5	5611.4	10479.4
Lambda (return input)	1297.5	5567.3	10284.4
AssertLambda (return True)	1300.8	5567.1	10235.6
AssertShape (None, H, W, None)	1271.6	5431.1	10001.8
ChannelShuffle (0.5)	1351.2	5589.4	10447.6
Add	1360.6	5590.4	10432.7
AddElementwise	1361.6	5640.5	10443.3
AdditiveGaussianNoise	1351.0	5616.6	10528.6
AdditiveLaplaceNoise	1362.5	5568.4	10364.0

Table 3 – continued from previous page

AdditivePoissonNoise	1359.3	5620.1	10447.9
Multiply	1349.5	5619.9	10354.5
MultiplyElementwise	1351.1	5580.9	10203.0
Cutout (1 iter, constant fill)	1349.3	5584.4	10434.1
Dropout (1-5%)	1371.0	5635.7	10357.0
CoarseDropout (1-5%, size=1-10%)	1344.5	5618.2	10530.4
Dropout2d (10%)	859.7	5068.6	7527.2
TotalDropout (10%)	965.1	5453.2	8083.1
ReplaceElementwise	1359.0	5603.3	10451.3
ImpulseNoise	1359.5	5602.0	10516.4
SaltAndPepper	1352.8	5635.6	10447.3
CoarseSaltAndPepper	1356.7	5614.1	10310.9
Salt	1348.9	5576.4	10497.3
CoarseSalt	1339.5	5584.0	10521.0
Pepper	1365.9	5583.9	10475.6
CoarsePepper	1352.6	6526.4	10546.9
Invert (10%)	1364.9	6532.4	10436.1
JpegCompression (50-99%)	1366.4	6563.1	10513.6
Cartoon	1340.7	6543.2	10352.2
BlendAlpha (Identity)	662.3	2880.4	5086.3
BlendAlphaElementwise (Identity)	298.6	451.2	1251.4
BlendAlphaSimplexNoise (Identity)	100.7	122.9	800.4
BlendAlphaFrequencyNoise (Identity)	125.3	147.0	882.5
BlendAlphaSomeColors (Identity)	n/a	n/a	n/a
BlendAlphaHorizontalL.Grad. (Identity)	259.0	534.5	1243.2
BlendAlphaVerticalL.Grad. (Identity)	253.6	491.5	1219.2
BlendAlphaRegularGrid (Identity)	242.7	479.7	1200.4
BlendAlphaCheckerboard (Identity)	246.2	461.2	1212.2
GaussianBlur (sigma=(1,5))	1274.8	6491.3	10490.6
AverageBlur	1266.5	6470.4	10485.1
MedianBlur	1274.3	6516.3	10496.8
BilateralBlur	1277.1	6503.2	10455.4
MotionBlur	1261.7	6515.3	10305.2
MeanShiftBlur	1285.3	6645.9	10363.5
RandAugment (n=2, m=(6,12))	n/a	n/a	n/a
WithColorspace (HSV, Identity)	1185.0	6766.3	10008.5
WithBrightnessChannels (Identity)	1192.7	6776.9	10012.8
MultiplyAndAddToBrightness	1072.6	6750.4	9290.9
MultiplyBrightness	1176.3	6767.6	9863.5
AddToBrightness	1117.8	6759.0	9902.6
WithHueAndSaturation	1152.0	6768.5	10050.7
MultiplyHueAndSaturation	1170.8	6808.5	9967.6
MultiplyHue	868.0	3343.3	6367.6
MultiplySaturation	861.3	3353.0	6426.9
RemoveSaturation	875.9	3337.8	6363.9
AddToHueAndSaturation	1304.3	6766.2	10446.2
AddToHue	1339.0	6785.4	10526.2
AddToSaturation	1336.5	6799.2	10456.1
ChangeColorspace (HSV)	1278.9	6776.7	10554.9
Grayscale	1298.5	6832.6	10475.2

Table 3 – continued from previous page

KMeansColorQuantization (2-16 colors)	1285.0	6794.2	10472.9
UniformColorQuantization (2-16 colors)	1286.6	6813.7	10553.4
UniformColorQuant.NBits (1-7 bits)	1309.5	6819.9	10343.9
Posterize (1-7 bits)	1357.7	6843.1	10515.6
GammaContrast	1337.8	6800.4	10495.7
SigmoidContrast	1337.0	6806.4	10466.6
LogContrast	1344.0	6762.3	10463.7
LinearContrast	1349.4	6793.7	10413.8
AllChannelsHistogramEqualization	1365.5	6783.6	10488.7
HistogramEqualization	1364.7	6801.1	10477.2
AllChannelsCLAHE	1330.2	6800.5	10415.1
CLAHE	1372.3	6833.8	10539.7
Convolve (3x3)	1356.2	6806.5	10419.5
Sharpen	1339.8	6799.7	10392.9
Emboss	1354.7	6767.9	10413.8
EdgeDetect	1368.3	6769.3	10518.3
DirectedEdgeDetect	1358.6	6795.3	10466.6
Canny	1344.4	6789.5	10480.2
Fliplr (p=100%)	1064.6	6145.1	8835.8
Flipud (p=100%)	1080.3	6538.2	8888.0
Affine (order=0, constant)	255.7	654.9	1134.6
Affine (order=1, constant)	264.3	654.2	1139.9
Affine (order=3, constant)	262.5	654.2	1129.1
Affine (order=1, edge)	260.1	654.5	1137.4
Affine (order=1, constant, skimage)	161.4	267.1	372.0
PiecewiseAffine (4x4, order=1, constant)	19.3	20.6	51.5
PiecewiseAffine (4x4, order=0, constant)	19.5	20.4	51.0
PiecewiseAffine (4x4, order=1, edge)	19.4	20.4	51.3
PiecewiseAffine (8x8, order=1, constant)	5.6	5.6	31.3
PerspectiveTransform	258.5	925.1	1411.7
PerspectiveTransform (keep_size)	225.5	603.7	1094.5
ElasticTransformation (order=0, constant)	69.4	97.6	1217.6
ElasticTransformation (order=1, constant)	70.5	97.6	1204.3
ElasticTransformation (order=1, nearest)	70.0	96.1	1211.0
ElasticTransformation (order=1, reflect)	70.3	96.0	1215.7
Rot90	748.1	5247.6	7075.9
Rot90 (keep_size)	574.6	2012.9	3971.2
WithPolarWarping (Identity)	622.5	1869.3	2621.8
Jigsaw (rows/cols=(3,8), 1 step)	117.7	171.8	1500.0
AveragePooling	516.2	1850.0	4782.4
AveragePooling (keep_size)	1243.4	6670.4	10415.7
MaxPooling	528.5	1844.9	4879.3
MaxPooling (keep_size)	1274.8	6704.8	10582.7
MinPooling	521.7	1871.9	4895.7
MinPooling (keep_size)	1273.7	6721.3	10444.5
MedianPooling	523.9	1870.4	4888.3
MedianPooling (keep_size)	1265.8	6728.4	10477.3
imgcorruptlike.GaussianNoise((1,5))	1279.1	6750.4	10470.0
imgcorruptlike.ShotNoise((1,5))	1277.7	6764.5	10524.5
imgcorruptlike.ImpulseNoise((1,5))	1252.8	6745.3	10458.3

Table 3 – continued from previous page

imgcorruptlike.SpeckleNoise((1,5))	1262.7	6708.5	10519.9
imgcorruptlike.GaussianBlur((1,5))	1271.5	6695.1	10375.6
imgcorruptlike.GlassBlur((1,5))	1279.0	6740.8	10482.6
imgcorruptlike.DefocusBlur((1,5))	1264.1	6725.7	10457.5
imgcorruptlike.MotionBlur((1,5))	1285.1	6735.1	10389.7
imgcorruptlike.ZoomBlur((1,5))	1281.2	6754.2	10308.4
imgcorruptlike.Fog((1,5))	1277.6	6663.9	10386.9
imgcorruptlike.Frost((1,5))	1280.2	6737.7	10449.3
imgcorruptlike.Snow((1,5))	1264.4	6732.5	10613.1
imgcorruptlike.Spatter((1,5))	1248.5	6735.6	10394.4
imgcorruptlike.Contrast((1,5))	1251.7	6737.4	10483.9
imgcorruptlike.Brightness((1,5))	1266.7	6748.6	10519.9
imgcorruptlike.Saturate((1,5))	1270.4	6771.8	10566.0
imgcorruptlike.JpegCompression((1,5))	1256.7	6735.9	10589.2
imgcorruptlike.Pixelate((1,5))	1282.2	6739.7	10463.6
imgcorruptlike.ElasticTransform((1,5))	1255.9	6715.8	10436.7
pillike.Solarize (p=1.0)	1274.0	6747.8	10415.5
pillike.Posterize (1-7 bits)	1278.4	6738.6	10475.7
pillike.Equalize	1274.1	6758.1	10450.9
pillike.Autocontrast	1259.0	6723.7	10419.1
pillike.EnhanceColor	1272.0	6719.8	10397.4
pillike.EnhanceContrast	1257.2	6720.5	10479.4
pillike.EnhanceBrightness	1267.2	6732.6	10512.5
pillike.EnhanceSharpness	1241.9	6735.0	10433.6
pillike.FilterBlur	1259.5	6661.7	10495.7
pillike.FilterSmooth	1247.8	6718.1	10483.4
pillike.FilterSmoothMore	1240.1	6690.4	10318.6
pillike.FilterEdgeEnhance	1265.3	6717.4	10390.9
pillike.FilterEdgeEnhanceMore	1270.1	6728.4	10098.6
pillike.FilterFindEdges	1266.7	6751.5	10353.5
pillike.FilterContour	1249.3	6694.9	10333.1
pillike.FilterEmboss	1267.5	6767.2	10476.8
pillike.FilterSharpen	1271.6	6737.6	10398.5
pillike.FilterDetail	1281.0	6780.6	10496.7
pillike.Affine	n/a	n/a	n/a
Superpixels (max_size=64, cubic)	1295.0	6561.1	10527.2
Superpixels (max_size=64, linear)	1257.2	6569.5	10486.4
Superpixels (max_size=128, linear)	1260.1	6512.3	10486.7
Superpixels (max_size=224, linear)	1280.0	6589.3	10361.3
UniformVoronoi (250-1000k points, linear)	1289.5	6565.9	10375.7
RegularGridVoronoi (16-31 rows/cols)	1284.1	6565.6	10567.1
RelativeRegularGridVoronoi (7%-14% rows/cols)	1282.7	6580.6	10424.6
Resize (nearest)	485.9	1426.2	3312.5
Resize (linear)	500.5	1344.2	3064.6
Resize (cubic)	492.0	1267.1	2730.4
CropAndPad	417.4	1488.9	2927.0
CropAndPad (edge)	414.1	1486.4	2919.0
CropAndPad (keep_size)	337.4	855.3	1956.8
Crop	529.6	2379.2	5026.1
Crop (keep_size)	406.8	1097.0	2568.4

Table 3 – continued from previous page

Pad	405.3	1367.0	2601.9
Pad (edge)	402.4	1365.7	2615.3
Pad (keep_size)	329.1	801.9	1703.2
PadToFixedSize	441.9	1683.2	2947.0
CropToFixedSize	630.4	2987.1	5543.8
KeepSizeByResize (CropToFixedSize(nearest))	402.9	1293.7	2947.5
KeepSizeByResize (CropToFixedSize(linear))	398.9	1223.9	2707.1
KeepSizeByResize (CropToFixedSize(cubic))	395.7	1160.9	2440.0
FastSnowyLandscape	1266.6	6468.7	10481.2
Clouds	719.1	5982.0	6979.2
Fog	1284.1	6527.3	10390.2
CloudLayer	1268.0	6527.0	10328.0
Snowflakes	717.6	5974.3	6852.2
SnowflakesLayer	1282.2	6503.1	10732.7
Rain	723.3	5951.5	6901.8
RainLayer	1189.0	6438.2	10457.1

Number of heatmap instances per sec for 64x64x5 or 224x224x5 heatmaps on 224x224x3 images:

	64x64x5 on 224x224x3		224x224x5
Augmenter	B=1	B=128	B=1
Sequential (2xIdentity)	9479.6	44643.9	6686.2
Sequential (2xIdentity, random_order)	8838.1	44350.4	6157.9
SomeOf (1-3, 3xIdentity)	5761.4	41319.8	4476.3
SomeOf (1-3, 3xIdentity, random_order)	5650.8	41282.8	4403.5
OneOf (3xIdentity)	8207.1	43109.2	5939.9
Sometimes (Identity)	6648.8	43839.3	5117.4
WithChannels ([1,2], Identity)	7209.9	20981.2	4221.6
Identity	10634.7	44851.8	6913.0
Noop	10812.0	44891.4	6843.7
Lambda (return input)	10379.6	44538.3	6716.4
AssertLambda (return True)	10406.7	44536.9	6684.5
AssertShape (None, H, W, None)	10172.8	43448.8	6531.8
ChannelShuffle (0.5)	10809.8	44715.3	6822.9
Add	10884.9	44723.0	6813.2
AddElementwise	10892.6	45124.3	6820.1
AdditiveGaussianNoise	10807.8	44932.5	6875.8
AdditiveLaplaceNoise	10899.7	44547.2	6768.3
AdditivePoissonNoise	10874.6	44960.6	6823.1
Multiply	10796.3	44959.0	6762.1
MultiplyElementwise	10808.7	44647.6	6663.2
Cutout (1 iter, constant fill)	10794.6	44675.1	6814.1
Dropout (1-5%)	10967.9	45085.7	6763.8
CoarseDropout (1-5%, size=1-10%)	10756.2	44945.9	6877.0
Dropout2d (10%)	6877.6	40548.7	4915.7
TotalDropout (10%)	7721.0	43625.5	5278.8
ReplaceElementwise	10871.7	44826.6	6825.3
ImpulseNoise	10875.7	44816.0	6867.9
SaltAndPepper	10822.3	45085.0	6822.7
CoarseSaltAndPepper	10853.6	44912.4	6733.7

Table 4 – continued from previous page

Salt	10791.0	44611.6	6855.3
CoarseSalt	10716.2	44671.6	6870.8
Pepper	10927.1	44671.0	6841.2
CoarsePepper	10821.1	52210.9	6887.8
Invert (10%)	10919.3	52259.6	6815.4
JpegCompression (50-99%)	10931.5	52505.0	6866.0
Cartoon	10725.2	52345.4	6760.6
BlendAlpha (Identity)	5298.2	23043.2	3321.7
BlendAlphaElementwise (Identity)	2388.4	3609.9	817.2
BlendAlphaSimplexNoise (Identity)	805.7	983.0	522.7
BlendAlphaFrequencyNoise (Identity)	1002.4	1175.6	576.3
BlendAlphaSomeColors (Identity)	n/a	n/a	n/a
BlendAlphaHorizontalL.Grad. (Identity)	2072.0	4275.7	811.9
BlendAlphaVerticalL.Grad. (Identity)	2029.2	3931.6	796.2
BlendAlphaRegularGrid (Identity)	1941.8	3837.9	784.0
BlendAlphaCheckerboard (Identity)	1969.7	3689.2	791.6
GaussianBlur (sigma=(1,5))	10198.4	51930.5	6851.0
AverageBlur	10132.1	51763.5	6847.4
MedianBlur	10194.0	52130.6	6855.1
BilateralBlur	10216.7	52025.7	6828.0
MotionBlur	10093.6	52122.1	6730.0
MeanShiftBlur	10282.4	53167.2	6768.0
RandAugment (n=2, m=(6,12))	n/a	n/a	n/a
WithColorspace (HSV, Identity)	9479.7	54130.4	6536.1
WithBrightnessChannels (Identity)	9541.3	54215.5	6539.0
MultiplyAndAddToBrightness	8581.1	54003.0	6067.5
MultiplyBrightness	9410.6	54140.9	6441.5
AddToBrightness	8942.4	54072.0	6467.0
WithHueAndSaturation	9215.7	54148.2	6563.7
MultiplyHueAndSaturation	9366.5	54467.8	6509.5
MultiplyHue	6943.9	26746.8	4158.4
MultiplySaturation	6890.2	26823.7	4197.1
RemoveSaturation	7006.9	26702.7	4156.0
AddToHueAndSaturation	10434.8	54129.4	6822.0
AddToHue	10711.8	54283.0	6874.2
AddToSaturation	10692.0	54393.9	6828.5
ChangeColorspace (HSV)	10231.2	54213.2	6893.0
Grayscale	10387.7	54661.1	6841.0
KMeansColorQuantization (2-16 colors)	10280.3	54353.4	6839.4
UniformColorQuantization (2-16 colors)	10292.9	54509.9	6892.0
UniformColorQuant.NBits (1-7 bits)	10476.1	54559.5	6755.2
Posterize (1-7 bits)	10861.3	54744.7	6867.4
GammaContrast	10702.7	54403.1	6854.3
SigmoidContrast	10696.1	54451.1	6835.3
LogContrast	10751.8	54098.5	6833.4
LinearContrast	10795.1	54349.9	6800.9
AllChannelsHistogramEqualization	10924.1	54268.6	6849.7
HistogramEqualization	10917.2	54409.1	6842.2
AllChannelsCLAHE	10641.3	54403.8	6801.7
CLAHE	10978.3	54670.0	6883.1

Table 4 – continued from previous page

Convolve (3x3)	10849.8	54451.7	6804.5
Sharpen	10718.5	54397.6	6787.2
Emboss	10838.0	54143.0	6800.8
EdgeDetect	10946.7	54154.6	6869.1
DirectedEdgeDetect	10868.6	54362.6	6835.3
Canny	10755.1	54316.1	6844.2
Fliplr (p=100%)	8516.7	49161.0	5770.3
Flipud (p=100%)	8642.8	52305.2	5804.4
Affine (order=0, constant)	2045.6	5238.8	741.0
Affine (order=1, constant)	2114.6	5233.6	744.4
Affine (order=3, constant)	2100.4	5233.6	737.4
Affine (order=1, edge)	2080.8	5236.2	742.8
Affine (order=1, constant, skimage)	1291.6	2137.1	242.9
PiecewiseAffine (4x4, order=1, constant)	154.3	164.8	33.6
PiecewiseAffine (4x4, order=0, constant)	155.8	163.5	33.3
PiecewiseAffine (4x4, order=1, edge)	155.2	163.4	33.5
PiecewiseAffine (8x8, order=1, constant)	44.8	44.9	20.4
PerspectiveTransform	2068.0	7401.0	921.9
PerspectiveTransform (keep_size)	1804.2	4829.3	714.8
ElasticTransformation (order=0, constant)	555.2	780.9	795.2
ElasticTransformation (order=1, constant)	564.2	780.9	786.5
ElasticTransformation (order=1, nearest)	559.9	769.1	790.9
ElasticTransformation (order=1, reflect)	562.3	768.3	793.9
Rot90	5985.0	41980.6	4621.0
Rot90 (keep_size)	4596.8	16103.2	2593.4
WithPolarWarping (Identity)	4979.9	14954.7	1712.2
Jigsaw (rows/cols=(3,8), 1 step)	941.8	1374.1	979.6
AveragePooling	4129.7	14799.8	3123.2
AveragePooling (keep_size)	9947.3	53363.1	6802.1
MaxPooling	4227.7	14759.2	3186.5
MaxPooling (keep_size)	10198.2	53638.1	6911.1
MinPooling	4173.4	14975.5	3197.2
MinPooling (keep_size)	10189.7	53770.8	6820.9
MedianPooling	4191.0	14963.4	3192.3
MedianPooling (keep_size)	10126.2	53827.5	6842.3
imgcorruptlike.GaussianNoise((1,5))	10233.1	54003.5	6837.6
imgcorruptlike.ShotNoise((1,5))	10221.4	54116.1	6873.2
imgcorruptlike.ImpulseNoise((1,5))	10022.6	53962.0	6829.9
imgcorruptlike.SpeckleNoise((1,5))	10101.6	53668.1	6870.1
imgcorruptlike.GaussianBlur((1,5))	10172.3	53560.8	6775.9
imgcorruptlike.GlassBlur((1,5))	10231.6	53926.2	6845.8
imgcorruptlike.DefocusBlur((1,5))	10112.5	53805.7	6829.4
imgcorruptlike.MotionBlur((1,5))	10280.8	53881.1	6785.1
imgcorruptlike.ZoomBlur((1,5))	10249.5	54033.7	6732.0
imgcorruptlike.Fog((1,5))	10220.8	53311.2	6783.3
imgcorruptlike.Frost((1,5))	10241.5	53901.9	6824.1
imgcorruptlike.Snow((1,5))	10115.0	53859.8	6931.0
imgcorruptlike.Spatter((1,5))	9988.2	53884.5	6788.2
imgcorruptlike.Contrast((1,5))	10013.3	53899.3	6846.6
imgcorruptlike.Brightness((1,5))	10133.9	53989.0	6870.2

Table 4 – continued from previous page

imgcorruptlike.Saturate((1,5))	10163.1	54174.7	6900.2
imgcorruptlike.JpegCompression((1,5))	10053.6	53886.8	6915.4
imgcorruptlike.Pixelate((1,5))	10257.9	53917.3	6833.4
imgcorruptlike.ElasticTransform((1,5))	n/a	n/a	n/a
pillike.Solarize (p=1.0)	10192.1	53982.3	6802.0
pillike.Posterize (1-7 bits)	10227.5	53909.2	6841.3
pillike.Equalize	10192.7	54064.7	6825.1
pillike.Autocontrast	10071.7	53789.6	6804.3
pillike.EnhanceColor	10176.0	53758.3	6790.1
pillike.EnhanceContrast	10057.3	53764.1	6843.7
pillike.EnhanceBrightness	10137.6	53860.6	6865.3
pillike.EnhanceSharpness	9935.5	53879.9	6813.8
pillike.FilterBlur	10075.9	53293.9	6854.4
pillike.FilterSmooth	9982.0	53744.7	6846.3
pillike.FilterSmoothMore	9920.8	53523.5	6738.7
pillike.FilterEdgeEnhance	10122.4	53739.0	6785.9
pillike.FilterEdgeEnhanceMore	10161.0	53826.9	6595.0
pillike.FilterFindEdges	10133.6	54011.7	6761.5
pillike.FilterContour	9994.2	53559.4	6748.1
pillike.FilterEmboss	10139.8	54137.7	6842.0
pillike.FilterSharpen	10172.6	53900.8	6790.9
pillike.FilterDetail	10248.3	54244.5	6855.0
pillike.Affine	n/a	n/a	n/a
Superpixels (max_size=64, cubic)	10359.8	52488.8	6874.9
Superpixels (max_size=64, linear)	10057.4	52556.2	6848.3
Superpixels (max_size=128, linear)	10080.4	52098.7	6848.5
Superpixels (max_size=224, linear)	10239.8	52714.2	6766.6
UniformVoronoi (250-1000k points, linear)	10315.7	52527.3	6775.9
RegularGridVoronoi (16-31 rows/cols)	10273.0	52524.5	6901.0
RelativeRegularGridVoronoi (7%-14% rows/cols)	10261.5	52644.6	6807.9
Resize (nearest)	3887.2	11409.5	2163.3
Resize (linear)	4004.2	10753.6	2001.4
Resize (cubic)	3936.0	10136.8	1783.1
CropAndPad	3339.4	11911.4	1911.5
CropAndPad (edge)	3312.9	11891.2	1906.3
CropAndPad (keep_size)	2699.4	6842.7	1277.9
Crop	4236.7	19033.3	3282.4
Crop (keep_size)	3254.2	8776.2	1677.3
Pad	3242.5	10935.9	1699.2
Pad (edge)	3219.3	10925.4	1708.0
Pad (keep_size)	2633.0	6415.0	1112.3
PadToFixedSize	3535.1	13465.6	1924.5
CropToFixedSize	5043.4	23896.6	3620.4
KeepSizeByResize (CropToFixedSize(nearest))	3223.5	10349.5	1924.9
KeepSizeByResize (CropToFixedSize(linear))	3191.6	9791.4	1767.9
KeepSizeByResize (CropToFixedSize(cubic))	3165.5	9287.2	1593.5
FastSnowyLandscape	10132.9	51749.5	6844.8
Clouds	5753.1	47856.2	4557.8
Fog	10272.6	52218.0	6785.4
CloudLayer	10143.9	52216.4	6744.8

Table 4 – continued from previous page

Snowflakes	5740.8	47794.5	4474.9
SnowflakesLayer	10257.8	52025.2	7009.1
Rain	5786.5	47612.3	4507.3
RainLayer	9512.3	51505.4	6829.1

10.4 Keypoints and Bounding Boxes

Numbers below are for keypoints on small and large images. Each `KeypointsOnImage` instance contained `10 Keypoint` instances. `B=1` denotes a batch size of `1`, `B=128` one of `128`.

The numbers for bounding boxes can be derived by dividing each value by 4.

Number of augmented Keypoint instances per sec (divide by 10 for KeypointsOnImage instances):

	10 KPs on 224x224x3	
Augmenter	B=1	B=128
Sequential (2xIdentity)	37012.5	1082118.9
Sequential (2xIdentity, random_order)	29576.2	1018050.4
SomeOf (1-3, 3xIdentity)	14592.0	524940.8
SomeOf (1-3, 3xIdentity, random_order)	14223.8	518189.2
OneOf (3xIdentity)	26817.8	734837.0
Sometimes (Identity)	18553.2	818134.7
WithChannels ([1,2], Identity)	33973.0	569703.3
Identity	42157.2	1072697.7
Noop	42000.7	1069077.8
Lambda (return input)	38945.8	1036321.0
AssertLambda (return True)	37957.5	1032700.9
AssertShape (None, H, W, None)	34536.2	805366.0
ChannelShuffle (0.5)	41079.3	1064524.5
Add	41118.5	1065848.9
AddElementwise	40655.1	1059295.9
AdditiveGaussianNoise	40977.6	1048931.1
AdditiveLaplaceNoise	40870.1	1064200.9
AdditivePoissonNoise	40427.3	1056247.5
Multiply	40921.6	1054364.0
MultiplyElementwise	40417.9	1065871.4
Cutout (1 iter, constant fill)	41902.8	1063992.8
Dropout (1-5%)	40515.3	1052003.2
CoarseDropout (1-5%, size=1-10%)	40723.5	1057345.9
Dropout2d (10%)	15761.6	460892.5
TotalDropout (10%)	19089.1	817259.6
ReplaceElementwise	40863.7	996760.1
ImpulseNoise	41210.6	983126.4
SaltAndPepper	42135.2	993933.0
CoarseSaltAndPepper	41175.5	1003854.8
Salt	41061.6	1001330.9
CoarseSalt	40894.0	988559.6
Pepper	41267.4	1006313.5
CoarsePepper	41782.6	1005808.2

Continued on next page

Table 5 – continued from previous page

Invert (10%)	41179.3	1013027.7
JpegCompression (50-99%)	41222.5	1002183.3
Cartoon	41742.7	1016170.7
BlendAlpha (Identity)	13638.3	231006.1
BlendAlphaElementwise (Identity)	7628.0	15529.5
BlendAlphaSimplexNoise (Identity)	1805.8	2281.7
BlendAlphaFrequencyNoise (Identity)	2417.1	2942.4
BlendAlphaSomeColors (Identity)	n/a	n/a
BlendAlphaHorizontalL.Grad. (Identity)	6762.3	24122.5
BlendAlphaVerticalL.Grad. (Identity)	6420.3	19841.8
BlendAlphaRegularGrid (Identity)	6007.6	18590.6
BlendAlphaCheckerboard (Identity)	6039.8	16668.2
GaussianBlur (sigma=(1,5))	40927.0	1086402.0
AverageBlur	41359.6	1067578.3
MedianBlur	41554.1	1071704.2
BilateralBlur	41461.5	1077808.8
MotionBlur	41110.4	1076676.1
MeanShiftBlur	41129.7	1062169.5
RandAugment (n=2, m=(6,12))	n/a	n/a
WithColorspace (HSV, Identity)	36815.8	1046132.3
WithBrightnessChannels (Identity)	36561.2	1039062.1
MultiplyAndAddToBrightness	28022.1	964196.3
MultiplyBrightness	35559.8	1017404.7
AddToBrightness	35527.7	1022234.8
WithHueAndSaturation	35777.0	1024870.5
MultiplyHueAndSaturation	35910.5	1021646.0
MultiplyHue	29659.9	592502.1
MultiplySaturation	28923.3	600158.0
RemoveSaturation	29378.7	598766.6
AddToHueAndSaturation	40986.2	1119725.2
AddToHue	41381.9	1103111.1
AddToSaturation	41058.4	1114935.0
ChangeColorspace (HSV)	40768.4	1117990.4
Grayscale	41099.1	1100846.1
KMeansColorQuantization (2-16 colors)	42215.3	1103010.2
UniformColorQuantization (2-16 colors)	41109.3	1112418.4
UniformColorQuant.NBits (1-7 bits)	41959.8	1096977.3
Posterize (1-7 bits)	41053.5	1095550.6
GammaContrast	41231.7	1080049.3
SigmoidContrast	40382.1	1100256.5
LogContrast	40291.1	1106150.6
LinearContrast	41104.0	1107891.5
AllChannelsHistogramEqualization	40714.5	1080249.2
HistogramEqualization	41734.9	1090530.2
AllChannelsCLAHE	41319.9	1100346.7
CLAHE	41405.8	1114551.1
Convolve (3x3)	41350.9	1102244.4
Sharpen	40292.7	1104893.9
Emboss	41969.9	1100852.1
EdgeDetect	40670.9	1103836.8

Continued on next page

Table 5 – continued from previous page

DirectedEdgeDetect	40743.0	1106445.4
Canny	40905.1	1091215.8
Fliplr (p=100%)	23140.1	706230.7
Flipud (p=100%)	23031.6	699545.0
Affine (order=0, constant)	6562.1	33646.8
Affine (order=1, constant)	6531.6	33854.2
Affine (order=3, constant)	6551.7	33967.8
Affine (order=1, edge)	6539.0	34790.0
Affine (order=1, constant, skimage)	6545.8	34800.9
PiecewiseAffine (4x4, order=1, constant)	124.5	128.0
PiecewiseAffine (4x4, order=0, constant)	124.9	125.9
PiecewiseAffine (4x4, order=1, edge)	124.5	126.0
PiecewiseAffine (8x8, order=1, constant)	56.8	57.1
PerspectiveTransform	7538.8	65378.6
PerspectiveTransform (keep_size)	6866.2	34643.2
ElasticTransformation (order=0, constant)	1443.7	1698.9
ElasticTransformation (order=1, constant)	1403.2	1699.2
ElasticTransformation (order=1, nearest)	1387.6	1703.0
ElasticTransformation (order=1, reflect)	1438.2	1724.1
Rot90	15168.5	332321.6
Rot90 (keep_size)	14909.9	326491.3
WithPolarWarping (Identity)	15354.0	103429.7
Jigsaw (rows/cols=(3,8), 1 step)	5414.4	11859.6
AveragePooling	10527.7	58352.0
AveragePooling (keep_size)	40770.0	1077687.6
MaxPooling	10744.3	57821.0
MaxPooling (keep_size)	40547.7	1083825.7
MinPooling	10595.1	57142.4
MinPooling (keep_size)	40991.0	1081313.9
MedianPooling	10725.2	57799.0
MedianPooling (keep_size)	40305.6	1084803.9
imgcorruptlike.GaussianNoise((1,5))	40730.9	1087593.4
imgcorruptlike.ShotNoise((1,5))	41105.6	1074974.5
imgcorruptlike.ImpulseNoise((1,5))	40888.7	1092987.1
imgcorruptlike.SpeckleNoise((1,5))	40711.9	1070581.5
imgcorruptlike.GaussianBlur((1,5))	40999.5	1082354.5
imgcorruptlike.GlassBlur((1,5))	41458.0	1080668.9
imgcorruptlike.DefocusBlur((1,5))	40992.6	1078547.8
imgcorruptlike.MotionBlur((1,5))	42025.0	1043556.9
imgcorruptlike.ZoomBlur((1,5))	40322.1	1090790.9
imgcorruptlike.Fog((1,5))	41586.4	1074558.5
imgcorruptlike.Frost((1,5))	40853.3	1090004.7
imgcorruptlike.Snow((1,5))	40003.9	1086178.6
imgcorruptlike.Spatter((1,5))	41532.1	1076336.5
imgcorruptlike.Contrast((1,5))	40690.3	1089199.7
imgcorruptlike.Brightness((1,5))	41673.2	1078498.0
imgcorruptlike.Saturate((1,5))	40142.6	1082613.8
imgcorruptlike.JpegCompression((1,5))	41298.2	1090813.8
imgcorruptlike.Pixelate((1,5))	40576.9	1078943.8
imgcorruptlike.ElasticTransform((1,5))	n/a	n/a

Continued on next page

Table 5 – continued from previous page

pillike.Solarize (p=1.0)	40884.9	1050872.0
pillike.Posterize (1-7 bits)	41180.4	1079403.7
pillike.Equalize	40595.3	1093551.1
pillike.Autocontrast	40986.2	1083557.4
pillike.EnhanceColor	41340.6	1095687.7
pillike.EnhanceContrast	41217.6	1082945.4
pillike.EnhanceBrightness	41036.4	1084161.3
pillike.EnhanceSharpness	41666.4	1102461.7
pillike.FilterBlur	40530.9	1093928.5
pillike.FilterSmooth	41571.7	1092077.1
pillike.FilterSmoothMore	40467.9	1099664.6
pillike.FilterEdgeEnhance	40796.9	1084941.3
pillike.FilterEdgeEnhanceMore	41247.9	1092474.1
pillike.FilterFindEdges	41696.2	1096042.6
pillike.FilterContour	40493.4	1092456.3
pillike.FilterEmboss	41637.7	1099112.3
pillike.FilterSharpen	41646.0	1098764.4
pillike.FilterDetail	40681.9	1104490.8
pillike.Affine	n/a	n/a
Superpixels (max_size=64, cubic)	42058.7	1086970.3
Superpixels (max_size=64, linear)	40705.6	1090370.0
Superpixels (max_size=128, linear)	41916.8	1093384.5
Superpixels (max_size=224, linear)	41026.1	1056588.1
UniformVoronoi (250-1000k points, linear)	41827.0	1070845.2
RegularGridVoronoi (16-31 rows/cols)	40995.0	1074150.0
RelativeRegularGridVoronoi (7%-14% rows/cols)	41378.2	1109759.1
Resize (nearest)	11282.7	53402.6
Resize (linear)	11297.8	53518.4
Resize (cubic)	11210.7	52994.5
CropAndPad	8939.7	79988.7
CropAndPad (edge)	8870.1	79660.3
CropAndPad (keep_size)	7565.8	33852.5
Crop	10042.0	80279.6
Crop (keep_size)	8376.0	34519.8
Pad	8946.2	79155.3
Pad (edge)	8950.1	79899.8
Pad (keep_size)	7589.9	34261.0
PadToFixedSize	10459.6	255851.2
CropToFixedSize	13165.0	269638.4
KeepSizeByResize (CropToFixedSize(nearest))	8923.0	55861.6
KeepSizeByResize (CropToFixedSize(linear))	8876.8	55959.4
KeepSizeByResize (CropToFixedSize(cubic))	8877.9	56122.5
FastSnowyLandscape	39966.1	1080643.5
Clouds	14326.3	548492.1
Fog	41600.9	1069986.2
CloudLayer	40820.5	1085618.5
Snowflakes	14110.5	536123.0
SnowflakesLayer	40336.6	1086376.4
Rain	13720.2	533951.5
RainLayer	40263.1	1086607.3

CHAPTER 11

dtype support

The function `augment_images()`, which all augmenters in `imgaug` offer, works by default with numpy arrays. In most cases, these arrays will have the numpy dtype `uint8`, i.e. the images will have values in the range `[0, 255]`. This is the datatype returned by most image loading functions. Sometimes however you may want to augment other datatypes, e.g. `float64`. While all augmenters support `uint8`, the support for other datatypes varies. The tables further below show which datatype is supported by which augmenter (alongside the dtype support in some helper functions). The API of each augmenter may contain more details.

Note: Whenever possible it is suggested to use `uint8` as that is the most thoroughly tested dtype. In general, the use of large dtypes (i.e. `uint64`, `int64`, `float128`) is discouraged, even when they are marked as supported. That is because writing correct tests for these dtypes can be difficult as no larger dtypes are available to which values can be temporarily cast. Additionally, when using inputs for which precise discrete values are important (e.g. segmentation maps, where an accidental change by `1` would break the map), medium sized dtypes (`uint32`, `int32`) should be used with caution. This is because other libraries may convert temporarily to `float64`, which could lead to inaccuracies for some numbers.

11.1 Legend

Support level (color of table cells):

- Green: Datatype is considered supported by the augmenter.
- Yellow: Limited support for the datatype, e.g. due to inaccuracies around large values. See the API for the respective augmenter for more details.
- Red: Datatype is not supported by the augmenter.

Test level (symbols in table cells):

- +++: Datatype support is thoroughly tested (via unittests or integration tests).
- ++: Datatype support is tested, though not thoroughly.
- +: Datatype support is indirectly tested via tests for other augmenters.
- −: Datatype support is not tested.

- `?`: Unknown support level for the datatype.

11.2 imgaug helper functions

	uint8	uint16	uint32	uint64	int8	int16	int32	int64	float16	float32	float64	float128	bool
draw_text	+++									-			
imresize_many_images	+++	++			++	++	++		++	++	++		++
imresize_single_image	+++	++			++	++	++		++	++	++		++
pool	+++	++	++		++	++	++		++	++	++	++	++
avg_pool	+++	++	++		++	++	++		++	++	++	++	++
max_pool	+++	++	++		++	++	++		++	++	++	++	++
min_pool	+++	++	++		++	++	++		++	++	++	++	++
median_pool	+++	++	++		++	++	++		++	++	++	++	++
draw_grid	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++
show_grid	-	?	?	?	?	?	?	?	?	?	?	?	?
imshow	-	?	?	?	?	?	?	?	?	?	?	?	?
apply_lut	+++												
apply_lut_	+++												

Fig. 1: Dtype support of helper functions in `imgaug`, e.g. `import imgaug; imgaug.imresize_single_image(array, size)`.

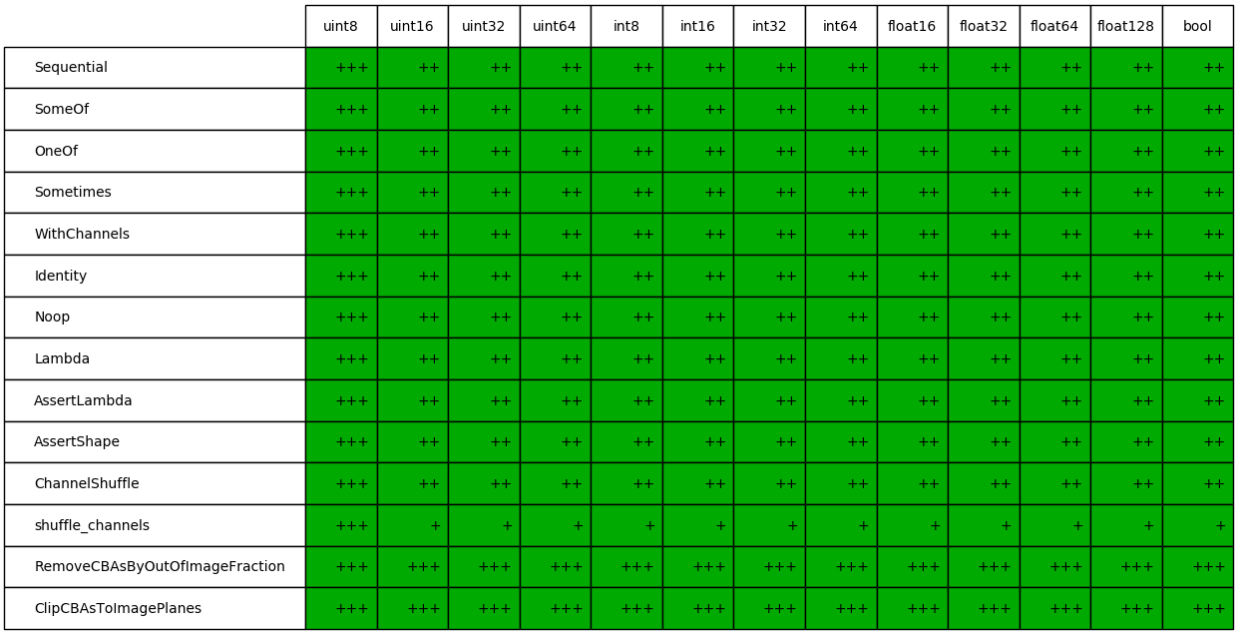

	uint8	uint16	uint32	uint64	int8	int16	int32	int64	float16	float32	float64	float128	bool
Sequential	+++	++	++	++	++	++	++	++	++	++	++	++	++
SomeOf	+++	++	++	++	++	++	++	++	++	++	++	++	++
OneOf	+++	++	++	++	++	++	++	++	++	++	++	++	++
Sometimes	+++	++	++	++	++	++	++	++	++	++	++	++	++
WithChannels	+++	++	++	++	++	++	++	++	++	++	++	++	++
Identity	+++	++	++	++	++	++	++	++	++	++	++	++	++
Noop	+++	++	++	++	++	++	++	++	++	++	++	++	++
Lambda	+++	++	++	++	++	++	++	++	++	++	++	++	++
AssertLambda	+++	++	++	++	++	++	++	++	++	++	++	++	++
AssertShape	+++	++	++	++	++	++	++	++	++	++	++	++	++
ChannelShuffle	+++	++	++	++	++	++	++	++	++	++	++	++	++
shuffle_channels	+++	+	+	+	+	+	+	+	+	+	+	+	+
RemoveCBAsByOutOfImageFraction	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++
ClipCBAsToImagePlanes	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++

Fig. 2: Image dtypes supported by augmenters and helper functions in `imgaug.augmenters.meta`.

	uint8	uint16	uint32	uint64	int8	int16	int32	int64	float16	float32	float64	float128	bool
add_scalar	+++	++			++	++			++	++			++
add_elementwise	+++	++			++	++			++	++			++
multiply_scalar	+++	++			++	++			++	++			++
multiply_elementwise	+++	++			++	++			++	++			++
cutout	+++	++	++	++	++	++	++	++	++	++	++	++	++
cutout_	+++	++	++	++	++	++	++	++	++	++	++	++	++
replace_elementwise_	+++	++	++		++	++	++		++	++	++		++
invert	+++	++	++		++	++	++		++	++			
invert_ (min_value=None and max_value=None)	+++	++	++	++	++	++	++	++	++	++	++	++	++
invert_ (min_value!=None or max_value!=None)	+++	++	++		++	++	++		++	++			
solarize	+++	++	++	++	++	++	++	++	++	++	++	++	++
solarize_	+++	++	++	++	++	++	++	++	++	++	++	++	++
compress_jpeg	+++	?	?	?	?	?	?	?	?	?	?	?	?
Add	+++	++			++	++			++	++			++
AddElementwise	+++	++			++	++			++	++			++
AdditiveGaussianNoise	+++	++			++	++			++	++			++
AdditiveLaplaceNoise	+++	++			++	++			++	++			++
AdditivePoissonNoise	+++	++			++	++			++	++			++
Multiply	+++	++			++	++			++	++			++
MultiplyElementwise	+++	++			++	++			++	++			++
Cutout	+++	++	++	++	++	++	++	++	++	++	++	++	++
Dropout	+++	++			++	++			++	++			++
CoarseDropout	+++	++			++	++			++	++			++
Dropout2d	+++	++	++	++	++	++	++	++	++	++	++	++	++
TotalDropout	+++	++	++	++	++	++	++	++	++	++	++	++	++

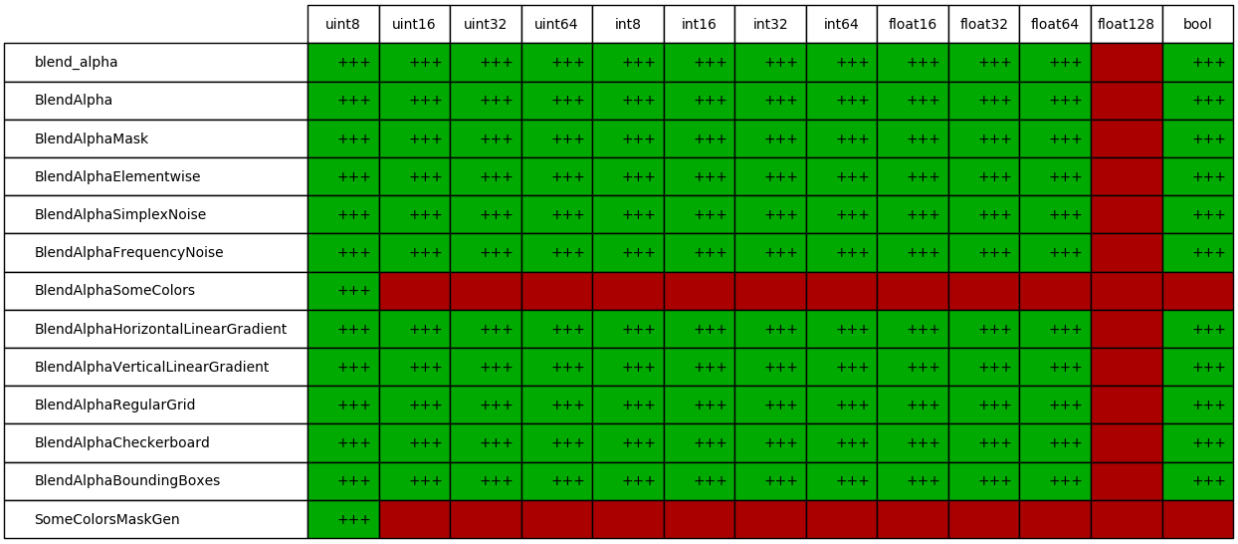

	uint8	uint16	uint32	uint64	int8	int16	int32	int64	float16	float32	float64	float128	bool
blend_alpha	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++		+++
BlendAlpha	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++		+++
BlendAlphaMask	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++		+++
BlendAlphaElementwise	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++		+++
BlendAlphaSimplexNoise	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++		+++
BlendAlphaFrequencyNoise	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++		+++
BlendAlphaSomeColors	+++												
BlendAlphaHorizontalLinearGradient	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++		+++
BlendAlphaVerticalLinearGradient	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++		+++
BlendAlphaRegularGrid	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++		+++
BlendAlphaCheckerboard	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++		+++
BlendAlphaBoundingBoxes	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++		+++
SomeColorsMaskGen	+++												

Fig. 4: Image dtypes supported by augmenters and helper functions in `imgaug.augmenters.blend`.

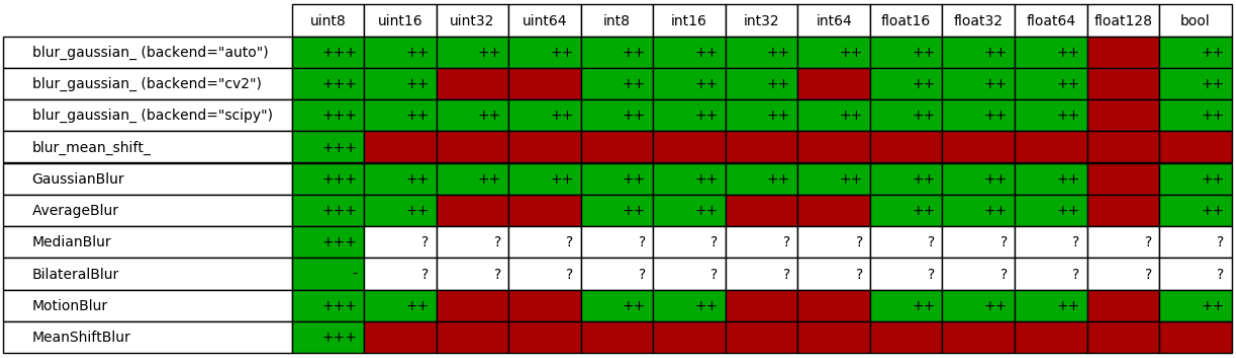

	uint8	uint16	uint32	uint64	int8	int16	int32	int64	float16	float32	float64	float128	bool
blur_gaussian_ (backend="auto")	+++	++	++	++	++	++	++	++	++	++	++		++
blur_gaussian_ (backend="cv2")	+++	++			++	++	++		++	++	++		++
blur_gaussian_ (backend="scipy")	+++	++	++	++	++	++	++	++	++	++	++		++
blur_mean_shift_	+++												
GaussianBlur	+++	++	++	++	++	++	++	++	++	++	++		++
AverageBlur	+++	++			++	++			++	++	++		++
MedianBlur	+++	?	?	?	?	?	?	?	?	?	?	?	?
BilateralBlur	-	?	?	?	?	?	?	?	?	?	?	?	?
MotionBlur	+++	++			++	++			++	++	++		++
MeanShiftBlur	+++												

Fig. 5: Image dtypes supported by augmenters and helper functions in `imgaug.augmenters.blur`.

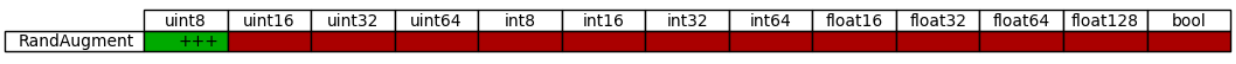

	uint8	uint16	uint32	uint64	int8	int16	int32	int64	float16	float32	float64	float128	bool
RandAugment	+++												

Fig. 6: Image dtypes supported by augmenters and helper functions in `imgaug.augmenters.collections`.

	uint8	uint16	uint32	uint64	int8	int16	int32	int64	float16	float32	float64	float128	bool
change_colorspace_	+++												
change_colorspaces_	+++												
change_color_temperatures_	+++												
change_color_temperature	+++												
WithColorspace	+++												
WithBrightnessChannels	+++												
MultiplyAndAddToBrightness	+++												
MultiplyBrightness	+++												
AddToBrightness	+++												
WithHueAndSaturation	+++												
MultiplyHueAndSaturation	+++												
MultiplyHue	+++												
MultiplySaturation	+++												
RemoveSaturation	+++												
AddToHueAndSaturation	+++												
AddToHue	+++												
AddToSaturation	+++												
ChangeColorspace	+++												
Grayscale	+++												
ChangeColorTemperature	+++												
KMeansColorQuantization (image size <= max_size)	+++												
KMeansColorQuantization (image size > max_size)	+++												
quantize_kmeans	+++												
UniformColorQuantization (image size <= max_size)	+++												
UniformColorQuantization (image size > max_size)	+++												
UniformColorQuantizationToNBits (image size <= max_size)	+++												
UniformColorQuantizationToNBits (image size > max_size)	+++												
Posterize	+++												
quantize_uniform	+++												
quantize_uniform_	+++												
quantize_uniform_to_n_bits	+++												

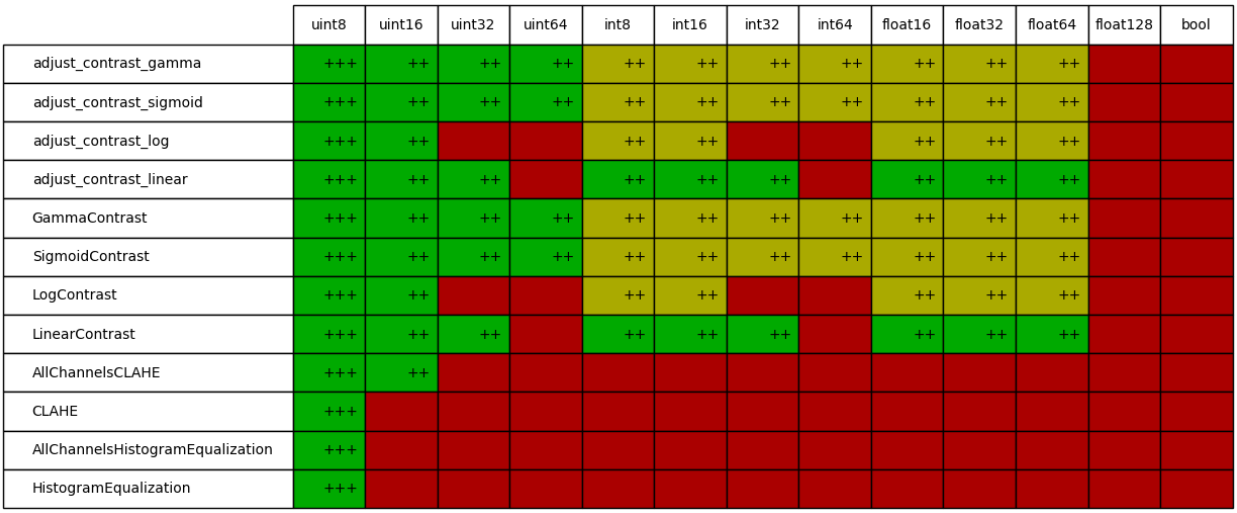

	uint8	uint16	uint32	uint64	int8	int16	int32	int64	float16	float32	float64	float128	bool
adjust_contrast_gamma	+++	++	++	++	++	++	++	++	++	++	++		
adjust_contrast_sigmoid	+++	++	++	++	++	++	++	++	++	++	++		
adjust_contrast_log	+++	++			++	++			++	++	++		
adjust_contrast_linear	+++	++	++		++	++	++		++	++	++		
GammaContrast	+++	++	++	++	++	++	++	++	++	++	++		
SigmoidContrast	+++	++	++	++	++	++	++	++	++	++	++		
LogContrast	+++	++			++	++			++	++	++		
LinearContrast	+++	++	++		++	++	++		++	++	++		
AllChannelsCLAHE	+++	++											
CLAHE	+++												
AllChannelsHistogramEqualization	+++												
HistogramEqualization	+++												

Fig. 8: Image dtypes supported by augmenters and helper functions in `imgaug.augmenters.contrast`.

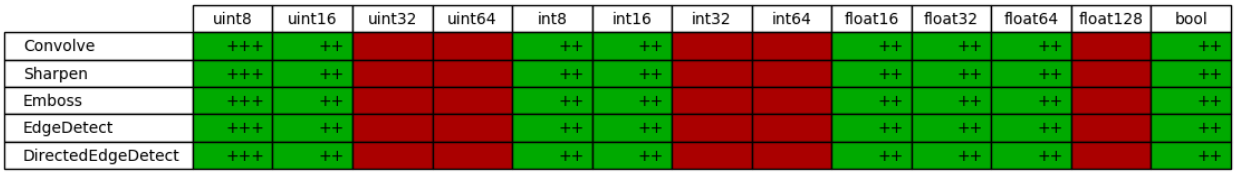

	uint8	uint16	uint32	uint64	int8	int16	int32	int64	float16	float32	float64	float128	bool
Convolve	+++	++			++	++			++	++	++		++
Sharpen	+++	++			++	++			++	++	++		++
Emboss	+++	++			++	++			++	++	++		++
EdgeDetect	+++	++			++	++			++	++	++		++
DirectedEdgeDetect	+++	++			++	++			++	++	++		++

Fig. 9: Image dtypes supported by augmenters and helper functions in `imgaug.augmenters.convolutional`.

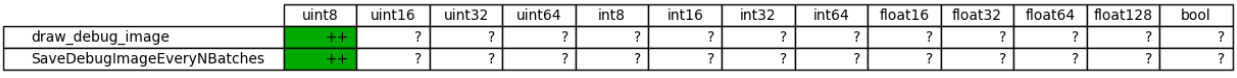

	uint8	uint16	uint32	uint64	int8	int16	int32	int64	float16	float32	float64	float128	bool
draw_debug_image	++	?	?	?	?	?	?	?	?	?	?	?	?
SaveDebugImageEveryNBatches	++	?	?	?	?	?	?	?	?	?	?	?	?

Fig. 10: Image dtypes supported by augmenters and helper functions in `imgaug.augmenters.debug`.

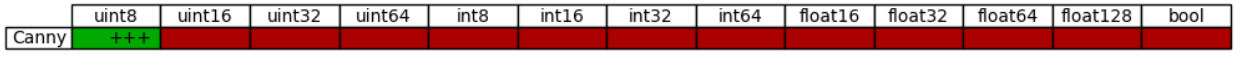

	uint8	uint16	uint32	uint64	int8	int16	int32	int64	float16	float32	float64	float128	bool
Canny	+++												

Fig. 11: Image dtypes supported by augmenters and helper functions in `imgaug.augmenters.edges`.

11.3 imgaug.augmenters.meta

11.4 imgaug.augmenters.arithmetic

11.5 imgaug.augmenters.blend

11.6 imgaug.augmenters.blur

11.7 imgaug.augmenters.collections

11.8 imgaug.augmenters.color

11.9 imgaug.augmenters.contrast

11.10 imgaug.augmenters.convolutional

11.11 imgaug.augmenters.debug

11.12 imgaug.augmenters.edges

11.13 imgaug.augmenters.flip

	uint8	uint16	uint32	uint64	int8	int16	int32	int64	float16	float32	float64	float128	bool
fliplr	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++
flipud	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++
Fliplr	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++
Flipud	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++

Fig. 12: Image dtypes supported by augmenters and helper functions in `imgaug.augmenters.flip`.

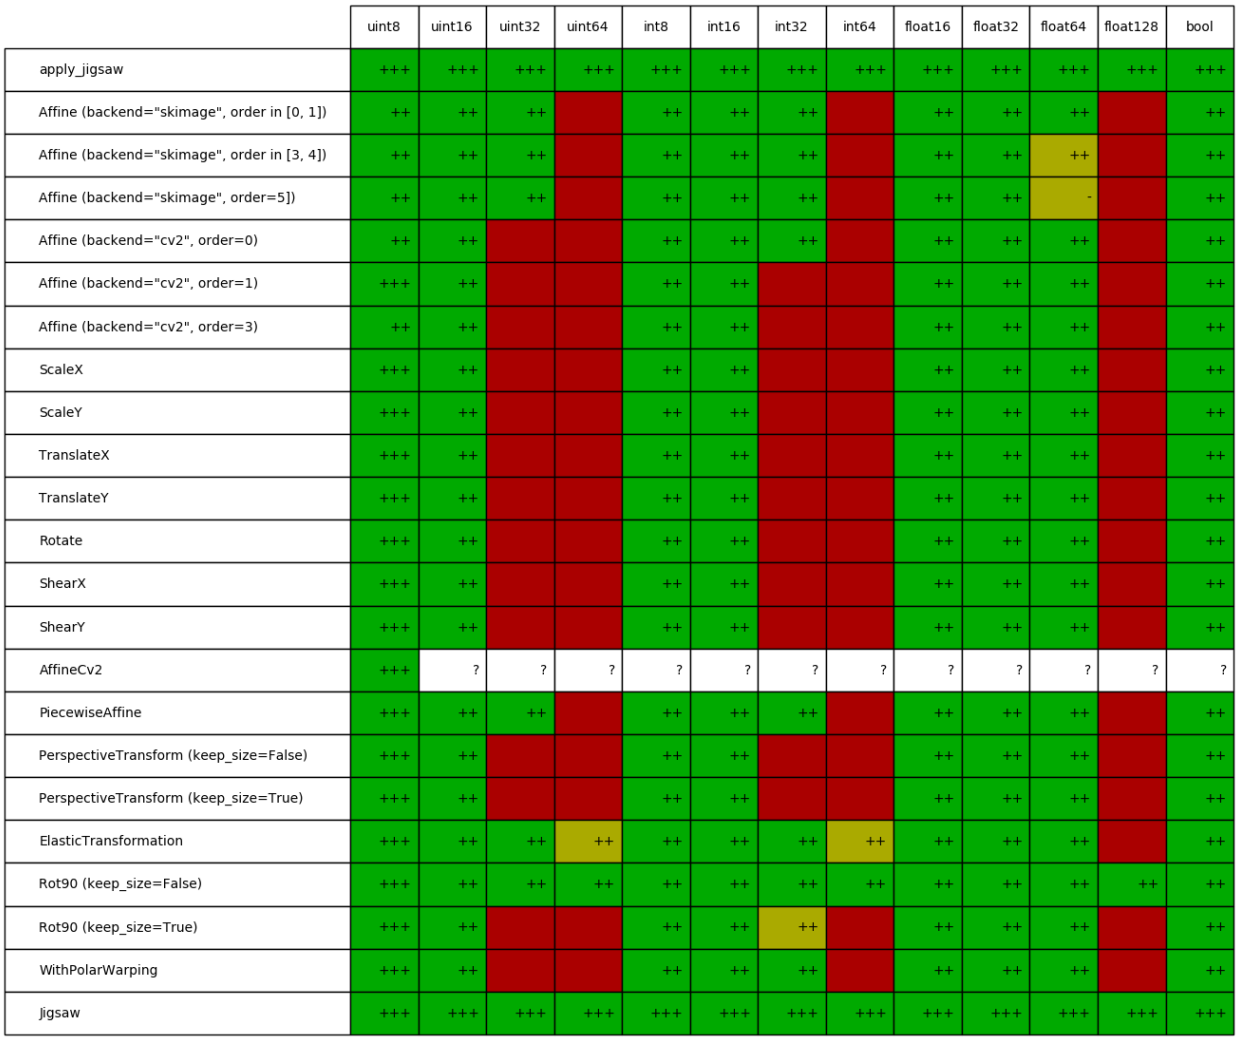

Fig. 13: Image dtypes supported by augmenters and helper functions in `imgaug.augmenters.geometric`.

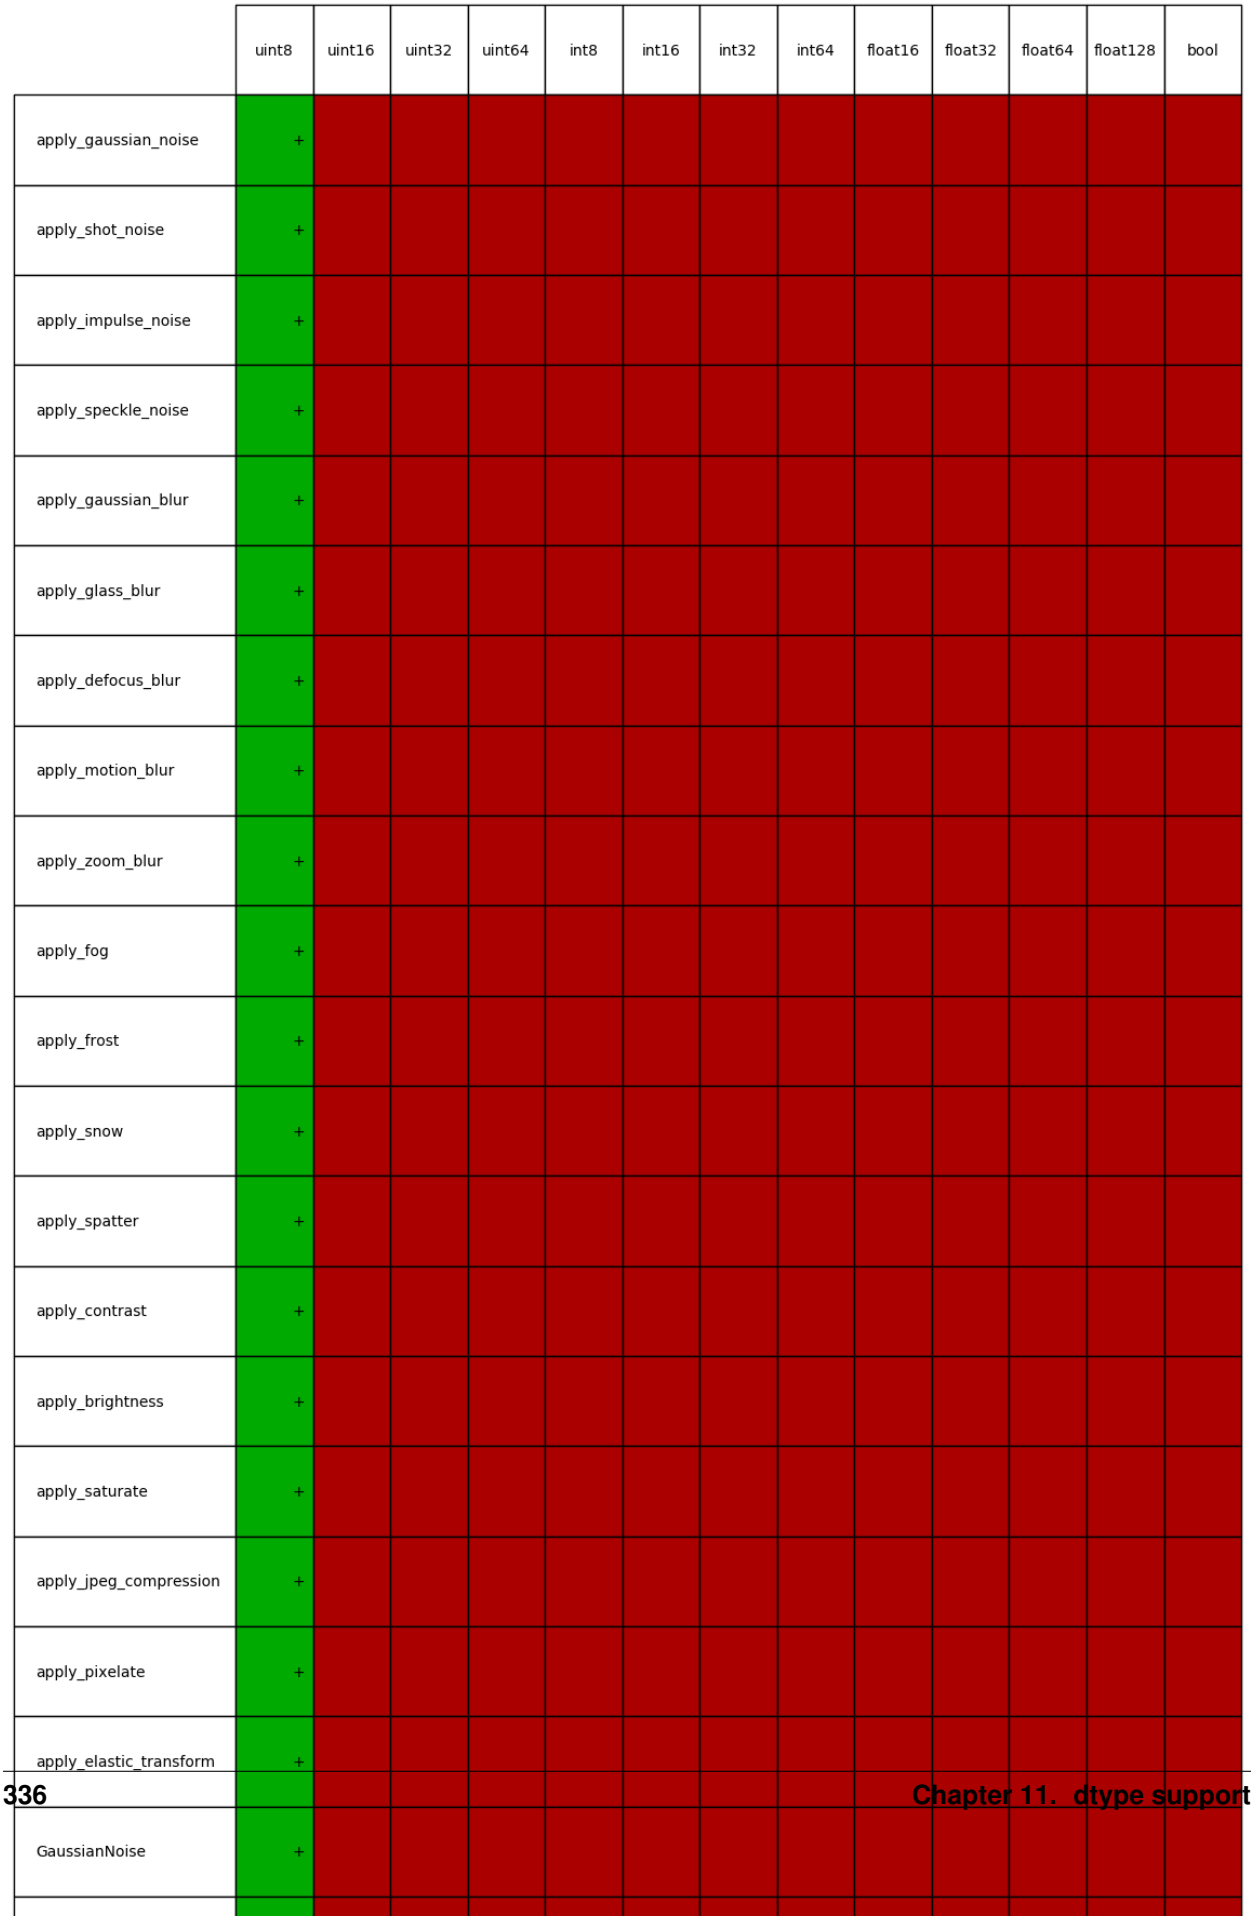
uint8
uint16
uint32
uint64
int8
int16
int32
int64
float16
float32
float64
float128
bool
apply_gaussian_noise
apply_shot_noise
apply_impulse_noise
apply_speckle_noise
apply_gaussian_blur
apply_glass_blur
apply_defocus_blur
apply_motion_blur
apply_zoom_blur
apply_fog
apply_frost
apply_snow
apply_spatter
apply_contrast
apply_brightness
apply_saturate
apply_jpeg_compression
apply_pixelate
apply_elastic_transform
GaussianNoise

	uint8	uint16	uint32	uint64	int8	int16	int32	int64	float16	float32	float64	float128	bool
solarize_	+++	++	++	++	++	++	++	++	++	++	++	++	++
solarize	+++	++	++	++	++	++	++	++	++	++	++	++	++
posterize_	+++												
posterize	+++												
equalize	+++												
equalize_	+++												
autocontrast	+++												
enhance_color	+++												
enhance_contrast	+++												
enhance_brightness	+++												
enhance_sharpness	+++												
filter_blur	+++												
filter_smooth	+++												
filter_smooth_more	+++												
filter_edge_enhance	+++												
filter_edge_enhance_more	+++												
filter_find_edges	+++												
filter_contour	+++												
filter_emboss	+++												
filter_sharpen	+++												

	uint8	uint16	uint32	uint64	int8	int16	int32	int64	float16	float32	float64	float128	bool
Superpixels (image size <= max_size)	+++	++	++	-	++	++	++	-					++
Superpixels (image size > max_size)	+++	++			++	++	++						++
Voronoi (image size <= max_size)	+++												
Voronoi (image size > max_size)	+++												
UniformVoronoi	+++												
RegularGridVoronoi	+++												
RelativeRegularGridVoronoi	+++												

Fig. 16: Image dtypes supported by augmenters and helper functions in `imgaug.augmenters.segmentation`.

11.14 imgaug.augmenters.geometric

11.15 imgaug.augmenters.imgcorruptlike

11.16 imgaug.augmenters.pillike

11.17 imgaug.augmenters.segmentation

11.18 imgaug.augmenters.size

11.19 imgaug.augmenters.weather

	uint8	uint16	uint32	uint64	int8	int16	int32	int64	float16	float32	float64	float128	bool
pad	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	++
pad_to_aspect_ratio	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	++
pad_to_multiples_of	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	++
Resize	+++	++			++	++	++		++	++	++		++
CropAndPad (keep_size=False)	+++	++	++	++	++	++	++	++	++	++	++	++	++
CropAndPad (keep_size=True)	+++	++			++	++	++		++	++	++		++
Pad	+++	++			++	++	++		++	++	++		++
Crop	+++	++			++	++	++		++	++	++		++
PadToFixedSize	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	++
CenterPadToFixedSize	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	++
CropToFixedSize	+++	++	++	++	++	++	++	++	++	++	++	++	++
CenterCropToFixedSize	+++	++	++	++	++	++	++	++	++	++	++	++	++
CropToMultiplesOf	+++	++	++	++	++	++	++	++	++	++	++	++	++
CenterCropToMultiplesOf	+++	++	++	++	++	++	++	++	++	++	++	++	++
PadToMultiplesOf	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	++
CenterPadToMultiplesOf	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	++
CropToPowersOf	+++	++	++	++	++	++	++	++	++	++	++	++	++
CenterCropToPowersOf	+++	++	++	++	++	++	++	++	++	++	++	++	++
PadToPowersOf	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	++
CenterPadToPowersOf	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	++
CropToAspectRatio	+++	++	++	++	++	++	++	++	++	++	++	++	++
CenterCropToAspectRatio	+++	++	++	++	++	++	++	++	++	++	++	++	++
PadToAspectRatio	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	++
CenterPadToAspectRatio	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	++
CropToSquare	+++	++	++	++	++	++	++	++	++	++	++	++	++
CenterCropToSquare	+++	++	++	++	++	++	++	++	++	++	++	++	++
PadToSquare	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	++

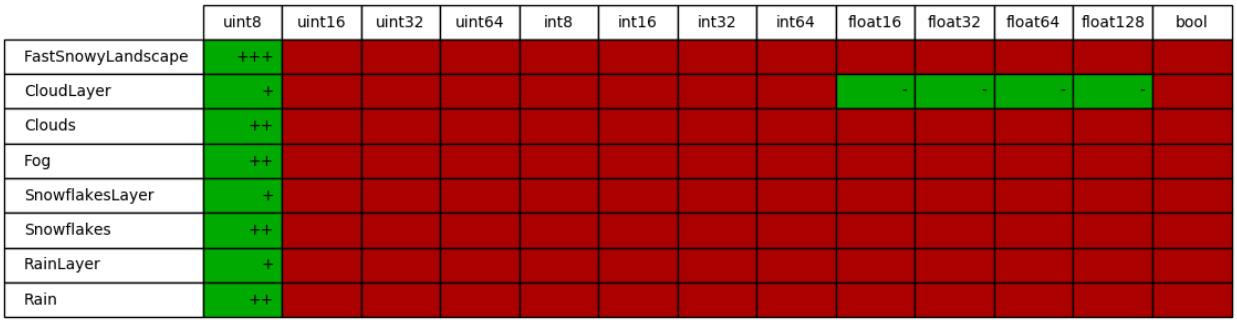

	uint8	uint16	uint32	uint64	int8	int16	int32	int64	float16	float32	float64	float128	bool
FastSnowyLandscape	+++												
CloudLayer	+								-	-	-	-	
Clouds	++												
Fog	++												
SnowflakesLayer	+												
Snowflakes	++												
RainLayer	+												
Rain	++												

Fig. 18: Image dtypes supported by augmenters and helper functions in `imgaug.augmenters.weather`.

CHAPTER 12

Jupyter Notebooks

Several jupyter notebooks are available that provide tutorials about *imgaug*'s functionality. They are hosted at imgaug-doc/notebooks. The notebooks can be downloaded to interactively modify the examples.

List of Notebooks:

- A01 - Load and Augment an Image
- A03 - Multicore Augmentation
- B01 - Augment Keypoints (aka Landmarks)
- B02 - Augment Bounding Boxes
- B03 - Augment Polygons
- B06 - Augment Line Strings
- B04 - Augment Heatmaps
- B05 - Augment Segmentation Maps
- C01 - Using Probability Distributions as Parameters
- C02 - Using imgaug with more Control Flow
- C03 - Stochastic and Deterministic Augmentation
- C04 - Copying Random States and Using Multiple Augmentation Sequences

CHAPTER 13

API

13.1 imgaug

Collection of basic functions used throughout imgaug.

imgaug.imgaug.**BackgroundAugmenter**(**args*, ***kwargs*)

imgaug.imgaug.**Batch**(**args*, ***kwargs*)

imgaug.imgaug.**BatchLoader**(**args*, ***kwargs*)

imgaug.imgaug.**BoundingBox**(**args*, ***kwargs*)

imgaug.imgaug.**BoundingBoxesOnImage**(**args*, ***kwargs*)

exception imgaug.imgaug.**DeprecationWarning**

Bases: Warning

Warning for deprecated calls.

Since python 2.7 DeprecatedWarning is silent by default. So we define our own DeprecatedWarning here so that it is not silent by default.

imgaug.imgaug.**HeatmapsOnImage**(**args*, ***kwargs*)

class imgaug.imgaug.**HooksHeatmaps**(*activator=None*, *propagator=None*, *preprocessor=None*, *postprocessor=None*)

Bases: *imgaug.imgaug.HooksImages*

Class to intervene with heatmap augmentation runs.

This is e.g. useful to dynamically deactivate some augmenters.

This class is currently the same as the one for images. This may or may not change in the future.

Methods

is_activated(self, images, augmenter, ...)	Estimate whether an augmenter may be executed.
is_propagating(self, images, augmenter, ...)	Estimate whether an augmenter may call its children.
postprocess(self, images, augmenter, parents)	Postprocess input data per augmenter after augmentation.
preprocess(self, images, augmenter, parents)	Preprocess input data per augmenter before augmentation.

class imgaug.imgaug.**HooksImages**(*activator=None*, *propagator=None*, *preprocessor=None*, *postprocessor=None*)

Bases: object

Class to intervene with image augmentation runs.

This is e.g. useful to dynamically deactivate some augmenters.

Parameters

- **activator** (*None or callable, optional*) – A function that gives permission to execute an augmenter. The expected interface is:

```
``f(images, augmenter, parents, default)``
```

  where images are the input images to augment, augmenter is the instance of the augmenter to execute, parents are previously executed augmenters and default is an expected default value to be returned if the activator function does not plan to make a decision for the given inputs.

- **propagator** (*None or callable, optional*) – A function that gives permission to propagate the augmentation further to the children of an augmenter. This happens after the activator. In theory, an augmenter may augment images itself (if allowed by the activator) and then execute child augmenters afterwards (if allowed by the propagator). If the activator returned False, the propagation step will never be executed. The expected interface is:

```
``f(images, augmenter, parents, default)``
```

  with all arguments having identical meaning to the activator.

- **preprocessor** (*None or callable, optional*) – A function to call before an augmenter performed any augmentations. The interface is:

```
f(images, augmenter, parents)
```

  with all arguments having identical meaning to the activator. It is expected to return the input images, optionally modified.

- **postprocessor** (*None or callable, optional*) – A function to call after an augmenter performed augmentations. The interface is the same as for the *preprocessor*.

Examples

```
>>> import numpy as np
>>> import imgaug as ia
>>> import imgaug.augmenters as iaa
>>> seq = iaa.Sequential([
>>>     iaa.GaussianBlur(3.0, name="blur"),
>>>     iaa.Dropout(0.05, name="dropout"),
>>>     iaa.Affine(translate_px=-5, name="affine")
```

(continues on next page)

(continued from previous page)

```
>>> ])
>>> images = [np.zeros((10, 10), dtype=np.uint8)]
>>>
>>> def activator(images, augmenter, parents, default):
>>>     return False if augmenter.name in ["blur", "dropout"] else default
>>>
>>> seq_det = seq.to_deterministic()
>>> images_aug = seq_det.augment_images(images)
>>> heatmaps = [np.random.rand(*(3, 10, 10))]
>>> heatmaps_aug = seq_det.augment_images(
>>>     heatmaps,
>>>     hooks=ia.HooksImages(activator=activator)
>>> )
```

This augments images and their respective heatmaps in the same way. The heatmaps however are only modified by `Affine`, not by `GaussianBlur` or `Dropout`.

Methods

is_activated(self, images, augmenter, ...)	Estimate whether an augmenter may be executed.
is_propagating(self, images, augmenter, ...)	Estimate whether an augmenter may call its children.
postprocess(self, images, augmenter, parents)	Postprocess input data per augmenter after augmentation.
preprocess(self, images, augmenter, parents)	Preprocess input data per augmenter before augmentation.

is_activated (*self*, *images*, *augmenter*, *parents*, *default*)
Estimate whether an augmenter may be executed.

This also affects propagation of data to child augmenters.

Returns If `True`, the augmenter may be executed. Otherwise `False`.

Return type bool

is_propagating (*self*, *images*, *augmenter*, *parents*, *default*)
Estimate whether an augmenter may call its children.

This function decides whether an augmenter with children is allowed to call these in order to further augment the inputs. Note that if the augmenter itself performs augmentations (before/after calling its children), these may still be executed, even if this method returns `False`.

Returns If `True`, the augmenter may propagate data to its children. Otherwise `False`.

Return type bool

postprocess (*self*, *images*, *augmenter*, *parents*)
Postprocess input data per augmenter after augmentation.

Returns The input images, optionally modified.

Return type (N,H,W,C) ndarray or (N,H,W) ndarray or list of (H,W,C) ndarray or list of (H,W) ndarray

preprocess (*self*, *images*, *augmenter*, *parents*)
Preprocess input data per augmenter before augmentation.

Returns The input images, optionally modified.

Return type (N,H,W,C) ndarray or (N,H,W) ndarray or list of (H,W,C) ndarray or list of (H,W) ndarray

class imgaug.imgaug.**HooksKeypoints**(*activator=None*, *propagator=None*, *preprocessor=None*, *postprocessor=None*)

Bases: *imgaug.imgaug.HooksImages*

Class to intervene with keypoint augmentation runs.

This is e.g. useful to dynamically deactivate some augmenters.

This class is currently the same as the one for images. This may or may not change in the future.

Methods

is_activated(self, images, augmenter, ...)	Estimate whether an augmenter may be executed.
is_propagating(self, images, augmenter, ...)	Estimate whether an augmenter may call its children.
postprocess(self, images, augmenter, parents)	Postprocess input data per augmenter after augmentation.
preprocess(self, images, augmenter, parents)	Preprocess input data per augmenter before augmentation.

imgaug.imgaug.**Keypoint**(**args*, ***kwargs*)

imgaug.imgaug.**KeypointsOnImage**(**args*, ***kwargs*)

imgaug.imgaug.**MultiPolygon**(**args*, ***kwargs*)

imgaug.imgaug.**Polygon**(**args*, ***kwargs*)

imgaug.imgaug.**PolygonsOnImage**(**args*, ***kwargs*)

imgaug.imgaug.**SegmentationMapsOnImage**(**args*, ***kwargs*)

imgaug.imgaug.**angle_between_vectors**(*v1*, *v2*)

Calculcate the angle in radians between vectors *v1* and *v2*.

From http://stackoverflow.com/questions/2827393/angles-between-two-n-dimensional-vectors-in-python

Parameters

- **v1** (*(N,) ndarray*) – First vector.
- **v2** (*(N,) ndarray*) – Second vector.

Returns Angle in radians.

Return type float

Examples

```
>>> angle_between_vectors(np.float32([1, 0, 0]), np.float32([0, 1, 0]))
1.570796...
```

```
>>> angle_between_vectors(np.float32([1, 0, 0]), np.float32([1, 0, 0]))
0.0
```

```
>>> angle_between_vectors(np.float32([1, 0, 0]), np.float32([-1, 0, 0]))
3.141592...
```

imgaug.imgaug.**apply_lut**(*image*, *table*)
Map an input image to a new one using a lookup table.

Added in 0.4.0.

Supported dtypes:

See *apply_lut_()*.

Parameters

- **image** (*ndarray*) – See *apply_lut_()*.
- **table** (*ndarray or list of ndarray*) – See *apply_lut_()*.

Returns Image after mapping via lookup table.

Return type ndarray

imgaug.imgaug.**apply_lut_**(*image*, *table*)
Map an input image in-place to a new one using a lookup table.

Added in 0.4.0.

Supported dtypes:

- `uint8`: yes; fully tested
- `uint16`: no
- `uint32`: no
- `uint64`: no
- `int8`: no
- `int16`: no
- `int32`: no
- `int64`: no
- `float16`: no
- `float32`: no
- `float64`: no
- `float128`: no
- `bool`: no

Parameters

- **image** (*ndarray*) – Image of dtype `uint8` and shape `(H,W)` or `(H,W,C)`.
- **table** (*ndarray or list of ndarray*) – Table of dtype `uint8` containing the mapping from old to new values. Either a `list` of `C` `(256,)` arrays or a single array of shape `(256,)` or `(256, C)` or `(1, 256, C)`. In case of `(256,)` the same table is used for all channels, otherwise a channelwise table is used and `C` is expected to match the number of channels.

Returns Image after mapping via lookup table. This *might* be the same array instance as provided via *image*.

Return type ndarray

imgaug.imgaug.**avg_pool**(*arr*, *block_size*, *pad_mode='reflect'*, *pad_cval=128*, *preserve_dtype=True*, *cval=None*)

Resize an array using average pooling.

Defaults to `pad_mode="reflect"` to ensure that padded values do not affect the average.

Supported dtypes:

See *pool()*.

Parameters

- **arr** (*(H,W) ndarray or (H,W,C) ndarray*) – Image-like array to pool. See *pool()* for details.
- **block_size** (*int or tuple of int or tuple of int*) – Size of each block of values to pool. See *pool()* for details.
- **pad_mode** (*str, optional*) – Padding mode to use if the array cannot be divided by *block_size* without remainder. See *pad()* for details.
- **pad_cval** (*number, optional*) – Padding value. See *pool()* for details.
- **preserve_dtype** (*bool, optional*) – Whether to preserve the input array dtype. See *pool()* for details.
- **cval** (*None or number, optional*) – Deprecated. Old name for *pad_cval*.

Returns Array after average pooling.

Return type (H',W') ndarray or (H',W',C') ndarray

imgaug.imgaug.**caller_name**()

Return the name of the caller, e.g. a function.

Returns The name of the caller as a string

Return type str

imgaug.imgaug.**compute_geometric_median**(**args*, ***kwargs*)

imgaug.imgaug.**compute_line_intersection_point**(*x1*, *y1*, *x2*, *y2*, *x3*, *y3*, *x4*, *y4*)

Compute the intersection point of two lines.

Taken from https://stackoverflow.com/a/20679579 .

Parameters

- **x1** (*number*) – x coordinate of the first point on line 1. (The lines extends beyond this point.)
- **y1** (*number*) – y coordinate of the first point on line 1. (The lines extends beyond this point.)
- **x2** (*number*) – x coordinate of the second point on line 1. (The lines extends beyond this point.)
- **y2** (*number*) – y coordinate of the second point on line 1. (The lines extends beyond this point.)
- **x3** (*number*) – x coordinate of the first point on line 2. (The lines extends beyond this point.)
- **y3** (*number*) – y coordinate of the first point on line 2. (The lines extends beyond this point.)
- **x4** (*number*) – x coordinate of the second point on line 2. (The lines extends beyond this point.)
- **y4** (*number*) – y coordinate of the second point on line 2. (The lines extends beyond this point.)

Returns The coordinate of the intersection point as a `tuple (x, y)`. If the lines are parallel (no intersection point or an infinite number of them), the result is `False`.

Return type tuple of number or bool

`imgaug.imgaug.`**`compute_paddings_for_aspect_ratio`**`(*args, **kwargs)`

`imgaug.imgaug.`**`compute_paddings_to_reach_exponents_of`**`(*args, **kwargs)`

`imgaug.imgaug.`**`compute_paddings_to_reach_multiples_of`**`(*args, **kwargs)`

`imgaug.imgaug.`**`copy_random_state`**`(random_state, force_copy=False)`

Deprecated. Use `imgaug.random.copy_generator_unless_global_rng` instead.

Copy an existing numpy (random number) generator.

Parameters

- **random_state** (*numpy.random.Generator or numpy.random.RandomState*) – The generator to copy.
- **force_copy** (*bool, optional*) – If `True`, this function will always create a copy of every random state. If `False`, it will not copy numpy's default random state, but all other random states.

Returns **rs_copy** – The copied random state.

Return type numpy.random.RandomState

`imgaug.imgaug.`**`current_random_state`**`()`

Deprecated. Use `imgaug.random.get_global_rng` instead.

Get or create the current global RNG of imgaug.

Note that the first call to this function will create a global RNG.

Returns The global RNG to use.

Return type *imgaug.random.RNG*

`class` `imgaug.imgaug.`**`deprecated`**`(alt_func=None, behavior='warn', removed_version=None, comment=None)`

Bases: `object`

Decorator to mark deprecated functions with warning.

Adapted from <https://github.com/scikit-image/scikit-image/blob/master/skimage/_shared/utils.py>.

Parameters

- **alt_func** (*None or str, optional*) – If given, tell user what function to use instead.
- **behavior** (*{'warn', 'raise'}, optional*) – Behavior during call to deprecated function: `warn` means that the user is warned that the function is deprecated; `raise` means that an error is raised.
- **removed_version** (*None or str, optional*) – The package version in which the deprecated function will be removed.
- **comment** (*None or str, optional*) – An optional comment that will be appended to the warning message.

Methods

`__call__`(self, func)	Call self as a function.

imgaug.imgaug.**derive_random_state**(*random_state*)

Deprecated. Use `imgaug.random.derive_generator_` instead.

Derive a child numpy random generator from another one.

Parameters **random_state** (*numpy.random.Generator or numpy.random.RandomState*) – The generator from which to derive a new child generator.

Returns In numpy <=1.16 a `RandomState`, in 1.17+ a `Generator`. In both cases a derived child generator.

Return type numpy.random.Generator or numpy.random.RandomState

imgaug.imgaug.**derive_random_states**(*random_state*, *n=1*)

Deprecated. Use `imgaug.random.derive_generators_` instead.

Derive child numpy random generators from another one.

Parameters

- **random_state** (*numpy.random.Generator or numpy.random.RandomState*) – The generator from which to derive new child generators.
- **n** (*int, optional*) – Number of child generators to derive.

Returns In numpy <=1.16 a `list` of `RandomState` s, in 1.17+ a `list` of `Generator` s. In both cases lists of derived child generators.

Return type list of numpy.random.Generator or list of numpy.random.RandomState

imgaug.imgaug.**do_assert**(*condition*, *message='Assertion failed.'*)

Assert that a `condition` holds or raise an `Exception` otherwise.

This was added because *assert* statements are removed in optimized code. It replaced *assert* statements throughout the library, but that was reverted again for readability and performance reasons.

Parameters

- **condition** (*bool*) – If `False`, an exception is raised.
- **message** (*str, optional*) – Error message.

imgaug.imgaug.**draw_grid**(*images*, *rows=None*, *cols=None*)

Combine multiple images into a single grid-like image.

Calling this function with four images of the same shape and `rows=2`, `cols=2` will combine the four images to a single image array of shape `(2*H, 2*W, C)`, where `H` is the height of any of the images (analogous `W`) and `C` is the number of channels of any image.

Calling this function with four images of the same shape and `rows=4`, `cols=1` is analogous to calling `numpy.vstack()` on the images.

Supported dtypes:

- `uint8`: yes; fully tested
- `uint16`: yes; fully tested
- `uint32`: yes; fully tested
- `uint64`: yes; fully tested
- `int8`: yes; fully tested

- int16: yes; fully tested
- int32: yes; fully tested
- int64: yes; fully tested
- float16: yes; fully tested
- float32: yes; fully tested
- float64: yes; fully tested
- float128: yes; fully tested
- bool: yes; fully tested

Parameters

- **images** (*(N,H,W,3) ndarray or iterable of (H,W,3) array*) – The input images to convert to a grid.
- **rows** (*None or int, optional*) – The number of rows to show in the grid. If None, it will be automatically derived.
- **cols** (*None or int, optional*) – The number of cols to show in the grid. If None, it will be automatically derived.

Returns Image of the generated grid.

Return type (H',W',3) ndarray

imgaug.imgaug.**draw_text**(*img*, *y*, *x*, *text*, *color=(0*, *255*, *0)*, *size=25*)

Draw text on an image.

This uses by default DejaVuSans as its font, which is included in this library.

Supported dtypes:

- uint8: yes; fully tested
- uint16: no
- uint32: no
- uint64: no
- int8: no
- int16: no
- int32: no
- int64: no
- float16: no
- float32: yes; not tested
- float64: no
- float128: no
- bool: no

TODO check if other dtypes could be enabled

Parameters

- **img** (*(H,W,3) ndarray*) – The image array to draw text on. Expected to be of dtype `uint8` or `float32` (expected value range is `[0.0, 255.0]`).
- **y** (*int*) – x-coordinate of the top left corner of the text.
- **x** (*int*) – y- coordinate of the top left corner of the text.
- **text** (*str*) – The text to draw.
- **color** (*iterable of int, optional*) – Color of the text to draw. For RGB-images this is expected to be an RGB color.
- **size** (*int, optional*) – Font size of the text to draw.

Returns Input image with text drawn on it.

Return type (H,W,3) ndarray

`imgaug.imgaug.`**`dummy_random_state`**`()`

Deprecated. Use `imgaug.random.convert_seed_to_rng` instead.

Create a dummy random state using a seed of `1`.

Returns The new random state.

Return type *imgaug.random.RNG*

`imgaug.imgaug.`**`flatten`**`(`*nested_iterable*`)`

Flatten arbitrarily nested lists/tuples.

Code partially taken from https://stackoverflow.com/a/10824420.

Parameters **nested_iterable** – A `list` or `tuple` of arbitrarily nested values.

Yields *any* – All values in *nested_iterable*, flattened.

`imgaug.imgaug.`**`forward_random_state`**`(`*random_state*`)`

Deprecated. Use `imgaug.random.advance_generator_` instead.

Advance a numpy random generator's internal state.

Parameters **random_state** (*numpy.random.Generator or numpy.random.RandomState*) – Generator of which to advance the internal state.

`imgaug.imgaug.`**`imresize_many_images`**`(`*images*, *sizes=None*, *interpolation=None*`)`

Resize each image in a list or array to a specified size.

Supported dtypes:

- `uint8`: yes; fully tested
- `uint16`: yes; tested
- `uint32`: no (1)
- `uint64`: no (2)
- `int8`: yes; tested (3)
- `int16`: yes; tested
- `int32`: limited; tested (4)
- `int64`: no (2)
- `float16`: yes; tested (5)
- `float32`: yes; tested

- `float64`: yes; tested
- `float128`: no (1)
- `bool`: yes; tested (6)
- (1) rejected by `cv2.imresize`
- (2) results too inaccurate
- (3) mapped internally to `int16` when interpolation!="nearest"
- (4) only supported for interpolation="nearest", other interpolations lead to cv2 error
- (5) mapped internally to `float32`
- (6) mapped internally to `uint8`

Parameters

- **images** (*(N,H,W,[C]) ndarray or list of (H,W,[C]) ndarray*) – Array of the images to resize. Usually recommended to be of dtype `uint8`.
- **sizes** (*float or iterable of int or iterable of float*) – The new size of the images, given either as a fraction (a single float) or as a `(height, width) tuple` of two integers or as a `(height fraction, width fraction) tuple` of two floats.
- **interpolation** (*None or str or int, optional*) – The interpolation to use during resize. If `int`, then expected to be one of:
  - `cv2.INTER_NEAREST` (nearest neighbour interpolation)
  - `cv2.INTER_LINEAR` (linear interpolation)
  - `cv2.INTER_AREA` (area interpolation)
  - `cv2.INTER_CUBIC` (cubic interpolation)

  If `str`, then expected to be one of:
  - `nearest` (identical to `cv2.INTER_NEAREST`)
  - `linear` (identical to `cv2.INTER_LINEAR`)
  - `area` (identical to `cv2.INTER_AREA`)
  - `cubic` (identical to `cv2.INTER_CUBIC`)

  If `None`, the interpolation will be chosen automatically. For size increases, `area` interpolation will be picked and for size decreases, `linear` interpolation will be picked.

Returns Array of the resized images.

Return type (N,H',W',[C]) ndarray

Examples

```
>>> import imgaug as ia
>>> images = np.zeros((2, 8, 16, 3), dtype=np.uint8)
>>> images_resized = ia.imresize_many_images(images, 2.0)
>>> images_resized.shape
(2, 16, 32, 3)
```

Convert two RGB images of height `8` and width `16` to images of height `2*8=16` and width `2*16=32`.

```
>>> images_resized = ia.imresize_many_images(images, (2.0, 4.0))
>>> images_resized.shape
(2, 16, 64, 3)
```

Convert two RGB images of height `8` and width `16` to images of height `2*8=16` and width `4*16=64`.

```
>>> images_resized = ia.imresize_many_images(images, (16, 32))
>>> images_resized.shape
(2, 16, 32, 3)
```

Converts two RGB images of height `8` and width `16` to images of height `16` and width `32`.

imgaug.imgaug.**imresize_single_image**(*image*, *sizes*, *interpolation=None*)

Resize a single image.

Supported dtypes:

See *imresize_many_images()*.

Parameters

- **image** (*(H,W,C) ndarray or (H,W) ndarray*) – Array of the image to resize. Usually recommended to be of dtype `uint8`.
- **sizes** (*float or iterable of int or iterable of float*) – See *imresize_many_images()*.
- **interpolation** (*None or str or int, optional*) – See *imresize_many_images()*.

Returns The resized image.

Return type (H',W',C) ndarray or (H',W') ndarray

imgaug.imgaug.**imshow**(*image*, *backend='matplotlib'*)

Show an image in a window.

Supported dtypes:

- `uint8`: yes; not tested
- `uint16`: ?
- `uint32`: ?
- `uint64`: ?
- `int8`: ?
- `int16`: ?
- `int32`: ?
- `int64`: ?
- `float16`: ?
- `float32`: ?
- `float64`: ?
- `float128`: ?
- `bool`: ?

Parameters

- **image** (*(H,W,3) ndarray*) – Image to show.

- **backend** (*{'matplotlib', 'cv2'}, optional*) – Library to use to show the image. May be either matplotlib or OpenCV ('cv2'). OpenCV tends to be faster, but apparently causes more technical issues.

imgaug.imgaug.**is_callable**(*val*)
Check whether a variable is a callable, e.g. a function.

Parameters **val** – The variable to check.

Returns `True` if the variable is a callable. Otherwise `False`.

Return type bool

imgaug.imgaug.**is_float_array**(*val*)
Check whether a variable is a numpy float array.

Parameters **val** – The variable to check.

Returns `True` if the variable is a numpy float array. Otherwise `False`.

Return type bool

imgaug.imgaug.**is_generator**(*val*)
Check whether a variable is a generator.

Parameters **val** – The variable to check.

Returns `True` is the variable is a generator. Otherwise `False`.

Return type bool

imgaug.imgaug.**is_integer_array**(*val*)
Check whether a variable is a numpy integer array.

Parameters **val** – The variable to check.

Returns `True` if the variable is a numpy integer array. Otherwise `False`.

Return type bool

imgaug.imgaug.**is_iterable**(*val*)
Checks whether a variable is iterable.

Parameters **val** – The variable to check.

Returns `True` if the variable is an iterable. Otherwise `False`.

Return type bool

imgaug.imgaug.**is_np_array**(*val*)
Check whether a variable is a numpy array.

Parameters **val** – The variable to check.

Returns `True` if the variable is a numpy array. Otherwise `False`.

Return type bool

imgaug.imgaug.**is_np_scalar**(*val*)
Check whether a variable is a numpy scalar.

Parameters **val** – The variable to check.

Returns `True` if the variable is a numpy scalar. Otherwise `False`.

Return type bool

imgaug.imgaug.**is_single_bool**(*val*)
Check whether a variable is a `bool`.

Parameters **val** – The variable to check.

Returns `True` if the variable is a `bool`. Otherwise `False`.

Return type bool

imgaug.imgaug.**is_single_float**(*val*)
Check whether a variable is a `float`.

Parameters **val** – The variable to check.

Returns `True` if the variable is a `float`. Otherwise `False`.

Return type bool

imgaug.imgaug.**is_single_integer**(*val*)
Check whether a variable is an `int`.

Parameters **val** – The variable to check.

Returns `True` if the variable is an `int`. Otherwise `False`.

Return type bool

imgaug.imgaug.**is_single_number**(*val*)
Check whether a variable is a `number`, i.e. an `int` or `float`.

Parameters **val** – The variable to check.

Returns `True` if the variable is a `number`. Otherwise `False`.

Return type bool

imgaug.imgaug.**is_string**(*val*)
Check whether a variable is a string.

Parameters **val** – The variable to check.

Returns `True` if the variable is a string. Otherwise `False`.

Return type bool

imgaug.imgaug.**max_pool**(*arr*, *block_size*, *pad_mode='edge'*, *pad_cval=0*, *preserve_dtype=True*, *cval=None*)
Resize an array using max-pooling.

Defaults to `pad_mode="edge"` to ensure that padded values do not affect the maximum, even if the dtype was something else than `uint8`.

Supported dtypes:

See *pool()*.

Parameters

- **arr** (*(H,W) ndarray or (H,W,C) ndarray*) – Image-like array to pool. See *pool()* for details.
- **block_size** (*int or tuple of int or tuple of int*) – Size of each block of values to pool. See *pool()* for details.
- **pad_mode** (*str, optional*) – Padding mode to use if the array cannot be divided by *block_size* without remainder. See *pad()* for details.
- **pad_cval** (*number, optional*) – Padding value. See *pool()* for details.

- **preserve_dtype** (*bool, optional*) – Whether to preserve the input array dtype. See *pool()* for details.
- **cval** (*None or number, optional*) – Deprecated. Old name for *pad_cval*.

Returns Array after max-pooling.

Return type (H',W') ndarray or (H',W',C') ndarray

imgaug.imgaug.**median_pool**(*arr*, *block_size*, *pad_mode='reflect'*, *pad_cval=128*, *preserve_dtype=True*)

Resize an array using median-pooling.

Defaults to `pad_mode="reflect"` to ensure that padded values do not affect the average.

Supported dtypes:

See *pool()*.

Parameters

- **arr** (*(H,W) ndarray or (H,W,C) ndarray*) – Image-like array to pool. See *pool()* for details.
- **block_size** (*int or tuple of int or tuple of int*) – Size of each block of values to pool. See *pool()* for details.
- **pad_mode** (*str, optional*) – Padding mode to use if the array cannot be divided by *block_size* without remainder. See *pad()* for details.
- **pad_cval** (*number, optional*) – Padding value. See *pool()* for details.
- **preserve_dtype** (*bool, optional*) – Whether to preserve the input array dtype. See *pool()* for details.

Returns Array after min-pooling.

Return type (H',W') ndarray or (H',W',C') ndarray

imgaug.imgaug.**min_pool**(*arr*, *block_size*, *pad_mode='edge'*, *pad_cval=255*, *preserve_dtype=True*)

Resize an array using min-pooling.

Defaults to `pad_mode="edge"` to ensure that padded values do not affect the minimum, even if the dtype was something else than `uint8`.

Supported dtypes:

See *pool()*.

Parameters

- **arr** (*(H,W) ndarray or (H,W,C) ndarray*) – Image-like array to pool. See *pool()* for details.
- **block_size** (*int or tuple of int or tuple of int*) – Size of each block of values to pool. See *pool()* for details.
- **pad_mode** (*str, optional*) – Padding mode to use if the array cannot be divided by *block_size* without remainder. See *pad()* for details.
- **pad_cval** (*number, optional*) – Padding value. See *pool()* for details.
- **preserve_dtype** (*bool, optional*) – Whether to preserve the input array dtype. See *pool()* for details.

Returns Array after min-pooling.

Return type (H',W') ndarray or (H',W',C') ndarray

imgaug.imgaug.**new_random_state**(*seed=None*, *fully_random=False*)

Deprecated. Use `imgaug.random.convert_seed_to_rng` instead.

Create a new numpy random number generator.

Parameters

- **seed** (*None or int, optional*) – The seed value to use. If `None` and *fully_random* is `False`, the seed will be derived from the global RNG. If *fully_random* is `True`, the seed will be provided by the OS.
- **fully_random** (*bool, optional*) – Whether the seed will be provided by the OS.

Returns In numpy <=1.16 a `RandomState`, in 1.17+ a `Generator`. Both are initialized with the provided seed.

Return type numpy.random.Generator or numpy.random.RandomState

imgaug.imgaug.**normalize_random_state**(*random_state*)

Deprecated. Use `imgaug.random.normalize_generator` instead.

Normalize various inputs to a numpy random generator.

Parameters **random_state** (*None or int or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.bit_generator.SeedSequence or numpy.random.RandomState*) – See *normalize_generator()*.

Returns In numpy <=1.16 a `RandomState`, in 1.17+ a `Generator` (even if the input was a `RandomState`).

Return type numpy.random.Generator or numpy.random.RandomState

imgaug.imgaug.**pad**(**args*, ***kwargs*)

imgaug.imgaug.**pad_to_aspect_ratio**(**args*, ***kwargs*)

imgaug.imgaug.**pad_to_multiples_of**(**args*, ***kwargs*)

imgaug.imgaug.**pool**(*arr*, *block_size*, *func*, *pad_mode='constant'*, *pad_cval=0*, *preserve_dtype=True*, *cval=None*)

Resize an array by pooling values within blocks.

Supported dtypes:

- `uint8`: yes; fully tested
- `uint16`: yes; tested
- `uint32`: yes; tested (2)
- `uint64`: no (1)
- `int8`: yes; tested
- `int16`: yes; tested
- `int32`: yes; tested (2)
- `int64`: no (1)
- `float16`: yes; tested
- `float32`: yes; tested

- `float64`: yes; tested
- `float128`: yes; tested (2)
- `bool`: yes; tested
- (1) results too inaccurate (at least when using np.average as func)
- (2) Note that scikit-image documentation says that the wrapped pooling function converts inputs to `float64`. Actual tests showed no indication of that happening (at least when using preserve_dtype=True).

Parameters

- **arr** (*(H,W) ndarray or (H,W,C) ndarray*) – Image-like array to pool. Ideally of datatype `float64`.
- **block_size** (*int or tuple of int*) –

  Spatial size of each group of values to pool, aka kernel size.

  - If a single `int`, then a symmetric block of that size along height and width will be used.
  - If a `tuple` of two values, it is assumed to be the block size along height and width of the image-like, with pooling happening per channel.
  - If a `tuple` of three values, it is assumed to be the block size along height, width and channels.
- **func** (*callable*) – Function to apply to a given block in order to convert it to a single number, e.g. `numpy.average()`, `numpy.min()`, `numpy.max()`.
- **pad_mode** (*str, optional*) – Padding mode to use if the array cannot be divided by *block_size* without remainder. See *pad()* for details.
- **pad_cval** (*number, optional*) – Value to use for padding if *mode* is `constant`. See `numpy.pad()` for details.
- **preserve_dtype** (*bool, optional*) – Whether to convert the array back to the input datatype if it is changed away from that in the pooling process.
- **cval** (*None or number, optional*) – Deprecated. Old name for *pad_cval*.

Returns Array after pooling.

Return type (H',W') ndarray or (H',W',C') ndarray

`imgaug.imgaug.`**`quokka`**(*size=None*, *extract=None*)

Return an image of a quokka as a numpy array.

Parameters

- **size** (*None or float or tuple of int, optional*) – Size of the output image. Input into *imresize_single_image()*. Usually expected to be a `tuple (H, W)`, where `H` is the desired height and `W` is the width. If `None`, then the image will not be resized.
- **extract** (*None or 'square' or tuple of number or imgaug.augmentables.bbs.BoundingBox or imgaug.augmentables.bbs.BoundingBoxesOnImage*) –

  Subarea of the quokka image to extract:

  - If `None`, then the whole image will be used.
  - If `str square`, then a squared area `(x: 0 to max 643, y: 0 to max 643)` will be extracted from the image.

- If a `tuple`, then expected to contain four `number` s denoting `(x1, y1, x2, y2)`.
- If a *`BoundingBox`*, then that bounding box's area will be extracted from the image.
- If a *`BoundingBoxesOnImage`*, then expected to contain exactly one bounding box and a shape matching the full image dimensions (i.e. `(643, 960, *)`). Then the one bounding box will be used similar to `BoundingBox` above.

Returns The image array of dtype `uint8`.

Return type (H,W,3) ndarray

imgaug.imgaug.**quokka_bounding_boxes**(*size=None*, *extract=None*)

Return example bounding boxes on the standard example quokke image.

Currently only a single bounding box is returned that covers the quokka.

Parameters

- **size** (*None or float or tuple of int or tuple of float, optional*) – Size of the output image on which the BBs are placed. If `None`, then the BBs are not projected to any new size (positions on the original image are used). `float` s lead to relative size changes, `int` s to absolute sizes in pixels.
- **extract** (*None or 'square' or tuple of number or imgaug.augmentables.bbs.BoundingBox or imgaug.augmentables.bbs.BoundingBoxesOnImage*) – Subarea to extract from the image. See *`quokka()`*.

Returns Example BBs on the quokka image.

Return type *imgaug.augmentables.bbs.BoundingBoxesOnImage*

imgaug.imgaug.**quokka_heatmap**(*size=None*, *extract=None*)

Return a heatmap (here: depth map) for the standard example quokka image.

Parameters

- **size** (*None or float or tuple of int, optional*) – See *`quokka()`*.
- **extract** (*None or 'square' or tuple of number or imgaug.augmentables.bbs.BoundingBox or imgaug.augmentables.bbs.BoundingBoxesOnImage*) – See *`quokka()`*.

Returns Depth map as an heatmap object. Values close to `0.0` denote objects that are close to the camera. Values close to `1.0` denote objects that are furthest away (among all shown objects).

Return type *imgaug.augmentables.heatmaps.HeatmapsOnImage*

imgaug.imgaug.**quokka_keypoints**(*size=None*, *extract=None*)

Return example keypoints on the standard example quokke image.

The keypoints cover the eyes, ears, nose and paws.

Parameters

- **size** (*None or float or tuple of int or tuple of float, optional*) – Size of the output image on which the keypoints are placed. If `None`, then the keypoints are not projected to any new size (positions on the original image are used). `float` s lead to relative size changes, `int` s to absolute sizes in pixels.
- **extract** (*None or 'square' or tuple of number or imgaug.augmentables.bbs.BoundingBox or imgaug.augmentables.bbs.BoundingBoxesOnImage*) – Subarea to extract from the image. See *`quokka()`*.

Returns Example keypoints on the quokka image.

Return type *imgaug.augmentables.kps.KeypointsOnImage*

imgaug.imgaug.**quokka_polygons**(*size=None*, *extract=None*)
Returns example polygons on the standard example quokke image.

The result contains one polygon, covering the quokka's outline.

Parameters

- **size** (*None or float or tuple of int or tuple of float, optional*) – Size of the output image on which the polygons are placed. If `None`, then the polygons are not projected to any new size (positions on the original image are used). `float` s lead to relative size changes, `int` s to absolute sizes in pixels.
- **extract** (*None or 'square' or tuple of number or imgaug.augmentables.bbs.BoundingBox or imgaug.augmentables.bbs.BoundingBoxesOnImage*) – Subarea to extract from the image. See *quokka()*.

Returns Example polygons on the quokka image.

Return type *imgaug.augmentables.polys.PolygonsOnImage*

imgaug.imgaug.**quokka_segmentation_map**(*size=None*, *extract=None*)
Return a segmentation map for the standard example quokka image.

Parameters

- **size** (*None or float or tuple of int, optional*) – See *quokka()*.
- **extract** (*None or 'square' or tuple of number or imgaug.augmentables.bbs.BoundingBox or imgaug.augmentables.bbs.BoundingBoxesOnImage*) – See *quokka()*.

Returns Segmentation map object.

Return type *imgaug.augmentables.segmaps.SegmentationMapsOnImage*

imgaug.imgaug.**quokka_square**(*size=None*)
Return an (square) image of a quokka as a numpy array.

Parameters **size** (*None or float or tuple of int, optional*) – Size of the output image. Input into *imresize_single_image()*. Usually expected to be a `tuple (H, W)`, where `H` is the desired height and `W` is the width. If `None`, then the image will not be resized.

Returns The image array of dtype `uint8`.

Return type (H,W,3) ndarray

imgaug.imgaug.**seed**(*entropy=None*, *seedval=None*)
Set the seed of imgaug's global RNG.

The global RNG controls most of the "randomness" in imgaug.

The global RNG is the default one used by all augmenters. Under special circumstances (e.g. when an augmenter is switched to deterministic mode), the global RNG is replaced with a local one. The state of that replacement may be dependent on the global RNG's state at the time of creating the child RNG.

Note: This function is not yet marked as deprecated, but might be in the future. The preferred way to seed *imgaug* is via *seed()*.

Parameters

- **entropy** (*int*) – The seed value to use.
- **seedval** (*None or int, optional*) – Deprecated since 0.4.0.

imgaug.imgaug.**show_grid**(*images*, *rows=None*, *cols=None*)
Combine multiple images into a single image and plot the result.

This will show a window of the results of *draw_grid()*.

Supported dtypes:

minimum of (*draw_grid()*, *imshow()*

)

Parameters

- **images** (*(N,H,W,3) ndarray or iterable of (H,W,3) array*) – See *draw_grid()*.
- **rows** (*None or int, optional*) – See *draw_grid()*.
- **cols** (*None or int, optional*) – See *draw_grid()*.

imgaug.imgaug.**warn**(*msg*, *category=<class 'UserWarning'>*, *stacklevel=2*)
Generate a a warning with stacktrace.

Parameters

- **msg** (*str*) – The message of the warning.
- **category** (*class*) – The class of the warning to produce.
- **stacklevel** (*int, optional*) – How many steps above this function to "jump" in the stacktrace when displaying file and line number of the error message. Usually `2`.

imgaug.imgaug.**warn_deprecated**(*msg*, *stacklevel=2*)
Generate a non-silent deprecation warning with stacktrace.

The used warning is `imgaug.imgaug.DeprecationWarning`.

Parameters

- **msg** (*str*) – The message of the warning.
- **stacklevel** (*int, optional*) – How many steps above this function to "jump" in the stacktrace when displaying file and line number of the error message. Usually `2`

13.2 imgaug.parameters

Classes and methods to use for parameters of augmenters.

This module contains e.g. classes representing probability distributions (guassian, poisson etc.), classes representing noise sources and methods to normalize parameter-related user inputs.

class imgaug.parameters.**Absolute**(*other_param*)
Bases: *imgaug.parameters.StochasticParameter*

Convert the samples of another parameter to their absolute values.

Parameters **other_param** (*imgaug.parameters.StochasticParameter*) – Other parameter which's sampled values are to be modified.

Examples

```
>>> import imgaug.parameters as iap
>>> param = iap.Absolute(iap.Uniform(-1.0, 1.0))
```

Convert a uniform distribution from `[-1.0, 1.0)` to `[0.0, 1.0]`.

Methods

`copy`(self)	Create a shallow copy of this parameter.
`deepcopy`(self)	Create a deep copy of this parameter.
`draw_distribution_graph`(self[, title, size, ...])	Generate an image visualizing the parameter's sample distribution.
`draw_sample`(self[, random_state])	Draws a single sample value from this parameter.
`draw_samples`(self, size[, random_state])	Draw one or more samples from the parameter.

class imgaug.parameters.**Add**(*other_param*, *val*, *elementwise=False*)

Bases: *imgaug.parameters.StochasticParameter*

Add to the samples of another stochastic parameter.

Parameters

- **other_param** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – Samples of *val* will be added to samples of this parameter. Let `S` be the requested shape of samples, then the datatype behaviour is as follows:
  - If a single `number`, this `number` will be used as a constant value to fill an array of shape `S`.
  - If a `tuple` of two `number`s `(a, b)`, an array of shape `S` will be filled with uniformly sampled values from the continuous interval `[a, b)`.
  - If a `list` of `number`, an array of shape `S` will be filled with randomly picked values from the `list`.
  - If a *StochasticParameter*, that parameter will be queried once per call to generate an array of shape `S`.

  "per call" denotes a call of `Add.draw_sample()` or `Add.draw_samples()`.
- **val** (*number or tuple of two number or list of number or imgaug.parameters.StochasticParameter*) – Value to add to the samples of *other_param*. Datatype behaviour is analogous to *other_param*, though if `elementwise=False` (the default), only a single sample will be generated per call instead of `S`.
- **elementwise** (*bool, optional*) – Controls the sampling behaviour of *val*. If set to `False`, a single samples will be requested from *val* and used as the constant multiplier. If set to `True`, samples of shape `S` will be requested from *val* and added elementwise with the samples of *other_param*.

Examples

```
>>> import imgaug.parameters as iap
>>> param = iap.Add(Uniform(0.0, 1.0), 1.0)
```

Convert a uniform distribution from `[0.0, 1.0)` to `[1.0, 2.0)`.

Methods

`copy`(self)	Create a shallow copy of this parameter.
`deepcopy`(self)	Create a deep copy of this parameter.
`draw_distribution_graph`(self[, title, size, ...])	Generate an image visualizing the parameter's sample distribution.
`draw_sample`(self[, random_state])	Draws a single sample value from this parameter.
`draw_samples`(self, size[, random_state])	Draw one or more samples from the parameter.

class `imgaug.parameters.`**Beta**(*alpha*, *beta*, *epsilon=0.0001*)

Bases: *imgaug.parameters.StochasticParameter*

Parameter that resembles a (continuous) beta distribution.

Parameters

- **alpha** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – alpha parameter of the beta distribution. Expected value range is `(0, inf)`. Values below `0` are automatically clipped to `0+epsilon`.
  - If a single `number`, this `number` will be used as a constant value.
  - If a `tuple` of two `number`s `(a, b)`, the value will be sampled from the continuous interval `[a, b)` once per call.
  - If a `list` of `number`, a random value will be picked from the `list` once per call.
  - If a *StochasticParameter*, that parameter will be queried once per call.

  "per call" denotes a call of `Beta.draw_sample()` or `Beta.draw_samples()`.
- **beta** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – Beta parameter of the beta distribution. Analogous to *alpha*.
- **epsilon** (*number*) – Clipping parameter. If *alpha* or *beta* end up `<=0`, they are clipped to `0+epsilon`.

Examples

```
>>> import imgaug.parameters as iap
>>> param = iap.Beta(0.4, 0.6)
```

Create a beta distribution with `alpha=0.4` and `beta=0.6`.

Methods

`copy`(self)	Create a shallow copy of this parameter.
`deepcopy`(self)	Create a deep copy of this parameter.
`draw_distribution_graph`(self[, title, size, ...])	Generate an image visualizing the parameter's sample distribution.
`draw_sample`(self[, random_state])	Draws a single sample value from this parameter.

Continued on next page

Table 7 – continued from previous page

draw_samples(self, size[, random_state])	Draw one or more samples from the parameter.

class imgaug.parameters.**Binomial**(*p*)

Bases: *imgaug.parameters.StochasticParameter*

Binomial distribution.

Parameters p (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – Probability of the binomial distribution. Expected to be in the interval `[0.0, 1.0]`.

- If a single `number`, this `number` will be used as a constant value.
- If a `tuple` of two `number` s `(a, b)`, the value will be sampled from the continuous interval `[a, b)` once per call.
- If a `list` of `number`, a random value will be picked from the `list` once per call.
- If a *StochasticParameter*, that parameter will be queried once per call.

"per call" denotes a call of `Binomial.draw_sample()` or `Binomial.draw_samples()`.

Examples

```
>>> import imgaug.parameters as iap
>>> param = iap.Binomial(Uniform(0.01, 0.2))
```

Create a binomial distribution that uses a varying probability between `0.01` and `0.2`, randomly and uniformly estimated once per sampling call.

Methods

copy(self)	Create a shallow copy of this parameter.
deepcopy(self)	Create a deep copy of this parameter.
draw_distribution_graph(self[, title, size, ...])	Generate an image visualizing the parameter's sample distribution.
draw_sample(self[, random_state])	Draws a single sample value from this parameter.
draw_samples(self, size[, random_state])	Draw one or more samples from the parameter.

class imgaug.parameters.**ChiSquare**(*df*)

Bases: *imgaug.parameters.StochasticParameter*

Parameter that resembles a (continuous) chi-square distribution.

This is a wrapper around numpy's `numpy.random.chisquare()`.

Parameters df (*int or tuple of two int or list of int or imgaug.parameters.StochasticParameter*) –

Degrees of freedom. Expected value range is `[1, inf)`.

- If a single `int`, this `int` will be used as a constant value.
- If a `tuple` of two `int` s `(a, b)`, the value will be sampled from the discrete interval `[a..b]` once per call.
- If a `list` of `int`, a random value will be picked from the `list` once per call.

- If a *StochasticParameter*, that parameter will be queried once per call.

"per call" denotes a call of `ChiSquare.draw_sample()` or `ChiSquare.draw_samples()`.

Examples

```
>>> import imgaug.parameters as iap
>>> param = iap.ChiSquare(df=2)
```

Create a chi-square distribution with two degrees of freedom.

Methods

copy(self)	Create a shallow copy of this parameter.
deepcopy(self)	Create a deep copy of this parameter.
draw_distribution_graph(self[, title, size, ...])	Generate an image visualizing the parameter's sample distribution.
draw_sample(self[, random_state])	Draws a single sample value from this parameter.
draw_samples(self, size[, random_state])	Draw one or more samples from the parameter.

class imgaug.parameters.**Choice**(*a*, *replace=True*, *p=None*)

Bases: *imgaug.parameters.StochasticParameter*

Parameter that samples value from a list of allowed values.

Parameters

- **a** (*iterable*) – List of allowed values. Usually expected to be `int` s, `float` s or `str` s. May also contain `StochasticParameter` s. Each `StochasticParameter` that is randomly picked will automatically be replaced by a sample of itself (or by `N` samples if the parameter was picked `N` times).
- **replace** (*bool, optional*) – Whether to perform sampling with or without replacing.
- **p** (*None or iterable of number, optional*) – Probabilities of each element in *a*. Must have the same length as *a* (if provided).

Examples

```
>>> import imgaug.parameters as iap
>>> param = iap.Choice([5, 17, 25], p=[0.25, 0.5, 0.25])
>>> sample = param.draw_sample()
>>> assert sample in [5, 17, 25]
```

Create and sample from a parameter, which will produce with `50%` probability the sample `17` and in the other `50%` of all cases the sample `5` or `25`..

Methods

copy(self)	Create a shallow copy of this parameter.

Continued on next page

Table 10 – continued from previous page

deepcopy(self)	Create a deep copy of this parameter.
draw_distribution_graph(self[, title, size, ...])	Generate an image visualizing the parameter's sample distribution.
draw_sample(self[, random_state])	Draws a single sample value from this parameter.
draw_samples(self, size[, random_state])	Draw one or more samples from the parameter.

class imgaug.parameters.**Clip**(*other_param*, *minval=None*, *maxval=None*)

Bases: *imgaug.parameters.StochasticParameter*

Clip another parameter to a defined value range.

Parameters

- **other_param** (*imgaug.parameters.StochasticParameter*) – The other parameter, which's values are to be clipped.
- **minval** (*None or number, optional*) – The minimum value to use. If None, no minimum will be used.
- **maxval** (*None or number, optional*) – The maximum value to use. If None, no maximum will be used.

Examples

```
>>> import imgaug.parameters as iap
>>> param = iap.Clip(Normal(0, 1.0), minval=-2.0, maxval=2.0)
```

Create a standard gaussian distribution, which's values never go below -2.0 or above 2.0. Note that this will lead to small "bumps" of higher probability at -2.0 and 2.0, as values below/above these will be clipped to them. For smoother limitations on gaussian distributions, see *TruncatedNormal*.

Methods

copy(self)	Create a shallow copy of this parameter.
deepcopy(self)	Create a deep copy of this parameter.
draw_distribution_graph(self[, title, size, ...])	Generate an image visualizing the parameter's sample distribution.
draw_sample(self[, random_state])	Draws a single sample value from this parameter.
draw_samples(self, size[, random_state])	Draw one or more samples from the parameter.

class imgaug.parameters.**Deterministic**(*value*)

Bases: *imgaug.parameters.StochasticParameter*

Parameter that is a constant value.

If N values are sampled from this parameter, it will return N times V, where V is the constant value.

Parameters **value** (*number or str or imgaug.parameters.StochasticParameter*) – A constant value to use. A string may be provided to generate arrays of strings. If this is a StochasticParameter, a single value will be sampled from it exactly once and then used as the constant value.

Examples

```
>>> import imgaug.parameters as iap
>>> param = iap.Deterministic(10)
>>> param.draw_sample()
10
```

Will always sample the value 10.

Methods

copy(self)	Create a shallow copy of this parameter.
deepcopy(self)	Create a deep copy of this parameter.
draw_distribution_graph(self[, title, size, ...])	Generate an image visualizing the parameter's sample distribution.
draw_sample(self[, random_state])	Draws a single sample value from this parameter.
draw_samples(self, size[, random_state])	Draw one or more samples from the parameter.

class imgaug.parameters.**DeterministicList**(*values*)

Bases: *imgaug.parameters.StochasticParameter*

Parameter that repeats elements from a list in the given order.

E.g. of samples of shape `(A, B, C)` are requested, this parameter will return the first `A*B*C` elements, reshaped to `(A, B, C)` from the provided list. If the list contains less than `A*B*C` elements, it will (by default) be tiled until it is long enough (i.e. the sampling will start again at the first element, if necessary multiple times).

Added in 0.4.0.

Parameters **values** (*ndarray or iterable of number*) – An iterable of values to sample from in the order within the iterable.

Methods

copy(self)	Create a shallow copy of this parameter.
deepcopy(self)	Create a deep copy of this parameter.
draw_distribution_graph(self[, title, size, ...])	Generate an image visualizing the parameter's sample distribution.
draw_sample(self[, random_state])	Draws a single sample value from this parameter.
draw_samples(self, size[, random_state])	Draw one or more samples from the parameter.

class imgaug.parameters.**DiscreteUniform**(*a*, *b*)

Bases: *imgaug.parameters.StochasticParameter*

Uniform distribution over the discrete interval `[a..b]`.

Parameters

- **a** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter*) – Lower bound of the interval. If `a>b`, *a* and *b* will automatically be flipped. If `a==b`, all generated values will be identical to *a*.
  - If a single `int`, this `int` will be used as a constant value.

- If a `tuple` of two `int` s `(a, b)`, the value will be sampled from the discrete interval `[a..b]` once per call.
- If a `list` of `int`, a random value will be picked from the `list` once per call.
- If a *`StochasticParameter`*, that parameter will be queried once per call.

"per call" denotes a call of `DiscreteUniform.draw_sample()` or `DiscreteUniform.draw_samples()`.

- **b** (*int or imgaug.parameters.StochasticParameter*) – Upper bound of the interval. Analogous to *a*.

Examples

```
>>> import imgaug.parameters as iap
>>> param = iap.DiscreteUniform(10, Choice([20, 30, 40]))
>>> sample = param.draw_sample()
>>> assert 10 <= sample <= 40
```

Create a discrete uniform distribution which's interval differs between calls and can be `[10..20]`, `[10..30]` or `[10..40]`.

Methods

`copy`(self)	Create a shallow copy of this parameter.
`deepcopy`(self)	Create a deep copy of this parameter.
`draw_distribution_graph`(self[, title, size, ...])	Generate an image visualizing the parameter's sample distribution.
`draw_sample`(self[, random_state])	Draws a single sample value from this parameter.
`draw_samples`(self, size[, random_state])	Draw one or more samples from the parameter.

class imgaug.parameters.**Discretize**(*other_param*, *round=True*)

Bases: *imgaug.parameters.StochasticParameter*

Convert a continuous distribution to a discrete one.

This will round the values and then cast them to integers. Values sampled from already discrete distributions are not changed.

Parameters

- **other_param** (*imgaug.parameters.StochasticParameter*) – The other parameter, which's values are to be discretized.
- **round** (*bool, optional*) – Whether to round before converting to integer dtype.

  Added in 0.4.0.

Examples

```
>>> import imgaug.parameters as iap
>>> param = iap.Discretize(iap.Normal(0, 1.0))
```

Create a discrete standard gaussian distribution.

Methods

copy(self)	Create a shallow copy of this parameter.
deepcopy(self)	Create a deep copy of this parameter.
draw_distribution_graph(self[, title, size, ...])	Generate an image visualizing the parameter's sample distribution.
draw_sample(self[, random_state])	Draws a single sample value from this parameter.
draw_samples(self, size[, random_state])	Draw one or more samples from the parameter.

class imgaug.parameters.**Divide**(*other_param*, *val*, *elementwise=False*)

Bases: *imgaug.parameters.StochasticParameter*

Divide the samples of another stochastic parameter.

This parameter will automatically prevent division by zero (uses 1.0) as the denominator in these cases.

Parameters

- **other_param** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – Other parameter which's sampled values are to be divided by *val*. Let `S` be the requested shape of samples, then the datatype behaviour is as follows:
  - If a single `number`, this `number` will be used as a constant value to fill an array of shape `S`.
  - If a `tuple` of two `number` s `(a, b)`, an array of shape `S` will be filled with uniformly sampled values from the continuous interval `[a, b)`.
  - If a `list` of `number`, an array of shape `S` will be filled with randomly picked values from the `list`.
  - If a *StochasticParameter*, that parameter will be queried once per call to generate an array of shape `S`.

  "per call" denotes a call of `Divide.draw_sample()` or `Divide.draw_samples()`.
- **val** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – Denominator to use. Datatype behaviour is analogous to *other_param*, though if `elementwise=False` (the default), only a single sample will be generated per call instead of `S`.
- **elementwise** (*bool, optional*) – Controls the sampling behaviour of *val*. If set to `False`, a single samples will be requested from *val* and used as the constant denominator. If set to `True`, samples of shape `S` will be requested from *val* and used to divide the samples of *other_param* elementwise.

Examples

```
>>> import imgaug.parameters as iap
>>> param = iap.Divide(iap.Uniform(0.0, 1.0), 2)
```

Convert a uniform distribution `[0.0, 1.0)` to `[0, 0.5)`.

Methods

copy(self)	Create a shallow copy of this parameter.
deepcopy(self)	Create a deep copy of this parameter.
draw_distribution_graph(self[, title, size, ...])	Generate an image visualizing the parameter's sample distribution.
draw_sample(self[, random_state])	Draws a single sample value from this parameter.
draw_samples(self, size[, random_state])	Draw one or more samples from the parameter.

class imgaug.parameters.**ForceSign**(*other_param*, *positive*, *mode='invert'*, *reroll_count_max=2*)

Bases: *imgaug.parameters.StochasticParameter*

Convert a parameter's samples to either positive or negative values.

Parameters

- **other_param** (*imgaug.parameters.StochasticParameter*) – Other parameter which's sampled values are to be modified.
- **positive** (*bool*) – Whether to force all signs to be positive (`True`) or negative (`False`).
- **mode** (*{'invert', 'reroll'}, optional*) – Method to change the signs. Valid values are `invert` and `reroll`. `invert` means that wrong signs are simply flipped. `reroll` means that all samples with wrong signs are sampled again, optionally many times, until they randomly end up having the correct sign.
- **reroll_count_max** (*int, optional*) – If *mode* is set to `reroll`, this determines how often values may be rerolled before giving up and simply flipping the sign (as in `mode="invert"`). This shouldn't be set too high, as rerolling is expensive.

Examples

```
>>> import imgaug.parameters as iap
>>> param = iap.ForceSign(iap.Poisson(1), positive=False)
```

Create a poisson distribution with `alpha=1` that is flipped towards negative values.

Methods

copy(self)	Create a shallow copy of this parameter.
deepcopy(self)	Create a deep copy of this parameter.
draw_distribution_graph(self[, title, size, ...])	Generate an image visualizing the parameter's sample distribution.
draw_sample(self[, random_state])	Draws a single sample value from this parameter.
draw_samples(self, size[, random_state])	Draw one or more samples from the parameter.

class imgaug.parameters.**FrequencyNoise**(*exponent=(-4, 4), size_px_max=(4, 32), upscale_method=['linear', 'nearest']*)

Bases: *imgaug.parameters.StochasticParameter*

Parameter to generate noise of varying frequencies.

This parameter expects to sample noise for 2d planes, i.e. for sizes `(H, W, [C])` and will return a value in

the range `[0.0, 1.0]` per spatial location in that plane.

The exponent controls the frequencies and therefore noise patterns. Small values (around `-4.0`) will result in large blobs. Large values (around `4.0`) will result in small, repetitive patterns.

The noise is sampled from low resolution planes and upscaled to the requested height and width. The size of the low resolution plane may be defined (high values can be slow) and the interpolation method for upscaling can be set.

Parameters

- **exponent** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Exponent to use when scaling in the frequency domain. Sane values are in the range `-4` (large blobs) to `4` (small patterns). To generate cloud-like structures, use roughly `-2`.
  - If a single `number`, this `number` will be used as a constant value.
  - If a `tuple` of two `number`s `(a, b)`, the value will be sampled from the continuous interval `[a, b)` once per call.
  - If a `list` of `number`, a random value will be picked from the `list` once per call.
  - If a *`StochasticParameter`*, that parameter will be queried once per call.
- **size_px_max** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter, optional*) – Maximum height and width in pixels of the low resolution plane. Upon any sampling call, the requested shape will be downscaled until the height or width (whichever is larger) does not exceed this maximum value anymore. Then the noise will be sampled at that shape and later upscaled back to the requested shape.
  - If a single `int`, this `int` will be used as a constant value.
  - If a `tuple` of two `int`s `(a, b)`, the value will be sampled from the discrete interval `[a..b]` once per call.
  - If a `list` of `int`, a random value will be picked from the `list` once per call.
  - If a *`StochasticParameter`*, that parameter will be queried once per call.

  "per call" denotes a call of `FrequencyNoise.draw_sample()` or `FrequencyNoise.draw_samples()`.
- **upscale_method** (*imgaug.ALL or str or list of str or imgaug.parameters.StochasticParameter, optional*) – After generating the noise maps in low resolution environments, they have to be upscaled to the originally requested shape (i.e. usually the image size). This parameter controls the interpolation method to use. See also *`imresize_many_images()`* for a description of possible values.
  - If `imgaug.ALL`, then either `nearest` or `linear` or `area` or `cubic` is picked per iteration (all same probability).
  - If `str`, then that value will always be used as the method (must be `nearest` or `linear` or `area` or `cubic`).
  - If `list` of `str`, then a random value will be picked from that list per call.
  - If *`StochasticParameter`*, then a random value will be sampled from that parameter per call.

Examples

```
>>> import imgaug.parameters as iap
>>> param = iap.FrequencyNoise(
>>>     exponent=-2,
>>>     size_px_max=(16, 32),
>>>     upscale_method="linear")
```

Create a parameter that produces noise with cloud-like patterns.

Methods

copy(self)	Create a shallow copy of this parameter.
deepcopy(self)	Create a deep copy of this parameter.
draw_distribution_graph(self[, title, size, ...])	Generate an image visualizing the parameter's sample distribution.
draw_sample(self[, random_state])	Draws a single sample value from this parameter.
draw_samples(self, size[, random_state])	Draw one or more samples from the parameter.

class imgaug.parameters.**FromLowerResolution**(*other_param*, *size_percent=None*, *size_px=None*, *method='nearest'*, *min_size=1*)

Bases: *imgaug.parameters.StochasticParameter*

Parameter to sample from other parameters at lower image resolutions.

This parameter is intended to be used with parameters that would usually sample one value per pixel (or one value per pixel and channel). Instead of sampling from the other parameter at full resolution, it samples at lower resolution, e.g. `0.5*H x 0.5*W` with `H` being the height and `W` being the width. After the low-resolution sampling this parameter then upscales the result to `HxW`.

This parameter is intended to produce coarse samples. E.g. combining this with *Binomial* can lead to large rectangular areas of `1` s and `0` s.

Parameters

- **other_param** (*imgaug.parameters.StochasticParameter*) – The other parameter which is to be sampled on a coarser image.
- **size_percent** (*None or number or iterable of number or imgaug.parameters.StochasticParameter, optional*) – Size of the 2d sampling plane in percent of the requested size. I.e. this is relative to the size provided in the call to `draw_samples(size)`. Lower values will result in smaller sampling planes, which are then upsampled to *size*. This means that lower values will result in larger rectangles. The size may be provided as a constant value or a tuple `(a, b)`, which will automatically be converted to the continuous uniform range `[a, b)` or a *StochasticParameter*, which will be queried per call to `FromLowerResolution.draw_sample()` and `FromLowerResolution.draw_samples()`.
- **size_px** (*None or number or iterable of numbers or imgaug.parameters.StochasticParameter, optional*) – Size of the 2d sampling plane in pixels. Lower values will result in smaller sampling planes, which are then upsampled to the input *size* of `draw_samples(size)`. This means that lower values will result in larger rectangles. The size may be provided as a constant value or a tuple `(a, b)`, which will automatically be converted to the discrete uniform range `[a..b]` or a *StochasticParameter*,

which will be queried once per call to `FromLowerResolution.draw_sample()` and `FromLowerResolution.draw_samples()`.

- **method** (*str or int or imgaug.parameters.StochasticParameter, optional*) – Upsampling/interpolation method to use. This is used after the sampling is finished and the low resolution plane has to be upsampled to the requested *size* in `draw_samples(size, ...)`. The method may be the same as in *imresize_many_images()*. Usually `nearest` or `linear` are good choices. `nearest` will result in rectangles with sharp edges and `linear` in rectangles with blurry and round edges. The method may be provided as a *StochasticParameter*, which will be queried once per call to `FromLowerResolution.draw_sample()` and `FromLowerResolution.draw_samples()`.
- **min_size** (*int, optional*) – Minimum size in pixels of the low resolution sampling plane.

Examples

```
>>> import imgaug.parameters as iap
>>> param = iap.FromLowerResolution(
>>>     Binomial(0.05),
>>>     size_px=(2, 16),
>>>     method=Choice(["nearest", "linear"]))
```

Samples from a binomial distribution with `p=0.05`. The sampling plane will always have a size HxWxC with H and W being independently sampled from `[2..16]` (i.e. it may range from `2x2xC` up to `16x16xC` max, but may also be e.g. `4x8xC`). The upsampling method will be `nearest` in `50%` of all cases and `linear` in the other 50 percent. The result will sometimes be rectangular patches of sharp `1` s surrounded by `0` s and sometimes blurry blobs of `1``s, surrounded by values ``<1.0`.

Methods

`copy`(self)	Create a shallow copy of this parameter.
`deepcopy`(self)	Create a deep copy of this parameter.
`draw_distribution_graph`(self[, title, size, ...])	Generate an image visualizing the parameter's sample distribution.
`draw_sample`(self[, random_state])	Draws a single sample value from this parameter.
`draw_samples`(self, size[, random_state])	Draw one or more samples from the parameter.

class imgaug.parameters.**IterativeNoiseAggregator**(*other_param, iterations=(1, 3), aggregation_method=['max', 'avg']*)

Bases: *imgaug.parameters.StochasticParameter*

Aggregate multiple iterations of samples from another parameter.

This is supposed to be used in conjunction with *SimplexNoise* or *FrequencyNoise*. If a shape `S` is requested, it will request `I` times `S` samples from the underlying parameter, where `I` is the number of iterations. The `I` arrays will be combined to a single array of shape `S` using an aggregation method, e.g. simple averaging.

Parameters

- **other_param** (*StochasticParameter*) – The other parameter from which to sample one or more times.
- **iterations** (*int or iterable of int or list of int or imgaug.parameters.StochasticParameter, optional*) –

The number of iterations.

- If a single `int`, this `int` will be used as a constant value.
- If a `tuple` of two `int` s `(a, b)`, the value will be sampled from the discrete interval `[a..b]` once per call.
- If a `list` of `int`, a random value will be picked from the `list` once per call.
- If a *StochasticParameter*, that parameter will be queried once per call.

"per call" denotes a call of `IterativeNoiseAggregator.draw_sample()` or `IterativeNoiseAggregator.draw_samples()`.

- **aggregation_method** (*imgaug.ALL or {'min', 'avg', 'max'} or list of str or imgaug.parameters.StochasticParameter, optional*) – The method to use to aggregate the samples of multiple iterations to a single output array. All methods combine several arrays of shape `S` each to a single array of shape `S` and hence work elementwise. Known methods are `min` (take the minimum over all iterations), `max` (take the maximum) and `avg` (take the average).
  - If an `str`, it must be one of the described methods and will be used for all calls..
  - If a `list` of `str`, it must contain one or more of the described methods and a random one will be samples once per call.
  - If `imgaug.ALL`, then equivalent to the `list ["min", "max", "avg"]`.
  - If *StochasticParameter*, a value will be sampled from that parameter once per call and must be one of the described methods..

  "per call" denotes a call of `IterativeNoiseAggregator.draw_sample()` or `IterativeNoiseAggregator.draw_samples()`.

Examples

```
>>> import imgaug.parameters as iap
>>> noise = iap.IterativeNoiseAggregator(
>>>     iap.SimplexNoise(),
>>>     iterations=(2, 5),
>>>     aggregation_method="max")
```

Create a parameter that – upon each call – generates `2` to `5` arrays of simplex noise with the same shape. Then it combines these noise maps to a single map using elementwise maximum.

Methods

`copy`(self)	Create a shallow copy of this parameter.
`deepcopy`(self)	Create a deep copy of this parameter.
`draw_distribution_graph`(self[, title, size, ...])	Generate an image visualizing the parameter's sample distribution.
`draw_sample`(self[, random_state])	Draws a single sample value from this parameter.
`draw_samples`(self, size[, random_state])	Draw one or more samples from the parameter.

class `imgaug.parameters.`**Laplace**(*loc*, *scale*)

Bases: *imgaug.parameters.StochasticParameter*

Parameter that resembles a (continuous) laplace distribution.

This is a wrapper around numpy's `numpy.random.laplace()`.

Parameters

- **loc** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – The position of the distribution peak, similar to the mean in normal distributions.
  - If a single `number`, this `number` will be used as a constant value.
  - If a `tuple` of two `number`s `(a, b)`, the value will be sampled from the continuous interval `[a, b)` once per call.
  - If a `list` of `number`, a random value will be picked from the `list` once per call.
  - If a *StochasticParameter*, that parameter will be queried once per call.

  "per call" denotes a call of `Laplace.draw_sample()` or `Laplace.draw_samples()`.
- **scale** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – The exponential decay factor, similar to the standard deviation in gaussian distributions. If this parameter reaches `0`, the output array will be filled with *loc*. Datatype behaviour is the analogous to *loc*.

Examples

```
>>> import imgaug.parameters as iap
>>> param = iap.Laplace(0, 1.0)
```

Create a laplace distribution, which's peak is at `0` and decay is `1.0`.

Methods

`copy`(self)	Create a shallow copy of this parameter.
`deepcopy`(self)	Create a deep copy of this parameter.
`draw_distribution_graph`(self[, title, size, ...])	Generate an image visualizing the parameter's sample distribution.
`draw_sample`(self[, random_state])	Draws a single sample value from this parameter.
`draw_samples`(self, size[, random_state])	Draw one or more samples from the parameter.

class imgaug.parameters.**Multiply**(*other_param*, *val*, *elementwise=False*)

Bases: *imgaug.parameters.StochasticParameter*

Multiply the samples of another stochastic parameter.

Parameters

- **other_param** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – Other parameter which's sampled values are to be multiplied with *val*. Let `S` be the requested shape of samples, then the datatype behaviour is as follows:
  - If a single `number`, this `number` will be used as a constant value to fill an array of shape `S`.

  - If a `tuple` of two `number`s `(a, b)`, an array of shape `S` will be filled with uniformly sampled values from the continuous interval `[a, b)`.
  - If a `list` of `number`, an array of shape `S` will be filled with randomly picked values from the `list`.
  - If a *StochasticParameter*, that parameter will be queried once per call to generate an array of shape `S`.

  "per call" denotes a call of `Multiply.draw_sample()` or `Multiply.draw_samples()`.

- **val** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – Multiplier to use. Datatype behaviour is analogous to *other_param*, though if `elementwise=False` (the default), only a single sample will be generated per call instead of `S`.
- **elementwise** (*bool, optional*) – Controls the sampling behaviour of *val*. If set to `False`, a single samples will be requested from *val* and used as the constant multiplier. If set to `True`, samples of shape `S` will be requested from *val* and multiplied elementwise with the samples of *other_param*.

Examples

```
>>> import imgaug.parameters as iap
>>> param = iap.Multiply(iap.Uniform(0.0, 1.0), -1)
```

Convert a uniform distribution from `[0.0, 1.0)` to `(-1.0, 0.0]`.

Methods

`copy`(self)	Create a shallow copy of this parameter.
`deepcopy`(self)	Create a deep copy of this parameter.
`draw_distribution_graph`(self[, title, size, ...])	Generate an image visualizing the parameter's sample distribution.
`draw_sample`(self[, random_state])	Draws a single sample value from this parameter.
`draw_samples`(self, size[, random_state])	Draw one or more samples from the parameter.

`imgaug.parameters.`**`Negative`**(*other_param*, *mode='invert'*, *reroll_count_max=2*)

Convert another parameter's results to negative values.

Parameters

- **other_param** (*imgaug.parameters.StochasticParameter*) – Other parameter which's sampled values are to be modified.
- **mode** (*{'invert', 'reroll'}, optional*) – How to change the signs. Valid values are `invert` and `reroll`. `invert` means that wrong signs are simply flipped. `reroll` means that all samples with wrong signs are sampled again, optionally many times, until they randomly end up having the correct sign.
- **reroll_count_max** (*int, optional*) – If *mode* is set to `reroll`, this determines how often values may be rerolled before giving up and simply flipping the sign (as in `mode="invert"`). This shouldn't be set too high, as rerolling is expensive.

Examples

```
>>> import imgaug.parameters as iap
>>> param = iap.Negative(iap.Normal(0, 1), mode="reroll")
```

Create a gaussian distribution that has only negative values. If any positive value is sampled in the process, that sample is resampled up to two times to get a negative one. If it isn't negative after the second resampling step, the sign is simply flipped.

class imgaug.parameters.**Normal**(*loc*, *scale*)
Bases: *imgaug.parameters.StochasticParameter*

Parameter that resembles a normal/gaussian distribution.

Parameters

- **loc** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) –

  The mean of the normal distribution.

  - If a single `number`, this `number` will be used as a constant value.
  - If a `tuple` of two `number` s `(a, b)`, the value will be sampled from the continuous interval `[a, b)` once per call.
  - If a `list` of `number`, a random value will be picked from the `list` once per call.
  - If a *StochasticParameter*, that parameter will be queried once per call.

  "per call" denotes a call of `Laplace.draw_sample()` or `Laplace.draw_samples()`.

- **scale** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – The standard deviation of the normal distribution. If this parameter reaches `0`, the output array will be filled with *loc*. Datatype behaviour is the analogous to *loc*.

Examples

```
>>> import imgaug.parameters as iap
>>> param = iap.Normal(Choice([-1.0, 1.0]), 1.0)
```

Create a gaussian distribution with a mean that differs by call. Samples values may sometimes follow `N(-1.0, 1.0)` and sometimes `N(1.0, 1.0)`.

Methods

copy(self)	Create a shallow copy of this parameter.
deepcopy(self)	Create a deep copy of this parameter.
draw_distribution_graph(self[, title, size, ...])	Generate an image visualizing the parameter's sample distribution.
draw_sample(self[, random_state])	Draws a single sample value from this parameter.
draw_samples(self, size[, random_state])	Draw one or more samples from the parameter.

class imgaug.parameters.**Poisson**(*lam*)
Bases: *imgaug.parameters.StochasticParameter*

Parameter that resembles a poisson distribution.

A poisson distribution with `lambda=0` has its highest probability at point `0` and decreases quickly from there. Poisson distributions are discrete and never negative.

Parameters **lam** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) –

Lambda parameter of the poisson distribution.

- If a single `number`, this `number` will be used as a constant value.
- If a `tuple` of two `number` s `(a, b)`, the value will be sampled from the continuous interval `[a, b)` once per call.
- If a `list` of `number`, a random value will be picked from the `list` once per call.
- If a *StochasticParameter*, that parameter will be queried once per call.

"per call" denotes a call of `Poisson.draw_sample()` or `Poisson.draw_samples()`.

Examples

```
>>> import imgaug.parameters as iap
>>> param = iap.Poisson(1)
>>> sample = param.draw_sample()
>>> assert sample >= 0
```

Create a poisson distribution with `lambda=1` and sample a value from it.

Methods

copy(self)	Create a shallow copy of this parameter.
deepcopy(self)	Create a deep copy of this parameter.
draw_distribution_graph(self[, title, size, ...])	Generate an image visualizing the parameter's sample distribution.
draw_sample(self[, random_state])	Draws a single sample value from this parameter.
draw_samples(self, size[, random_state])	Draw one or more samples from the parameter.

imgaug.parameters.**Positive**(*other_param*, *mode='invert'*, *reroll_count_max=2*)
Convert another parameter's results to positive values.

Parameters

- **other_param** (*imgaug.parameters.StochasticParameter*) – Other parameter which's sampled values are to be modified.
- **mode** (*{'invert', 'reroll'}, optional*) – How to change the signs. Valid values are `invert` and `reroll`. `invert` means that wrong signs are simply flipped. `reroll` means that all samples with wrong signs are sampled again, optionally many times, until they randomly end up having the correct sign.
- **reroll_count_max** (*int, optional*) – If *mode* is set to `reroll`, this determines how often values may be rerolled before giving up and simply flipping the sign (as in `mode="invert"`). This shouldn't be set too high, as rerolling is expensive.

Examples

```
>>> import imgaug.parameters as iap
>>> param = iap.Positive(iap.Normal(0, 1), mode="reroll")
```

Create a gaussian distribution that has only positive values. If any negative value is sampled in the process, that sample is resampled up to two times to get a positive one. If it isn't positive after the second resampling step, the sign is simply flipped.

class imgaug.parameters.**Power**(*other_param*, *val*, *elementwise=False*)
Bases: *imgaug.parameters.StochasticParameter*

Exponentiate the samples of another stochastic parameter.

Parameters

- **other_param** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – Other parameter which's sampled values are to be exponentiated by *val*. Let `S` be the requested shape of samples, then the datatype behaviour is as follows:
  - If a single `number`, this `number` will be used as a constant value to fill an array of shape `S`.
  - If a `tuple` of two `number` s `(a, b)`, an array of shape `S` will be filled with uniformly sampled values from the continuous interval `[a, b)`.
  - If a `list` of `number`, an array of shape `S` will be filled with randomly picked values from the `list`.
  - If a *StochasticParameter*, that parameter will be queried once per call to generate an array of shape `S`.

  "per call" denotes a call of `Power.draw_sample()` or `Power.draw_samples()`.
- **val** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – Value to use exponentiate the samples of *other_param*. Datatype behaviour is analogous to *other_param*, though if `elementwise=False` (the default), only a single sample will be generated per call instead of `S`.
- **elementwise** (*bool, optional*) – Controls the sampling behaviour of *val*. If set to `False`, a single samples will be requested from *val* and used as the constant multiplier. If set to `True`, samples of shape `S` will be requested from *val* and used to exponentiate elementwise the samples of *other_param*.

Examples

```
>>> import imgaug.parameters as iap
>>> param = iap.Power(iap.Uniform(0.0, 1.0), 2)
```

Converts a uniform range `[0.0, 1.0)` to a distribution that is peaked towards 1.0.

Methods

copy(self)	Create a shallow copy of this parameter.

Continued on next page

Table 25 – continued from previous page

deepcopy(self)	Create a deep copy of this parameter.
draw_distribution_graph(self[, title, size, ...])	Generate an image visualizing the parameter's sample distribution.
draw_sample(self[, random_state])	Draws a single sample value from this parameter.
draw_samples(self, size[, random_state])	Draw one or more samples from the parameter.

class imgaug.parameters.**RandomSign**(*other_param*, *p_positive=0.5*)

Bases: *imgaug.parameters.StochasticParameter*

Convert a parameter's samples randomly to positive or negative values.

Parameters

- **other_param** (*imgaug.parameters.StochasticParameter*) – Other parameter which's sampled values are to be modified.
- **p_positive** (*number*) – Fraction of values that are supposed to be turned to positive values.

Examples

```
>>> import imgaug.parameters as iap
>>> param = iap.RandomSign(iap.Poisson(1))
```

Create a poisson distribution with `alpha=1` that is mirrored/copied (not flipped) at the y-axis.

Methods

copy(self)	Create a shallow copy of this parameter.
deepcopy(self)	Create a deep copy of this parameter.
draw_distribution_graph(self[, title, size, ...])	Generate an image visualizing the parameter's sample distribution.
draw_sample(self[, random_state])	Draws a single sample value from this parameter.
draw_samples(self, size[, random_state])	Draw one or more samples from the parameter.

class imgaug.parameters.**Sigmoid**(*other_param*, *threshold=(-10, 10)*, *activated=True*, *mul=1*, *add=0*)

Bases: *imgaug.parameters.StochasticParameter*

Apply a sigmoid function to the outputs of another parameter.

This is intended to be used in combination with *SimplexNoise* or *FrequencyNoise*. It pushes the noise values away from ~`0.5` and towards `0.0` or `1.0`, making the noise maps more binary.

Parameters

- **other_param** (*imgaug.parameters.StochasticParameter*) – The other parameter to which the sigmoid will be applied.
- **threshold** (*number or tuple of number or iterable of number or imgaug.parameters.StochasticParameter, optional*) – Sets the value of the sigmoid's saddle point, i.e. where values start to quickly shift from `0.0` to `1.0`.
  - If a single `number`, this `number` will be used as a constant value.
  - If a `tuple` of two `number`s `(a, b)`, the value will be sampled from the continuous interval `[a, b)` once per call.

  - If a `list` of `number`, a random value will be picked from the `list` once per call.
  - If a *`StochasticParameter`*, that parameter will be queried once per call.

  "per call" denotes a call of `Sigmoid.draw_sample()` or `Sigmoid.draw_samples()`.

- **activated** (*bool or number, optional*) – Defines whether the sigmoid is activated. If this is `False`, the results of *other_param* will not be altered. This may be set to a `float p` in value range''[0.0, 1.0]'', which will result in *activated* being `True` in `p` percent of all calls.
- **mul** (*number, optional*) – The results of *other_param* will be multiplied with this value before applying the sigmoid. For noise values (range `[0.0, 1.0]`) this should be set to about `20`.
- **add** (*number, optional*) – This value will be added to the results of *other_param* before applying the sigmoid. For noise values (range `[0.0, 1.0]`) this should be set to about `-10.0`, provided *mul* was set to `20`.

Examples

```
>>> import imgaug.parameters as iap
>>> param = iap.Sigmoid(
>>>     iap.SimplexNoise(),
>>>     activated=0.5,
>>>     mul=20,
>>>     add=-10)
```

Applies a sigmoid to simplex noise in `50%` of all calls. The noise results are modified to match the sigmoid's expected value range. The sigmoid's outputs are in the range `[0.0, 1.0]`.

Methods

`copy`(self)	Create a shallow copy of this parameter.
`create_for_noise`(other_param[, threshold, ...])	Create a Sigmoid adjusted for noise parameters.
`deepcopy`(self)	Create a deep copy of this parameter.
`draw_distribution_graph`(self[, title, size, ...])	Generate an image visualizing the parameter's sample distribution.
`draw_sample`(self[, random_state])	Draws a single sample value from this parameter.
`draw_samples`(self, size[, random_state])	Draw one or more samples from the parameter.

static create_for_noise (*other_param*, *threshold=(-10, 10)*, *activated=True*)

Create a Sigmoid adjusted for noise parameters.

"noise" here denotes *`SimplexNoise`* and *`FrequencyNoise`*.

Parameters

- **other_param** (*imgaug.parameters.StochasticParameter*) – See `__init__()`.
- **threshold** (*number or tuple of number or iterable of number or imgaug.parameters.StochasticParameter, optional*) – See `__init__()`.
- **activated** (*bool or number, optional*) – See `__init__()`.

Returns A sigmoid adjusted to be used with noise.

Return type *Sigmoid*

class imgaug.parameters.**SimplexNoise**(*size_px_max=(2, 16), upscale_method=['linear', 'nearest']*)

Bases: *imgaug.parameters.StochasticParameter*

Parameter that generates simplex noise of varying resolutions.

This parameter expects to sample noise for 2d planes, i.e. for sizes `(H, W, [C])` and will return a value in the range `[0.0, 1.0]` per spatial location in that plane.

The noise is sampled from low resolution planes and upscaled to the requested height and width. The size of the low resolution plane may be defined (large values can be slow) and the interpolation method for upscaling can be set.

Parameters

- **size_px_max** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter, optional*) – Maximum height and width in pixels of the low resolution plane. Upon any sampling call, the requested shape will be downscaled until the height or width (whichever is larger) does not exceed this maximum value anymore. Then the noise will be sampled at that shape and later upscaled back to the requested shape.
  - If a single `int`, this `int` will be used as a constant value.
  - If a `tuple` of two `int` s `(a, b)`, the value will be sampled from the discrete interval `[a..b]` once per call.
  - If a `list` of `int`, a random value will be picked from the `list` once per call.
  - If a *StochasticParameter*, that parameter will be queried once per call.

  "per call" denotes a call of `SimplexNoise.draw_sample()` or `SimplexNoise.draw_samples()`.
- **upscale_method** (*str or int or list of str or list of int or imgaug.parameters.StochasticParameter, optional*) – After generating the noise maps in low resolution environments, they have to be upscaled to the originally requested shape (i.e. usually the image size). This parameter controls the interpolation method to use. See also *imresize_many_images()* for a description of possible values.
  - If `imgaug.ALL`, then either `nearest` or `linear` or `area` or `cubic` is picked per iteration (all same probability).
  - If `str`, then that value will always be used as the method (must be `nearest` or `linear` or `area` or `cubic`).
  - If `list` of `str`, then a random value will be picked from that list per call.
  - If *StochasticParameter*, then a random value will be sampled from that parameter per call.

Examples

```
>>> import imgaug.parameters as iap
>>> param = iap.SimplexNoise(upscale_method="linear")
```

Create a parameter that produces smooth simplex noise of varying sizes.

```
>>> param = iap.SimplexNoise(
>>>     size_px_max=(8, 16),
>>>     upscale_method="nearest")
```

Create a parameter that produces rectangular simplex noise of rather high detail.

Methods

copy(self)	Create a shallow copy of this parameter.
deepcopy(self)	Create a deep copy of this parameter.
draw_distribution_graph(self[, title, size, ...])	Generate an image visualizing the parameter's sample distribution.
draw_sample(self[, random_state])	Draws a single sample value from this parameter.
draw_samples(self, size[, random_state])	Draw one or more samples from the parameter.

class imgaug.parameters.**StochasticParameter**

Bases: object

Abstract parent class for all stochastic parameters.

Stochastic parameters are here all parameters from which values are supposed to be sampled. Usually the sampled values are to a degree random. E.g. a stochastic parameter may be the uniform distribution over the interval `[-10, 10]`. Samples from that distribution (and therefore the stochastic parameter) could be `5.2`, `-3.7`, `-9.7`, `6.4`, etc.

Methods

copy(self)	Create a shallow copy of this parameter.
deepcopy(self)	Create a deep copy of this parameter.
draw_distribution_graph(self[, title, size, ...])	Generate an image visualizing the parameter's sample distribution.
draw_sample(self[, random_state])	Draws a single sample value from this parameter.
draw_samples(self, size[, random_state])	Draw one or more samples from the parameter.

copy(*self*)

Create a shallow copy of this parameter.

Returns Shallow copy.

Return type *imgaug.parameters.StochasticParameter*

deepcopy(*self*)

Create a deep copy of this parameter.

Returns Deep copy.

Return type *imgaug.parameters.StochasticParameter*

draw_distribution_graph(*self*, *title=None*, *size=(1000, 1000)*, *bins=100*)

Generate an image visualizing the parameter's sample distribution.

Parameters

- **title** (*None or False or str, optional*) – Title of the plot. `None` is automatically replaced by a title derived from `str(param)`. If set to `False`, no title will be shown.

- **size** (*tuple of int*) – Number of points to sample. This is always expected to have at least two values. The first defines the number of sampling runs, the second (and further) dimensions define the size assigned to each `draw_samples()` call. E.g. `(10, 20, 15)` will lead to `10` calls of `draw_samples(size=(20, 15))`. The results will be merged to a single 1d array.
- **bins** (*int*) – Number of bins in the plot histograms.

Returns **data** – Image of the plot.

Return type (H,W,3) ndarray

draw_sample(*self*, *random_state=None*)
Draws a single sample value from this parameter.

Parameters **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – A seed or random number generator to use during the sampling process. If `None`, the global RNG will be used. See also `__init__()` for a similar parameter with more details.

Returns A single sample value.

Return type any

draw_samples(*self*, *size*, *random_state=None*)
Draw one or more samples from the parameter.

Parameters

- **size** (*tuple of int or int*) – Number of samples by dimension.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – A seed or random number generator to use during the sampling process. If `None`, the global RNG will be used. See also `__init__()` for a similar parameter with more details.

Returns Sampled values. Usually a numpy ndarray of basically any dtype, though not strictly limited to numpy arrays. Its shape is expected to match *size*.

Return type ndarray

class imgaug.parameters.**Subtract**(*other_param*, *val*, *elementwise=False*)
Bases: *imgaug.parameters.StochasticParameter*

Subtract from the samples of another stochastic parameter.

Parameters

- **other_param** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – Samples of *val* will be subtracted from samples of this parameter. Let `S` be the requested shape of samples, then the datatype behaviour is as follows:
  - If a single `number`, this `number` will be used as a constant value to fill an array of shape `S`.
  - If a `tuple` of two `number` s `(a, b)`, an array of shape `S` will be filled with uniformly sampled values from the continuous interval `[a, b)`.
  - If a `list` of `number`, an array of shape `S` will be filled with randomly picked values from the `list`.

– If a *StochasticParameter*, that parameter will be queried once per call to generate an array of shape `S`.

"per call" denotes a call of `Subtract.draw_sample()` or `Subtract.draw_samples()`.

- **val** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – Value to subtract from the other parameter. Datatype behaviour is analogous to *other_param*, though if `elementwise=False` (the default), only a single sample will be generated per call instead of `S`.
- **elementwise** (*bool, optional*) – Controls the sampling behaviour of *val*. If set to `False`, a single samples will be requested from *val* and used as the constant multiplier. If set to `True`, samples of shape `S` will be requested from *val* and subtracted elementwise from the samples of *other_param*.

Examples

```
>>> import imgaug.parameters as iap
>>> param = iap.Subtract(iap.Uniform(0.0, 1.0), 1.0)
```

Convert a uniform distribution from `[0.0, 1.0)` to `[-1.0, 0.0)`.

Methods

`copy`(self)	Create a shallow copy of this parameter.
`deepcopy`(self)	Create a deep copy of this parameter.
`draw_distribution_graph`(self[, title, size, ...])	Generate an image visualizing the parameter's sample distribution.
`draw_sample`(self[, random_state])	Draws a single sample value from this parameter.
`draw_samples`(self, size[, random_state])	Draw one or more samples from the parameter.

class imgaug.parameters.**TruncatedNormal**(*loc*, *scale*, *low=-inf*, *high=inf*)

Bases: *imgaug.parameters.StochasticParameter*

Parameter that resembles a truncated normal distribution.

A truncated normal distribution is similar to a normal distribution, except the domain is smoothly bounded to a min and max value.

This is a wrapper around `scipy.stats.truncnorm()`.

Parameters

- **loc** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) –

  The mean of the normal distribution.

  – If a single `number`, this `number` will be used as a constant value.

  – If a `tuple` of two `number` s `(a, b)`, the value will be sampled from the continuous interval `[a, b)` once per call.

  – If a `list` of `number`, a random value will be picked from the `list` once per call.

  – If a *StochasticParameter*, that parameter will be queried once per call.

"per call" denotes a call of `TruncatedNormal.draw_sample()` or `TruncatedNormal.draw_samples()`.

- **scale** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – The standard deviation of the normal distribution. If this parameter reaches `0`, the output array will be filled with *loc*. Datatype behaviour is the same as for *loc*.
- **low** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – The minimum value of the truncated normal distribution. Datatype behaviour is the same as for *loc*.
- **high** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – The maximum value of the truncated normal distribution. Datatype behaviour is the same as for *loc*.

Examples

```
>>> import imgaug.parameters as iap
>>> param = iap.TruncatedNormal(0, 5.0, low=-10, high=10)
>>> samples = param.draw_samples(100, random_state=0)
>>> assert np.all(samples >= -10)
>>> assert np.all(samples <= 10)
```

Create a truncated normal distribution with its minimum at `-10.0` and its maximum at `10.0`.

Methods

`copy`(self)	Create a shallow copy of this parameter.
`deepcopy`(self)	Create a deep copy of this parameter.
`draw_distribution_graph`(self[, title, size, ...])	Generate an image visualizing the parameter's sample distribution.
`draw_sample`(self[, random_state])	Draws a single sample value from this parameter.
`draw_samples`(self, size[, random_state])	Draw one or more samples from the parameter.

class imgaug.parameters.**Uniform**(*a*, *b*)

Bases: *imgaug.parameters.StochasticParameter*

Parameter that resembles a uniform distribution over `[a, b)`.

Parameters

- **a** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – Lower bound of the interval. If `a>b`, *a* and *b* will automatically be flipped. If `a==b`, all generated values will be identical to *a*.
  - If a single `number`, this `number` will be used as a constant value.
  - If a `tuple` of two `number` s `(a, b)`, the value will be sampled from the continuous interval `[a, b)` once per call.
  - If a `list` of `number`, a random value will be picked from the `list` once per call.
  - If a *StochasticParameter*, that parameter will be queried once per call.

  "per call" denotes a call of `Uniform.draw_sample()` or `Uniform.draw_samples()`.

- **b** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – Upper bound of the interval. Analogous to *a*.

Examples

```
>>> import imgaug.parameters as iap
>>> param = iap.Uniform(0, 10.0)
>>> sample = param.draw_sample()
>>> assert 0 <= sample < 10.0
```

Create and sample from a uniform distribution over `[0, 10.0)`.

Methods

`copy`(self)	Create a shallow copy of this parameter.
`deepcopy`(self)	Create a deep copy of this parameter.
`draw_distribution_graph`(self[, title, size, ...])	Generate an image visualizing the parameter's sample distribution.
`draw_sample`(self[, random_state])	Draws a single sample value from this parameter.
`draw_samples`(self, size[, random_state])	Draw one or more samples from the parameter.

class imgaug.parameters.**Weibull**(*a*)

Bases: *imgaug.parameters.StochasticParameter*

Parameter that resembles a (continuous) weibull distribution.

This is a wrapper around numpy's `numpy.random.weibull()`.

Parameters **a** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) –

Shape parameter of the distribution.

- If a single `number`, this `number` will be used as a constant value.
- If a `tuple` of two `number` s `(a, b)`, the value will be sampled from the continuous interval `[a, b)` once per call.
- If a `list` of `number`, a random value will be picked from the `list` once per call.
- If a *StochasticParameter*, that parameter will be queried once per call.

"per call" denotes a call of `Weibull.draw_sample()` or `Weibull.draw_samples()`.

Examples

```
>>> import imgaug.parameters as iap
>>> param = iap.Weibull(a=0.5)
```

Create a weibull distribution with shape 0.5.

Methods

copy(self)	Create a shallow copy of this parameter.
deepcopy(self)	Create a deep copy of this parameter.
draw_distribution_graph(self[, title, size, ...])	Generate an image visualizing the parameter's sample distribution.
draw_sample(self[, random_state])	Draws a single sample value from this parameter.
draw_samples(self, size[, random_state])	Draw one or more samples from the parameter.

imgaug.parameters.**both_np_float_if_one_is_float**(*a*, *b*)

imgaug.parameters.**draw_distributions_grid**(*params*, *rows=None*, *cols=None*, *graph_sizes=(350, 350)*, *sample_sizes=None*, *titles=None*)

imgaug.parameters.**force_np_float_dtype**(*val*)

imgaug.parameters.**handle_categorical_string_param**(*param*, *name*, *valid_values=None*)

imgaug.parameters.**handle_continuous_param**(*param*, *name*, *value_range=None*, *tuple_to_uniform=True*, *list_to_choice=True*)

imgaug.parameters.**handle_discrete_kernel_size_param**(*param*, *name*, *value_range=(1, None)*, *allow_floats=True*)

imgaug.parameters.**handle_discrete_param**(*param*, *name*, *value_range=None*, *tuple_to_uniform=True*, *list_to_choice=True*, *allow_floats=True*)

imgaug.parameters.**handle_probability_param**(*param*, *name*, *tuple_to_uniform=False*, *list_to_choice=False*)

imgaug.parameters.**show_distributions_grid**(*params*, *rows=None*, *cols=None*, *graph_sizes=(350, 350)*, *sample_sizes=None*, *titles=None*)

13.3 imgaug.multicore

Classes and functions dealing with augmentation on multiple CPU cores.

class imgaug.multicore.**BackgroundAugmenter**(*batch_loader*, *augseq*, *queue_size=50*, *nb_workers='auto'*)

Bases: object

Deprecated. Augment batches in the background processes.

Deprecated. Use imgaug.multicore.Pool instead.

This is a wrapper around the multiprocessing module.

Parameters

- **batch_loader** (*BatchLoader or multiprocessing.Queue*) – BatchLoader object that loads the data fed into the BackgroundAugmenter, or alternatively a Queue. If a Queue, then it must be made sure that a final None in the Queue signals that the loading is finished and no more batches will follow. Otherwise the BackgroundAugmenter will wait forever for the next batch.
- **augseq** (*Augmenter*) – An augmenter to apply to all loaded images. This may be e.g. a Sequential to apply multiple augmenters.

- **queue_size** (*int*) – Size of the queue that is used to temporarily save the augmentation results. Larger values offer the background processes more room to save results when the main process doesn't load much, i.e. they can lead to smoother and faster training. For large images, high values can block a lot of RAM though.
- **nb_workers** (*'auto' or int*) – Number of background workers to spawn. If `auto`, it will be set to `C-1`, where `C` is the number of CPU cores.

Methods

get_batch(self)	Returns a batch from the queue of augmented batches.
terminate(self)	Terminates all background processes immediately.

all_finished	

all_finished(*self*)

get_batch(*self*)

Returns a batch from the queue of augmented batches.

If workers are still running and there are no batches in the queue, it will automatically wait for the next batch.

Returns **out** – One batch or None if all workers have finished.

Return type None or imgaug.Batch

terminate(*self*)

Terminates all background processes immediately.

This will also free their RAM.

class imgaug.multicore.**BatchLoader**(*load_batch_func*, *queue_size=50*, *nb_workers=1*, *threaded=True*)

Bases: `object`

Deprecated. Load batches in the background.

Deprecated. Use `imgaug.multicore.Pool` instead.

Loaded batches can be accesses using `imgaug.BatchLoader.queue`.

Parameters

- **load_batch_func** (*callable or generator*) – Generator or generator function (i.e. function that yields Batch objects) or a function that returns a list of Batch objects. Background loading automatically stops when the last batch was yielded or the last batch in the list was reached.
- **queue_size** (*int, optional*) – Maximum number of batches to store in the queue. May be set higher for small images and/or small batches.
- **nb_workers** (*int, optional*) – Number of workers to run in the background.
- **threaded** (*bool, optional*) – Whether to run the background processes using threads (True) or full processes (False).

Methods

all_finished(self)	Determine whether the workers have finished the loading process.
terminate(self)	Stop all workers.

count_workers_alive	

all_finished(*self*)
Determine whether the workers have finished the loading process.

Returns **out** – True if all workers have finished. Else False.

Return type bool

count_workers_alive(*self*)

terminate(*self*)
Stop all workers.

class imgaug.multicore.**Pool**(*augseq*, *processes=None*, *maxtasksperchild=None*, *seed=None*)
Bases: object

Wrapper around multiprocessing.Pool for multicore augmentation.

Parameters

- **augseq** (*imgaug.augmenters.meta.Augmenter*) – The augmentation sequence to apply to batches.
- **processes** (*None or int, optional*) – The number of background workers, similar to the same parameter in multiprocessing.Pool. If None, the number of the machine's CPU cores will be used (this counts hyperthreads as CPU cores). If this is set to a negative value p, then P - abs(p) will be used, where P is the number of CPU cores. E.g. -1 would use all cores except one (this is useful to e.g. reserve one core to feed batches to the GPU).
- **maxtasksperchild** (*None or int, optional*) – The number of tasks done per worker process before the process is killed and restarted, similar to the same parameter in multiprocessing.Pool. If None, worker processes will not be automatically restarted.
- **seed** (*None or int, optional*) – The seed to use for child processes. If None, a random seed will be used.

Attributes

pool	Return or create the multiprocessing.Pool instance.

Methods

close(self)	Close the pool gracefully.
imap_batches(self, batches[, chunksize, ...])	Augment batches from a generator.
imap_batches_unordered(self, batches[, ...])	Augment batches from a generator (without preservation of order).
join(self)	Wait for the workers to exit.
map_batches(self, batches[, chunksize])	Augment a list of batches.
map_batches_async(self, batches[, ...])	Augment batches asynchonously.

Continued on next page

Table 36 – continued from previous page

terminate(self)	Terminate the pool immediately.

close (*self*)
Close the pool gracefully.

imap_batches (*self*, *batches*, *chunksize=1*, *output_buffer_size=None*)
Augment batches from a generator.

Pattern for output buffer constraint is from https://stackoverflow.com/a/47058399.

Parameters

- **batches** (*generator of imgaug.augmentables.batches.Batch*) – The batches to augment, provided as a generator. Each call to the generator should yield exactly one batch.
- **chunksize** (*None or int, optional*) – Rough indicator of how many tasks should be sent to each worker. Increasing this number can improve performance.
- **output_buffer_size** (*None or int, optional*) – Max number of batches to handle *at the same time* in the *whole* pipeline (including already augmented batches that are waiting to be requested). If the buffer size is reached, no new batches will be loaded from *batches* until a produced (i.e. augmented) batch is consumed (i.e. requested from this method). The buffer is unlimited if this is set to `None`. For large datasets, this should be set to an integer value to avoid filling the whole RAM if loading+augmentation happens faster than training.

  New in version 0.3.0.

Yields *imgaug.augmentables.batches.Batch* – Augmented batch.

imap_batches_unordered (*self*, *batches*, *chunksize=1*, *output_buffer_size=None*)
Augment batches from a generator (without preservation of order).

Pattern for output buffer constraint is from https://stackoverflow.com/a/47058399.

Parameters

- **batches** (*generator of imgaug.augmentables.batches.Batch*) – The batches to augment, provided as a generator. Each call to the generator should yield exactly one batch.
- **chunksize** (*None or int, optional*) – Rough indicator of how many tasks should be sent to each worker. Increasing this number can improve performance.
- **output_buffer_size** (*None or int, optional*) – Max number of batches to handle *at the same time* in the *whole* pipeline (including already augmented batches that are waiting to be requested). If the buffer size is reached, no new batches will be loaded from *batches* until a produced (i.e. augmented) batch is consumed (i.e. requested from this method). The buffer is unlimited if this is set to `None`. For large datasets, this should be set to an integer value to avoid filling the whole RAM if loading+augmentation happens faster than training.

  New in version 0.3.0.

Yields *imgaug.augmentables.batches.Batch* – Augmented batch.

join (*self*)
Wait for the workers to exit.

This may only be called after first calling *close()* or *terminate()*.

map_batches (*self*, *batches*, *chunksize=None*)
Augment a list of batches.

Parameters

- **batches** (*list of imgaug.augmentables.batches.Batch*) – The batches to augment.
- **chunksize** (*None or int, optional*) – Rough indicator of how many tasks should be sent to each worker. Increasing this number can improve performance.

Returns Augmented batches.

Return type list of imgaug.augmentables.batches.Batch

map_batches_async(*self*, *batches*, *chunksize=None*, *callback=None*, *error_callback=None*)
Augment batches asynchonously.

Parameters

- **batches** (*list of imgaug.augmentables.batches.Batch*) – The batches to augment.
- **chunksize** (*None or int, optional*) – Rough indicator of how many tasks should be sent to each worker. Increasing this number can improve performance.
- **callback** (*None or callable, optional*) – Function to call upon finish. See `multiprocessing.Pool`.
- **error_callback** (*None or callable, optional*) – Function to call upon errors. See `multiprocessing.Pool`.

Returns Asynchonous result. See `multiprocessing.Pool`.

Return type multiprocessing.MapResult

pool
Return or create the `multiprocessing.Pool` instance.

This creates a new instance upon the first call and afterwards returns that instance (until the property `_pool` is set to `None` again).

Returns The `multiprocessing.Pool` used internally by this `imgaug.multicore.Pool`.

Return type multiprocessing.Pool

terminate(*self*)
Terminate the pool immediately.

13.4 imgaug.dtypes

Functions to interact/analyze with numpy dtypes.

imgaug.dtypes.**change_dtype_**(*arr*, *dtype*, *clip=True*, *round=True*)

imgaug.dtypes.**change_dtypes_**(*images*, *dtypes*, *clip=True*, *round=True*)

imgaug.dtypes.**clip_**(*array*, *min_value*, *max_value*)

imgaug.dtypes.**clip_to_dtype_value_range_**(*array*, *dtype*, *validate=True*, *validate_values=None*)

imgaug.dtypes.**copy_dtypes_for_restore**(*images*, *force_list=False*)

imgaug.dtypes.**gate_dtypes**(*dtypes*, *allowed*, *disallowed*, *augmenter=None*)

imgaug.dtypes.**get_minimal_dtype**(*arrays*, *increase_itemsize_factor=1*)

imgaug.dtypes.**get_value_range_of_dtype**(*dtype*)

imgaug.dtypes.**increase_array_resolutions_**(*arrays*, *factor*)

imgaug.dtypes.**increase_itemsize_of_dtype**(*dtype*, *factor*)

imgaug.dtypes.**normalize_dtype**(*dtype*)

imgaug.dtypes.**normalize_dtypes**(*dtypes*)

imgaug.dtypes.**promote_array_dtypes_**(*arrays*, *dtypes=None*, *increase_itemsize_factor=1*)

imgaug.dtypes.**restore_dtypes_**(*images*, *dtypes*, *clip=True*, *round=True*)

13.5 imgaug.random

Classes and functions related to pseudo-random number generation.

This module deals with the generation of pseudo-random numbers. It provides the *RNG* class, which is the primary random number generator in `imgaug`. It also provides various utility functions related random number generation, such as copying random number generators or setting their state.

The main benefit of this module is to hide the actually used random number generation classes and methods behin imgaug-specific classes and methods. This allows to deal with numpy using two different interfaces (one old interface in numpy <=1.16 and a new one in numpy 1.17+). It also allows to potentially switch to a different framework/library in the future.

13.5.1 Definitions

- *numpy generator* or *numpy random number generator*: Usually an instance of `numpy.random.Generator`. Can often also denote an instance of `numpy.random.RandomState` as both have almost the same interface.
- *RandomState*: An instance of *numpy.random.RandomState*. Note that outside of this module, the term "random state" often roughly translates to "any random number generator with numpy-like interface in a given state", i.e. it can then include instances of `numpy.random.Generator` or *RNG*.
- *RNG*: An instance of *RNG*.

Examples

```
>>> import imgaug.random as iarandom
>>> rng = iarandom.RNG(1234)
>>> rng.integers(0, 1000)
```

Initialize a random number generator with seed `1234`, then sample a single integer from the discrete interval `[0, 1000)`. This will use a `numpy.random.Generator` in numpy 1.17+ and automatically fall back to `numpy.random.RandomState` in numpy <=1.16.

class imgaug.random.**RNG**(*generator*)

Bases: `object`

Random number generator for imgaug.

This class is a wrapper around `numpy.random.Generator` and automatically falls back to `numpy.random.RandomState` in case of numpy version 1.16 or lower. It allows to use numpy 1.17's sampling functions in 1.16 too and supports a variety of useful functions on the wrapped sampler, e.g. gettings its state or copying it.

Not supported sampling functions of numpy <=1.16:

- `numpy.random.RandomState.rand()`
- `numpy.random.RandomState.randn()`
- `numpy.random.RandomState.randint()`
- `numpy.random.RandomState.random_integers()`
- `numpy.random.RandomState.random_sample()`
- `numpy.random.RandomState.ranf()`
- `numpy.random.RandomState.sample()`
- `numpy.random.RandomState.seed()`
- `numpy.random.RandomState.get_state()`
- `numpy.random.RandomState.set_state()`

In *choice()*, the *axis* argument is not yet supported.

Parameters generator (*None or int or RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState*) – The numpy random number generator to use. In case of numpy version 1.17 or later, this shouldn't be a `RandomState` as that class is outdated. Behaviour for different datatypes:

- If `None`: The global RNG is wrapped by this RNG (they are then effectively identical, any sampling on this RNG will affect the global RNG).
- If `int`: In numpy 1.17+, the value is used as a seed for a `Generator` wrapped by this RNG. I.e. it will be provided as the entropy to a `SeedSequence`, which will then be used for an `SFC64` bit generator and wrapped by a `Generator`. In numpy <=1.16, the value is used as a seed for a `RandomState`, which is then wrapped by this RNG.
- If *RNG*: That RNG's `generator` attribute will be used as the generator for this RNG, i.e. the same as `RNG(other_rng.generator)`.
- If `numpy.random.Generator`: That generator will be wrapped.
- If `numpy.random.BitGenerator`: A numpy generator will be created (and wrapped by this RNG) that contains the bit generator.
- If `numpy.random.SeedSequence`: A numpy generator will be created (and wrapped by this RNG) that contains an `SFC64` bit generator initialized with the given `SeedSequence`.
- If `numpy.random.RandomState`: In numpy <=1.16, this `RandomState` will be wrapped and used to sample random values. In numpy 1.17+, a seed will be derived from this `RandomState` and a new `numpy.generator.Generator` based on an `SFC64` bit generator will be created and wrapped by this RNG.

Attributes

state	Get the state of this RNG.

Methods

*advance_*(self)	Advance the RNG's internal state in-place by one step.
beta(self, a, b[, size])	Call `numpy.random.Generator.beta()`.

Continued on next page

Table 37 – continued from previous page

binomial(self, n, p[, size])	Call `numpy.random.Generator.binomial()`.
bytes(self, length)	Call `numpy.random.Generator.bytes()`.
chisquare(self, df[, size])	Call `numpy.random.Generator.chisquare()`.
choice(self, a[, size, replace, p])	Call `numpy.random.Generator.choice()`.
copy(self)	Create a copy of this RNG.
copy_unless_global_rng(self)	Create a copy of this RNG unless it is the global RNG.
create_fully_random()	Create a new RNG, based on entropy provided from the OS.
*create_pseudo_random_*()	Create a new RNG in pseudo-random fashion.
*derive_rng_*(self)	Create a child RNG.
*derive_rngs_*(self, n)	Create *n* child RNGs.
dirichlet(self, alpha[, size])	Call `numpy.random.Generator.dirichlet()`.
duplicate(self, n)	Create a list containing *n* times this RNG.
equals(self, other)	Estimate whether this RNG and *other* have the same state.
equals_global_rng(self)	Estimate whether this RNG has the same state as the global RNG.
exponential(self[, scale, size])	Call `numpy.random.Generator.exponential()`.
f(self, dfnum, dfden[, size])	Call `numpy.random.Generator.f()`.
gamma(self, shape[, scale, size])	Call `numpy.random.Generator.gamma()`.
*generate_seed_*(self)	Sample a random seed.
*generate_seeds_*(self, n)	Generate *n* random seed values.
geometric(self, p[, size])	Call `numpy.random.Generator.geometric()`.
gumbel(self[, loc, scale, size])	Call `numpy.random.Generator.gumbel()`.
hypergeometric(self, ngood, nbad, nsample[, ...])	Call `numpy.random.Generator.hypergeometric()`.
integers(self, low[, high, size, dtype, ...])	Call numpy's `integers()` or `randint()`.
is_global_rng(self)	Estimate whether this RNG is identical to the global RNG.
laplace(self[, loc, scale, size])	Call `numpy.random.Generator.laplace()`.
logistic(self[, loc, scale, size])	Call `numpy.random.Generator.logistic()`.
lognormal(self[, mean, sigma, size])	Call `numpy.random.Generator.lognormal()`.
logseries(self, p[, size])	Call `numpy.random.Generator.logseries()`.
multinomial(self, n, pvals[, size])	Call `numpy.random.Generator.multinomial()`.
multivariate_normal(self, mean, cov[, size, ...])	Call `numpy.random.Generator.multivariate_normal()`.
negative_binomial(self, n, p[, size])	Call `numpy.random.Generator.negative_binomial()`.
noncentral_chisquare(self, df, nonc[, size])	Call `numpy.random.Generator.noncentral_chisquare()`.

Continued on next page

Table 37 – continued from previous page

noncentral_f(self, dfnum, dfden, nonc[, size])	Call `numpy.random.Generator.noncentral_f()`.
normal(self[, loc, scale, size])	Call `numpy.random.Generator.normal()`.
pareto(self, a[, size])	Call `numpy.random.Generator.pareto()`.
permutation(self, x)	Call `numpy.random.Generator.permutation()`.
poisson(self[, lam, size])	Call `numpy.random.Generator.poisson()`.
power(self, a[, size])	Call `numpy.random.Generator.power()`.
rand(self, *args)	Call `numpy.random.RandomState.rand()`.
randint(self, low[, high, size, dtype])	Call `numpy.random.RandomState.randint()`.
randn(self, *args)	Call `numpy.random.RandomState.randn()`.
random(self, size[, dtype, out])	Call numpy's `random()` or `random_sample()`.
random_integers(self, low[, high, size])	Call `numpy.random.RandomState.random_integers()`.
random_sample(self, size)	Call `numpy.random.RandomState.random_sample()`.
rayleigh(self[, scale, size])	Call `numpy.random.Generator.rayleigh()`.
*reset_cache_*(self)	Reset all cache of this RNG.
*set_state_*(self, value)	Set the state if the RNG in-place.
shuffle(self, x)	Call `numpy.random.Generator.shuffle()`.
standard_cauchy(self[, size])	Call `numpy.random.Generator.standard_cauchy()`.
standard_exponential(self[, size, dtype, ...])	Call `numpy.random.Generator.standard_exponential()`.
standard_gamma(self, shape[, size, dtype, out])	Call `numpy.random.Generator.standard_gamma()`.
standard_normal(self[, size, dtype, out])	Call `numpy.random.Generator.standard_normal()`.
standard_t(self, df[, size])	Call `numpy.random.Generator.standard_t()`.
tomaxint(self[, size])	Call `numpy.random.RandomState.tomaxint()`.
triangular(self, left, mode, right[, size])	Call `numpy.random.Generator.triangular()`.
uniform(self[, low, high, size])	Call `numpy.random.Generator.uniform()`.
*use_state_of_*(self, other)	Copy and use (in-place) the state of another RNG.
vonmises(self, mu, kappa[, size])	Call `numpy.random.Generator.vonmises()`.
wald(self, mean, scale[, size])	Call `numpy.random.Generator.wald()`.
weibull(self, a[, size])	Call `numpy.random.Generator.weibull()`.
zipf(self, a[, size])	Call `numpy.random.Generator.zipf()`.

advance_(*self*)

Advance the RNG's internal state in-place by one step.

This advances the underlying generator's state.

Note: This simply samples one or more random values. This means that a call of this method will not completely change the outputs of the next called sampling method. To achieve more drastic output changes, call *derive_rng_()*.

Returns The RNG itself.

Return type *RNG*

beta(*self*, *a*, *b*, *size=None*)
Call `numpy.random.Generator.beta()`.

binomial(*self*, *n*, *p*, *size=None*)
Call `numpy.random.Generator.binomial()`.

bytes(*self*, *length*)
Call `numpy.random.Generator.bytes()`.

chisquare(*self*, *df*, *size=None*)
Call `numpy.random.Generator.chisquare()`.

choice(*self*, *a*, *size=None*, *replace=True*, *p=None*)
Call `numpy.random.Generator.choice()`.

copy(*self*)
Create a copy of this RNG.

Returns Copy of this RNG. The copy will produce the same random samples.

Return type *RNG*

copy_unless_global_rng(*self*)
Create a copy of this RNG unless it is the global RNG.

Returns Copy of this RNG unless it is the global RNG. In the latter case the RNG instance itself will be returned without any changes.

Return type *RNG*

classmethod create_fully_random()
Create a new RNG, based on entropy provided from the OS.

Returns A new RNG. It is not derived from any other previously created RNG, nor does it depend on the seeding of imgaug or numpy.

Return type *RNG*

classmethod create_pseudo_random_()
Create a new RNG in pseudo-random fashion.

A seed will be sampled from the current global RNG and used to initialize the new RNG.

This advandes the global RNG's state.

Returns A new RNG, derived from the current global RNG.

Return type *RNG*

derive_rng_(*self*)
Create a child RNG.

This advances the underlying generator's state.

Returns A child RNG.

Return type *RNG*

derive_rngs_(*self*, *n*)
Create *n* child RNGs.

This advances the underlying generator's state.

Parameters **n** (*int*) – Number of child RNGs to derive.

Returns Child RNGs.

Return type list of RNG

dirichlet(*self*, *alpha*, *size=None*)
Call `numpy.random.Generator.dirichlet()`.

duplicate(*self*, *n*)
Create a list containing *n* times this RNG.

This method was mainly introduced as a replacement for previous calls of *derive_rngs_()*. These calls turned out to be very slow in numpy 1.17+ and were hence replaced by simple duplication (except for the cases where child RNGs absolutely *had* to be created). This RNG duplication method doesn't help very much against code repetition, but it does *mark* the points where it would be desirable to create child RNGs for various reasons. Once deriving child RNGs is somehow sped up in the future, these calls can again be easily found and replaced.

Parameters **n** (*int*) – Length of the output list.

Returns List containing *n* times this RNG (same instances, no copies).

Return type list of RNG

equals(*self*, *other*)
Estimate whether this RNG and *other* have the same state.

Returns `True` if this RNG's generator and the generator of *other* have equal internal states. `False` otherwise.

Return type bool

equals_global_rng(*self*)
Estimate whether this RNG has the same state as the global RNG.

Returns `True` is this RNG has the same state as the global RNG, i.e. it will lead to the same sampled values given the same sampling method calls. The RNGs *don't* have to be identical object instances, which protects against e.g. copy effects. `False` otherwise.

Return type bool

exponential(*self*, *scale=1.0*, *size=None*)
Call `numpy.random.Generator.exponential()`.

f(*self*, *dfnum*, *dfden*, *size=None*)
Call `numpy.random.Generator.f()`.

gamma(*self*, *shape*, *scale=1.0*, *size=None*)
Call `numpy.random.Generator.gamma()`.

generate_seed_(*self*)
Sample a random seed.

This advances the underlying generator's state.

See `SEED_MIN_VALUE` and `SEED_MAX_VALUE` for the seed's value range.

Returns The sampled seed.

Return type int

generate_seeds_(*self*, *n*)
Generate *n* random seed values.

This advances the underlying generator's state.

See `SEED_MIN_VALUE` and `SEED_MAX_VALUE` for the seed's value range.

Parameters **n** (*int*) – Number of seeds to sample.

Returns 1D-array of `int32` seeds.

Return type ndarray

geometric(*self*, *p*, *size=None*)
Call `numpy.random.Generator.geometric()`.

gumbel(*self*, *loc=0.0*, *scale=1.0*, *size=None*)
Call `numpy.random.Generator.gumbel()`.

hypergeometric(*self*, *ngood*, *nbad*, *nsample*, *size=None*)
Call `numpy.random.Generator.hypergeometric()`.

integers(*self*, *low*, *high=None*, *size=None*, *dtype='int32'*, *endpoint=False*)
Call numpy's `integers()` or `randint()`.

Note: Changed *dtype* argument default value from numpy's `int64` to `int32`.

---

is_global_rng(*self*)
Estimate whether this RNG is identical to the global RNG.

Returns `True` is this RNG's underlying generator is identical to the global RNG's underlying generator. The RNGs themselves may be different, only the wrapped generator matters. `False` otherwise.

Return type bool

laplace(*self*, *loc=0.0*, *scale=1.0*, *size=None*)
Call `numpy.random.Generator.laplace()`.

logistic(*self*, *loc=0.0*, *scale=1.0*, *size=None*)
Call `numpy.random.Generator.logistic()`.

lognormal(*self*, *mean=0.0*, *sigma=1.0*, *size=None*)
Call `numpy.random.Generator.lognormal()`.

logseries(*self*, *p*, *size=None*)
Call `numpy.random.Generator.logseries()`.

multinomial(*self*, *n*, *pvals*, *size=None*)
Call `numpy.random.Generator.multinomial()`.

multivariate_normal(*self*, *mean*, *cov*, *size=None*, *check_valid='warn'*, *tol=1e-08*)
Call `numpy.random.Generator.multivariate_normal()`.

negative_binomial(*self*, *n*, *p*, *size=None*)
Call `numpy.random.Generator.negative_binomial()`.

noncentral_chisquare(*self*, *df*, *nonc*, *size=None*)
Call `numpy.random.Generator.noncentral_chisquare()`.

noncentral_f (*self*, *dfnum*, *dfden*, *nonc*, *size=None*)
Call `numpy.random.Generator.noncentral_f()`.

normal (*self*, *loc=0.0*, *scale=1.0*, *size=None*)
Call `numpy.random.Generator.normal()`.

pareto (*self*, *a*, *size=None*)
Call `numpy.random.Generator.pareto()`.

permutation (*self*, *x*)
Call `numpy.random.Generator.permutation()`.

poisson (*self*, *lam=1.0*, *size=None*)
Call `numpy.random.Generator.poisson()`.

power (*self*, *a*, *size=None*)
Call `numpy.random.Generator.power()`.

rand (*self*, **args*)
Call `numpy.random.RandomState.rand()`.

> **Warning:** This method is outdated in numpy. Use *RNG.random()* instead.

Added in 0.4.0.

randint (*self*, *low*, *high=None*, *size=None*, *dtype='int32'*)
Call `numpy.random.RandomState.randint()`.

> **Note:** Changed *dtype* argument default value from numpy's `I` to `int32`.

> **Warning:** This method is outdated in numpy. Use *RNG.integers()* instead.

Added in 0.4.0.

randn (*self*, **args*)
Call `numpy.random.RandomState.randn()`.

> **Warning:** This method is outdated in numpy. Use *RNG.standard_normal()* instead.

Added in 0.4.0.

random (*self*, *size*, *dtype='float32'*, *out=None*)
Call numpy's `random()` or `random_sample()`.

> **Note:** Changed *dtype* argument default value from numpy's `d` to `float32`.

random_integers (*self*, *low*, *high=None*, *size=None*)
Call `numpy.random.RandomState.random_integers()`.

> **Warning:** This method is outdated in numpy. Use *RNG.integers()* instead.

Added in 0.4.0.

random_sample (*self*, *size*)
Call `numpy.random.RandomState.random_sample()`.

> **Warning:** This method is outdated in numpy. Use *RNG.uniform()* instead.

Added in 0.4.0.

rayleigh (*self*, *scale=1.0*, *size=None*)
Call `numpy.random.Generator.rayleigh()`.

reset_cache_ (*self*)
Reset all cache of this RNG.

Returns The RNG itself.

Return type *RNG*

set_state_ (*self*, *value*)
Set the state if the RNG in-place.

Parameters **value** (*tuple or dict*) – The new state of the RNG. Should correspond to the output of the `state` property.

Returns The RNG itself.

Return type *RNG*

shuffle (*self*, *x*)
Call `numpy.random.Generator.shuffle()`.

standard_cauchy (*self*, *size=None*)
Call `numpy.random.Generator.standard_cauchy()`.

standard_exponential (*self*, *size=None*, *dtype='float32'*, *method='zig'*, *out=None*)
Call `numpy.random.Generator.standard_exponential()`.

> **Note:** Changed *dtype* argument default value from numpy's `d` to `float32`.

standard_gamma (*self*, *shape*, *size=None*, *dtype='float32'*, *out=None*)
Call `numpy.random.Generator.standard_gamma()`.

> **Note:** Changed *dtype* argument default value from numpy's `d` to `float32`.

standard_normal (*self*, *size=None*, *dtype='float32'*, *out=None*)
Call `numpy.random.Generator.standard_normal()`.

> **Note:** Changed *dtype* argument default value from numpy's `d` to `float32`.

standard_t (*self*, *df*, *size=None*)
Call `numpy.random.Generator.standard_t()`.

state
Get the state of this RNG.

Returns The state of the RNG. In numpy 1.17+, the bit generator's state will be returned. In numpy <=1.16, the `RandomState` 's state is returned. In both cases the state is a copy. In-place changes will not affect the RNG.

Return type tuple or dict

tomaxint (*self*, *size=None*)
Call `numpy.random.RandomState.tomaxint()`.

Warning: This method is outdated in numpy. Use *RNG.integers()* instead.

Added in 0.4.0.

triangular (*self*, *left*, *mode*, *right*, *size=None*)
Call `numpy.random.Generator.triangular()`.

uniform (*self*, *low=0.0*, *high=1.0*, *size=None*)
Call `numpy.random.Generator.uniform()`.

use_state_of_ (*self*, *other*)
Copy and use (in-place) the state of another RNG.

Note: It is often sensible to first verify that neither this RNG nor *other* are identical to the global RNG.

Parameters **other** (*RNG*) – The other RNG, which's state will be copied.

Returns The RNG itself.

Return type *RNG*

vonmises (*self*, *mu*, *kappa*, *size=None*)
Call `numpy.random.Generator.vonmises()`.

wald (*self*, *mean*, *scale*, *size=None*)
Call `numpy.random.Generator.wald()`.

weibull (*self*, *a*, *size=None*)
Call `numpy.random.Generator.weibull()`.

zipf (*self*, *a*, *size=None*)
Call `numpy.random.Generator.zipf()`.

imgaug.random.**advance_generator_** (*generator*)
Advance a numpy random generator's internal state in-place by one step.

This advances the generator's state.

Note: This simply samples one or more random values. This means that a call of this method will not completely change the outputs of the next called sampling method. To achieve more drastic output changes, call *derive_generator_()*.

Parameters **generator** (*numpy.random.Generator or numpy.random.RandomState*) – Generator of which to advance the internal state.

imgaug.random.**convert_seed_sequence_to_generator**(*seed_sequence*)
Convert a seed sequence to a numpy (random number) generator.

Parameters **seed_sequence** (*numpy.random.SeedSequence*) – The seed value to use.

Returns Generator initialized with the provided seed sequence.

Return type numpy.random.Generator

imgaug.random.**convert_seed_to_generator**(*entropy*)
Convert a seed value to a numpy (random number) generator.

Parameters **entropy** (*int*) – The seed value to use.

Returns In numpy <=1.16 a `RandomState`, in 1.17+ a `Generator`. Both are initialized with the provided seed.

Return type numpy.random.Generator or numpy.random.RandomState

imgaug.random.**copy_generator**(*generator*)
Copy an existing numpy (random number) generator.

Parameters **generator** (*numpy.random.Generator or numpy.random.RandomState*) – The generator to copy.

Returns In numpy <=1.16 a `RandomState`, in 1.17+ a `Generator`. Both are copies of the input argument.

Return type numpy.random.Generator or numpy.random.RandomState

imgaug.random.**copy_generator_unless_global_generator**(*generator*)
Copy a numpy generator unless it is the current global generator.

"global generator" here denotes the generator contained in the global RNG's `.generator` attribute.

Parameters **generator** (*numpy.random.Generator or numpy.random.RandomState*) – The generator to copy.

Returns In numpy <=1.16 a `RandomState`, in 1.17+ a `Generator`. Both are copies of the input argument, unless that input is identical to the global generator. If it is identical, the instance itself will be returned without copying it.

Return type numpy.random.Generator or numpy.random.RandomState

imgaug.random.**create_fully_random_generator**()
Create a new numpy (random) generator, derived from OS's entropy.

Returns In numpy <=1.16 a `RandomState`, in 1.17+ a `Generator`. Both are initialized with entropy requested from the OS. They are hence independent of entered seeds or the library's global RNG.

Return type numpy.random.Generator or numpy.random.RandomState

imgaug.random.**create_pseudo_random_generator_**()
Create a new numpy (random) generator, derived from the global RNG.

This function advances the global RNG's state.

Returns In numpy <=1.16 a `RandomState`, in 1.17+ a `Generator`. Both are initialized with a seed sampled from the global RNG.

Return type numpy.random.Generator or numpy.random.RandomState

imgaug.random.**derive_generator_**(*generator*)
Create a child numpy (random number) generator from an existing one.

This advances the generator's state.

Parameters **generator** (*numpy.random.Generator or numpy.random.RandomState*) – The generator from which to derive a new child generator.

Returns In numpy <=1.16 a `RandomState`, in 1.17+ a `Generator`. In both cases a derived child generator.

Return type numpy.random.Generator or numpy.random.RandomState

`imgaug.random.`**`derive_generators_`**(*generator*, *n*)
Create child numpy (random number) generators from an existing one.

Parameters

- **generator** (*numpy.random.Generator or numpy.random.RandomState*) – The generator from which to derive new child generators.
- **n** (*int*) – Number of child generators to derive.

Returns In numpy <=1.16 a list of `RandomState` s, in 1.17+ a list of `Generator` s. In both cases lists of derived child generators.

Return type list of numpy.random.Generator or list of numpy.random.RandomState

`imgaug.random.`**`generate_seed_`**(*generator*)
Sample a seed from the provided generator.

This function advances the generator's state.

See `SEED_MIN_VALUE` and `SEED_MAX_VALUE` for the seed's value range.

Parameters **generator** (*numpy.random.Generator or numpy.random.RandomState*) – The generator from which to sample the seed.

Returns The sampled seed.

Return type int

`imgaug.random.`**`generate_seeds_`**(*generator*, *n*)
Sample *n* seeds from the provided generator.

This function advances the generator's state.

Parameters

- **generator** (*numpy.random.Generator or numpy.random.RandomState*) – The generator from which to sample the seed.
- **n** (*int*) – Number of seeds to sample.

Returns 1D-array of `int32` seeds.

Return type ndarray

`imgaug.random.`**`get_generator_state`**(*generator*)
Get the state of this provided generator.

Parameters **generator** (*numpy.random.Generator or numpy.random.RandomState*) – The generator, which's state is supposed to be extracted.

Returns The state of the generator. In numpy 1.17+, the bit generator's state will be returned. In numpy <=1.16, the `RandomState` 's state is returned. In both cases the state is a copy. In-place changes will not affect the RNG.

Return type tuple or dict

imgaug.random.**get_global_rng**()
Get or create the current global RNG of imgaug.

Note that the first call to this function will create a global RNG.

Returns The global RNG to use.

Return type *RNG*

imgaug.random.**is_generator_equal_to**(*generator*, *other_generator*)
Estimate whether two generator have the same class and state.

Parameters

- **generator** (*numpy.random.Generator or numpy.random.RandomState*) – First generator used in the comparison.
- **other_generator** (*numpy.random.Generator or numpy.random.RandomState*) – Second generator used in the comparison.

Returns `True` if *generator* 's class and state are the same as the class and state of *other_generator*. `False` otherwise.

Return type bool

imgaug.random.**normalize_generator**(*generator*)
Normalize various inputs to a numpy (random number) generator.

This function will first copy the provided argument, i.e. it never returns a provided instance itself.

Parameters **generator** (*None or int or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState*) – The numpy random number generator to normalize. In case of numpy version 1.17 or later, this shouldn't be a `RandomState` as that class is outdated. Behaviour for different datatypes:

- If `None`: The global RNG's generator is returned.
- If `int`: In numpy 1.17+, the value is used as a seed for a `Generator`, i.e. it will be provided as the entropy to a `SeedSequence`, which will then be used for an `SFC64` bit generator and wrapped by a `Generator`, which is then returned. In numpy <=1.16, the value is used as a seed for a `RandomState`, which will then be returned.
- If `numpy.random.Generator`: That generator will be returned.
- If `numpy.random.BitGenerator`: A numpy generator will be created and returned that contains the bit generator.
- If `numpy.random.SeedSequence`: A numpy generator will be created and returned that contains an `SFC64` bit generator initialized with the given `SeedSequence`.
- If `numpy.random.RandomState`: In numpy <=1.16, this `RandomState` will be returned. In numpy 1.17+, a seed will be derived from this `RandomState` and a new `numpy.generator.Generator` based on an `SFC64` bit generator will be created and returned.

Returns In numpy <=1.16 a `RandomState`, in 1.17+ a `Generator` (even if the input was a `RandomState`).

Return type numpy.random.Generator or numpy.random.RandomState

imgaug.random.**normalize_generator_**(*generator*)
Normalize in-place various inputs to a numpy (random number) generator.

This function will try to return the provided instance itself.

Parameters **generator** (*None or int or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState*) – See `normalize_generator()`.

Returns In numpy <=1.16 a `RandomState`, in 1.17+ a `Generator` (even if the input was a `RandomState`).

Return type numpy.random.Generator or numpy.random.RandomState

imgaug.random.**polyfill_integers**(*generator*, *low*, *high=None*, *size=None*, *dtype='int32'*, *endpoint=False*)

Sample integers from a generator in different numpy versions.

Parameters

- **generator** (*numpy.random.Generator or numpy.random.RandomState*) – The generator to sample from. If it is a `RandomState`, `numpy.random.RandomState.randint()` will be called, otherwise `numpy.random.Generator.integers()`.
- **low** (*int or array-like of ints*) – See `numpy.random.Generator.integers()`.
- **high** (*int or array-like of ints, optional*) – See `numpy.random.Generator.integers()`.
- **size** (*int or tuple of ints, optional*) – See `numpy.random.Generator.integers()`.
- **dtype** (*{str, dtype}, optional*) – See `numpy.random.Generator.integers()`.
- **endpoint** (*bool, optional*) – See `numpy.random.Generator.integers()`.

Returns See `numpy.random.Generator.integers()`.

Return type int or ndarray of ints

imgaug.random.**polyfill_random**(*generator*, *size*, *dtype='float32'*, *out=None*)

Sample random floats from a generator in different numpy versions.

Parameters

- **generator** (*numpy.random.Generator or numpy.random.RandomState*) – The generator to sample from. Both `RandomState` and `Generator` support `random()`, but with different interfaces.
- **size** (*int or tuple of ints, optional*) – See `numpy.random.Generator.random()`.
- **dtype** (*{str, dtype}, optional*) – See `numpy.random.Generator.random()`.
- **out** (*ndarray, optional*) – See `numpy.random.Generator.random()`.

Returns See `numpy.random.Generator.random()`.

Return type float or ndarray of floats

imgaug.random.**reset_generator_cache_**(*generator*)

Reset a numpy (random number) generator's internal cache.

This function modifies the generator's state in-place.

Parameters **generator** (*numpy.random.Generator or numpy.random.RandomState*) – The generator of which to reset the cache.

Returns In numpy <=1.16 a `RandomState`, in 1.17+ a `Generator`. In both cases the input argument itself.

Return type numpy.random.Generator or numpy.random.RandomState

imgaug.random.**seed**(*entropy*)
Set the seed of imgaug's global RNG (in-place).

The global RNG controls most of the "randomness" in imgaug.

The global RNG is the default one used by all augmenters. Under special circumstances (e.g. when an augmenter is switched to deterministic mode), the global RNG is replaced with a local one. The state of that replacement may be dependent on the global RNG's state at the time of creating the child RNG.

Parameters **entropy** (*int*) – The seed value to use.

imgaug.random.**set_generator_state_**(*generator*, *state*)
Set the state of a numpy (random number) generator in-place.

Parameters

- **generator** (*numpy.random.Generator or numpy.random.RandomState*) – The generator, which's state is supposed to be modified.
- **state** (*tuple or dict*) – The new state of the generator. Should correspond to the output of *get_generator_state()*.

imgaug.random.**supports_new_numpy_rng_style**()
Determine whether numpy supports the new `random` interface (v1.17+).

Returns `True` if the new `random` interface is supported by numpy, i.e. if numpy has version 1.17 or later. Otherwise `False`, i.e. numpy has version 1.16 or older and `numpy.random.RandomState` should be used instead.

Return type bool

class imgaug.random.**temporary_numpy_seed**(*entropy=None*)
Bases: `object`

Context to temporarily alter the random state of `numpy.random`.

The random state's internal state will be set back to the original one once the context finishes.

Added in 0.4.0.

Parameters **entropy** (*None or int*) – The seed value to use. If *None* then the seed will not be altered and the internal state of `numpy.random` will not be reset back upon context exit (i.e. this context will do nothing).

13.6 imgaug.validation

Helper functions to validate input data and produce error messages.

imgaug.validation.**assert_is_iterable_of**(*iterable_var*, *classes*)
Assert that *iterable_var* only contains instances of given classes.

Parameters

- **iterable_var** (*iterable*) – See *is_iterable_of()*.
- **classes** (*type or iterable of type*) – See *is_iterable_of()*.

imgaug.validation.**convert_iterable_to_string_of_types**(*iterable_var*)
Convert an iterable of values to a string of their types.

Parameters **iterable_var** (*iterable*) – An iterable of variables, e.g. a list of integers.

Returns String representation of the types in *iterable_var*. One per item in *iterable_var*. Separated by commas.

Return type str

imgaug.validation.**is_iterable_of**(*iterable_var*, *classes*)

Check whether *iterable_var* contains only instances of given classes.

Parameters

- **iterable_var** (*iterable*) – An iterable of items that will be matched against *classes*.
- **classes** (*type or iterable of type*) – One or more classes that each item in *var* must be an instanceof. If this is an iterable, a single match per item is enough.

Returns Whether *var* only contains instances of *classes*. If *var* was empty, `True` will be returned.

Return type bool

13.7 imgaug.augmentables.base

Interfaces used by augmentable objects.

Added in 0.4.0.

class imgaug.augmentables.base.**IAugmentable**

Bases: `object`

Interface of augmentable objects.

This interface is right now only used to "mark" augmentable objects. It does not enforce any methods yet (but will probably in the future).

Currently, only `*OnImage` clases are marked as augmentable. Non-OnImage objects are normalized to OnImage-objects. Batches are not yet marked as augmentable, but might be in the future.

Added in 0.4.0.

13.8 imgaug.augmentables.batches

Classes representing batches of normalized or unnormalized data.

class imgaug.augmentables.batches.**Batch**(*images=None*, *heatmaps=None*, *segmentation_maps=None*, *keypoints=None*, *bounding_boxes=None*, *polygons=None*, *line_strings=None*, *data=None*)

Bases: `object`

Class encapsulating a batch before and after augmentation.

Parameters

- **images** (*None or (N,H,W,C) ndarray or list of (H,W,C) ndarray*) – The images to augment.
- **heatmaps** (*None or list of imgaug.augmentables.heatmaps.HeatmapsOnImage*) – The heatmaps to augment.
- **segmentation_maps** (*None or list of imgaug.augmentables.segmaps.SegmentationMapsOnImage*) – The segmentation maps to augment.

- **keypoints** (*None or list of imgaug.augmentables.kps.KeypointOnImage*) – The keypoints to augment.
- **bounding_boxes** (*None or list of imgaug.augmentables.bbs.BoundingBoxesOnImage*) – The bounding boxes to augment.
- **polygons** (*None or list of imgaug.augmentables.polys.PolygonsOnImage*) – The polygons to augment.
- **line_strings** (*None or list of imgaug.augmentables.lines.LineStringsOnImage*) – The line strings to augment.
- **data** – Additional data that is saved in the batch and may be read out after augmentation. This could e.g. contain filepaths to each image in *images*. As this object is usually used for background augmentation with multiple processes, the augmented Batch objects might not be returned in the original order, making this information useful.

Attributes

bounding_boxes **Deprecated**. Use `Batch.bounding_boxes_unaug` instead.

heatmaps **Deprecated**. Use `Batch.heatmaps_unaug` instead.

images **Deprecated**. Use `Batch.images_unaug` instead.

keypoints **Deprecated**. Use `Batch.keypoints_unaug` instead.

segmentation_maps **Deprecated**. Use `Batch.segmentation_maps_unaug` instead.

Methods

deepcopy(self[, images_unaug, images_aug, ...])	Copy this batch and all of its column values.
*fill_from_batch_in_augmentation_*(self, ...)	Set the columns in this batch to the column values of another batch.
get_column_names(self)	Get the names of types of augmentables that contain data.
to_batch_in_augmentation(self)	Convert this batch to a `_BatchInAugmentation` instance.
to_normalized_batch(self)	Return this batch.

bounding_boxes

Deprecated. Use `Batch.bounding_boxes_unaug` instead.

Get unaugmented bounding boxes.

deepcopy(*self*, *images_unaug='DEFAULT'*, *images_aug='DEFAULT'*, *heatmaps_unaug='DEFAULT'*, *heatmaps_aug='DEFAULT'*, *segmentation_maps_unaug='DEFAULT'*, *segmentation_maps_aug='DEFAULT'*, *keypoints_unaug='DEFAULT'*, *keypoints_aug='DEFAULT'*, *bounding_boxes_unaug='DEFAULT'*, *bounding_boxes_aug='DEFAULT'*, *polygons_unaug='DEFAULT'*, *polygons_aug='DEFAULT'*, *line_strings_unaug='DEFAULT'*, *line_strings_aug='DEFAULT'*)

Copy this batch and all of its column values.

Parameters

- **images_unaug** (*imgaug.augmentables.batches.DEFAULT or None or (N,H,W,C) ndarray or list of (H,W,C) ndarray*) – Copies the current attribute value without changes if set

to `imgaug.augmentables.batches.DEFAULT`. Otherwise same as in `Batch.__init__()`.

- **images_aug** (*imgaug.augmentables.batches.DEFAULT or None or (N,H,W,C) ndarray or list of (H,W,C) ndarray*) – Copies the current attribute value without changes if set to `imgaug.augmentables.batches.DEFAULT`. Otherwise same as in `Batch.__init__()`.
- **heatmaps_unaug** (*imgaug.augmentables.batches.DEFAULT or None or list of imgaug.augmentables.heatmaps.HeatmapsOnImage*) – Copies the current attribute value without changes if set to `imgaug.augmentables.batches.DEFAULT`. Otherwise same as in `Batch.__init__()`.
- **heatmaps_aug** (*imgaug.augmentables.batches.DEFAULT or None or list of imgaug.augmentables.heatmaps.HeatmapsOnImage*) – Copies the current attribute value without changes if set to `imgaug.augmentables.batches.DEFAULT`. Otherwise same as in `Batch.__init__()`.
- **segmentation_maps_unaug** (*imgaug.augmentables.batches.DEFAULT or None or list of imgaug.augmentables.segmaps.SegmentationMapsOnImage*) – Copies the current attribute value without changes if set to `imgaug.augmentables.batches.DEFAULT`. Otherwise same as in `Batch.__init__()`.
- **segmentation_maps_aug** (*imgaug.augmentables.batches.DEFAULT or None or list of imgaug.augmentables.segmaps.SegmentationMapsOnImage*) – Copies the current attribute value without changes if set to `imgaug.augmentables.batches.DEFAULT`. Otherwise same as in `Batch.__init__()`.
- **keypoints_unaug** (*imgaug.augmentables.batches.DEFAULT or None or list of imgaug.augmentables.kps.KeypointOnImage*) – Copies the current attribute value without changes if set to `imgaug.augmentables.batches.DEFAULT`. Otherwise same as in `Batch.__init__()`.
- **keypoints_aug** (*imgaug.augmentables.batches.DEFAULT or None or list of imgaug.augmentables.kps.KeypointOnImage*) – Copies the current attribute value without changes if set to `imgaug.augmentables.batches.DEFAULT`. Otherwise same as in `Batch.__init__()`.
- **bounding_boxes_unaug** (*imgaug.augmentables.batches.DEFAULT or None or list of imgaug.augmentables.bbs.BoundingBoxesOnImage*) – Copies the current attribute value without changes if set to `imgaug.augmentables.batches.DEFAULT`. Otherwise same as in `Batch.__init__()`.
- **bounding_boxes_aug** (*imgaug.augmentables.batches.DEFAULT or None or list of imgaug.augmentables.bbs.BoundingBoxesOnImage*) – Copies the current attribute value without changes if set to `imgaug.augmentables.batches.DEFAULT`. Otherwise same as in `Batch.__init__()`.
- **polygons_unaug** (*imgaug.augmentables.batches.DEFAULT or None or list of imgaug.augmentables.polys.PolygonsOnImage*) – Copies the current attribute value without changes if set to `imgaug.augmentables.batches.DEFAULT`. Otherwise same as in `Batch.__init__()`.
- **polygons_aug** (*imgaug.augmentables.batches.DEFAULT or None or list of imgaug.augmentables.polys.PolygonsOnImage*) – Copies the current attribute value without changes if set to `imgaug.augmentables.batches.DEFAULT`. Otherwise same as in `Batch.__init__()`.
- **line_strings_unaug** (*imgaug.augmentables.batches.DEFAULT or None or list of imgaug.augmentables.lines.LineStringsOnImage*) – Copies the current attribute value with-

out changes if set to `imgaug.augmentables.batches.DEFAULT`. Otherwise same as in `Batch.__init__()`.

- **line_strings_aug** (*imgaug.augmentables.batches.DEFAULT or None or list of imgaug.augmentables.lines.LineStringsOnImage*) – Copies the current attribute value without changes if set to `imgaug.augmentables.batches.DEFAULT`. Otherwise same as in `Batch.__init__()`.

Returns Deep copy of the batch, optionally with new attributes.

Return type *Batch*

fill_from_batch_in_augmentation_(*self*, *batch_in_augmentation*)
Set the columns in this batch to the column values of another batch.

This method works in-place.

Added in 0.4.0.

Parameters **batch_in_augmentation** (*_BatchInAugmentation*) – Batch of which to use the column values. The values are *not* copied. Only their references are used.

Returns The updated batch. (Modified in-place.)

Return type *Batch*

get_column_names(*self*)
Get the names of types of augmentables that contain data.

This method is intended for situations where one wants to know which data is contained in the batch that has to be augmented, visualized or something similar.

Added in 0.4.0.

Returns Names of types of augmentables. E.g. `["images", "polygons"]`.

Return type list of str

heatmaps
Deprecated. Use `Batch.heatmaps_unaug` instead.

Get unaugmented heatmaps.

images
Deprecated. Use `Batch.images_unaug` instead.

Get unaugmented images.

keypoints
Deprecated. Use `Batch.keypoints_unaug` instead.

Get unaugmented keypoints.

segmentation_maps
Deprecated. Use `Batch.segmentation_maps_unaug` instead.

Get unaugmented segmentation maps.

to_batch_in_augmentation(*self*)
Convert this batch to a `_BatchInAugmentation` instance.

Added in 0.4.0.

Returns The converted batch.

Return type imgaug.augmentables.batches._BatchInAugmentation

to_normalized_batch(*self*)
Return this batch.

This method does nothing and only exists to simplify interfaces that accept both *UnnormalizedBatch* and *Batch*.

Added in 0.4.0.

Returns This batch (not copied).

Return type *imgaug.augmentables.batches.Batch*

class imgaug.augmentables.batches.**UnnormalizedBatch**(*images=None*, *heatmaps=None*, *segmentation_maps=None*, *keypoints=None*, *bounding_boxes=None*, *polygons=None*, *line_strings=None*, *data=None*)

Bases: object

Class for batches of unnormalized data before and after augmentation.

Parameters

- **images** (*None or (N,H,W,C) ndarray or (N,H,W) ndarray or iterable of (H,W,C) ndarray or iterable of (H,W) ndarray*) – The images to augment.
- **heatmaps** (*None or (N,H,W,C) ndarray or imgaug.augmentables.heatmaps.HeatmapsOnImage or iterable of (H,W,C) ndarray or iterable of imgaug.augmentables.heatmaps.HeatmapsOnImage*) – The heatmaps to augment. If anything else than HeatmapsOnImage, then the number of heatmaps must match the number of images provided via parameter *images*. The number is contained either in N or the first iterable's size.
- **segmentation_maps** (*None or (N,H,W) ndarray or imgaug.augmentables.segmaps.SegmentationMapsOnImage or iterable of (H,W) ndarray or iterable of imgaug.augmentables.segmaps.SegmentationMapsOnImage*) – The segmentation maps to augment. If anything else than SegmentationMapsOnImage, then the number of segmaps must match the number of images provided via parameter *images*. The number is contained either in N or the first iterable's size.
- **keypoints** (*None or list of (N,K,2) ndarray or tuple of number or imgaug.augmentables.kps.Keypoint or iterable of (K,2) ndarray or iterable of tuple of number or iterable of imgaug.augmentables.kps.Keypoint or iterable of imgaug.augmentables.kps.KeypointOnImage or iterable of iterable of tuple of number or iterable of iterable of imgaug.augmentables.kps.Keypoint*) – The keypoints to augment. If a tuple (or iterable(s) of tuple), then iterpreted as (x,y) coordinates and must hence contain two numbers. A single tuple represents a single coordinate on one image, an iterable of tuples the coordinates on one image and an iterable of iterable of tuples the coordinates on several images. Analogous if Keypoint objects are used instead of tuples. If an ndarray, then N denotes the number of images and K the number of keypoints on each image. If anything else than KeypointsOnImage is provided, then the number of keypoint groups must match the number of images provided via parameter *images*. The number is contained e.g. in N or in case of "iterable of iterable of tuples" in the first iterable's size.
- **bounding_boxes** (*None or (N,B,4) ndarray or tuple of number or imgaug.augmentables.bbs.BoundingBox or imgaug.augmentables.bbs.BoundingBoxesOnImage or iterable of (B,4) ndarray or iterable of tuple of number or iterable of imgaug.augmentables.bbs.BoundingBox or iterable of imgaug.augmentables.bbs.BoundingBoxesOnImage or iterable of iterable of tuple of*

number or iterable of iterable imgaug.augmentables.bbs.BoundingBox) – The bounding boxes to augment. This is analogous to the *keypoints* parameter. However, each tuple – and also the last index in case of arrays – has size 4, denoting the bounding box coordinates `x1`, `y1`, `x2` and `y2`.

- **polygons** (*None or (N,#polys,#points,2) ndarray or imgaug.augmentables.polys.Polygon or imgaug.augmentables.polys.PolygonsOnImage or iterable of (#polys,#points,2) ndarray or iterable of tuple of number or iterable of imgaug.augmentables.kps.Keypoint or iterable of imgaug.augmentables.polys.Polygon or iterable of imgaug.augmentables.polys.PolygonsOnImage or iterable of iterable of (#points,2) ndarray or iterable of iterable of tuple of number or iterable of iterable of imgaug.augmentables.kps.Keypoint or iterable of iterable of imgaug.augmentables.polys.Polygon or iterable of iterable of iterable of tuple of number or iterable of iterable of iterable of tuple of imgaug.augmentables.kps.Keypoint*) – The polygons to augment. This is similar to the *keypoints* parameter. However, each polygon may be made up of several `(x,y)` coordinates (three or more are required for valid polygons). The following datatypes will be interpreted as a single polygon on a single image:

  – `imgaug.augmentables.polys.Polygon`

  – `iterable of tuple of number`

  – `iterable of imgaug.augmentables.kps.Keypoint`

  The following datatypes will be interpreted as multiple polygons on a single image:

  – `imgaug.augmentables.polys.PolygonsOnImage`

  – `iterable of imgaug.augmentables.polys.Polygon`

  – `iterable of iterable of tuple of number`

  – `iterable of iterable of imgaug.augmentables.kps.Keypoint`

  – `iterable of iterable of imgaug.augmentables.polys.Polygon`

  The following datatypes will be interpreted as multiple polygons on multiple images:

  – `(N,#polys,#points,2) ndarray`

  – `iterable of (#polys,#points,2) ndarray`

  – `iterable of iterable of (#points,2) ndarray`

  – `iterable of iterable of iterable of tuple of number`

  – `iterable of iterable of iterable of tuple of imgaug.augmentables.kps.Keypoint`

- **line_strings** (*None or (N,#lines,#points,2) ndarray or imgaug.augmentables.lines.LineString or imgaug.augmentables.lines.LineStringOnImage or iterable of (#lines,#points,2) ndarray or iterable of tuple of number or iterable of imgaug.augmentables.kps.Keypoint or iterable of imgaug.augmentables.lines.LineString or iterable of imgaug.augmentables.lines.LineStringOnImage or iterable of iterable of (#points,2) ndarray or iterable of iterable of tuple of number or iterable of iterable of imgaug.augmentables.kps.Keypoint or iterable of iterable of imgaug.augmentables.polys.LineString or iterable of iterable of iterable of tuple of number or iterable of iterable of iterable of tuple of imgaug.augmentables.kps.Keypoint*) – The line strings to augment. See *polygons* for more details as polygons follow a similar structure to line strings.

- **data** – Additional data that is saved in the batch and may be read out after augmentation. This could e.g. contain filepaths to each image in *images*. As this object is usually used for background augmentation with multiple processes, the augmented Batch objects might not be returned in the original order, making this information useful.

Methods

fill_from_augmented_normalized_batch(self, ...)	Fill this batch with (normalized) augmentation results.
*fill_from_augmented_normalized_batch_*(self, ...)	Fill this batch with (normalized) augmentation results in-place.
get_column_names(self)	Get the names of types of augmentables that contain data.
to_normalized_batch(self)	Convert this unnormalized batch to an instance of Batch.

fill_from_augmented_normalized_batch(*self*, *batch_aug_norm*)
Fill this batch with (normalized) augmentation results.

This method receives a (normalized) Batch instance, takes all `*_aug` attributes out if it and assigns them to this batch *in unnormalized form*. Hence, the datatypes of all `*_aug` attributes will match the datatypes of the `*_unaug` attributes.

Parameters **batch_aug_norm** (*imgaug.augmentables.batches.Batch*) – Batch after normalization and augmentation.

Returns New UnnormalizedBatch instance. All `*_unaug` attributes are taken from the old UnnormalizedBatch (without deepcopying them) and all `*_aug` attributes are taken from *batch_normalized*, converted to unnormalized form.

Return type *imgaug.augmentables.batches.UnnormalizedBatch*

fill_from_augmented_normalized_batch_(*self*, *batch_aug_norm*)
Fill this batch with (normalized) augmentation results in-place.

This method receives a (normalized) Batch instance, takes all `*_aug` attributes out if it and assigns them to this batch *in unnormalized form*. Hence, the datatypes of all `*_aug` attributes will match the datatypes of the `*_unaug` attributes.

Added in 0.4.0.

Parameters **batch_aug_norm** (*imgaug.augmentables.batches.Batch*) – Batch after normalization and augmentation.

Returns This instance itself. All `*_unaug` attributes are unchanged. All `*_aug` attributes are taken from *batch_normalized*, converted to unnormalized form.

Return type *imgaug.augmentables.batches.UnnormalizedBatch*

get_column_names(*self*)
Get the names of types of augmentables that contain data.

This method is intended for situations where one wants to know which data is contained in the batch that has to be augmented, visualized or something similar.

Added in 0.4.0.

Returns Names of types of augmentables. E.g. `["images", "polygons"]`.

Return type list of str

to_normalized_batch(*self*)
Convert this unnormalized batch to an instance of Batch.

As this method is intended to be called before augmentation, it assumes that none of the `*_aug` attributes is yet set. It will produce an AssertionError otherwise.

The newly created Batch's `*_unaug` attributes will match the ones in this batch, just in normalized form.

Returns The batch, with `*_unaug` attributes being normalized.

Return type *imgaug.augmentables.batches.Batch*

13.9 imgaug.augmentables.bbs

Classes representing bounding boxes.

class imgaug.augmentables.bbs.**BoundingBox**(*x1*, *y1*, *x2*, *y2*, *label=None*)
Bases: object

Class representing bounding boxes.

Each bounding box is parameterized by its top left and bottom right corners. Both are given as x and y-coordinates. The corners are intended to lie inside the bounding box area. As a result, a bounding box that lies completely inside the image but has maximum extensions would have coordinates `(0.0, 0.0)` and `(W - epsilon, H - epsilon)`. Note that coordinates are saved internally as floats.

Parameters

- **x1** (*number*) – X-coordinate of the top left of the bounding box.
- **y1** (*number*) – Y-coordinate of the top left of the bounding box.
- **x2** (*number*) – X-coordinate of the bottom right of the bounding box.
- **y2** (*number*) – Y-coordinate of the bottom right of the bounding box.
- **label** (*None or str, optional*) – Label of the bounding box, e.g. a string representing the class.

Attributes

area Estimate the area of the bounding box.

center_x Estimate the x-coordinate of the center point of the bounding box.

center_y Estimate the y-coordinate of the center point of the bounding box.

coords Get the top-left and bottom-right coordinates as one array.

height Estimate the height of the bounding box.

width Estimate the width of the bounding box.

x1_int Get the x-coordinate of the top left corner as an integer.

x2_int Get the x-coordinate of the bottom left corner as an integer.

y1_int Get the y-coordinate of the top left corner as an integer.

y2_int Get the y-coordinate of the bottom left corner as an integer.

Methods

almost_equals(self, other[, max_distance])	Compare this and another BB's label and coordinates.
clip_out_of_image(self, image)	Clip off all parts of the BB box that are outside of the image.
*clip_out_of_image_*(self, image)	Clip off parts of the BB box that are outside of the image in-place.
compute_out_of_image_area(self, image)	Compute the area of the BB that is outside of the image plane.
compute_out_of_image_fraction(self, image)	Compute fraction of BB area outside of the image plane.
contains(self, other)	Estimate whether the bounding box contains a given point.
coords_almost_equals(self, other[, max_distance])	Estimate if this and another BB have almost identical coordinates.
copy(self[, x1, y1, x2, y2, label])	Create a shallow copy of this BoundingBox instance.
cut_out_of_image(self, *args, **kwargs)	**Deprecated**.
deepcopy(self[, x1, y1, x2, y2, label])	Create a deep copy of the BoundingBox object.
draw_box_on_image(self, image[, color, ...])	Draw the rectangle of the bounding box on an image.
draw_label_on_image(self, image[, color, ...])	Draw a box showing the BB's label.
draw_on_image(self, image[, color, alpha, ...])	Draw the bounding box on an image.
extend(self[, all_sides, top, right, ...])	Extend the size of the bounding box along its sides.
*extend_*(self[, all_sides, top, right, ...])	Extend the size of the bounding box along its sides in-place.
extract_from_image(self, image[, pad, ...])	Extract the image pixels within the bounding box.
from_point_soup(xy)	Convert a `(2P,) or (P,2) ndarray` to a BB instance.
intersection(self, other[, default])	Compute the intersection BB between this BB and another BB.
iou(self, other)	Compute the IoU between this bounding box and another one.
is_fully_within_image(self, image)	Estimate whether the bounding box is fully inside the image area.
is_out_of_image(self, image[, fully, partly])	Estimate whether the BB is partially/fully outside of the image area.
is_partly_within_image(self, image)	Estimate whether the BB is at least partially inside the image area.
project(self, from_shape, to_shape)	Project the bounding box onto a differently shaped image.
*project_*(self, from_shape, to_shape)	Project the bounding box onto a differently shaped image in-place.
shift(self[, x, y, top, right, bottom, left])	Move this bounding box along the x/y-axis.
*shift_*(self[, x, y])	Move this bounding box along the x/y-axis in-place.
to_keypoints(self)	Convert the BB's corners to keypoints (clockwise, from top left).
to_polygon(self)	Convert this bounding box to a polygon covering the same area.
union(self, other)	Compute the union BB between this BB and another BB.

almost_equals(*self*, *other*, *max_distance=0.0001*)
Compare this and another BB's label and coordinates.

This is the same as `coords_almost_equals()` but additionally compares the labels.

Added in 0.4.0.

Parameters

- **other** (*imgaug.augmentables.bbs.BoundingBox or iterable*) – The other object to compare against. Expected to be a `BoundingBox`.
- **max_distance** (*number, optional*) – See `coords_almost_equals()`.

Returns `True` if the coordinates are almost equal and additionally the labels are equal. Otherwise `False`.

Return type bool

area
Estimate the area of the bounding box.

Returns Area of the bounding box, i.e. `height * width`.

Return type number

center_x
Estimate the x-coordinate of the center point of the bounding box.

Returns X-coordinate of the center point of the bounding box.

Return type number

center_y
Estimate the y-coordinate of the center point of the bounding box.

Returns Y-coordinate of the center point of the bounding box.

Return type number

clip_out_of_image(*self*, *image*)
Clip off all parts of the BB box that are outside of the image.

Parameters **image** (*(H,W,...) ndarray or tuple of int*) – Image dimensions to use for the clipping of the bounding box. If an `ndarray`, its shape will be used. If a `tuple`, it is assumed to represent the image shape and must contain at least two integers.

Returns Bounding box, clipped to fall within the image dimensions.

Return type *imgaug.augmentables.bbs.BoundingBox*

clip_out_of_image_(*self*, *image*)
Clip off parts of the BB box that are outside of the image in-place.

Added in 0.4.0.

Parameters **image** (*(H,W,...) ndarray or tuple of int*) – Image dimensions to use for the clipping of the bounding box. If an `ndarray`, its shape will be used. If a `tuple`, it is assumed to represent the image shape and must contain at least two integers.

Returns Bounding box, clipped to fall within the image dimensions. The object may have been modified in-place.

Return type *imgaug.augmentables.bbs.BoundingBox*

compute_out_of_image_area(*self*, *image*)
Compute the area of the BB that is outside of the image plane.

Added in 0.4.0.

Parameters **image** (*(H,W,...) ndarray or tuple of int*) – Image dimensions to use. If an `ndarray`, its shape will be used. If a `tuple`, it is assumed to represent the image shape and must contain at least two integers.

Returns Total area of the bounding box that is outside of the image plane. Can be `0.0`.

Return type float

compute_out_of_image_fraction(*self*, *image*)
Compute fraction of BB area outside of the image plane.

This estimates `f = A_ooi / A`, where `A_ooi` is the area of the bounding box that is outside of the image plane, while `A` is the total area of the bounding box.

Added in 0.4.0.

Parameters **image** (*(H,W,...) ndarray or tuple of int*) – Image dimensions to use. If an `ndarray`, its shape will be used. If a `tuple`, it is assumed to represent the image shape and must contain at least two integers.

Returns Fraction of the bounding box area that is outside of the image plane. Returns `0.0` if the bounding box is fully inside of the image plane. If the bounding box has an area of zero, the result is `1.0` if its coordinates are outside of the image plane, otherwise `0.0`.

Return type float

contains(*self*, *other*)
Estimate whether the bounding box contains a given point.

Parameters **other** (*tuple of number or imgaug.augmentables.kps.Keypoint*) – Point to check for.

Returns `True` if the point is contained in the bounding box, `False` otherwise.

Return type bool

coords
Get the top-left and bottom-right coordinates as one array.

Added in 0.4.0.

Returns A `(N, 2)` numpy array with `N=2` containing the top-left and bottom-right coordinates.

Return type ndarray

coords_almost_equals(*self*, *other*, *max_distance=0.0001*)
Estimate if this and another BB have almost identical coordinates.

Added in 0.4.0.

Parameters

- **other** (*imgaug.augmentables.bbs.BoundingBox or iterable*) – The other bounding box with which to compare this one. If this is an `iterable`, it is assumed to represent the top-left and bottom-right coordinates of that bounding box, given as e.g. an `(2,2)` ndarray or an `(4,)` ndarray or as a similar list.
- **max_distance** (*number, optional*) – The maximum euclidean distance between a corner on one bounding box and the closest corner on the other bounding box. If the distance is exceeded for any such pair, the two BBs are not viewed as equal.

Returns Whether the two bounding boxes have almost identical corner coordinates.

Return type bool

copy(*self*, *x1=None*, *y1=None*, *x2=None*, *y2=None*, *label=None*)
Create a shallow copy of this BoundingBox instance.

Parameters

- **x1** (*None or number*) – If not `None`, then the `x1` coordinate of the copied object will be set to this value.
- **y1** (*None or number*) – If not `None`, then the `y1` coordinate of the copied object will be set to this value.
- **x2** (*None or number*) – If not `None`, then the `x2` coordinate of the copied object will be set to this value.
- **y2** (*None or number*) – If not `None`, then the `y2` coordinate of the copied object will be set to this value.
- **label** (*None or string*) – If not `None`, then the `label` of the copied object will be set to this value.

Returns Shallow copy.

Return type *imgaug.augmentables.bbs.BoundingBox*

cut_out_of_image(*self*, **args*, ***kwargs*)
Deprecated. Use `BoundingBox.clip_out_of_image()` instead. clip_out_of_image() has the exactly same interface.

Clip off all parts of the BB box that are outside of the image.

deepcopy(*self*, *x1=None*, *y1=None*, *x2=None*, *y2=None*, *label=None*)
Create a deep copy of the BoundingBox object.

Parameters

- **x1** (*None or number*) – If not `None`, then the `x1` coordinate of the copied object will be set to this value.
- **y1** (*None or number*) – If not `None`, then the `y1` coordinate of the copied object will be set to this value.
- **x2** (*None or number*) – If not `None`, then the `x2` coordinate of the copied object will be set to this value.
- **y2** (*None or number*) – If not `None`, then the `y2` coordinate of the copied object will be set to this value.
- **label** (*None or string*) – If not `None`, then the `label` of the copied object will be set to this value.

Returns Deep copy.

Return type *imgaug.augmentables.bbs.BoundingBox*

draw_box_on_image(*self*, *image*, *color=(0, 255, 0)*, *alpha=1.0*, *size=1*, *copy=True*, *raise_if_out_of_image=False*, *thickness=None*)
Draw the rectangle of the bounding box on an image.

This method does not draw the label.

Added in 0.4.0.

Parameters

- **image** (*(H,W,C) ndarray*) – The image onto which to draw the bounding box rectangle. Currently expected to be `uint8`.
- **color** (*iterable of int, optional*) – The color to use, corresponding to the channel layout of the image. Usually RGB.
- **alpha** (*float, optional*) – The transparency of the drawn bounding box, where `1.0` denotes no transparency and `0.0` is invisible.
- **size** (*int, optional*) – The thickness of the bounding box in pixels. If the value is larger than `1`, then additional pixels will be added around the bounding box (i.e. extension towards the outside).
- **copy** (*bool, optional*) – Whether to copy the input image or change it in-place.
- **raise_if_out_of_image** (*bool, optional*) – Whether to raise an error if the bounding box is fully outside of the image. If set to `False`, no error will be raised and only the parts inside the image will be drawn.
- **thickness** (*None or int, optional*) – Deprecated.

Returns Image with bounding box drawn on it.

Return type (H,W,C) ndarray(uint8)

draw_label_on_image(*self*, *image*, *color=(0*, *255*, *0)*, *color_text=None*, *color_bg=None*, *alpha=1.0*, *size=1*, *size_text=20*, *height=30*, *copy=True*, *raise_if_out_of_image=False*)

Draw a box showing the BB's label.

The box is placed right above the BB's rectangle.

Added in 0.4.0.

Parameters

- **image** (*(H,W,C) ndarray*) – The image onto which to draw the label. Currently expected to be `uint8`.
- **color** (*None or iterable of int, optional*) – The color to use, corresponding to the channel layout of the image. Usually RGB. Text and background colors will be derived from this.
- **color_text** (*None or iterable of int, optional*) – The text color to use. If `None`, derived from *color_bg*.
- **color_bg** (*None or iterable of int, optional*) – The background color of the label box. If `None`, derived from *color*.
- **alpha** (*float, optional*) – The transparency of the drawn bounding box, where `1.0` denotes no transparency and `0.0` is invisible.
- **size** (*int, optional*) – The thickness of the bounding box in pixels. If the value is larger than `1`, then additional pixels will be added around the bounding box (i.e. extension towards the outside).
- **size_text** (*int, optional*) – Font size to use.
- **height** (*int, optional*) – Height of the label box in pixels.
- **copy** (*bool, optional*) – Whether to copy the input image or change it in-place.
- **raise_if_out_of_image** (*bool, optional*) – Whether to raise an error if the bounding box is fully outside of the image. If set to `False`, no error will be raised and only the parts inside the image will be drawn.

Returns Image with bounding box drawn on it.

Return type (H,W,C) ndarray(uint8)

draw_on_image(*self*, *image*, *color=(0, 255, 0)*, *alpha=1.0*, *size=1*, *copy=True*, *raise_if_out_of_image=False*, *thickness=None*)

Draw the bounding box on an image.

This will automatically also draw the label, unless it is `None`. To only draw the box rectangle use *draw_box_on_image()*. To draw the label even if it is `None` or to configure e.g. its color, use *draw_label_on_image()*.

Parameters

- **image** (*((H,W,C) ndarray*) – The image onto which to draw the bounding box. Currently expected to be `uint8`.
- **color** (*iterable of int, optional*) – The color to use, corresponding to the channel layout of the image. Usually RGB.
- **alpha** (*float, optional*) – The transparency of the drawn bounding box, where `1.0` denotes no transparency and `0.0` is invisible.
- **size** (*int, optional*) – The thickness of the bounding box in pixels. If the value is larger than `1`, then additional pixels will be added around the bounding box (i.e. extension towards the outside).
- **copy** (*bool, optional*) – Whether to copy the input image or change it in-place.
- **raise_if_out_of_image** (*bool, optional*) – Whether to raise an error if the bounding box is fully outside of the image. If set to `False`, no error will be raised and only the parts inside the image will be drawn.
- **thickness** (*None or int, optional*) – Deprecated.

Returns Image with bounding box drawn on it.

Return type (H,W,C) ndarray(uint8)

extend(*self*, *all_sides=0*, *top=0*, *right=0*, *bottom=0*, *left=0*)

Extend the size of the bounding box along its sides.

Parameters

- **all_sides** (*number, optional*) – Value by which to extend the bounding box size along all sides.
- **top** (*number, optional*) – Value by which to extend the bounding box size along its top side.
- **right** (*number, optional*) – Value by which to extend the bounding box size along its right side.
- **bottom** (*number, optional*) – Value by which to extend the bounding box size along its bottom side.
- **left** (*number, optional*) – Value by which to extend the bounding box size along its left side.

Returns Extended bounding box.

Return type imgaug.BoundingBox

extend_(*self*, *all_sides=0*, *top=0*, *right=0*, *bottom=0*, *left=0*)

Extend the size of the bounding box along its sides in-place.

Added in 0.4.0.

Parameters

- **all_sides** (*number, optional*) – Value by which to extend the bounding box size along all sides.
- **top** (*number, optional*) – Value by which to extend the bounding box size along its top side.
- **right** (*number, optional*) – Value by which to extend the bounding box size along its right side.
- **bottom** (*number, optional*) – Value by which to extend the bounding box size along its bottom side.
- **left** (*number, optional*) – Value by which to extend the bounding box size along its left side.

Returns Extended bounding box. The object may have been modified in-place.

Return type imgaug.BoundingBox

extract_from_image(*self*, *image*, *pad=True*, *pad_max=None*, *prevent_zero_size=True*)

Extract the image pixels within the bounding box.

This function will zero-pad the image if the bounding box is partially/fully outside of the image.

Parameters

- **image** (*(H,W) ndarray or (H,W,C) ndarray*) – The image from which to extract the pixels within the bounding box.
- **pad** (*bool, optional*) – Whether to zero-pad the image if the object is partially/fully outside of it.
- **pad_max** (*None or int, optional*) – The maximum number of pixels that may be zero-padded on any side, i.e. if this has value `N` the total maximum of added pixels is `4*N`. This option exists to prevent extremely large images as a result of single points being moved very far away during augmentation.
- **prevent_zero_size** (*bool, optional*) – Whether to prevent the height or width of the extracted image from becoming zero. If this is set to `True` and the height or width of the bounding box is below `1`, the height/width will be increased to `1`. This can be useful to prevent problems, e.g. with image saving or plotting. If it is set to `False`, images will be returned as `(H', W')` or `(H', W', 3)` with `H` or `W` potentially being 0.

Returns Pixels within the bounding box. Zero-padded if the bounding box is partially/fully outside of the image. If *prevent_zero_size* is activated, it is guarantueed that `H'>0` and `W'>0`, otherwise only `H'>=0` and `W'>=0`.

Return type (H',W') ndarray or (H',W',C) ndarray

classmethod from_point_soup(*xy*)

Convert a `(2P,) or (P,2) ndarray` to a BB instance.

This is the inverse of `to_xyxy_array()`.

Added in 0.4.0.

Parameters **xy** (*(2P,) ndarray or (P, 2) array or iterable of number or iterable of iterable of number*) – Array containing `P` points in xy-form denoting a soup of points around which to place a bounding box. The array should usually be of dtype `float32`.

Returns Bounding box around the points.

Return type *imgaug.augmentables.bbs.BoundingBox*

height
Estimate the height of the bounding box.

Returns Height of the bounding box.

Return type number

intersection (*self*, *other*, *default=None*)
Compute the intersection BB between this BB and another BB.

Note that in extreme cases, the intersection can be a single point. In that case the intersection bounding box exists and it will be returned, but it will have a height and width of zero.

Parameters

- **other** (*imgaug.augmentables.bbs.BoundingBox*) – Other bounding box with which to generate the intersection.
- **default** (*any, optional*) – Default value to return if there is no intersection.

Returns Intersection bounding box of the two bounding boxes if there is an intersection. If there is no intersection, the default value will be returned, which can by anything.

Return type *imgaug.augmentables.bbs.BoundingBox* or any

iou (*self*, *other*)
Compute the IoU between this bounding box and another one.

IoU is the intersection over union, defined as:

```
``area(intersection(A, B)) / area(union(A, B))``
``= area(intersection(A, B))
    / (area(A) + area(B) - area(intersection(A, B)))``
```

Parameters **other** (*imgaug.augmentables.bbs.BoundingBox*) – Other bounding box with which to compare.

Returns IoU between the two bounding boxes.

Return type float

is_fully_within_image (*self*, *image*)
Estimate whether the bounding box is fully inside the image area.

Parameters **image** (*(H,W,...) ndarray or tuple of int*) – Image dimensions to use. If an `ndarray`, its shape will be used. If a `tuple`, it is assumed to represent the image shape and must contain at least two integers.

Returns `True` if the bounding box is fully inside the image area. `False` otherwise.

Return type bool

is_out_of_image (*self*, *image*, *fully=True*, *partly=False*)
Estimate whether the BB is partially/fully outside of the image area.

Parameters

- **image** (*(H,W,...) ndarray or tuple of int*) – Image dimensions to use. If an `ndarray`, its shape will be used. If a `tuple`, it is assumed to represent the image shape and must contain at least two integers.
- **fully** (*bool, optional*) – Whether to return `True` if the bounding box is fully outside of the image area.

- **partly** (*bool, optional*) – Whether to return `True` if the bounding box is at least partially outside fo the image area.

Returns `True` if the bounding box is partially/fully outside of the image area, depending on defined parameters. `False` otherwise.

Return type bool

is_partly_within_image(*self*, *image*)
Estimate whether the BB is at least partially inside the image area.

Parameters **image** (*(H,W,...) ndarray or tuple of int*) – Image dimensions to use. If an `ndarray`, its shape will be used. If a `tuple`, it is assumed to represent the image shape and must contain at least two integers.

Returns `True` if the bounding box is at least partially inside the image area. `False` otherwise.

Return type bool

project(*self*, *from_shape*, *to_shape*)
Project the bounding box onto a differently shaped image.

E.g. if the bounding box is on its original image at `x1=(10 of 100 pixels)` and `y1=(20 of 100 pixels)` and is projected onto a new image with size `(width=200, height=200)`, its new position will be `(x1=20, y1=40)`. (Analogous for `x2`/`y2`.)

This is intended for cases where the original image is resized. It cannot be used for more complex changes (e.g. padding, cropping).

Parameters

- **from_shape** (*tuple of int or ndarray*) – Shape of the original image. (Before resize.)
- **to_shape** (*tuple of int or ndarray*) – Shape of the new image. (After resize.)

Returns `BoundingBox` instance with new coordinates.

Return type *imgaug.augmentables.bbs.BoundingBox*

project_(*self*, *from_shape*, *to_shape*)
Project the bounding box onto a differently shaped image in-place.

E.g. if the bounding box is on its original image at `x1=(10 of 100 pixels)` and `y1=(20 of 100 pixels)` and is projected onto a new image with size `(width=200, height=200)`, its new position will be `(x1=20, y1=40)`. (Analogous for `x2`/`y2`.)

This is intended for cases where the original image is resized. It cannot be used for more complex changes (e.g. padding, cropping).

Added in 0.4.0.

Parameters

- **from_shape** (*tuple of int or ndarray*) – Shape of the original image. (Before resize.)
- **to_shape** (*tuple of int or ndarray*) – Shape of the new image. (After resize.)

Returns `BoundingBox` instance with new coordinates. The object may have been modified in-place.

Return type *imgaug.augmentables.bbs.BoundingBox*

shift(*self*, *x=0*, *y=0*, *top=None*, *right=None*, *bottom=None*, *left=None*)
Move this bounding box along the x/y-axis.

The origin `(0, 0)` is at the top left of the image.

Parameters

- **x** (*number, optional*) – Value to be added to all x-coordinates. Positive values shift towards the right images.
- **y** (*number, optional*) – Value to be added to all y-coordinates. Positive values shift towards the bottom images.
- **top** (*None or int, optional*) – Deprecated since 0.4.0. Amount of pixels by which to shift this object *from* the top (towards the bottom).
- **right** (*None or int, optional*) – Deprecated since 0.4.0. Amount of pixels by which to shift this object *from* the right (towards the left).
- **bottom** (*None or int, optional*) – Deprecated since 0.4.0. Amount of pixels by which to shift this object *from* the bottom (towards the top).
- **left** (*None or int, optional*) – Deprecated since 0.4.0. Amount of pixels by which to shift this object *from* the left (towards the right).

Returns Shifted bounding box.

Return type *imgaug.augmentables.bbs.BoundingBox*

shift_ (*self*, *x=0*, *y=0*)

Move this bounding box along the x/y-axis in-place.

The origin `(0, 0)` is at the top left of the image.

Added in 0.4.0.

Parameters

- **x** (*number, optional*) – Value to be added to all x-coordinates. Positive values shift towards the right images.
- **y** (*number, optional*) – Value to be added to all y-coordinates. Positive values shift towards the bottom images.

Returns Shifted bounding box. The object may have been modified in-place.

Return type *imgaug.augmentables.bbs.BoundingBox*

to_keypoints (*self*)

Convert the BB's corners to keypoints (clockwise, from top left).

Returns Corners of the bounding box as keypoints.

Return type list of imgaug.augmentables.kps.Keypoint

to_polygon (*self*)

Convert this bounding box to a polygon covering the same area.

Added in 0.4.0.

Returns The bounding box converted to a polygon.

Return type *imgaug.augmentables.polys.Polygon*

union (*self*, *other*)

Compute the union BB between this BB and another BB.

This is equivalent to drawing a bounding box around all corner points of both bounding boxes.

Parameters **other** (*imgaug.augmentables.bbs.BoundingBox*) – Other bounding box with which to generate the union.

Returns Union bounding box of the two bounding boxes.

Return type *imgaug.augmentables.bbs.BoundingBox*

width
Estimate the width of the bounding box.

Returns Width of the bounding box.

Return type number

x1_int
Get the x-coordinate of the top left corner as an integer.

Returns X-coordinate of the top left corner, rounded to the closest integer.

Return type int

x2_int
Get the x-coordinate of the bottom left corner as an integer.

Returns X-coordinate of the bottom left corner, rounded to the closest integer.

Return type int

y1_int
Get the y-coordinate of the top left corner as an integer.

Returns Y-coordinate of the top left corner, rounded to the closest integer.

Return type int

y2_int
Get the y-coordinate of the bottom left corner as an integer.

Returns Y-coordinate of the bottom left corner, rounded to the closest integer.

Return type int

class imgaug.augmentables.bbs.**BoundingBoxesOnImage**(*bounding_boxes*, *shape*)
Bases: *imgaug.augmentables.base.IAugmentable*

Container for the list of all bounding boxes on a single image.

Parameters

- **bounding_boxes** (*list of imgaug.augmentables.bbs.BoundingBox*) – List of bounding boxes on the image.
- **shape** (*tuple of int or ndarray*) – The shape of the image on which the objects are placed. Either an image with shape `(H,W,[C])` or a `tuple` denoting such an image shape.

Examples

```
>>> import numpy as np
>>> from imgaug.augmentables.bbs import BoundingBox, BoundingBoxesOnImage
>>>
>>> image = np.zeros((100, 100))
>>> bbs = [
>>>     BoundingBox(x1=10, y1=20, x2=20, y2=30),
>>>     BoundingBox(x1=25, y1=50, x2=30, y2=70)
>>> ]
>>> bbs_oi = BoundingBoxesOnImage(bbs, shape=image.shape)
```

Attributes

empty Determine whether this instance contains zero bounding boxes.

height Get the height of the image on which the bounding boxes fall.

items Get the bounding boxes in this container.

width Get the width of the image on which the bounding boxes fall.

Methods

clip_out_of_image(self)	Clip off all parts from all BBs that are outside of the image.
*clip_out_of_image_*(self)	Clip off in-place all parts from all BBs that are outside of the image.
copy(self[, bounding_boxes, shape])	Create a shallow copy of the `BoundingBoxesOnImage` instance.
cut_out_of_image(self)	**Deprecated**.
deepcopy(self[, bounding_boxes, shape])	Create a deep copy of the `BoundingBoxesOnImage` object.
draw_on_image(self, image[, color, alpha, ...])	Draw all bounding boxes onto a given image.
*fill_from_xy_array_*(self, xy)	Modify the BB coordinates of this instance in-place.
*fill_from_xyxy_array_*(self, xyxy)	Modify the BB coordinates of this instance in-place.
from_point_soups(xy, shape)	Convert an `(N, 2P) or (N, P, 2) ndarray` to a BBsOI instance.
from_xyxy_array(xyxy, shape)	Convert an `(N, 4) or (N, 2, 2) ndarray` to a BBsOI instance.
*invert_to_keypoints_on_image_*(self, kpsoi)	Invert the output of `to_keypoints_on_image()` in-place.
on(self, image)	Project bounding boxes from one image (shape) to a another one.
*on_*(self, image)	Project BBs from one image (shape) to a another one in-place.
remove_out_of_image(self[, fully, partly])	Remove all BBs that are fully/partially outside of the image.
*remove_out_of_image_*(self[, fully, partly])	Remove in-place all BBs that are fully/partially outside of the image.
remove_out_of_image_fraction(self, fraction)	Remove all BBs with an out of image fraction of at least *fraction*.
*remove_out_of_image_fraction_*(self, fraction)	Remove in-place all BBs with an OOI fraction of at least *fraction*.
shift(self[, x, y, top, right, bottom, left])	Move all BBs along the x/y-axis.
*shift_*(self[, x, y])	Move all BBs along the x/y-axis in-place.
to_keypoints_on_image(self)	Convert the bounding boxes to one `KeypointsOnImage` instance.
to_polygons_on_image(self)	Convert the bounding boxes to one `PolygonsOnImage` instance.
to_xy_array(self)	Convert the `BoundingBoxesOnImage` object to an `(N,2) ndarray`.
to_xyxy_array(self[, dtype])	Convert the `BoundingBoxesOnImage` object to an `(N,4) ndarray`.

clip_out_of_image(*self*)
Clip off all parts from all BBs that are outside of the image.

Returns Bounding boxes, clipped to fall within the image dimensions.

Return type *imgaug.augmentables.bbs.BoundingBoxesOnImage*

clip_out_of_image_(*self*)
Clip off in-place all parts from all BBs that are outside of the image.

Added in 0.4.0.

Returns Bounding boxes, clipped to fall within the image dimensions. The object and its items may have been modified in-place.

Return type *imgaug.augmentables.bbs.BoundingBoxesOnImage*

copy(*self*, *bounding_boxes=None*, *shape=None*)
Create a shallow copy of the `BoundingBoxesOnImage` instance.

Parameters

- **bounding_boxes** (*None or list of imgaug.augmntables.bbs.BoundingBox, optional*) – List of bounding boxes on the image. If `None`, the instance's bounding boxes will be copied.
- **shape** (*tuple of int, optional*) – The shape of the image on which the bounding boxes are placed. If `None`, the instance's shape will be copied.

Returns Shallow copy.

Return type *imgaug.augmentables.bbs.BoundingBoxesOnImage*

cut_out_of_image(*self*)
Deprecated. Use `BoundingBoxesOnImage.clip_out_of_image()` instead. clip_out_of_image() has the exactly same interface.

Clip off all parts from all BBs that are outside of the image.

deepcopy(*self*, *bounding_boxes=None*, *shape=None*)
Create a deep copy of the `BoundingBoxesOnImage` object.

Parameters

- **bounding_boxes** (*None or list of imgaug.augmntables.bbs.BoundingBox, optional*) – List of bounding boxes on the image. If `None`, the instance's bounding boxes will be copied.
- **shape** (*tuple of int, optional*) – The shape of the image on which the bounding boxes are placed. If `None`, the instance's shape will be copied.

Returns Deep copy.

Return type *imgaug.augmentables.bbs.BoundingBoxesOnImage*

draw_on_image(*self*, *image*, *color=(0, 255, 0)*, *alpha=1.0*, *size=1*, *copy=True*, *raise_if_out_of_image=False*, *thickness=None*)
Draw all bounding boxes onto a given image.

Parameters

- **image** (*(H,W,3) ndarray*) – The image onto which to draw the bounding boxes. This image should usually have the same shape as set in `BoundingBoxesOnImage.shape`.
- **color** (*int or list of int or tuple of int or (3,) ndarray, optional*) – The RGB color of all bounding boxes. If a single `int C`, then that is equivalent to `(C,C,C)`.
- **alpha** (*float, optional*) – Alpha/transparency of the bounding box.

- **size** (*int, optional*) – Thickness in pixels.
- **copy** (*bool, optional*) – Whether to copy the image before drawing the bounding boxes.
- **raise_if_out_of_image** (*bool, optional*) – Whether to raise an exception if any bounding box is outside of the image.
- **thickness** (*None or int, optional*) – Deprecated.

Returns Image with drawn bounding boxes.

Return type (H,W,3) ndarray

empty
Determine whether this instance contains zero bounding boxes.

Returns True if this object contains zero bounding boxes.

Return type bool

fill_from_xy_array_ (*self*, *xy*)
Modify the BB coordinates of this instance in-place.

See *fill_from_xyxy_array_()*.

Added in 0.4.0.

Parameters **xy** (*(2*B, 2) ndarray or iterable of iterable of number*) – Coordinates of `B` bounding boxes on an image, given as a `(2*B,2)` array of two corner xy-coordinates per bounding box. `B` must match the number of bounding boxes in this instance.

Returns This instance itself, with updated bounding box coordinates. Note that the instance was modified in-place.

Return type *BoundingBoxesOnImage*

fill_from_xyxy_array_ (*self*, *xyxy*)
Modify the BB coordinates of this instance in-place.

Note: This currently expects exactly one entry in *xyxy* per bounding in this instance. (I.e. two corner coordinates per instance.) Otherwise, an `AssertionError` will be raised.

Note: This method will automatically flip x-coordinates if `x1>x2` for a bounding box. (Analogous for y-coordinates.)

Added in 0.4.0.

Parameters **xyxy** (*(N, 4) ndarray or iterable of iterable of number*) – Coordinates of `N` bounding boxes on an image, given as a `(N,4)` array of two corner xy-coordinates per bounding box. `N` must match the number of bounding boxes in this instance.

Returns This instance itself, with updated bounding box coordinates. Note that the instance was modified in-place.

Return type *BoundingBoxesOnImage*

classmethod from_point_soups (*xy*, *shape*)
Convert an `(N, 2P) or (N, P, 2) ndarray` to a BBsOI instance.

Added in 0.4.0.

Parameters

- **xy** (*(N, 2P) ndarray or (N, P, 2) array or iterable of iterable of number or iterable of iterable of iterable of number*) – Array containing the corner coordinates of `N` bounding boxes. Each bounding box is represented by a soup of `P` points. If `(N, P)` then the second axis is expected to be in xy-form (e.g. `x1`, `y1`, `x2`, `y2`, ...). The final bounding box coordinates will be derived using `min` and `max` operations on the xy-values. The array should usually be of dtype `float32`.
- **shape** (*tuple of int*) – Shape of the image on which the bounding boxes are placed. Should usually be `(H, W, C)` or `(H, W)`.

Returns Object containing a list of `BoundingBox` instances derived from the provided point soups.

Return type *imgaug.augmentables.bbs.BoundingBoxesOnImage*

classmethod from_xyxy_array(*xyxy*, *shape*)

Convert an `(N, 4) or (N, 2, 2) ndarray` to a BBsOI instance.

This is the inverse of `to_xyxy_array()`.

Parameters

- **xyxy** (*(N, 4) ndarray or (N, 2, 2) array*) – Array containing the corner coordinates of `N` bounding boxes. Each bounding box is represented by its top-left and bottom-right coordinates. The array should usually be of dtype `float32`.
- **shape** (*tuple of int*) – Shape of the image on which the bounding boxes are placed. Should usually be `(H, W, C)` or `(H, W)`.

Returns Object containing a list of `BoundingBox` instances derived from the provided corner coordinates.

Return type *imgaug.augmentables.bbs.BoundingBoxesOnImage*

height

Get the height of the image on which the bounding boxes fall.

Returns Image height.

Return type int

invert_to_keypoints_on_image_(*self*, *kpsoi*)

Invert the output of `to_keypoints_on_image()` in-place.

This function writes in-place into this `BoundingBoxesOnImage` instance.

Added in 0.4.0.

Parameters **kpsoi** (*imgaug.augmentables.kps.KeypointsOnImages*) – Keypoints to convert back to bounding boxes, i.e. the outputs of `to_keypoints_on_image()`.

Returns Bounding boxes container with updated coordinates. Note that the instance is also updated in-place.

Return type *BoundingBoxesOnImage*

items

Get the bounding boxes in this container.

Added in 0.4.0.

Returns Bounding boxes within this container.

Return type list of BoundingBox

on (*self*, *image*)
Project bounding boxes from one image (shape) to a another one.

Parameters **image** (*ndarray or tuple of int*) – New image onto which the bounding boxes are to be projected. May also simply be that new image's shape tuple.

Returns Object containing the same bounding boxes after projection to the new image shape.

Return type *imgaug.augmentables.bbs.BoundingBoxesOnImage*

on_ (*self*, *image*)
Project BBs from one image (shape) to a another one in-place.

Added in 0.4.0.

Parameters **image** (*ndarray or tuple of int*) – New image onto which the bounding boxes are to be projected. May also simply be that new image's shape tuple.

Returns Object containing the same bounding boxes after projection to the new image shape. The object and its items may have been modified in-place.

Return type *imgaug.augmentables.bbs.BoundingBoxesOnImage*

remove_out_of_image (*self*, *fully=True*, *partly=False*)
Remove all BBs that are fully/partially outside of the image.

Parameters

- **fully** (*bool, optional*) – Whether to remove bounding boxes that are fully outside of the image.
- **partly** (*bool, optional*) – Whether to remove bounding boxes that are partially outside of the image.

Returns Reduced set of bounding boxes, with those that were fully/partially outside of the image being removed.

Return type *imgaug.augmentables.bbs.BoundingBoxesOnImage*

remove_out_of_image_ (*self*, *fully=True*, *partly=False*)
Remove in-place all BBs that are fully/partially outside of the image.

Added in 0.4.0.

Parameters

- **fully** (*bool, optional*) – Whether to remove bounding boxes that are fully outside of the image.
- **partly** (*bool, optional*) – Whether to remove bounding boxes that are partially outside of the image.

Returns Reduced set of bounding boxes, with those that were fully/partially outside of the image being removed. The object and its items may have been modified in-place.

Return type *imgaug.augmentables.bbs.BoundingBoxesOnImage*

remove_out_of_image_fraction (*self*, *fraction*)
Remove all BBs with an out of image fraction of at least *fraction*.

Added in 0.4.0.

Parameters **fraction** (*number*) – Minimum out of image fraction that a bounding box has to have in order to be removed. A fraction of `1.0` removes only bounding boxes that are `100%` outside of the image. A fraction of `0.0` removes all bounding boxes.

Returns Reduced set of bounding boxes, with those that had an out of image fraction greater or equal the given one removed.

Return type *imgaug.augmentables.bbs.BoundingBoxesOnImage*

remove_out_of_image_fraction_(*self*, *fraction*)
Remove in-place all BBs with an OOI fraction of at least *fraction*.

'OOI' is the abbreviation for 'out of image'.

Added in 0.4.0.

Parameters **fraction** (*number*) – Minimum out of image fraction that a bounding box has to have in order to be removed. A fraction of `1.0` removes only bounding boxes that are `100%` outside of the image. A fraction of `0.0` removes all bounding boxes.

Returns Reduced set of bounding boxes, with those that had an out of image fraction greater or equal the given one removed. The object and its items may have been modified in-place.

Return type *imgaug.augmentables.bbs.BoundingBoxesOnImage*

shift(*self*, *x=0*, *y=0*, *top=None*, *right=None*, *bottom=None*, *left=None*)
Move all BBs along the x/y-axis.

The origin `(0, 0)` is at the top left of the image.

Parameters

- **x** (*number, optional*) – Value to be added to all x-coordinates. Positive values shift towards the right images.
- **y** (*number, optional*) – Value to be added to all y-coordinates. Positive values shift towards the bottom images.
- **top** (*None or int, optional*) – Deprecated since 0.4.0. Amount of pixels by which to shift all objects *from* the top (towards the bottom).
- **right** (*None or int, optional*) – Deprecated since 0.4.0. Amount of pixels by which to shift all objects *from* the right (towads the left).
- **bottom** (*None or int, optional*) – Deprecated since 0.4.0. Amount of pixels by which to shift all objects *from* the bottom (towards the top).
- **left** (*None or int, optional*) – Deprecated since 0.4.0. Amount of pixels by which to shift all objects *from* the left (towards the right).

Returns Shifted bounding boxes.

Return type *imgaug.augmentables.bbs.BoundingBoxesOnImage*

shift_(*self*, *x=0*, *y=0*)
Move all BBs along the x/y-axis in-place.

The origin `(0, 0)` is at the top left of the image.

Added in 0.4.0.

Parameters

- **x** (*number, optional*) – Value to be added to all x-coordinates. Positive values shift towards the right images.
- **y** (*number, optional*) – Value to be added to all y-coordinates. Positive values shift towards the bottom images.

Returns Shifted bounding boxes. The object and its items may have been modified in-place.

Return type *imgaug.augmentables.bbs.BoundingBoxesOnImage*

to_keypoints_on_image(*self*)
Convert the bounding boxes to one `KeypointsOnImage` instance.

Added in 0.4.0.

Returns A keypoints instance containing `N*4` coordinates for `N` bounding boxes. Order matches the order in `bounding_boxes`.

Return type *imgaug.augmentables.kps.KeypointsOnImage*

to_polygons_on_image(*self*)
Convert the bounding boxes to one `PolygonsOnImage` instance.

Added in 0.4.0.

Returns A `PolygonsOnImage` containing polygons. Each polygon covers the same area as the corresponding bounding box.

Return type *imgaug.augmentables.polys.PolygonsOnImage*

to_xy_array(*self*)
Convert the `BoundingBoxesOnImage` object to an `(N,2) ndarray`.

Added in 0.4.0.

Returns `(2*B,2) ndarray` of xy-coordinates, where `B` denotes the number of bounding boxes.

Return type ndarray

to_xyxy_array(*self*, *dtype=<class 'numpy.float32'>*)
Convert the `BoundingBoxesOnImage` object to an `(N,4) ndarray`.

This is the inverse of `from_xyxy_array()`.

Parameters **dtype** (*numpy.dtype, optional*) – Desired output datatype of the ndarray.

Returns `(N,4) ndarray`, where `N` denotes the number of bounding boxes and `4` denotes the top-left and bottom-right bounding box corner coordinates in form `(x1, y1, x2, y2)`.

Return type ndarray

width
Get the width of the image on which the bounding boxes fall.

Returns Image width.

Return type int

13.10 imgaug.augmentables.heatmaps

Classes to represent heatmaps, i.e. float arrays of `[0.0, 1.0]`.

class imgaug.augmentables.heatmaps.**HeatmapsOnImage**(*arr*, *shape*, *min_value=0.0*, *max_value=1.0*)
Bases: *imgaug.augmentables.base.IAugmentable*

Object representing heatmaps on a single image.

Parameters

- **arr** (*(H,W) ndarray or (H,W,C) ndarray*) – Array representing the heatmap(s) on a single image. Multiple heatmaps may be provided, in which case `C` is expected to denote the heatmap index. The array must be of dtype `float32`.
- **shape** (*tuple of int*) – Shape of the image on which the heatmap(s) is/are placed. **Not** the shape of the heatmap(s) array, unless it is identical to the image shape (note the likely difference between the arrays in the number of channels). This is expected to be `(H, W)` or `(H, W, C)` with `C` usually being `3`. If there is no corresponding image, use `(H_arr, W_arr)` instead, where `H_arr` is the height of the heatmap(s) array (analogous `W_arr`).
- **min_value** (*float, optional*) – Minimum value for the heatmaps that *arr* represents. This will usually be `0.0`.
- **max_value** (*float, optional*) – Maximum value for the heatmaps that *arr* represents. This will usually be `1.0`.

Methods

avg_pool(self, block_size)	Average-pool the heatmap(s) array using a given block/kernel size.
change_normalization(arr, source, target)	Change the value range of a heatmap array.
copy(self)	Create a shallow copy of the heatmaps object.
deepcopy(self)	Create a deep copy of the heatmaps object.
draw(self[, size, cmap])	Render the heatmaps as RGB images.
draw_on_image(self, image[, alpha, cmap, resize])	Draw the heatmaps as overlays over an image.
from_0to1(arr_0to1, shape[, min_value, ...])	Create a heatmaps object from a `[0.0, 1.0]` float array.
from_uint8(arr_uint8, shape[, min_value, ...])	Create a `float`-based heatmaps object from an `uint8` array.
get_arr(self)	Get the heatmap's array in value range provided to `__init__()`.
invert(self)	Invert each component in the heatmap.
max_pool(self, block_size)	Max-pool the heatmap(s) array using a given block/kernel size.
pad(self[, top, right, bottom, left, mode, cval])	Pad the heatmaps at their top/right/bottom/left side.
pad_to_aspect_ratio(self, aspect_ratio[, ...])	Pad the heatmaps until they match a target aspect ratio.
resize(self, sizes[, interpolation])	Resize the heatmap(s) array given a target size and interpolation.
scale(self, *args, **kwargs)	**Deprecated**.
to_uint8(self)	Convert this heatmaps object to an `uint8` array.

avg_pool(*self*, *block_size*)

Average-pool the heatmap(s) array using a given block/kernel size.

Parameters **block_size** (*int or tuple of int*) – Size of each block of values to pool, aka kernel size. See *pool()* for details.

Returns Heatmaps after average pooling.

Return type *imgaug.augmentables.heatmaps.HeatmapsOnImage*

classmethod change_normalization(*arr*, *source*, *target*)

Change the value range of a heatmap array.

E.g. the value range may be changed from the interval `[0.0, 1.0]` to `[-1.0, 1.0]`.

Parameters

- **arr** (*ndarray*) – Heatmap array to modify.
- **source** (*tuple of float*) – Current value range of the input array, given as a tuple `(min, max)`, where both are `float` values.
- **target** (*tuple of float*) – Desired output value range of the array, given as a tuple `(min, max)`, where both are `float` values.

Returns Input array, with value range projected to the desired target value range.

Return type ndarray

copy (*self*)
Create a shallow copy of the heatmaps object.

Returns Shallow copy.

Return type *imgaug.augmentables.heatmaps.HeatmapsOnImage*

deepcopy (*self*)
Create a deep copy of the heatmaps object.

Returns Deep copy.

Return type *imgaug.augmentables.heatmaps.HeatmapsOnImage*

draw (*self*, *size=None*, *cmap='jet'*)
Render the heatmaps as RGB images.

Parameters

- **size** (*None or float or iterable of int or iterable of float, optional*) – Size of the rendered RGB image as `(height, width)`. See *imresize_single_image()* for details. If set to `None`, no resizing is performed and the size of the heatmaps array is used.
- **cmap** (*str or None, optional*) – Name of the `matplotlib` color map to use when convert the heatmaps to RGB images. If set to `None`, no color map will be used and the heatmaps will be converted to simple intensity maps.

Returns Rendered heatmaps as `uint8` arrays. Always a **list** containing one RGB image per heatmap array channel.

Return type list of (H,W,3) ndarray

draw_on_image (*self*, *image*, *alpha=0.75*, *cmap='jet'*, *resize='heatmaps'*)
Draw the heatmaps as overlays over an image.

Parameters

- **image** (*(H,W,3) ndarray*) – Image onto which to draw the heatmaps. Expected to be of dtype `uint8`.
- **alpha** (*float, optional*) – Alpha/opacity value to use for the mixing of image and heatmaps. Larger values mean that the heatmaps will be more visible and the image less visible.
- **cmap** (*str or None, optional*) – Name of the `matplotlib` color map to use. See *HeatmapsOnImage.draw()* for details.
- **resize** (*{'heatmaps', 'image'}, optional*) – In case of size differences between the image and heatmaps, either the image or the heatmaps can be resized. This parameter controls which of the two will be resized to the other's size.

Returns Rendered overlays as `uint8` arrays. Always a **list** containing one RGB image per heatmap array channel.

Return type list of (H,W,3) ndarray

static from_0to1(*arr_0to1*, *shape*, *min_value=0.0*, *max_value=1.0*)

Create a heatmaps object from a `[0.0, 1.0]` float array.

Parameters

- **arr_0to1** (*(H,W) or (H,W,C) ndarray*) – Heatmap(s) array, where `H` is the height, `W` is the width and `C` is the number of heatmap channels. Expected dtype is `float32`.
- **shape** (*tuple of ints*) – Shape of the image on which the heatmap(s) is/are placed. **Not** the shape of the heatmap(s) array, unless it is identical to the image shape (note the likely difference between the arrays in the number of channels). If there is not a corresponding image, use the shape of the heatmaps array.
- **min_value** (*float, optional*) – Minimum value of the float heatmaps that the input array represents. This will usually be 0.0. In most other cases it will be close to the interval `[0.0, 1.0]`. Calling `get_arr()`, will automatically convert the interval `[0.0, 1.0]` float array to this `[min, max]` interval.
- **max_value** (*float, optional*) – Minimum value of the float heatmaps that the input array represents. This will usually be 1.0. See parameter *min_value* for details.

Returns Heatmaps object.

Return type *imgaug.augmentables.heatmaps.HeatmapsOnImage*

static from_uint8(*arr_uint8*, *shape*, *min_value=0.0*, *max_value=1.0*)

Create a `float`-based heatmaps object from an `uint8` array.

Parameters

- **arr_uint8** (*(H,W) ndarray or (H,W,C) ndarray*) – Heatmap(s) array, where `H` is height, `W` is width and `C` is the number of heatmap channels. Expected dtype is `uint8`.
- **shape** (*tuple of int*) – Shape of the image on which the heatmap(s) is/are placed. **Not** the shape of the heatmap(s) array, unless it is identical to the image shape (note the likely difference between the arrays in the number of channels). If there is not a corresponding image, use the shape of the heatmaps array.
- **min_value** (*float, optional*) – Minimum value of the float heatmaps that the input array represents. This will usually be 0.0. In most other cases it will be close to the interval `[0.0, 1.0]`. Calling `get_arr()`, will automatically convert the interval `[0.0, 1.0]` float array to this `[min, max]` interval.
- **max_value** (*float, optional*) – Minimum value of the float heatmaps that the input array represents. This will usually be 1.0. See parameter *min_value* for details.

Returns Heatmaps object.

Return type *imgaug.augmentables.heatmaps.HeatmapsOnImage*

get_arr(*self*)

Get the heatmap's array in value range provided to `__init__()`.

The *`HeatmapsOnImage`* object saves heatmaps internally in the value range `[0.0, 1.0]`. This function converts the internal representation to `[min, max]`, where `min` and `max` are provided to `HeatmapsOnImage.__init__()` upon instantiation of the object.

Returns Heatmap array of dtype `float32`.

Return type (H,W) ndarray or (H,W,C) ndarray

invert(*self*)
Invert each component in the heatmap.

This shifts low values towards high values and vice versa.

This changes each value to:

```
v' = max - (v - min)
```

where `v` is the value at a spatial location, `min` is the minimum value in the heatmap and `max` is the maximum value. As the heatmap uses internally a `0.0` to `1.0` representation, this simply becomes `v' = 1.0 - v`.

This function can be useful e.g. when working with depth maps, where algorithms might have an easier time representing the furthest away points with zeros, requiring an inverted depth map.

Returns Inverted heatmap.

Return type *imgaug.augmentables.heatmaps.HeatmapsOnImage*

max_pool(*self*, *block_size*)
Max-pool the heatmap(s) array using a given block/kernel size.

Parameters **block_size** (*int or tuple of int*) – Size of each block of values to pool, aka kernel size. See *pool()* for details.

Returns Heatmaps after max-pooling.

Return type *imgaug.augmentables.heatmaps.HeatmapsOnImage*

pad(*self*, *top=0*, *right=0*, *bottom=0*, *left=0*, *mode='constant'*, *cval=0.0*)
Pad the heatmaps at their top/right/bottom/left side.

Parameters

- **top** (*int, optional*) – Amount of pixels to add at the top side of the heatmaps. Must be `0` or greater.
- **right** (*int, optional*) – Amount of pixels to add at the right side of the heatmaps. Must be `0` or greater.
- **bottom** (*int, optional*) – Amount of pixels to add at the bottom side of the heatmaps. Must be `0` or greater.
- **left** (*int, optional*) – Amount of pixels to add at the left side of the heatmaps. Must be `0` or greater.
- **mode** (*string, optional*) – Padding mode to use. See *pad()* for details.
- **cval** (*number, optional*) – Value to use for padding *mode* is `constant`. See *pad()* for details.

Returns Padded heatmaps of height `H'=H+top+bottom` and width `W'=W+left+right`.

Return type *imgaug.augmentables.heatmaps.HeatmapsOnImage*

pad_to_aspect_ratio(*self*, *aspect_ratio*, *mode='constant'*, *cval=0.0*, *return_pad_amounts=False*)
Pad the heatmaps until they match a target aspect ratio.

Depending on which dimension is smaller (height or width), only the corresponding sides (left/right or top/bottom) will be padded. In each case, both of the sides will be padded equally.

Parameters

- **aspect_ratio** (*float*) – Target aspect ratio, given as width/height. E.g. `2.0` denotes the image having twice as much width as height.
- **mode** (*str, optional*) – Padding mode to use. See *pad()* for details.
- **cval** (*number, optional*) – Value to use for padding if *mode* is `constant`. See *pad()* for details.
- **return_pad_amounts** (*bool, optional*) – If `False`, then only the padded instance will be returned. If `True`, a tuple with two entries will be returned, where the first entry is the padded instance and the second entry are the amounts by which each array side was padded. These amounts are again a tuple of the form `(top, right, bottom, left)`, with each value being an integer.

Returns

- *imgaug.augmentables.heatmaps.HeatmapsOnImage* – Padded heatmaps as *HeatmapsOnImage* instance.
- *tuple of int* – Amounts by which the instance's array was padded on each side, given as a tuple `(top, right, bottom, left)`. This tuple is only returned if *return_pad_amounts* was set to `True`.

resize(*self*, *sizes*, *interpolation='cubic'*)

Resize the heatmap(s) array given a target size and interpolation.

Parameters

- **sizes** (*float or iterable of int or iterable of float*) – New size of the array in `(height, width)`. See *imresize_single_image()* for details.
- **interpolation** (*None or str or int, optional*) – The interpolation to use during resize. See *imresize_single_image()* for details.

Returns Resized heatmaps object.

Return type *imgaug.augmentables.heatmaps.HeatmapsOnImage*

scale(*self*, **args*, ***kwargs*)

Deprecated. Use `HeatmapsOnImage.resize()` instead. resize() has the exactly same interface.

Resize the heatmap(s) array given a target size and interpolation.

to_uint8(*self*)

Convert this heatmaps object to an `uint8` array.

Returns Heatmap as an `uint8` array, i.e. with the discrete value range `[0, 255]`.

Return type (H,W,C) ndarray

13.11 imgaug.augmentables.kps

Classes to represent keypoints, i.e. points given as xy-coordinates.

class imgaug.augmentables.kps.**Keypoint**(*x*, *y*)

Bases: `object`

A single keypoint (aka landmark) on an image.

Parameters

- **x** (*number*) – Coordinate of the keypoint on the x axis.

- **y** (*number*) – Coordinate of the keypoint on the y axis.

Attributes

coords Get the xy-coordinates as an `(N,2)` ndarray.

x_int Get the keypoint's x-coordinate, rounded to the closest integer.

xy Get the keypoint's x- and y-coordinate as a single array.

xy_int Get the keypoint's xy-coord, rounded to closest integer.

y_int Get the keypoint's y-coordinate, rounded to the closest integer.

Methods

almost_equals(self, other[, max_distance])	Compare this and another KP's coordinates.
compute_out_of_image_fraction(self, image)	Compute fraction of the keypoint that is out of the image plane.
coords_almost_equals(self, other[, max_distance])	Estimate if this and another KP have almost identical coordinates.
copy(self[, x, y])	Create a shallow copy of the keypoint instance.
deepcopy(self[, x, y])	Create a deep copy of the keypoint instance.
draw_on_image(self, image[, color, alpha, ...])	Draw the keypoint onto a given image.
generate_similar_points_manhattan(self, ...)	Generate nearby points based on manhattan distance.
is_out_of_image(self, image)	Estimate whether this point is outside of the given image plane.
project(self, from_shape, to_shape)	Project the keypoint onto a new position on a new image.
*project_*(self, from_shape, to_shape)	Project in-place the keypoint onto a new position on a new image.
shift(self[, x, y])	Move the keypoint around on an image.
*shift_*(self[, x, y])	Move the keypoint around on an image in-place.

almost_equals (*self*, *other*, *max_distance=0.0001*)
Compare this and another KP's coordinates.

Note: This method is currently identical to `coords_almost_equals`. It exists for consistency with `BoundingBox` and `Polygons`.

Added in 0.4.0.

Parameters

- **other** (*imgaug.augmentables.kps.Keypoint or iterable*) – The other object to compare against. Expected to be a `Keypoint`.
- **max_distance** (*number, optional*) – See *coords_almost_equals()*.

Returns `True` if the coordinates are almost equal. Otherwise `False`.

Return type bool

compute_out_of_image_fraction (*self*, *image*)
Compute fraction of the keypoint that is out of the image plane.

The fraction is always either `1.0` (point is outside of the image plane) or `0.0` (point is inside the image plane). This method exists for consistency with other augmentables, e.g. bounding boxes.

Added in 0.4.0.

Parameters **image** (*(H,W,...) ndarray or tuple of int*) – Image dimensions to use. If an `ndarray`, its shape will be used. If a `tuple`, it is assumed to represent the image shape and must contain at least two integers.

Returns Either `1.0` (point is outside of the image plane) or `0.0` (point is inside of it).

Return type float

coords

Get the xy-coordinates as an `(N,2)` ndarray.

Added in 0.4.0.

Returns An `(N, 2) float32` ndarray with `N=1` containing the coordinates of this keypoints.

Return type ndarray

coords_almost_equals(*self*, *other*, *max_distance=0.0001*)

Estimate if this and another KP have almost identical coordinates.

Added in 0.4.0.

Parameters

- **other** (*imgaug.augmentables.kps.Keypoint or iterable*) – The other keypoint with which to compare this one. If this is an `iterable`, it is assumed to contain the xy-coordinates of a keypoint.
- **max_distance** (*number, optional*) – The maximum euclidean distance between a this keypoint and the other one. If the distance is exceeded, the two keypoints are not viewed as equal.

Returns Whether the two keypoints have almost identical coordinates.

Return type bool

copy(*self*, *x=None*, *y=None*)

Create a shallow copy of the keypoint instance.

Parameters

- **x** (*None or number, optional*) – Coordinate of the keypoint on the x axis. If `None`, the instance's value will be copied.
- **y** (*None or number, optional*) – Coordinate of the keypoint on the y axis. If `None`, the instance's value will be copied.

Returns Shallow copy.

Return type *imgaug.augmentables.kps.Keypoint*

deepcopy(*self*, *x=None*, *y=None*)

Create a deep copy of the keypoint instance.

Parameters

- **x** (*None or number, optional*) – Coordinate of the keypoint on the x axis. If `None`, the instance's value will be copied.
- **y** (*None or number, optional*) – Coordinate of the keypoint on the y axis. If `None`, the instance's value will be copied.

Returns Deep copy.

Return type *imgaug.augmentables.kps.Keypoint*

draw_on_image(*self*, *image*, *color=(0, 255, 0)*, *alpha=1.0*, *size=3*, *copy=True*, *raise_if_out_of_image=False*)

Draw the keypoint onto a given image.

The keypoint is drawn as a square.

Parameters

- **image** (*(H,W,3) ndarray*) – The image onto which to draw the keypoint.
- **color** (*int or list of int or tuple of int or (3,) ndarray, optional*) – The RGB color of the keypoint. If a single `int C`, then that is equivalent to `(C,C,C)`.
- **alpha** (*float, optional*) – The opacity of the drawn keypoint, where `1.0` denotes a fully visible keypoint and `0.0` an invisible one.
- **size** (*int, optional*) – The size of the keypoint. If set to `S`, each square will have size `S x S`.
- **copy** (*bool, optional*) – Whether to copy the image before drawing the keypoint.
- **raise_if_out_of_image** (*bool, optional*) – Whether to raise an exception if the keypoint is outside of the image.

Returns **image** – Image with drawn keypoint.

Return type (H,W,3) ndarray

generate_similar_points_manhattan(*self*, *nb_steps*, *step_size*, *return_array=False*)

Generate nearby points based on manhattan distance.

To generate the first neighbouring points, a distance of `S` (step size) is moved from the center point (this keypoint) to the top, right, bottom and left, resulting in four new points. From these new points, the pattern is repeated. Overlapping points are ignored.

The resulting points have a shape similar to a square rotated by `45` degrees.

Parameters

- **nb_steps** (*int*) – The number of steps to move from the center point. `nb_steps=1` results in a total of `5` output points (one center point + four neighbours).
- **step_size** (*number*) – The step size to move from every point to its neighbours.
- **return_array** (*bool, optional*) – Whether to return the generated points as a list of *Keypoint* or an array of shape `(N,2)`, where `N` is the number of generated points and the second axis contains the x-/y-coordinates.

Returns If *return_array* was `False`, then a list of *Keypoint*. Otherwise a numpy array of shape `(N,2)`, where `N` is the number of generated points and the second axis contains the x-/y-coordinates. The center keypoint (the one on which this function was called) is always included.

Return type list of imgaug.augmentables.kps.Keypoint or (N,2) ndarray

is_out_of_image(*self*, *image*)

Estimate whether this point is outside of the given image plane.

Added in 0.4.0.

Parameters **image** (*(H,W,...) ndarray or tuple of int*) – Image dimensions to use. If an `ndarray`, its shape will be used. If a `tuple`, it is assumed to represent the image shape and must contain at least two integers.

Returns `True` is the point is inside the image plane, `False` otherwise.

Return type bool

project (*self*, *from_shape*, *to_shape*)

Project the keypoint onto a new position on a new image.

E.g. if the keypoint is on its original image at `x=(10 of 100 pixels)` and `y=(20 of 100 pixels)` and is projected onto a new image with size `(width=200, height=200)`, its new position will be `(20, 40)`.

This is intended for cases where the original image is resized. It cannot be used for more complex changes (e.g. padding, cropping).

Parameters

- **from_shape** (*tuple of int*) – Shape of the original image. (Before resize.)
- **to_shape** (*tuple of int*) – Shape of the new image. (After resize.)

Returns Keypoint object with new coordinates.

Return type *imgaug.augmentables.kps.Keypoint*

project_ (*self*, *from_shape*, *to_shape*)

Project in-place the keypoint onto a new position on a new image.

E.g. if the keypoint is on its original image at `x=(10 of 100 pixels)` and `y=(20 of 100 pixels)` and is projected onto a new image with size `(width=200, height=200)`, its new position will be `(20, 40)`.

This is intended for cases where the original image is resized. It cannot be used for more complex changes (e.g. padding, cropping).

Added in 0.4.0.

Parameters

- **from_shape** (*tuple of int*) – Shape of the original image. (Before resize.)
- **to_shape** (*tuple of int*) – Shape of the new image. (After resize.)

Returns Keypoint object with new coordinates. The instance of the keypoint may have been modified in-place.

Return type *imgaug.augmentables.kps.Keypoint*

shift (*self*, *x=0*, *y=0*)

Move the keypoint around on an image.

Parameters

- **x** (*number, optional*) – Move by this value on the x axis.
- **y** (*number, optional*) – Move by this value on the y axis.

Returns Keypoint object with new coordinates.

Return type *imgaug.augmentables.kps.Keypoint*

shift_ (*self*, *x=0*, *y=0*)

Move the keypoint around on an image in-place.

Added in 0.4.0.

Parameters

- **x** (*number, optional*) – Move by this value on the x axis.
- **y** (*number, optional*) – Move by this value on the y axis.

Returns Keypoint object with new coordinates. The instance of the keypoint may have been modified in-place.

Return type *imgaug.augmentables.kps.Keypoint*

x_int

Get the keypoint's x-coordinate, rounded to the closest integer.

Returns **result** – Keypoint's x-coordinate, rounded to the closest integer.

Return type int

xy

Get the keypoint's x- and y-coordinate as a single array.

Added in 0.4.0.

Returns A `(2,)` `ndarray` denoting the xy-coordinate pair.

Return type ndarray

xy_int

Get the keypoint's xy-coord, rounded to closest integer.

Added in 0.4.0.

Returns A `(2,)` `ndarray` denoting the xy-coordinate pair.

Return type ndarray

y_int

Get the keypoint's y-coordinate, rounded to the closest integer.

Returns **result** – Keypoint's y-coordinate, rounded to the closest integer.

Return type int

class imgaug.augmentables.kps.**KeypointsOnImage**(*keypoints*, *shape*)

Bases: *imgaug.augmentables.base.IAugmentable*

Container for all keypoints on a single image.

Parameters

- **keypoints** (*list of imgaug.augmentables.kps.Keypoint*) – List of keypoints on the image.
- **shape** (*tuple of int or ndarray*) – The shape of the image on which the objects are placed. Either an image with shape `(H,W,[C])` or a `tuple` denoting such an image shape.

Examples

```
>>> import numpy as np
>>> from imgaug.augmentables.kps import Keypoint, KeypointsOnImage
>>>
>>> image = np.zeros((70, 70))
>>> kps = [Keypoint(x=10, y=20), Keypoint(x=34, y=60)]
>>> kps_oi = KeypointsOnImage(kps, shape=image.shape)
```

Attributes

empty Determine whether this object contains zero keypoints.

height Get the image height.

items Get the keypoints in this container.

width Get the image width.

Methods

clip_out_of_image(self)	Remove all KPs that are outside of the image plane.
*clip_out_of_image_*(self)	Remove all KPs that are outside of the image plane.
copy(self[, keypoints, shape])	Create a shallow copy of the `KeypointsOnImage` object.
deepcopy(self[, keypoints, shape])	Create a deep copy of the `KeypointsOnImage` object.
draw_on_image(self, image[, color, alpha, ...])	Draw all keypoints onto a given image.
*fill_from_xy_array_*(self, xy)	Modify the keypoint coordinates of this instance in-place.
from_coords_array(coords, shape)	**Deprecated**.
from_distance_maps(distance_maps[, ...])	Convert outputs of `to_distance_maps()` to `KeypointsOnImage`.
from_keypoint_image(image[, ...])	Convert `to_keypoint_image()` outputs to `KeypointsOnImage`.
from_xy_array(xy, shape)	Convert an `(N,2)` array to a `KeypointsOnImage` object.
get_coords_array(self)	**Deprecated**.
*invert_to_keypoints_on_image_*(self, kpsoi)	Invert the output of `to_keypoints_on_image()` in-place.
on(self, image)	Project all keypoints from one image shape to a new one.
*on_*(self, image)	Project all keypoints from one image shape to a new one in-place.
remove_out_of_image_fraction(self, fraction)	Remove all KPs with an out of image fraction of at least *fraction*.
*remove_out_of_image_fraction_*(self, fraction)	Remove all KPs with an OOI fraction of at least *fraction* in-place.
shift(self[, x, y])	Move the keypoints on the x/y-axis.
*shift_*(self[, x, y])	Move the keypoints on the x/y-axis in-place.
to_distance_maps(self[, inverted])	Generate a `(H,W,N)` array of distance maps for `N` keypoints.
to_keypoint_image(self[, size])	Create an `(H,W,N)` image with keypoint coordinates set to `255`.
to_keypoints_on_image(self)	Convert the keypoints to one `KeypointsOnImage` instance.
to_xy_array(self)	Convert all keypoint coordinates to an array of shape `(N,2)`.

clip_out_of_image(*self*)

Remove all KPs that are outside of the image plane.

This method exists for consistency with other augmentables, e.g. bounding boxes.

Added in 0.4.0.

Returns Keypoints that are inside the image plane.

Return type *imgaug.augmentables.kps.KeypointsOnImage*

clip_out_of_image_(*self*)
Remove all KPs that are outside of the image plane.

This method exists for consistency with other augmentables, e.g. bounding boxes.

Added in 0.4.0.

Returns Keypoints that are inside the image plane. The object may have been modified in-place.

Return type *imgaug.augmentables.kps.KeypointsOnImage*

copy(*self*, *keypoints=None*, *shape=None*)
Create a shallow copy of the `KeypointsOnImage` object.

Parameters

- **keypoints** (*None or list of imgaug.Keypoint, optional*) – List of keypoints on the image. If `None`, the instance's keypoints will be copied.
- **shape** (*tuple of int, optional*) – The shape of the image on which the keypoints are placed. If `None`, the instance's shape will be copied.

Returns Shallow copy.

Return type *imgaug.augmentables.kps.KeypointsOnImage*

deepcopy(*self*, *keypoints=None*, *shape=None*)
Create a deep copy of the `KeypointsOnImage` object.

Parameters

- **keypoints** (*None or list of imgaug.Keypoint, optional*) – List of keypoints on the image. If `None`, the instance's keypoints will be copied.
- **shape** (*tuple of int, optional*) – The shape of the image on which the keypoints are placed. If `None`, the instance's shape will be copied.

Returns Deep copy.

Return type *imgaug.augmentables.kps.KeypointsOnImage*

draw_on_image(*self*, *image*, *color=(0, 255, 0)*, *alpha=1.0*, *size=3*, *copy=True*, *raise_if_out_of_image=False*)
Draw all keypoints onto a given image.

Each keypoint is drawn as a square of provided color and size.

Parameters

- **image** (*(H,W,3) ndarray*) – The image onto which to draw the keypoints. This image should usually have the same shape as set in `KeypointsOnImage.shape`.
- **color** (*int or list of int or tuple of int or (3,) ndarray, optional*) – The RGB color of all keypoints. If a single `int C`, then that is equivalent to `(C,C,C)`.
- **alpha** (*float, optional*) – The opacity of the drawn keypoint, where `1.0` denotes a fully visible keypoint and `0.0` an invisible one.
- **size** (*int, optional*) – The size of each point. If set to `C`, each square will have size `C x C`.
- **copy** (*bool, optional*) – Whether to copy the image before drawing the points.

- **raise_if_out_of_image** (*bool, optional*) – Whether to raise an exception if any keypoint is outside of the image.

Returns Image with drawn keypoints.

Return type (H,W,3) ndarray

empty
Determine whether this object contains zero keypoints.

Returns `True` if this object contains zero keypoints.

Return type bool

fill_from_xy_array_(*self*, *xy*)
Modify the keypoint coordinates of this instance in-place.

Note: This currently expects that *xy* contains exactly as many coordinates as there are keypoints in this instance. Otherwise, an `AssertionError` will be raised.

Added in 0.4.0.

Parameters **xy** (*(N, 2) ndarray or iterable of iterable of number*) – Coordinates of `N` keypoints on an image, given as a `(N,2)` array of xy-coordinates. `N` must match the number of keypoints in this instance.

Returns This instance itself, with updated keypoint coordinates. Note that the instance was modified in-place.

Return type *KeypointsOnImage*

static from_coords_array(*coords*, *shape*)
Deprecated. Use `KeypointsOnImage.from_xy_array()` instead.

Convert an `(N,2)` array to a `KeypointsOnImage` object.

Parameters

coords [(N, 2) ndarray]

Coordinates of `N` keypoints on an image, given as a `(N,2)` array of xy-coordinates.

shape [tuple] The shape of the image on which the keypoints are placed.

Returns

imgaug.augmentables.kps.KeypointsOnImage *KeypointsOnImage* object containing the array's keypoints.

static from_distance_maps(*distance_maps*, *inverted=False*, *if_not_found_coords={'x': -1, 'y': -1}*, *threshold=None*, *nb_channels=None*)
Convert outputs of `to_distance_maps()` to `KeypointsOnImage`.

This is the inverse of *KeypointsOnImage.to_distance_maps()*.

Parameters

- **distance_maps** (*(H,W,N) ndarray*) – The distance maps. `N` is the number of keypoints.
- **inverted** (*bool, optional*) – Whether the given distance maps were generated in inverted mode (i.e. *KeypointsOnImage.to_distance_maps()* was called with `inverted=True`) or in non-inverted mode.

- **if_not_found_coords** (*tuple or list or dict or None, optional*) – Coordinates to use for keypoints that cannot be found in *distance_maps*.
  - If this is a `list`/`tuple`, it must contain two `int` values.
  - If it is a `dict`, it must contain the keys `x` and `y` with each containing one `int` value.
  - If this is `None`, then the keypoint will not be added to the final *KeypointsOnImage* object.
- **threshold** (*float, optional*) – The search for keypoints works by searching for the argmin (non-inverted) or argmax (inverted) in each channel. This parameters contains the maximum (non-inverted) or minimum (inverted) value to accept in order to view a hit as a keypoint. Use `None` to use no min/max.
- **nb_channels** (*None or int, optional*) – Number of channels of the image on which the keypoints are placed. Some keypoint augmenters require that information. If set to `None`, the keypoint's shape will be set to `(height, width)`, otherwise `(height, width, nb_channels)`.

Returns The extracted keypoints.

Return type *imgaug.augmentables.kps.KeypointsOnImage*

static from_keypoint_image(*image*, *if_not_found_coords={'x': -1, 'y': -1}*, *threshold=1*, *nb_channels=None*)

Convert `to_keypoint_image()` outputs to `KeypointsOnImage`.

This is the inverse of *KeypointsOnImage.to_keypoint_image()*.

Parameters

- **image** (*(H,W,N) ndarray*) – The keypoints image. N is the number of keypoints.
- **if_not_found_coords** (*tuple or list or dict or None, optional*) – Coordinates to use for keypoints that cannot be found in *image*.
  - If this is a `list`/`tuple`, it must contain two `int` values.
  - If it is a `dict`, it must contain the keys `x` and `y` with each containing one `int` value.
  - If this is `None`, then the keypoint will not be added to the final *KeypointsOnImage* object.
- **threshold** (*int, optional*) – The search for keypoints works by searching for the argmax in each channel. This parameters contains the minimum value that the max must have in order to be viewed as a keypoint.
- **nb_channels** (*None or int, optional*) – Number of channels of the image on which the keypoints are placed. Some keypoint augmenters require that information. If set to `None`, the keypoint's shape will be set to `(height, width)`, otherwise `(height, width, nb_channels)`.

Returns The extracted keypoints.

Return type *imgaug.augmentables.kps.KeypointsOnImage*

classmethod from_xy_array(*xy*, *shape*)

Convert an `(N,2)` array to a `KeypointsOnImage` object.

Parameters

- **xy** (*(N, 2) ndarray or iterable of iterable of number*) – Coordinates of `N` keypoints on an image, given as a `(N,2)` array of xy-coordinates.
- **shape** (*tuple of int or ndarray*) – The shape of the image on which the keypoints are placed.

Returns `KeypointsOnImage` object containing the array's keypoints.

Return type *imgaug.augmentables.kps.KeypointsOnImage*

get_coords_array(*self*)
Deprecated. Use `KeypointsOnImage.to_xy_array()` instead.

Convert all keypoint coordinates to an array of shape `(N,2)`.

Returns

(N, 2) ndarray Array containing the coordinates of all keypoints. `N` denotes the number of keypoints. The second axis denotes the x/y-coordinates.

height
Get the image height.

Returns Image height.

Return type int

invert_to_keypoints_on_image_(*self*, *kpsoi*)
Invert the output of `to_keypoints_on_image()` in-place.

This function writes in-place into this `KeypointsOnImage` instance.

Added in 0.4.0.

Parameters **kpsoi** (*imgaug.augmentables.kps.KeypointsOnImages*) – Keypoints to copy data from, i.e. the outputs of `to_keypoints_on_image()`.

Returns Keypoints container with updated coordinates. Note that the instance is also updated in-place.

Return type *KeypointsOnImage*

items
Get the keypoints in this container.

Added in 0.4.0.

Returns Keypoints within this container.

Return type list of Keypoint

on(*self*, *image*)
Project all keypoints from one image shape to a new one.

Parameters **image** (*ndarray or tuple of int*) – New image onto which the keypoints are to be projected. May also simply be that new image's shape tuple.

Returns Object containing all projected keypoints.

Return type *imgaug.augmentables.kps.KeypointsOnImage*

on_(*self*, *image*)
Project all keypoints from one image shape to a new one in-place.

Added in 0.4.0.

Parameters **image** (*ndarray or tuple of int*) – New image onto which the keypoints are to be projected. May also simply be that new image's shape tuple.

Returns Object containing all projected keypoints. The object may have been modified in-place.

Return type *imgaug.augmentables.kps.KeypointsOnImage*

remove_out_of_image_fraction (*self*, *fraction*)
Remove all KPs with an out of image fraction of at least *fraction*.

This method exists for consistency with other augmentables, e.g. bounding boxes.

Added in 0.4.0.

Parameters **fraction** (*number*) – Minimum out of image fraction that a keypoint has to have in order to be removed. Note that any keypoint can only have a fraction of either `1.0` (is outside) or `0.0` (is inside). Set this to `0.0+eps` to remove all points that are outside of the image. Setting this to `0.0` will remove all points.

Returns Reduced set of keypoints, with those thathad an out of image fraction greater or equal the given one removed.

Return type *imgaug.augmentables.kps.KeypointsOnImage*

remove_out_of_image_fraction_ (*self*, *fraction*)
Remove all KPs with an OOI fraction of at least *fraction* in-place.

'OOI' is the abbreviation for 'out of image'.

This method exists for consistency with other augmentables, e.g. bounding boxes.

Added in 0.4.0.

Parameters **fraction** (*number*) – Minimum out of image fraction that a keypoint has to have in order to be removed. Note that any keypoint can only have a fraction of either `1.0` (is outside) or `0.0` (is inside). Set this to `0.0+eps` to remove all points that are outside of the image. Setting this to `0.0` will remove all points.

Returns Reduced set of keypoints, with those thathad an out of image fraction greater or equal the given one removed. The object may have been modified in-place.

Return type *imgaug.augmentables.kps.KeypointsOnImage*

shift (*self*, *x=0*, *y=0*)
Move the keypoints on the x/y-axis.

Parameters

- **x** (*number, optional*) – Move each keypoint by this value on the x axis.
- **y** (*number, optional*) – Move each keypoint by this value on the y axis.

Returns Keypoints after moving them.

Return type *imgaug.augmentables.kps.KeypointsOnImage*

shift_ (*self*, *x=0*, *y=0*)
Move the keypoints on the x/y-axis in-place.

Added in 0.4.0.

Parameters

- **x** (*number, optional*) – Move each keypoint by this value on the x axis.
- **y** (*number, optional*) – Move each keypoint by this value on the y axis.

Returns Keypoints after moving them. The object and its items may have been modified in-place.

Return type *imgaug.augmentables.kps.KeypointsOnImage*

to_distance_maps(*self*, *inverted=False*)
Generate a `(H,W,N)` array of distance maps for `N` keypoints.

The `n`-th distance map contains at every location `(y, x)` the euclidean distance to the `n`-th keypoint.

This function can be used as a helper when augmenting keypoints with a method that only supports the augmentation of images.

Parameters **inverted** (*bool, optional*) – If `True`, inverted distance maps are returned where each distance value d is replaced by `d/(d+1)`, i.e. the distance maps have values in the range `(0.0, 1.0]` with `1.0` denoting exactly the position of the respective keypoint.

Returns A `float32` array containing `N` distance maps for `N` keypoints. Each location `(y, x, n)` in the array denotes the euclidean distance at `(y, x)` to the `n`-th keypoint. If *inverted* is `True`, the distance `d` is replaced by `d/(d+1)`. The height and width of the array match the height and width in `KeypointsOnImage.shape`.

Return type (H,W,N) ndarray

to_keypoint_image(*self*, *size=1*)
Create an `(H,W,N)` image with keypoint coordinates set to `255`.

This method generates a new `uint8` array of shape `(H,W,N)`, where `H` is the `.shape` height, `W` the `.shape` width and `N` is the number of keypoints. The array is filled with zeros. The coordinate of the `n`-th keypoint is set to `255` in the `n`-th channel.

This function can be used as a helper when augmenting keypoints with a method that only supports the augmentation of images.

Parameters **size** (*int*) – Size of each (squared) point.

Returns Image in which the keypoints are marked. `H` is the height, defined in `KeypointsOnImage.shape[0]` (analogous `W`). `N` is the number of keypoints.

Return type (H,W,N) ndarray

to_keypoints_on_image(*self*)
Convert the keypoints to one `KeypointsOnImage` instance.

This method exists for consistency with `BoundingBoxesOnImage`, `PolygonsOnImage` and `LineStringsOnImage`.

Added in 0.4.0.

Returns Copy of this keypoints instance.

Return type *imgaug.augmentables.kps.KeypointsOnImage*

to_xy_array(*self*)
Convert all keypoint coordinates to an array of shape `(N,2)`.

Returns Array containing the coordinates of all keypoints. `N` denotes the number of keypoints. The second axis denotes the x/y-coordinates.

Return type (N, 2) ndarray

width
Get the image width.

Returns Image width.

Return type int

imgaug.augmentables.kps.**compute_geometric_median**(*points=None*, *eps=1e-05*, *X=None*)
Estimate the geometric median of points in 2D.

Code from https://stackoverflow.com/a/30305181

Parameters

- **points** (*(N,2) ndarray*) – Points in 2D. Second axis must be given in xy-form.
- **eps** (*float, optional*) – Distance threshold when to return the median.
- **X** (*None or (N,2) ndarray, optional*) – Deprecated.

Returns Geometric median as xy-coordinate.

Return type (2,) ndarray

13.12 imgaug.augmentables.lines

Classes representing lines.

class imgaug.augmentables.lines.**LineString**(*coords*, *label=None*)
Bases: object

Class representing line strings.

A line string is a collection of connected line segments, each having a start and end point. Each point is given as its `(x, y)` absolute (sub-)pixel coordinates. The end point of each segment is also the start point of the next segment.

The line string is not closed, i.e. start and end point are expected to differ and will not be connected in drawings.

Parameters

- **coords** (*iterable of tuple of number or ndarray*) – The points of the line string.
- **label** (*None or str, optional*) – The label of the line string.

Attributes

height Compute the height of a bounding box encapsulating the line.

length Compute the total euclidean length of the line string.

width Compute the width of a bounding box encapsulating the line.

xx Get an array of x-coordinates of all points of the line string.

xx_int Get an array of discrete x-coordinates of all points.

yy Get an array of y-coordinates of all points of the line string.

yy_int Get an array of discrete y-coordinates of all points.

Methods

almost_equals(self, other[, max_distance, ...])	Compare this and another line string.
clip_out_of_image(self, image)	Clip off all parts of the line string that are outside of the image.
compute_distance(self, other[, default])	Compute the minimal distance between the line string and *other*.
compute_neighbour_distances(self)	Compute the euclidean distance between each two consecutive points.

Continued on next page

Table 45 – continued from previous page

compute_out_of_image_fraction(self, image)	Compute fraction of polygon area outside of the image plane.
compute_pointwise_distances(self, other[, ...])	Compute min distances between points of this and another line string.
concatenate(self, other)	Concatenate this line string with another one.
contains(self, other[, max_distance])	Estimate whether a point is on this line string.
coords_almost_equals(self, other[, ...])	Compare this and another LineString's coordinates.
copy(self[, coords, label])	Create a shallow copy of this line string.
deepcopy(self[, coords, label])	Create a deep copy of this line string.
draw_heatmap_array(self, image_shape[, ...])	Draw the line segments and points of the line string as a heatmap array.
draw_lines_heatmap_array(self, image_shape)	Draw the line segments of this line string as a heatmap array.
draw_lines_on_image(self, image[, color, ...])	Draw the line segments of this line string on a given image.
draw_mask(self, image_shape[, size_lines, ...])	Draw this line segment as a binary image mask.
draw_on_image(self, image[, color, ...])	Draw this line string onto an image.
draw_points_heatmap_array(self, image_shape)	Draw the points of this line string as a heatmap array.
draw_points_on_image(self, image[, color, ...])	Draw the points of this line string onto a given image.
extract_from_image(self, image[, size, pad, ...])	Extract all image pixels covered by the line string.
find_intersections_with(self, other)	Find all intersection points between this line string and *other*.
get_pointwise_inside_image_mask(self, image)	Determine per point whether it is inside of a given image plane.
is_fully_within_image(self, image[, default])	Estimate whether the line string is fully inside an image plane.
is_out_of_image(self, image[, fully, ...])	Estimate whether the line is partially/fully outside of the image area.
is_partly_within_image(self, image[, default])	Estimate whether the line string is at least partially inside the image.
project(self, from_shape, to_shape)	Project the line string onto a differently shaped image.
*project_*(self, from_shape, to_shape)	Project the line string onto a differently shaped image in-place.
shift(self[, x, y, top, right, bottom, left])	Move this line string along the x/y-axis.
*shift_*(self[, x, y])	Move this line string along the x/y-axis in-place.
subdivide(self, points_per_edge)	Derive a new line string with `N` interpolated points per edge.
to_bounding_box(self)	Generate a bounding box encapsulating the line string.
to_heatmap(self, image_shape[, size_lines, ...])	Generate a heatmap object from the line string.
to_keypoints(self)	Convert the line string points to keypoints.
to_polygon(self)	Generate a polygon from the line string points.
to_segmentation_map(self, image_shape[, ...])	Generate a segmentation map object from the line string.

almost_equals(*self*, *other*, *max_distance=0.0001*, *points_per_edge=8*)
Compare this and another line string.

Parameters

- **other** (*imgaug.augmentables.lines.LineString*) – The other object to compare against. Expected to be a `LineString`.
- **max_distance** (*float, optional*) – See *coords_almost_equals()*.
- **points_per_edge** (*int, optional*) – See *coords_almost_equals()*.

Returns `True` if the coordinates are almost equal and additionally the labels are equal. Otherwise `False`.

Return type bool

clip_out_of_image(*self*, *image*)

Clip off all parts of the line string that are outside of the image.

Parameters **image** (*ndarray or tuple of int*) – Either an image with shape `(H,W,[C])` or a `tuple` denoting such an image shape.

Returns Line strings, clipped to the image shape. The result may contain any number of line strins, including zero.

Return type list of imgaug.augmentables.lines.LineString

compute_distance(*self*, *other*, *default=None*)

Compute the minimal distance between the line string and *other*.

Parameters

- **other** (*tuple of number or imgaug.augmentables.kps.Keypoint or imgaug.augmentables.LineString*) – Other object to which to compute the distance.
- **default** (*any*) – Value to return if this line string or *other* contain no points.

Returns Minimal distance to *other* or *default* if no distance could be computed.

Return type float or any

compute_neighbour_distances(*self*)

Compute the euclidean distance between each two consecutive points.

Returns `(N-1,) float32` array of euclidean distances between point pairs. Same order as in *coords*.

Return type ndarray

compute_out_of_image_fraction(*self*, *image*)

Compute fraction of polygon area outside of the image plane.

This estimates `f = A_ooi / A`, where `A_ooi` is the area of the polygon that is outside of the image plane, while `A` is the total area of the bounding box.

Added in 0.4.0.

Parameters **image** (*(H,W,...) ndarray or tuple of int*) – Image dimensions to use. If an `ndarray`, its shape will be used. If a `tuple`, it is assumed to represent the image shape and must contain at least two integers.

Returns Fraction of the polygon area that is outside of the image plane. Returns `0.0` if the polygon is fully inside of the image plane. If the polygon has an area of zero, the polygon is treated similarly to a *LineString*, i.e. the fraction of the line that is inside the image plane is returned.

Return type float

compute_pointwise_distances(*self*, *other*, *default=None*)
Compute min distances between points of this and another line string.

Parameters

- **other** (*tuple of number or imgaug.augmentables.kps.Keypoint or imgaug.augmentables.LineString*) – Other object to which to compute the distances.
- **default** (*any*) – Value to return if *other* contains no points.

Returns For each coordinate of this line string, the distance to any closest location on *other*. *default* if no distance could be computed.

Return type list of float or any

concatenate(*self*, *other*)
Concatenate this line string with another one.

This will add a line segment between the end point of this line string and the start point of *other*.

Parameters **other** (*imgaug.augmentables.lines.LineString or ndarray or iterable of tuple of number*) – The points to add to this line string.

Returns New line string with concatenated points. The *label* of this line string will be kept.

Return type *imgaug.augmentables.lines.LineString*

contains(*self*, *other*, *max_distance=0.0001*)
Estimate whether a point is on this line string.

This method uses a maximum distance to estimate whether a point is on a line string.

Parameters

- **other** (*tuple of number or imgaug.augmentables.kps.Keypoint*) – Point to check for.
- **max_distance** (*float*) – Maximum allowed euclidean distance between the point and the closest point on the line. If the threshold is exceeded, the point is not considered to fall on the line.

Returns `True` if the point is on the line string, `False` otherwise.

Return type bool

coords_almost_equals(*self*, *other*, *max_distance=0.0001*, *points_per_edge=8*)
Compare this and another LineString's coordinates.

This is an approximate method based on pointwise distances and can in rare corner cases produce wrong outputs.

Parameters

- **other** (*imgaug.augmentables.lines.LineString or tuple of number or ndarray or list of ndarray or list of tuple of number*) – The other line string or its coordinates.
- **max_distance** (*float, optional*) – Max distance of any point from the other line string before the two line strings are evaluated to be unequal.
- **points_per_edge** (*int, optional*) – How many points to interpolate on each edge.

Returns Whether the two LineString's coordinates are almost identical, i.e. the max distance is below the threshold. If both have no coordinates, `True` is returned. If only one has no coordinates, `False` is returned. Beyond that, the number of points is not evaluated.

Return type bool

copy(*self*, *coords=None*, *label=None*)
Create a shallow copy of this line string.

Parameters

- **coords** (*None or iterable of tuple of number or ndarray*) – If not `None`, then the coords of the copied object will be set to this value.
- **label** (*None or str*) – If not `None`, then the label of the copied object will be set to this value.

Returns Shallow copy.

Return type *imgaug.augmentables.lines.LineString*

deepcopy(*self*, *coords=None*, *label=None*)
Create a deep copy of this line string.

Parameters

- **coords** (*None or iterable of tuple of number or ndarray*) – If not `None`, then the coords of the copied object will be set to this value.
- **label** (*None or str*) – If not `None`, then the label of the copied object will be set to this value.

Returns Deep copy.

Return type *imgaug.augmentables.lines.LineString*

draw_heatmap_array(*self*, *image_shape*, *alpha_lines=1.0*, *alpha_points=1.0*, *size_lines=1*, *size_points=0*, *antialiased=True*, *raise_if_out_of_image=False*)
Draw the line segments and points of the line string as a heatmap array.

Parameters

- **image_shape** (*tuple of int*) – The shape of the image onto which to draw the line mask.
- **alpha_lines** (*float, optional*) – Opacity of the line string. Higher values denote a more visible line string.
- **alpha_points** (*float, optional*) – Opacity of the line string points. Higher values denote a more visible points.
- **size_lines** (*int, optional*) – Thickness of the line segments.
- **size_points** (*int, optional*) – Size of the points in pixels.
- **antialiased** (*bool, optional*) – Whether to draw the line with anti-aliasing activated.
- **raise_if_out_of_image** (*bool, optional*) – Whether to raise an error if the line string is fully outside of the image. If set to `False`, no error will be raised and only the parts inside the image will be drawn.

Returns `float32` array of shape *image_shape* (no channel axis) with drawn line segments and points. All values are in the interval `[0.0, 1.0]`.

Return type ndarray

draw_lines_heatmap_array(*self*, *image_shape*, *alpha=1.0*, *size=1*, *antialiased=True*, *raise_if_out_of_image=False*)
Draw the line segments of this line string as a heatmap array.

Parameters

- **image_shape** (*tuple of int*) – The shape of the image onto which to draw the line mask.

- **alpha** (*float, optional*) – Opacity of the line string. Higher values denote a more visible line string.
- **size** (*int, optional*) – Thickness of the line segments.
- **antialiased** (*bool, optional*) – Whether to draw the line with anti-aliasing activated.
- **raise_if_out_of_image** (*bool, optional*) – Whether to raise an error if the line string is fully outside of the image. If set to `False`, no error will be raised and only the parts inside the image will be drawn.

Returns `float32` array of shape *image_shape* (no channel axis) with drawn line string. All values are in the interval `[0.0, 1.0]`.

Return type ndarray

draw_lines_on_image(*self*, *image*, *color=(0, 255, 0)*, *alpha=1.0*, *size=3*, *antialiased=True*, *raise_if_out_of_image=False*)

Draw the line segments of this line string on a given image.

Parameters

- **image** (*ndarray or tuple of int*) – The image onto which to draw. Expected to be `uint8` and of shape `(H, W, C)` with `C` usually being `3` (other values are not tested). If a tuple, expected to be `(H, W, C)` and will lead to a new `uint8` array of zeros being created.
- **color** (*int or iterable of int*) – Color to use as RGB, i.e. three values.
- **alpha** (*float, optional*) – Opacity of the line string. Higher values denote a more visible line string.
- **size** (*int, optional*) – Thickness of the line segments.
- **antialiased** (*bool, optional*) – Whether to draw the line with anti-aliasing activated.
- **raise_if_out_of_image** (*bool, optional*) – Whether to raise an error if the line string is fully outside of the image. If set to `False`, no error will be raised and only the parts inside the image will be drawn.

Returns *image* with line drawn on it.

Return type ndarray

draw_mask(*self*, *image_shape*, *size_lines=1*, *size_points=0*, *raise_if_out_of_image=False*)

Draw this line segment as a binary image mask.

Parameters

- **image_shape** (*tuple of int*) – The shape of the image onto which to draw the line mask.
- **size_lines** (*int, optional*) – Thickness of the line segments.
- **size_points** (*int, optional*) – Size of the points in pixels.
- **raise_if_out_of_image** (*bool, optional*) – Whether to raise an error if the line string is fully outside of the image. If set to `False`, no error will be raised and only the parts inside the image will be drawn.

Returns Boolean line mask of shape *image_shape* (no channel axis).

Return type ndarray

draw_on_image(*self*, *image*, *color=(0, 255, 0)*, *color_lines=None*, *color_points=None*, *alpha=1.0*, *alpha_lines=None*, *alpha_points=None*, *size=1*, *size_lines=None*, *size_points=None*, *antialiased=True*, *raise_if_out_of_image=False*)

Draw this line string onto an image.

Parameters

- **image** (*ndarray*) – The *(H,W,C) uint8* image onto which to draw the line string.
- **color** (*iterable of int, optional*) – Color to use as RGB, i.e. three values. The color of the line and points are derived from this value, unless they are set.
- **color_lines** (*None or iterable of int*) – Color to use for the line segments as RGB, i.e. three values. If `None`, this value is derived from *color*.
- **color_points** (*None or iterable of int*) – Color to use for the points as RGB, i.e. three values. If `None`, this value is derived from `0.5 * color`.
- **alpha** (*float, optional*) – Opacity of the line string. Higher values denote more visible points. The alphas of the line and points are derived from this value, unless they are set.
- **alpha_lines** (*None or float, optional*) – Opacity of the line string. Higher values denote more visible line string. If `None`, this value is derived from *alpha*.
- **alpha_points** (*None or float, optional*) – Opacity of the line string points. Higher values denote more visible points. If `None`, this value is derived from *alpha*.
- **size** (*int, optional*) – Size of the line string. The sizes of the line and points are derived from this value, unless they are set.
- **size_lines** (*None or int, optional*) – Thickness of the line segments. If `None`, this value is derived from *size*.
- **size_points** (*None or int, optional*) – Size of the points in pixels. If `None`, this value is derived from `3 * size`.
- **antialiased** (*bool, optional*) – Whether to draw the line with anti-aliasing activated. This does currently not affect the point drawing.
- **raise_if_out_of_image** (*bool, optional*) – Whether to raise an error if the line string is fully outside of the image. If set to `False`, no error will be raised and only the parts inside the image will be drawn.

Returns Image with line string drawn on it.

Return type ndarray

draw_points_heatmap_array(*self*, *image_shape*, *alpha=1.0*, *size=1*, *raise_if_out_of_image=False*)

Draw the points of this line string as a heatmap array.

Parameters

- **image_shape** (*tuple of int*) – The shape of the image onto which to draw the point mask.
- **alpha** (*float, optional*) – Opacity of the line string points. Higher values denote a more visible points.
- **size** (*int, optional*) – Size of the points in pixels.
- **raise_if_out_of_image** (*bool, optional*) – Whether to raise an error if the line string is fully outside of the image. If set to `False`, no error will be raised and only the parts inside the image will be drawn.

Returns `float32` array of shape *image_shape* (no channel axis) with drawn line string points. All values are in the interval `[0.0, 1.0]`.

Return type ndarray

draw_points_on_image(*self*, *image*, *color=(0, 128, 0)*, *alpha=1.0*, *size=3*, *copy=True*, *raise_if_out_of_image=False*)

Draw the points of this line string onto a given image.

Parameters

- **image** (*ndarray or tuple of int*) – The image onto which to draw. Expected to be `uint8` and of shape `(H, W, C)` with `C` usually being `3` (other values are not tested). If a tuple, expected to be `(H, W, C)` and will lead to a new `uint8` array of zeros being created.
- **color** (*iterable of int*) – Color to use as RGB, i.e. three values.
- **alpha** (*float, optional*) – Opacity of the line string points. Higher values denote a more visible points.
- **size** (*int, optional*) – Size of the points in pixels.
- **copy** (*bool, optional*) – Whether it is allowed to draw directly in the input array (`False`) or it has to be copied (`True`). The routine may still have to copy, even if `copy=False` was used. Always use the return value.
- **raise_if_out_of_image** (*bool, optional*) – Whether to raise an error if the line string is fully outside of the image. If set to `False`, no error will be raised and only the parts inside the image will be drawn.

Returns `float32` array of shape *image_shape* (no channel axis) with drawn line string points. All values are in the interval `[0.0, 1.0]`.

Return type ndarray

extract_from_image(*self*, *image*, *size=1*, *pad=True*, *pad_max=None*, *antialiased=True*, *prevent_zero_size=True*)

Extract all image pixels covered by the line string.

This will only extract pixels overlapping with the line string. As a rectangular image array has to be returned, non-overlapping pixels will be set to zero.

This function will by default zero-pad the image if the line string is partially/fully outside of the image. This is for consistency with the same methods for bounding boxes and polygons.

Parameters

- **image** (*ndarray*) – The image of shape *(H,W,[C])* from which to extract the pixels within the line string.
- **size** (*int, optional*) – Thickness of the line.
- **pad** (*bool, optional*) – Whether to zero-pad the image if the object is partially/fully outside of it.
- **pad_max** (*None or int, optional*) – The maximum number of pixels that may be zero-paded on any side, i.e. if this has value `N` the total maximum of added pixels is `4*N`. This option exists to prevent extremely large images as a result of single points being moved very far away during augmentation.
- **antialiased** (*bool, optional*) – Whether to apply anti-aliasing to the line string.
- **prevent_zero_size** (*bool, optional*) – Whether to prevent height or width of the extracted image from becoming zero. If this is set to `True` and height or width of the line string is below `1`, the height/width will be increased to `1`. This can be useful to prevent problems, e.g. with image saving or plotting. If it is set to `False`, images will be returned as `(H', W')` or `(H', W', 3)` with `H` or `W` potentially being `0`.

Returns Pixels overlapping with the line string. Zero-padded if the line string is partially/fully outside of the image and `pad=True`. If *prevent_zero_size* is activated, it is guarantueed that `H'>0` and `W'>0`, otherwise only `H'>=0` and `W'>=0`.

Return type (H',W') ndarray or (H',W',C) ndarray

find_intersections_with(*self*, *other*)
Find all intersection points between this line string and *other*.

Parameters **other** (*tuple of number or list of tuple of number or list of LineString or LineString*) – The other geometry to use during intersection tests.

Returns All intersection points. One list per pair of consecutive start and end point, i.e. *N-1* lists of *N* points. Each list may be empty or may contain multiple points.

Return type list of list of tuple of number

get_pointwise_inside_image_mask(*self*, *image*)
Determine per point whether it is inside of a given image plane.

Parameters **image** (*ndarray or tuple of int*) – Either an image with shape `(H,W,[C])` or a `tuple` denoting such an image shape.

Returns `(N,)` ``bool` array with one value for each of the `N` points indicating whether it is inside of the provided image plane (`True`) or not (`False`).

Return type ndarray

height
Compute the height of a bounding box encapsulating the line.

The height is computed based on the two points with lowest and largest y-coordinates.

Returns The height of the line string.

Return type float

is_fully_within_image(*self*, *image*, *default=False*)
Estimate whether the line string is fully inside an image plane.

Parameters

- **image** (*ndarray or tuple of int*) – Either an image with shape `(H,W,[C])` or a `tuple` denoting such an image shape.
- **default** (*any*) – Default value to return if the line string contains no points.

Returns `True` if the line string is fully inside the image area. `False` otherwise. Will return *default* if this line string contains no points.

Return type bool or any

is_out_of_image(*self*, *image*, *fully=True*, *partly=False*, *default=True*)
Estimate whether the line is partially/fully outside of the image area.

Parameters

- **image** (*ndarray or tuple of int*) – Either an image with shape `(H,W,[C])` or a tuple denoting such an image shape.
- **fully** (*bool, optional*) – Whether to return `True` if the line string is fully outside of the image area.
- **partly** (*bool, optional*) – Whether to return `True` if the line string is at least partially outside fo the image area.

- **default** (*any*) – Default value to return if the line string contains no points.

Returns `True` if the line string is partially/fully outside of the image area, depending on defined parameters. `False` otherwise. Will return *default* if this line string contains no points.

Return type bool or any

is_partly_within_image(*self*, *image*, *default=False*)
Estimate whether the line string is at least partially inside the image.

Parameters

- **image** (*ndarray or tuple of int*) – Either an image with shape `(H,W,[C])` or a `tuple` denoting such an image shape.
- **default** (*any*) – Default value to return if the line string contains no points.

Returns `True` if the line string is at least partially inside the image area. `False` otherwise. Will return *default* if this line string contains no points.

Return type bool or any

length
Compute the total euclidean length of the line string.

Returns The length based on euclidean distance, i.e. the sum of the lengths of each line segment.

Return type float

project(*self*, *from_shape*, *to_shape*)
Project the line string onto a differently shaped image.

E.g. if a point of the line string is on its original image at `x=(10 of 100 pixels)` and `y=(20 of 100 pixels)` and is projected onto a new image with size `(width=200, height=200)`, its new position will be `(x=20, y=40)`.

This is intended for cases where the original image is resized. It cannot be used for more complex changes (e.g. padding, cropping).

Parameters

- **from_shape** (*tuple of int or ndarray*) – Shape of the original image. (Before resize.)
- **to_shape** (*tuple of int or ndarray*) – Shape of the new image. (After resize.)

Returns Line string with new coordinates.

Return type *imgaug.augmentables.lines.LineString*

project_(*self*, *from_shape*, *to_shape*)
Project the line string onto a differently shaped image in-place.

E.g. if a point of the line string is on its original image at `x=(10 of 100 pixels)` and `y=(20 of 100 pixels)` and is projected onto a new image with size `(width=200, height=200)`, its new position will be `(x=20, y=40)`.

This is intended for cases where the original image is resized. It cannot be used for more complex changes (e.g. padding, cropping).

Added in 0.4.0.

Parameters

- **from_shape** (*tuple of int or ndarray*) – Shape of the original image. (Before resize.)
- **to_shape** (*tuple of int or ndarray*) – Shape of the new image. (After resize.)

Returns Line string with new coordinates. The object may have been modified in-place.

Return type *imgaug.augmentables.lines.LineString*

shift(*self*, *x=0*, *y=0*, *top=None*, *right=None*, *bottom=None*, *left=None*)
Move this line string along the x/y-axis.

The origin `(0, 0)` is at the top left of the image.

Parameters

- **x** (*number, optional*) – Value to be added to all x-coordinates. Positive values shift towards the right images.
- **y** (*number, optional*) – Value to be added to all y-coordinates. Positive values shift towards the bottom images.
- **top** (*None or int, optional*) – Deprecated since 0.4.0. Amount of pixels by which to shift this object *from* the top (towards the bottom).
- **right** (*None or int, optional*) – Deprecated since 0.4.0. Amount of pixels by which to shift this object *from* the right (towards the left).
- **bottom** (*None or int, optional*) – Deprecated since 0.4.0. Amount of pixels by which to shift this object *from* the bottom (towards the top).
- **left** (*None or int, optional*) – Deprecated since 0.4.0. Amount of pixels by which to shift this object *from* the left (towards the right).

Returns **result** – Shifted line string.

Return type *imgaug.augmentables.lines.LineString*

shift_(*self*, *x=0*, *y=0*)
Move this line string along the x/y-axis in-place.

The origin `(0, 0)` is at the top left of the image.

Added in 0.4.0.

Parameters

- **x** (*number, optional*) – Value to be added to all x-coordinates. Positive values shift towards the right images.
- **y** (*number, optional*) – Value to be added to all y-coordinates. Positive values shift towards the bottom images.

Returns **result** – Shifted line string. The object may have been modified in-place.

Return type *imgaug.augmentables.lines.LineString*

subdivide(*self*, *points_per_edge*)
Derive a new line string with `N` interpolated points per edge.

The interpolated points have (per edge) regular distances to each other.

For each edge between points `A` and `B` this adds points at `A + (i/(1+N)) * (B - A)`, where `i` is the index of the added point and `N` is the number of points to add per edge.

Calling this method two times will split each edge at its center and then again split each newly created edge at their center. It is equivalent to calling *subdivide(3)*.

Parameters **points_per_edge** (*int*) – Number of points to interpolate on each edge.

Returns Line string with subdivided edges.

Return type *imgaug.augmentables.lines.LineString*

to_bounding_box(*self*)
Generate a bounding box encapsulating the line string.

Returns Bounding box encapsulating the line string. `None` if the line string contained no points.

Return type None or *imgaug.augmentables.bbs.BoundingBox*

to_heatmap(*self*, *image_shape*, *size_lines=1*, *size_points=0*, *antialiased=True*, *raise_if_out_of_image=False*)
Generate a heatmap object from the line string.

This is similar to *draw_lines_heatmap_array()*, executed with `alpha=1.0`. The result is wrapped in a *HeatmapsOnImage* object instead of just an array. No points are drawn.

Parameters

- **image_shape** (*tuple of int*) – The shape of the image onto which to draw the line mask.
- **size_lines** (*int, optional*) – Thickness of the line.
- **size_points** (*int, optional*) – Size of the points in pixels.
- **antialiased** (*bool, optional*) – Whether to draw the line with anti-aliasing activated.
- **raise_if_out_of_image** (*bool, optional*) – Whether to raise an error if the line string is fully outside of the image. If set to `False`, no error will be raised and only the parts inside the image will be drawn.

Returns Heatmap object containing drawn line string.

Return type *imgaug.augmentables.heatmaps.HeatmapsOnImage*

to_keypoints(*self*)
Convert the line string points to keypoints.

Returns Points of the line string as keypoints.

Return type list of imgaug.augmentables.kps.Keypoint

to_polygon(*self*)
Generate a polygon from the line string points.

Returns Polygon with the same corner points as the line string. Note that the polygon might be invalid, e.g. contain less than `3` points or have self-intersections.

Return type *imgaug.augmentables.polys.Polygon*

to_segmentation_map(*self*, *image_shape*, *size_lines=1*, *size_points=0*, *raise_if_out_of_image=False*)
Generate a segmentation map object from the line string.

This is similar to *draw_mask()*. The result is wrapped in a `SegmentationMapsOnImage` object instead of just an array.

Parameters

- **image_shape** (*tuple of int*) – The shape of the image onto which to draw the line mask.
- **size_lines** (*int, optional*) – Thickness of the line.
- **size_points** (*int, optional*) – Size of the points in pixels.
- **raise_if_out_of_image** (*bool, optional*) – Whether to raise an error if the line string is fully outside of the image. If set to `False`, no error will be raised and only the parts inside the image will be drawn.

Returns Segmentation map object containing drawn line string.

Return type *imgaug.augmentables.segmaps.SegmentationMapsOnImage*

width
Compute the width of a bounding box encapsulating the line.

The width is computed based on the two points with lowest and largest x-coordinates.

Returns The width of the line string.

Return type float

xx
Get an array of x-coordinates of all points of the line string.

Returns `float32` x-coordinates of the line string points.

Return type ndarray

xx_int
Get an array of discrete x-coordinates of all points.

The conversion from `float32` coordinates to `int32` is done by first rounding the coordinates to the closest integer and then removing everything after the decimal point.

Returns `int32` x-coordinates of the line string points.

Return type ndarray

yy
Get an array of y-coordinates of all points of the line string.

Returns `float32` y-coordinates of the line string points.

Return type ndarray

yy_int
Get an array of discrete y-coordinates of all points.

The conversion from `float32` coordinates to `int32` is done by first rounding the coordinates to the closest integer and then removing everything after the decimal point.

Returns `int32` y-coordinates of the line string points.

Return type ndarray

class imgaug.augmentables.lines.**LineStringsOnImage**(*line_strings*, *shape*)
Bases: *imgaug.augmentables.base.IAugmentable*

Object that represents all line strings on a single image.

Parameters

- **line_strings** (*list of imgaug.augmentables.lines.LineString*) – List of line strings on the image.
- **shape** (*tuple of int or ndarray*) – The shape of the image on which the objects are placed. Either an image with shape `(H,W,[C])` or a `tuple` denoting such an image shape.

Examples

```
>>> import numpy as np
>>> from imgaug.augmentables.lines import LineString, LineStringsOnImage
>>>
>>> image = np.zeros((100, 100))
>>> lss = [
>>>     LineString([(0, 0), (10, 0)]),
>>>     LineString([(10, 20), (30, 30), (50, 70)])
>>> ]
>>> lsoi = LineStringsOnImage(lss, shape=image.shape)
```

Attributes

empty	Estimate whether this object contains zero line strings.
items	Get the line strings in this container.

Methods

clip_out_of_image(self)	Clip off all parts of the line strings that are outside of an image.
*clip_out_of_image_*(self)	Clip off all parts of the LSs that are outside of an image in-place.
copy(self[, line_strings, shape])	Create a shallow copy of this object.
deepcopy(self[, line_strings, shape])	Create a deep copy of the object.
draw_on_image(self, image[, color, ...])	Draw all line strings onto a given image.
*fill_from_xy_array_*(self, xy)	Modify the corner coordinates of all line strings in-place.
from_xy_arrays(xy, shape)	Convert an `(N,M,2)` ndarray to a `LineStringsOnImage` object.
*invert_to_keypoints_on_image_*(self, kpsoi)	Invert the output of `to_keypoints_on_image()` in-place.
on(self, image)	Project the line strings from one image shape to a new one.
*on_*(self, image)	Project the line strings from one image shape to a new one in-place.
remove_out_of_image(self[, fully, partly])	Remove all line strings that are fully/partially outside of an image.
*remove_out_of_image_*(self[, fully, partly])	Remove all LS that are fully/partially outside of an image in-place.
remove_out_of_image_fraction(self, fraction)	Remove all LS with an out of image fraction of at least *fraction*.
*remove_out_of_image_fraction_*(self, fraction)	Remove all LS with an OOI fraction of at least *fraction* in-place.
shift(self[, x, y, top, right, bottom, left])	Move the line strings along the x/y-axis.
*shift_*(self[, x, y])	Move the line strings along the x/y-axis in-place.
to_keypoints_on_image(self)	Convert the line strings to one `KeypointsOnImage` instance.
to_xy_array(self)	Convert all line string coordinates to one array of shape `(N,2)`.
to_xy_arrays(self[, dtype])	Convert this object to an iterable of `(M,2)` arrays of points.

clip_out_of_image(*self*)
Clip off all parts of the line strings that are outside of an image.

Note: The result can contain fewer line strings than the input did. That happens when a polygon is fully outside of the image plane.

Note: The result can also contain *more* line strings than the input did. That happens when distinct parts of a line string are only connected by line segments that are outside of the image plane and hence will be clipped off, resulting in two or more unconnected line string parts that are left in the image plane.

Returns Line strings, clipped to fall within the image dimensions. The count of output line strings may differ from the input count.

Return type *imgaug.augmentables.lines.LineStringsOnImage*

clip_out_of_image_(*self*)
Clip off all parts of the LSs that are outside of an image in-place.

Note: The result can contain fewer line strings than the input did. That happens when a polygon is fully outside of the image plane.

Note: The result can also contain *more* line strings than the input did. That happens when distinct parts of a line string are only connected by line segments that are outside of the image plane and hence will be clipped off, resulting in two or more unconnected line string parts that are left in the image plane.

Added in 0.4.0.

Returns Line strings, clipped to fall within the image dimensions. The count of output line strings may differ from the input count.

Return type *imgaug.augmentables.lines.LineStringsOnImage*

copy(*self*, *line_strings=None*, *shape=None*)
Create a shallow copy of this object.

Parameters

- **line_strings** (*None or list of imgaug.augmentables.lines.LineString, optional*) – List of line strings on the image. If not `None`, then the `line_strings` attribute of the copied object will be set to this value.
- **shape** (*None or tuple of int or ndarray, optional*) – The shape of the image on which the objects are placed. Either an image with shape `(H,W,[C])` or a tuple denoting such an image shape. If not `None`, then the `shape` attribute of the copied object will be set to this value.

Returns Shallow copy.

Return type *imgaug.augmentables.lines.LineStringsOnImage*

deepcopy(*self*, *line_strings=None*, *shape=None*)
Create a deep copy of the object.

Parameters

- **line_strings** (*None or list of imgaug.augmentables.lines.LineString, optional*) – List of line strings on the image. If not `None`, then the `line_strings` attribute of the copied object will be set to this value.
- **shape** (*None or tuple of int or ndarray, optional*) – The shape of the image on which the objects are placed. Either an image with shape `(H,W,[C])` or a tuple denoting such an image shape. If not `None`, then the `shape` attribute of the copied object will be set to this value.

Returns Deep copy.

Return type *imgaug.augmentables.lines.LineStringsOnImage*

draw_on_image(*self*, *image*, *color=(0, 255, 0)*, *color_lines=None*, *color_points=None*, *alpha=1.0*, *alpha_lines=None*, *alpha_points=None*, *size=1*, *size_lines=None*, *size_points=None*, *antialiased=True*, *raise_if_out_of_image=False*)

Draw all line strings onto a given image.

Parameters

- **image** (*ndarray*) – The `(H,W,C) uint8` image onto which to draw the line strings.
- **color** (*iterable of int, optional*) – Color to use as RGB, i.e. three values. The color of the lines and points are derived from this value, unless they are set.
- **color_lines** (*None or iterable of int*) – Color to use for the line segments as RGB, i.e. three values. If `None`, this value is derived from *color*.
- **color_points** (*None or iterable of int*) – Color to use for the points as RGB, i.e. three values. If `None`, this value is derived from `0.5 * color`.
- **alpha** (*float, optional*) – Opacity of the line strings. Higher values denote more visible points. The alphas of the line and points are derived from this value, unless they are set.
- **alpha_lines** (*None or float, optional*) – Opacity of the line strings. Higher values denote more visible line string. If `None`, this value is derived from *alpha*.
- **alpha_points** (*None or float, optional*) – Opacity of the line string points. Higher values denote more visible points. If `None`, this value is derived from *alpha*.
- **size** (*int, optional*) – Size of the line strings. The sizes of the line and points are derived from this value, unless they are set.
- **size_lines** (*None or int, optional*) – Thickness of the line segments. If `None`, this value is derived from *size*.
- **size_points** (*None or int, optional*) – Size of the points in pixels. If `None`, this value is derived from `3 * size`.
- **antialiased** (*bool, optional*) – Whether to draw the lines with anti-aliasing activated. This does currently not affect the point drawing.
- **raise_if_out_of_image** (*bool, optional*) – Whether to raise an error if a line string is fully outside of the image. If set to `False`, no error will be raised and only the parts inside the image will be drawn.

Returns Image with line strings drawn on it.

Return type ndarray

empty

Estimate whether this object contains zero line strings.

Returns `True` if this object contains zero line strings.

Return type bool

fill_from_xy_array_(*self*, *xy*)
Modify the corner coordinates of all line strings in-place.

Note: This currently expects that *xy* contains exactly as many coordinates as the line strings within this instance have corner points. Otherwise, an `AssertionError` will be raised.

Added in 0.4.0.

Parameters **xy** (*(N, 2) ndarray or iterable of iterable of number*) – XY-Coordinates of `N` corner points. `N` must match the number of corner points in all line strings within this instance.

Returns This instance itself, with updated coordinates. Note that the instance was modified in-place.

Return type *LineStringsOnImage*

classmethod from_xy_arrays(*xy*, *shape*)
Convert an `(N,M,2)` ndarray to a `LineStringsOnImage` object.

This is the inverse of *to_xy_array()*.

Parameters

- **xy** (*(N,M,2) ndarray or iterable of (M,2) ndarray*) – Array containing the point coordinates `N` line strings with each `M` points given as `(x,y)` coordinates. `M` may differ if an iterable of arrays is used. Each array should usually be of dtype `float32`.
- **shape** (*tuple of int*) – `(H,W,[C])` shape of the image on which the line strings are placed.

Returns Object containing a list of `LineString` objects following the provided point coordinates.

Return type *imgaug.augmentables.lines.LineStringsOnImage*

invert_to_keypoints_on_image_(*self*, *kpsoi*)
Invert the output of `to_keypoints_on_image()` in-place.

This function writes in-place into this `LineStringsOnImage` instance.

Added in 0.4.0.

Parameters **kpsoi** (*imgaug.augmentables.kps.KeypointsOnImages*) – Keypoints to convert back to line strings, i.e. the outputs of `to_keypoints_on_image()`.

Returns Line strings container with updated coordinates. Note that the instance is also updated in-place.

Return type *LineStringsOnImage*

items
Get the line strings in this container.

Added in 0.4.0.

Returns Line strings within this container.

Return type list of LineString

on(*self*, *image*)
Project the line strings from one image shape to a new one.

Parameters **image** (*ndarray or tuple of int*) – The new image onto which to project. Either an image with shape `(H,W,[C])` or a tuple denoting such an image shape.

Returns Object containing all projected line strings.

Return type imgaug.augmentables.lines.LineStrings

on_(*self*, *image*)
Project the line strings from one image shape to a new one in-place.

Added in 0.4.0.

Parameters **image** (*ndarray or tuple of int*) – The new image onto which to project. Either an image with shape `(H,W,[C])` or a tuple denoting such an image shape.

Returns Object containing all projected line strings. The object and its items may have been modified in-place.

Return type imgaug.augmentables.lines.LineStrings

remove_out_of_image(*self*, *fully=True*, *partly=False*)
Remove all line strings that are fully/partially outside of an image.

Parameters

- **fully** (*bool, optional*) – Whether to remove line strings that are fully outside of the image.
- **partly** (*bool, optional*) – Whether to remove line strings that are partially outside of the image.

Returns Reduced set of line strings. Those that are fully/partially outside of the given image plane are removed.

Return type *imgaug.augmentables.lines.LineStringsOnImage*

remove_out_of_image_(*self*, *fully=True*, *partly=False*)
Remove all LS that are fully/partially outside of an image in-place.

Added in 0.4.0.

Parameters

- **fully** (*bool, optional*) – Whether to remove line strings that are fully outside of the image.
- **partly** (*bool, optional*) – Whether to remove line strings that are partially outside of the image.

Returns Reduced set of line strings. Those that are fully/partially outside of the given image plane are removed. The object and its items may have been modified in-place.

Return type *imgaug.augmentables.lines.LineStringsOnImage*

remove_out_of_image_fraction(*self*, *fraction*)
Remove all LS with an out of image fraction of at least *fraction*.

Parameters **fraction** (*number*) – Minimum out of image fraction that a line string has to have in order to be removed. A fraction of `1.0` removes only line strings that are `100%` outside of the image. A fraction of `0.0` removes all line strings.

Returns Reduced set of line strings, with those that had an out of image fraction greater or equal the given one removed.

Return type *imgaug.augmentables.lines.LineStringsOnImage*

remove_out_of_image_fraction_(*self*, *fraction*)
Remove all LS with an OOI fraction of at least *fraction* in-place.

'OOI' is the abbreviation for 'out of image'.

Added in 0.4.0.

Parameters **fraction** (*number*) – Minimum out of image fraction that a line string has to have in order to be removed. A fraction of `1.0` removes only line strings that are `100%` outside of the image. A fraction of `0.0` removes all line strings.

Returns Reduced set of line strings, with those that had an out of image fraction greater or equal the given one removed. The object and its items may have been modified in-place.

Return type *imgaug.augmentables.lines.LineStringsOnImage*

shift(*self*, *x=0*, *y=0*, *top=None*, *right=None*, *bottom=None*, *left=None*)
Move the line strings along the x/y-axis.

The origin `(0, 0)` is at the top left of the image.

Parameters

- **x** (*number, optional*) – Value to be added to all x-coordinates. Positive values shift towards the right images.
- **y** (*number, optional*) – Value to be added to all y-coordinates. Positive values shift towards the bottom images.
- **top** (*None or int, optional*) – Deprecated since 0.4.0. Amount of pixels by which to shift all objects *from* the top (towards the bottom).
- **right** (*None or int, optional*) – Deprecated since 0.4.0. Amount of pixels by which to shift all objects *from* the right (towads the left).
- **bottom** (*None or int, optional*) – Deprecated since 0.4.0. Amount of pixels by which to shift all objects *from* the bottom (towards the top).
- **left** (*None or int, optional*) – Deprecated since 0.4.0. Amount of pixels by which to shift all objects *from* the left (towards the right).

Returns Shifted line strings.

Return type *imgaug.augmentables.lines.LineStringsOnImage*

shift_(*self*, *x=0*, *y=0*)
Move the line strings along the x/y-axis in-place.

The origin `(0, 0)` is at the top left of the image.

Added in 0.4.0.

Parameters

- **x** (*number, optional*) – Value to be added to all x-coordinates. Positive values shift towards the right images.
- **y** (*number, optional*) – Value to be added to all y-coordinates. Positive values shift towards the bottom images.

Returns Shifted line strings. The object and its items may have been modified in-place.

Return type *imgaug.augmentables.lines.LineStringsOnImage*

to_keypoints_on_image(*self*)
Convert the line strings to one `KeypointsOnImage` instance.

Added in 0.4.0.

Returns A keypoints instance containing `N` coordinates for a total of `N` points in the `coords` attributes of all line strings. Order matches the order in `line_strings` and `coords` attributes.

Return type *imgaug.augmentables.kps.KeypointsOnImage*

to_xy_array(*self*)
Convert all line string coordinates to one array of shape `(N,2)`.

Added in 0.4.0.

Returns Array containing all xy-coordinates of all line strings within this instance.

Return type (N, 2) ndarray

to_xy_arrays(*self*, *dtype=<class 'numpy.float32'>*)
Convert this object to an iterable of `(M,2)` arrays of points.

This is the inverse of `from_xy_array()`.

Parameters **dtype** (*numpy.dtype, optional*) – Desired output datatype of the ndarray.

Returns The arrays of point coordinates, each given as `(M,2)`.

Return type list of ndarray

13.13 imgaug.augmentables.normalization

Functions dealing with normalization of user input data to imgaug classes.

imgaug.augmentables.normalization.**estimate_bounding_boxes_norm_type**(*bounding_boxes*)

imgaug.augmentables.normalization.**estimate_heatmaps_norm_type**(*heatmaps*)

imgaug.augmentables.normalization.**estimate_keypoints_norm_type**(*keypoints*)

imgaug.augmentables.normalization.**estimate_line_strings_norm_type**(*line_strings*)

imgaug.augmentables.normalization.**estimate_normalization_type**(*inputs*)

imgaug.augmentables.normalization.**estimate_polygons_norm_type**(*polygons*)

imgaug.augmentables.normalization.**estimate_segmaps_norm_type**(*segmentation_maps*)

imgaug.augmentables.normalization.**find_first_nonempty**(*attr*, *parents=None*)

imgaug.augmentables.normalization.**invert_normalize_bounding_boxes**(*bounding_boxes*, *bounding_boxes_old*)

imgaug.augmentables.normalization.**invert_normalize_heatmaps**(*heatmaps*, *heatmaps_old*)

imgaug.augmentables.normalization.**invert_normalize_images**(*images*, *images_old*)

imgaug.augmentables.normalization.**invert_normalize_keypoints**(*keypoints*, *keypoints_old*)

imgaug.augmentables.normalization.**invert_normalize_line_strings**(*line_strings*, *line_strings_old*)

imgaug.augmentables.normalization.**invert_normalize_polygons**(*polygons*, *polygons_old*)

imgaug.augmentables.normalization.**invert_normalize_segmentation_maps**(*segmentation_maps*, *segmentation_maps_old*)

imgaug.augmentables.normalization.**normalize_bounding_boxes**(*inputs*, *shapes=None*)

imgaug.augmentables.normalization.**normalize_heatmaps**(*inputs*, *shapes=None*)

imgaug.augmentables.normalization.**normalize_images**(*images*)

imgaug.augmentables.normalization.**normalize_keypoints**(*inputs*, *shapes=None*)

imgaug.augmentables.normalization.**normalize_line_strings**(*inputs*, *shapes=None*)

imgaug.augmentables.normalization.**normalize_polygons**(*inputs*, *shapes=None*)

imgaug.augmentables.normalization.**normalize_segmentation_maps**(*inputs*, *shapes=None*)

imgaug.augmentables.normalization.**restore_dtype_and_merge**(*arr*, *input_dtype*)

13.14 imgaug.augmentables.polys

Classes dealing with polygons.

class imgaug.augmentables.polys.**MultiPolygon**(*geoms*)

Bases: object

Class that represents several polygons.

Parameters **geoms** (*list of imgaug.augmentables.polys.Polygon*) – List of the polygons.

Methods

from_shapely(geometry[, label])	Create a MultiPolygon from a shapely object.

static from_shapely(*geometry*, *label=None*)

Create a MultiPolygon from a shapely object.

This also creates all necessary Polygon s contained in this MultiPolygon.

Parameters

- **geometry** (*shapely.geometry.MultiPolygon or shapely.geometry.Polygon or shapely.geometry.collection.GeometryCollection*) – The object to convert to a MultiPolygon.
- **label** (*None or str, optional*) – A label assigned to all Polygons within the MultiPolygon.

Returns The derived MultiPolygon.

Return type *imgaug.augmentables.polys.MultiPolygon*

class imgaug.augmentables.polys.**Polygon**(*exterior*, *label=None*)

Bases: object

Class representing polygons.

Each polygon is parameterized by its corner points, given as absolute x- and y-coordinates with sub-pixel accuracy.

Parameters

- **exterior** (*list of imgaug.augmentables.kps.Keypoint or list of tuple of float or (N,2) ndarray*) – List of points defining the polygon. May be either a `list` of *Keypoint* objects or a `list` of `tuple` s in xy-form or a numpy array of shape (N,2) for `N` points in xy-form. All coordinates are expected to be the absolute subpixel-coordinates on the image, given as `float` s, e.g. `x=10.7` and `y=3.4` for a point at coordinates `(10.7, 3.4)`. Their order is expected to be clock-wise. They are expected to not be closed (i.e. first and last coordinate differ).
- **label** (*None or str, optional*) – Label of the polygon, e.g. a string representing the class.

Attributes

area Compute the area of the polygon.

coords Alias for attribute `exterior`.

height Compute the height of a bounding box encapsulating the polygon.

is_valid Estimate whether the polygon has a valid geometry.

width Compute the width of a bounding box encapsulating the polygon.

xx Get the x-coordinates of all points on the exterior.

xx_int Get the discretized x-coordinates of all points on the exterior.

yy Get the y-coordinates of all points on the exterior.

yy_int Get the discretized y-coordinates of all points on the exterior.

Methods

almost_equals(self, other[, max_distance, ...])	Estimate if this polygon's and another's geometry/labels are similar.
change_first_point_by_coords(self, x, y[, ...])	Reorder exterior points so that the point closest to given x/y is first.
change_first_point_by_index(self, point_idx)	Reorder exterior points so that the point with given index is first.
clip_out_of_image(self, image)	Cut off all parts of the polygon that are outside of an image.
compute_out_of_image_area(self, image)	Compute the area of the BB that is outside of the image plane.
compute_out_of_image_fraction(self, image)	Compute fraction of polygon area outside of the image plane.
coords_almost_equals(self, other[, ...])	Alias for *Polygon.exterior_almost_equals()*.
copy(self[, exterior, label])	Create a shallow copy of this object.
cut_out_of_image(self, image)	**Deprecated**.
deepcopy(self[, exterior, label])	Create a deep copy of this object.
draw_on_image(self, image[, color, ...])	Draw the polygon on an image.
exterior_almost_equals(self, other[, ...])	Estimate if this and another polygon's exterior are almost identical.

Continued on next page

Table 48 – continued from previous page

extract_from_image(self, image)	Extract all image pixels within the polygon area.
find_closest_point_index(self, x, y[, ...])	Find the index of the exterior point closest to given coordinates.
from_shapely(polygon_shapely[, label])	Create a polygon from a `Shapely Polygon`.
is_fully_within_image(self, image)	Estimate whether the polygon is fully inside an image plane.
is_out_of_image(self, image[, fully, partly])	Estimate whether the polygon is partially/fully outside of an image.
is_partly_within_image(self, image)	Estimate whether the polygon is at least partially inside an image.
project(self, from_shape, to_shape)	Project the polygon onto an image with different shape.
*project_*(self, from_shape, to_shape)	Project the polygon onto an image with different shape in-place.
shift(self[, x, y, top, right, bottom, left])	Move this polygon along the x/y-axis.
*shift_*(self[, x, y])	Move this polygon along the x/y-axis in-place.
subdivide(self, points_per_edge)	Derive a new polygon with `N` interpolated points per edge.
*subdivide_*(self, points_per_edge)	Derive a new poly with `N` interpolated points per edge in-place.
to_bounding_box(self)	Convert this polygon to a bounding box containing the polygon.
to_keypoints(self)	Convert this polygon's exterior to `Keypoint` instances.
to_line_string(self[, closed])	Convert this polygon's exterior to a `LineString` instance.
to_shapely_line_string(self[, closed, ...])	Convert this polygon to a `Shapely LineString` object.
to_shapely_polygon(self)	Convert this polygon to a `Shapely Polygon`.

almost_equals(*self*, *other*, *max_distance=0.0001*, *points_per_edge=8*)

Estimate if this polygon's and another's geometry/labels are similar.

This is the same as *exterior_almost_equals()* but additionally compares the labels.

Parameters

- **other** (*imgaug.augmentables.polys.Polygon*) – The other object to compare against. Expected to be a `Polygon`.
- **max_distance** (*float, optional*) – See *exterior_almost_equals()*.
- **points_per_edge** (*int, optional*) – See *exterior_almost_equals()*.

Returns `True` if the coordinates are almost equal and additionally the labels are equal. Otherwise `False`.

Return type bool

area

Compute the area of the polygon.

Returns Area of the polygon.

Return type number

change_first_point_by_coords(*self*, *x*, *y*, *max_distance=0.0001*, *raise_if_too_far_away=True*)

Reorder exterior points so that the point closest to given x/y is first.

This method takes a given `(x,y)` coordinate, finds the closest corner point on the exterior and reorders all exterior corner points so that the found point becomes the first one in the array.

If no matching points are found, an exception is raised.

Parameters

- **x** (*number*) – X-coordinate of the point.
- **y** (*number*) – Y-coordinate of the point.
- **max_distance** (*None or number, optional*) – Maximum distance past which possible matches are ignored. If `None` the distance limit is deactivated.
- **raise_if_too_far_away** (*bool, optional*) – Whether to raise an exception if the closest found point is too far away (`True`) or simply return an unchanged copy if this object (`False`).

Returns Copy of this polygon with the new point order.

Return type *imgaug.augmentables.polys.Polygon*

change_first_point_by_index(*self*, *point_idx*)

Reorder exterior points so that the point with given index is first.

This method takes a given index and reorders all exterior corner points so that the point with that index becomes the first one in the array.

An `AssertionError` will be raised if the index does not match any exterior point's index or the exterior does not contain any points.

Parameters **point_idx** (*int*) – Index of the desired starting point.

Returns Copy of this polygon with the new point order.

Return type *imgaug.augmentables.polys.Polygon*

clip_out_of_image(*self*, *image*)

Cut off all parts of the polygon that are outside of an image.

This operation may lead to new points being created. As a single polygon may be split into multiple new polygons, the result is always a list, which may contain more than one output polygon.

This operation will return an empty list if the polygon is completely outside of the image plane.

Parameters **image** (*(H,W,...) ndarray or tuple of int*) – Image dimensions to use for the clipping of the polygon. If an `ndarray`, its shape will be used. If a `tuple`, it is assumed to represent the image shape and must contain at least two `int` s.

Returns Polygon, clipped to fall within the image dimensions. Returned as a `list`, because the clipping can split the polygon into multiple parts. The list may also be empty, if the polygon was fully outside of the image plane.

Return type list of imgaug.augmentables.polys.Polygon

compute_out_of_image_area(*self*, *image*)

Compute the area of the BB that is outside of the image plane.

Added in 0.4.0.

Parameters **image** (*(H,W,...) ndarray or tuple of int*) – Image dimensions to use. If an `ndarray`, its shape will be used. If a `tuple`, it is assumed to represent the image shape and must contain at least two integers.

Returns Total area of the bounding box that is outside of the image plane. Can be `0.0`.

Return type float

compute_out_of_image_fraction(*self*, *image*)
Compute fraction of polygon area outside of the image plane.

This estimates `f = A_ooi / A`, where `A_ooi` is the area of the polygon that is outside of the image plane, while `A` is the total area of the bounding box.

Added in 0.4.0.

Parameters **image** (*(H,W,...) ndarray or tuple of int*) – Image dimensions to use. If an `ndarray`, its shape will be used. If a `tuple`, it is assumed to represent the image shape and must contain at least two integers.

Returns Fraction of the polygon area that is outside of the image plane. Returns `0.0` if the polygon is fully inside of the image plane or has zero points. If the polygon has an area of zero, the polygon is treated similarly to a `LineString`, i.e. the fraction of the line that is outside the image plane is returned.

Return type float

coords
Alias for attribute `exterior`.

Added in 0.4.0.

Returns An `(N, 2) float32` ndarray containing the coordinates of this polygon. This identical to the attribute `exterior`.

Return type ndarray

coords_almost_equals(*self*, *other*, *max_distance=0.0001*, *points_per_edge=8*)
Alias for *Polygon.exterior_almost_equals()*.

Parameters

- **other** (*imgaug.augmentables.polys.Polygon or (N,2) ndarray or list of tuple*) – See *exterior_almost_equals()*.
- **max_distance** (*number, optional*) – See *exterior_almost_equals()*.
- **points_per_edge** (*int, optional*) – See *exterior_almost_equals()*.

Returns Whether the two polygon's exteriors can be viewed as equal (approximate test).

Return type bool

copy(*self*, *exterior=None*, *label=None*)
Create a shallow copy of this object.

Parameters

- **exterior** (*list of imgaug.augmentables.kps.Keypoint or list of tuple or (N,2) ndarray, optional*) – List of points defining the polygon. See `__init__()` for details.
- **label** (*None or str, optional*) – If not `None`, the `label` of the copied object will be set to this value.

Returns Shallow copy.

Return type *imgaug.augmentables.polys.Polygon*

cut_out_of_image(*self*, *image*)

Deprecated. Use `Polygon.clip_out_of_image()` instead. clip_out_of_image() has the exactly same interface.

Cut off all parts of the polygon that are outside of an image.

deepcopy(*self*, *exterior=None*, *label=None*)

Create a deep copy of this object.

Parameters

- **exterior** (*list of Keypoint or list of tuple or (N,2) ndarray, optional*) – List of points defining the polygon. See *imgaug.augmentables.polys.Polygon.__init__* for details.
- **label** (*None or str*) – If not `None`, the `label` of the copied object will be set to this value.

Returns Deep copy.

Return type *imgaug.augmentables.polys.Polygon*

draw_on_image(*self*, *image*, *color=(0, 255, 0)*, *color_face=None*, *color_lines=None*, *color_points=None*, *alpha=1.0*, *alpha_face=None*, *alpha_lines=None*, *alpha_points=None*, *size=1*, *size_lines=None*, *size_points=None*, *raise_if_out_of_image=False*)

Draw the polygon on an image.

Parameters

- **image** (*(H,W,C) ndarray*) – The image onto which to draw the polygon. Usually expected to be of dtype `uint8`, though other dtypes are also handled.
- **color** (*iterable of int, optional*) – The color to use for the whole polygon. Must correspond to the channel layout of the image. Usually RGB. The values for *color_face*, *color_lines* and *color_points* will be derived from this color if they are set to `None`. This argument has no effect if *color_face*, *color_lines* and *color_points* are all set anything other than `None`.
- **color_face** (*None or iterable of int, optional*) – The color to use for the inner polygon area (excluding perimeter). Must correspond to the channel layout of the image. Usually RGB. If this is `None`, it will be derived from `color * 1.0`.
- **color_lines** (*None or iterable of int, optional*) – The color to use for the line (aka perimeter/border) of the polygon. Must correspond to the channel layout of the image. Usually RGB. If this is `None`, it will be derived from `color * 0.5`.
- **color_points** (*None or iterable of int, optional*) – The color to use for the corner points of the polygon. Must correspond to the channel layout of the image. Usually RGB. If this is `None`, it will be derived from `color * 0.5`.
- **alpha** (*float, optional*) – The opacity of the whole polygon, where `1.0` denotes a completely visible polygon and `0.0` an invisible one. The values for *alpha_face*, *alpha_lines* and *alpha_points* will be derived from this alpha value if they are set to `None`. This argument has no effect if *alpha_face*, *alpha_lines* and *alpha_points* are all set anything other than `None`.
- **alpha_face** (*None or number, optional*) – The opacity of the polygon's inner area (excluding the perimeter), where `1.0` denotes a completely visible inner area and `0.0` an invisible one. If this is `None`, it will be derived from `alpha * 0.5`.
- **alpha_lines** (*None or number, optional*) – The opacity of the polygon's line (aka perimeter/border), where `1.0` denotes a completely visible line and `0.0` an invisible one. If this is `None`, it will be derived from `alpha * 1.0`.

- **alpha_points** (*None or number, optional*) – The opacity of the polygon's corner points, where `1.0` denotes completely visible corners and `0.0` invisible ones. If this is `None`, it will be derived from `alpha * 1.0`.
- **size** (*int, optional*) – Size of the polygon. The sizes of the line and points are derived from this value, unless they are set.
- **size_lines** (*None or int, optional*) – Thickness of the polygon's line (aka perimeter/border). If `None`, this value is derived from *size*.
- **size_points** (*int, optional*) – Size of the points in pixels. If `None`, this value is derived from `3 * size`.
- **raise_if_out_of_image** (*bool, optional*) – Whether to raise an error if the polygon is fully outside of the image. If set to `False`, no error will be raised and only the parts inside the image will be drawn.

Returns Image with the polygon drawn on it. Result dtype is the same as the input dtype.

Return type (H,W,C) ndarray

exterior_almost_equals (*self*, *other*, *max_distance=0.0001*, *points_per_edge=8*)

Estimate if this and another polygon's exterior are almost identical.

The two exteriors can have different numbers of points, but any point randomly sampled on the exterior of one polygon should be close to the closest point on the exterior of the other polygon.

Note: This method works in an approximative way. One can come up with polygons with fairly different shapes that will still be estimated as equal by this method. In practice however this should be unlikely to be the case. The probability for something like that goes down as the interpolation parameter is increased.

Parameters

- **other** (*imgaug.augmentables.polys.Polygon or (N,2) ndarray or list of tuple*) – The other polygon with which to compare the exterior. If this is an `ndarray`, it is assumed to represent an exterior. It must then have dtype `float32` and shape `(N,2)` with the second dimension denoting xy-coordinates. If this is a `list` of `tuple` s, it is assumed to represent an exterior. Each tuple then must contain exactly two `number` s, denoting xy-coordinates.
- **max_distance** (*number, optional*) – The maximum euclidean distance between a point on one polygon and the closest point on the other polygon. If the distance is exceeded for any such pair, the two exteriors are not viewed as equal. The points are either the points contained in the polygon's exterior ndarray or interpolated points between these.
- **points_per_edge** (*int, optional*) – How many points to interpolate on each edge.

Returns Whether the two polygon's exteriors can be viewed as equal (approximate test).

Return type bool

extract_from_image (*self*, *image*)

Extract all image pixels within the polygon area.

This method returns a rectangular image array. All pixels within that rectangle that do not belong to the polygon area will be filled with zeros (i.e. they will be black). The method will also zero-pad the image if the polygon is partially/fully outside of the image.

Parameters **image** (*(H,W) ndarray or (H,W,C) ndarray*) – The image from which to extract the pixels within the polygon.

Returns Pixels within the polygon. Zero-padded if the polygon is partially/fully outside of the image.

Return type (H',W') ndarray or (H',W',C) ndarray

find_closest_point_index(*self*, *x*, *y*, *return_distance=False*)
Find the index of the exterior point closest to given coordinates.

"Closeness" is here defined based on euclidean distance. This method will raise an `AssertionError` if the exterior contains no points.

Parameters

- **x** (*number*) – X-coordinate around which to search for close points.
- **y** (*number*) – Y-coordinate around which to search for close points.
- **return_distance** (*bool, optional*) – Whether to also return the distance of the closest point.

Returns

- *int* – Index of the closest point.
- *number* – Euclidean distance to the closest point. This value is only returned if *return_distance* was set to `True`.

static from_shapely(*polygon_shapely*, *label=None*)
Create a polygon from a `Shapely Polygon`.

Note: This will remove any holes in the shapely polygon.

Parameters

- **polygon_shapely** (*shapely.geometry.Polygon*) – The shapely polygon.
- **label** (*None or str, optional*) – The label of the new polygon.

Returns A polygon with the same exterior as the `Shapely Polygon`.

Return type *imgaug.augmentables.polys.Polygon*

height
Compute the height of a bounding box encapsulating the polygon.

The height is computed based on the two exterior coordinates with lowest and largest x-coordinates.

Returns Height of the polygon.

Return type number

is_fully_within_image(*self*, *image*)
Estimate whether the polygon is fully inside an image plane.

Parameters **image** (*(H,W,...) ndarray or tuple of int*) – Image dimensions to use. If an `ndarray`, its shape will be used. If a `tuple`, it is assumed to represent the image shape and must contain at least two `int` s.

Returns `True` if the polygon is fully inside the image area. `False` otherwise.

Return type bool

is_out_of_image(*self*, *image*, *fully=True*, *partly=False*)
Estimate whether the polygon is partially/fully outside of an image.

Parameters

- **image** (*(H,W,...) ndarray or tuple of int*) – Image dimensions to use. If an `ndarray`, its shape will be used. If a `tuple`, it is assumed to represent the image shape and must contain at least two `int` s.
- **fully** (*bool, optional*) – Whether to return `True` if the polygon is fully outside of the image area.
- **partly** (*bool, optional*) – Whether to return `True` if the polygon is at least partially outside fo the image area.

Returns `True` if the polygon is partially/fully outside of the image area, depending on defined parameters. `False` otherwise.

Return type bool

is_partly_within_image(*self*, *image*)

Estimate whether the polygon is at least partially inside an image.

Parameters **image** (*(H,W,...) ndarray or tuple of int*) – Image dimensions to use. If an `ndarray`, its shape will be used. If a `tuple`, it is assumed to represent the image shape and must contain at least two `int` s.

Returns `True` if the polygon is at least partially inside the image area. `False` otherwise.

Return type bool

is_valid

Estimate whether the polygon has a valid geometry.

To to be considered valid, the polygon must be made up of at least `3` points and have a concave shape, i.e. line segments may not intersect or overlap. Multiple consecutive points are allowed to have the same coordinates.

Returns `True` if polygon has at least `3` points and is concave, otherwise `False`.

Return type bool

project(*self*, *from_shape*, *to_shape*)

Project the polygon onto an image with different shape.

The relative coordinates of all points remain the same. E.g. a point at `(x=20, y=20)` on an image `(width=100, height=200)` will be projected on a new image `(width=200, height=100)` to `(x=40, y=10)`.

This is intended for cases where the original image is resized. It cannot be used for more complex changes (e.g. padding, cropping).

Parameters

- **from_shape** (*tuple of int*) – Shape of the original image. (Before resize.)
- **to_shape** (*tuple of int*) – Shape of the new image. (After resize.)

Returns Polygon object with new coordinates.

Return type *imgaug.augmentables.polys.Polygon*

project_(*self*, *from_shape*, *to_shape*)

Project the polygon onto an image with different shape in-place.

The relative coordinates of all points remain the same. E.g. a point at `(x=20, y=20)` on an image `(width=100, height=200)` will be projected on a new image `(width=200, height=100)` to `(x=40, y=10)`.

This is intended for cases where the original image is resized. It cannot be used for more complex changes (e.g. padding, cropping).

Added in 0.4.0.

Parameters

- **from_shape** (*tuple of int*) – Shape of the original image. (Before resize.)
- **to_shape** (*tuple of int*) – Shape of the new image. (After resize.)

Returns Polygon object with new coordinates. The object may have been modified in-place.

Return type *imgaug.augmentables.polys.Polygon*

shift(*self*, *x=0*, *y=0*, *top=None*, *right=None*, *bottom=None*, *left=None*)
Move this polygon along the x/y-axis.

The origin `(0, 0)` is at the top left of the image.

Parameters

- **x** (*number, optional*) – Value to be added to all x-coordinates. Positive values shift towards the right images.
- **y** (*number, optional*) – Value to be added to all y-coordinates. Positive values shift towards the bottom images.
- **top** (*None or int, optional*) – Deprecated since 0.4.0. Amount of pixels by which to shift this object *from* the top (towards the bottom).
- **right** (*None or int, optional*) – Deprecated since 0.4.0. Amount of pixels by which to shift this object *from* the right (towards the left).
- **bottom** (*None or int, optional*) – Deprecated since 0.4.0. Amount of pixels by which to shift this object *from* the bottom (towards the top).
- **left** (*None or int, optional*) – Deprecated since 0.4.0. Amount of pixels by which to shift this object *from* the left (towards the right).

Returns Shifted polygon.

Return type *imgaug.augmentables.polys.Polygon*

shift_(*self*, *x=0*, *y=0*)
Move this polygon along the x/y-axis in-place.

The origin `(0, 0)` is at the top left of the image.

Added in 0.4.0.

Parameters

- **x** (*number, optional*) – Value to be added to all x-coordinates. Positive values shift towards the right images.
- **y** (*number, optional*) – Value to be added to all y-coordinates. Positive values shift towards the bottom images.

Returns Shifted polygon. The object may have been modified in-place.

Return type *imgaug.augmentables.polys.Polygon*

subdivide(*self*, *points_per_edge*)
Derive a new polygon with `N` interpolated points per edge.

See `subdivide()` for details.

Added in 0.4.0.

Parameters **points_per_edge** (*int*) – Number of points to interpolate on each edge.

Returns Polygon with subdivided edges.

Return type *imgaug.augmentables.polys.Polygon*

subdivide_(*self*, *points_per_edge*)
Derive a new poly with `N` interpolated points per edge in-place.

See *subdivide()* for details.

Added in 0.4.0.

Parameters **points_per_edge** (*int*) – Number of points to interpolate on each edge.

Returns Polygon with subdivided edges. The object may have been modified in-place.

Return type *imgaug.augmentables.polys.Polygon*

to_bounding_box(*self*)
Convert this polygon to a bounding box containing the polygon.

Returns Bounding box that tightly encapsulates the polygon.

Return type *imgaug.augmentables.bbs.BoundingBox*

to_keypoints(*self*)
Convert this polygon's exterior to `Keypoint` instances.

Returns Exterior vertices as *Keypoint* instances.

Return type list of imgaug.augmentables.kps.Keypoint

to_line_string(*self*, *closed=True*)
Convert this polygon's exterior to a `LineString` instance.

Parameters **closed** (*bool, optional*) – Whether to close the line string, i.e. to add the first point of the *exterior* also as the last point at the end of the line string. This has no effect if the polygon has a single point or zero points.

Returns Exterior of the polygon as a line string.

Return type *imgaug.augmentables.lines.LineString*

to_shapely_line_string(*self*, *closed=False*, *interpolate=0*)
Convert this polygon to a `Shapely` `LineString` object.

Parameters

- **closed** (*bool, optional*) – Whether to return the line string with the last point being identical to the first point.
- **interpolate** (*int, optional*) – Number of points to interpolate between any pair of two consecutive points. These points are added to the final line string.

Returns The `Shapely` `LineString` matching the polygon's exterior.

Return type shapely.geometry.LineString

to_shapely_polygon(*self*)
Convert this polygon to a `Shapely` `Polygon`.

Returns The `Shapely` `Polygon` matching this polygon's exterior.

Return type shapely.geometry.Polygon

width
Compute the width of a bounding box encapsulating the polygon.

The width is computed based on the two exterior coordinates with lowest and largest x-coordinates.

Returns Width of the polygon.

Return type number

xx
Get the x-coordinates of all points on the exterior.

Returns `float32` x-coordinates array of all points on the exterior.

Return type (N,2) ndarray

xx_int
Get the discretized x-coordinates of all points on the exterior.

The conversion from `float32` coordinates to `int32` is done by first rounding the coordinates to the closest integer and then removing everything after the decimal point.

Returns `int32` x-coordinates of all points on the exterior.

Return type (N,2) ndarray

yy
Get the y-coordinates of all points on the exterior.

Returns `float32` y-coordinates array of all points on the exterior.

Return type (N,2) ndarray

yy_int
Get the discretized y-coordinates of all points on the exterior.

The conversion from `float32` coordinates to `int32` is done by first rounding the coordinates to the closest integer and then removing everything after the decimal point.

Returns `int32` y-coordinates of all points on the exterior.

Return type (N,2) ndarray

class `imgaug.augmentables.polys.`**PolygonsOnImage**(*polygons*, *shape*)
Bases: *imgaug.augmentables.base.IAugmentable*

Container for all polygons on a single image.

Parameters

- **polygons** (*list of imgaug.augmentables.polys.Polygon*) – List of polygons on the image.
- **shape** (*tuple of int or ndarray*) – The shape of the image on which the objects are placed. Either an image with shape `(H,W,[C])` or a `tuple` denoting such an image shape.

Examples

```
>>> import numpy as np
>>> from imgaug.augmentables.polys import Polygon, PolygonsOnImage
>>> image = np.zeros((100, 100))
>>> polys = [
>>>     Polygon([(0.5, 0.5), (100.5, 0.5), (100.5, 100.5), (0.5, 100.5)]),
>>>     Polygon([(50.5, 0.5), (100.5, 50.5), (50.5, 100.5), (0.5, 50.5)])
```

(continues on next page)

(continued from previous page)

```
>>> ]
>>> polys_oi = PolygonsOnImage(polys, shape=image.shape)
```

Attributes

empty	Estimate whether this object contains zero polygons.
items	Get the polygons in this container.

Methods

clip_out_of_image(self)	Clip off all parts from all polygons that are outside of an image.
*clip_out_of_image_*(self)	Clip off all parts from all polygons that are OOI in-place.
copy(self[, polygons, shape])	Create a shallow copy of this object.
deepcopy(self[, polygons, shape])	Create a deep copy of this object.
draw_on_image(self, image[, color, ...])	Draw all polygons onto a given image.
*fill_from_xy_array_*(self, xy)	Modify the corner coordinates of all polygons in-place.
*invert_to_keypoints_on_image_*(self, kpsoi)	Invert the output of `to_keypoints_on_image()` in-place.
on(self, image)	Project all polygons from one image shape to a new one.
*on_*(self, image)	Project all polygons from one image shape to a new one in-place.
remove_out_of_image(self[, fully, partly])	Remove all polygons that are fully/partially outside of an image.
*remove_out_of_image_*(self[, fully, partly])	Remove all polygons that are fully/partially OOI in-place.
remove_out_of_image_fraction(self, fraction)	Remove all Polys with an out of image fraction of `>=fraction`.
*remove_out_of_image_fraction_*(self, fraction)	Remove all Polys with an OOI fraction of `>=fraction` in-place.
shift(self[, x, y, top, right, bottom, left])	Move the polygons along the x/y-axis.
*shift_*(self[, x, y])	Move the polygons along the x/y-axis in-place.
subdivide(self, points_per_edge)	Interpolate `N` points on each polygon.
*subdivide_*(self, points_per_edge)	Interpolate `N` points on each polygon.
to_keypoints_on_image(self)	Convert the polygons to one `KeypointsOnImage` instance.
to_xy_array(self)	Convert all polygon coordinates to one array of shape `(N,2)`.

clip_out_of_image(*self*)

Clip off all parts from all polygons that are outside of an image.

Note: The result can contain fewer polygons than the input did. That happens when a polygon is fully outside of the image plane.

Note: The result can also contain *more* polygons than the input did. That happens when distinct parts of a polygon are only connected by areas that are outside of the image plane and hence will be clipped off, resulting in two or more unconnected polygon parts that are left in the image plane.

Returns Polygons, clipped to fall within the image dimensions. The count of output polygons may differ from the input count.

Return type *imgaug.augmentables.polys.PolygonsOnImage*

clip_out_of_image_(*self*)
Clip off all parts from all polygons that are OOI in-place.

'OOI' is the abbreviation for 'out of image'.

Note: The result can contain fewer polygons than the input did. That happens when a polygon is fully outside of the image plane.

Note: The result can also contain *more* polygons than the input did. That happens when distinct parts of a polygon are only connected by areas that are outside of the image plane and hence will be clipped off, resulting in two or more unconnected polygon parts that are left in the image plane.

Added in 0.4.0.

Returns Polygons, clipped to fall within the image dimensions. The count of output polygons may differ from the input count. The object and its items may have been modified in-place.

Return type *imgaug.augmentables.polys.PolygonsOnImage*

copy(*self*, *polygons=None*, *shape=None*)
Create a shallow copy of this object.

Parameters

- **polygons** (*None or list of imgaug.augmentables.polys.Polygons, optional*) – List of polygons on the image. If not `None`, then the `polygons` attribute of the copied object will be set to this value.
- **shape** (*None or tuple of int or ndarray, optional*) – The shape of the image on which the objects are placed. Either an image with shape `(H,W,[C])` or a tuple denoting such an image shape. If not `None`, then the `shape` attribute of the copied object will be set to this value.

Returns Shallow copy.

Return type *imgaug.augmentables.polys.PolygonsOnImage*

deepcopy(*self*, *polygons=None*, *shape=None*)
Create a deep copy of this object.

Parameters

- **polygons** (*None or list of imgaug.augmentables.polys.Polygons, optional*) – List of polygons on the image. If not `None`, then the `polygons` attribute of the copied object will be set to this value.

- **shape** (*None or tuple of int or ndarray, optional*) – The shape of the image on which the objects are placed. Either an image with shape `(H,W,[C])` or a tuple denoting such an image shape. If not `None`, then the `shape` attribute of the copied object will be set to this value.

Returns Deep copy.

Return type *imgaug.augmentables.polys.PolygonsOnImage*

draw_on_image(*self*, *image*, *color=(0, 255, 0)*, *color_face=None*, *color_lines=None*, *color_points=None*, *alpha=1.0*, *alpha_face=None*, *alpha_lines=None*, *alpha_points=None*, *size=1*, *size_lines=None*, *size_points=None*, *raise_if_out_of_image=False*)

Draw all polygons onto a given image.

Parameters

- **image** (*(H,W,C) ndarray*) – The image onto which to draw the bounding boxes. This image should usually have the same shape as set in `PolygonsOnImage.shape`.
- **color** (*iterable of int, optional*) – The color to use for the whole polygons. Must correspond to the channel layout of the image. Usually RGB. The values for *color_face*, *color_lines* and *color_points* will be derived from this color if they are set to `None`. This argument has no effect if *color_face*, *color_lines* and *color_points* are all set anything other than `None`.
- **color_face** (*None or iterable of int, optional*) – The color to use for the inner polygon areas (excluding perimeters). Must correspond to the channel layout of the image. Usually RGB. If this is `None`, it will be derived from `color * 1.0`.
- **color_lines** (*None or iterable of int, optional*) – The color to use for the lines (aka perimeters/borders) of the polygons. Must correspond to the channel layout of the image. Usually RGB. If this is `None`, it will be derived from `color * 0.5`.
- **color_points** (*None or iterable of int, optional*) – The color to use for the corner points of the polygons. Must correspond to the channel layout of the image. Usually RGB. If this is `None`, it will be derived from `color * 0.5`.
- **alpha** (*float, optional*) – The opacity of the whole polygons, where `1.0` denotes completely visible polygons and `0.0` invisible ones. The values for *alpha_face*, *alpha_lines* and *alpha_points* will be derived from this alpha value if they are set to `None`. This argument has no effect if *alpha_face*, *alpha_lines* and *alpha_points* are all set anything other than `None`.
- **alpha_face** (*None or number, optional*) – The opacity of the polygon's inner areas (excluding the perimeters), where `1.0` denotes completely visible inner areas and `0.0` invisible ones. If this is `None`, it will be derived from `alpha * 0.5`.
- **alpha_lines** (*None or number, optional*) – The opacity of the polygon's lines (aka perimeters/borders), where `1.0` denotes completely visible perimeters and `0.0` invisible ones. If this is `None`, it will be derived from `alpha * 1.0`.
- **alpha_points** (*None or number, optional*) – The opacity of the polygon's corner points, where `1.0` denotes completely visible corners and `0.0` invisible ones. Currently this is an on/off choice, i.e. only `0.0` or `1.0` are allowed. If this is `None`, it will be derived from `alpha * 1.0`.
- **size** (*int, optional*) – Size of the polygons. The sizes of the line and points are derived from this value, unless they are set.
- **size_lines** (*None or int, optional*) – Thickness of the polygon lines (aka perimeter/border). If `None`, this value is derived from *size*.

- **size_points** (*int, optional*) – The size of all corner points. If set to `C`, each corner point will be drawn as a square of size `C x C`.
- **raise_if_out_of_image** (*bool, optional*) – Whether to raise an error if any polygon is fully outside of the image. If set to False, no error will be raised and only the parts inside the image will be drawn.

Returns Image with drawn polygons.

Return type (H,W,C) ndarray

`empty`
Estimate whether this object contains zero polygons.

Returns `True` if this object contains zero polygons.

Return type bool

`fill_from_xy_array_` (*self*, *xy*)
Modify the corner coordinates of all polygons in-place.

Note: This currently expects that *xy* contains exactly as many coordinates as the polygons within this instance have corner points. Otherwise, an `AssertionError` will be raised.

Warning: This does not validate the new coordinates or repair the resulting polygons. If bad coordinates are provided, the result will be invalid polygons (e.g. self-intersections).

Added in 0.4.0.

Parameters **xy** (*(N, 2) ndarray or iterable of iterable of number*) – XY-Coordinates of `N` corner points. `N` must match the number of corner points in all polygons within this instance.

Returns This instance itself, with updated coordinates. Note that the instance was modified in-place.

Return type *PolygonsOnImage*

`invert_to_keypoints_on_image_` (*self*, *kpsoi*)
Invert the output of `to_keypoints_on_image()` in-place.

This function writes in-place into this `PolygonsOnImage` instance.

Added in 0.4.0.

Parameters **kpsoi** (*imgaug.augmentables.kps.KeypointsOnImages*) – Keypoints to convert back to polygons, i.e. the outputs of `to_keypoints_on_image()`.

Returns Polygons container with updated coordinates. Note that the instance is also updated in-place.

Return type *PolygonsOnImage*

`items`
Get the polygons in this container.

Added in 0.4.0.

Returns Polygons within this container.

Return type list of Polygon

on (*self*, *image*)
Project all polygons from one image shape to a new one.

Parameters **image** (*ndarray or tuple of int*) – New image onto which the polygons are to be projected. May also simply be that new image's shape `tuple`.

Returns Object containing all projected polygons.

Return type *imgaug.augmentables.polys.PolygonsOnImage*

on_ (*self*, *image*)
Project all polygons from one image shape to a new one in-place.

Added in 0.4.0.

Parameters **image** (*ndarray or tuple of int*) – New image onto which the polygons are to be projected. May also simply be that new image's shape `tuple`.

Returns Object containing all projected polygons. The object and its items may have been modified in-place.

Return type *imgaug.augmentables.polys.PolygonsOnImage*

remove_out_of_image (*self*, *fully=True*, *partly=False*)
Remove all polygons that are fully/partially outside of an image.

Parameters

- **fully** (*bool, optional*) – Whether to remove polygons that are fully outside of the image.
- **partly** (*bool, optional*) – Whether to remove polygons that are partially outside of the image.

Returns Reduced set of polygons. Those that are fully/partially outside of the given image plane are removed.

Return type *imgaug.augmentables.polys.PolygonsOnImage*

remove_out_of_image_ (*self*, *fully=True*, *partly=False*)
Remove all polygons that are fully/partially OOI in-place.

'OOI' is the abbreviation for 'out of image'.

Added in 0.4.0.

Parameters

- **fully** (*bool, optional*) – Whether to remove polygons that are fully outside of the image.
- **partly** (*bool, optional*) – Whether to remove polygons that are partially outside of the image.

Returns Reduced set of polygons. Those that are fully/partially outside of the given image plane are removed. The object and its items may have been modified in-place.

Return type *imgaug.augmentables.polys.PolygonsOnImage*

remove_out_of_image_fraction (*self*, *fraction*)
Remove all Polys with an out of image fraction of `>=fraction`.

Added in 0.4.0.

Parameters **fraction** (*number*) – Minimum out of image fraction that a polygon has to have in order to be removed. A fraction of `1.0` removes only polygons that are `100%` outside of the image. A fraction of `0.0` removes all polygons.

Returns Reduced set of polygons, with those that had an out of image fraction greater or equal the given one removed.

Return type *imgaug.augmentables.polys.PolygonsOnImage*

remove_out_of_image_fraction_(*self*, *fraction*)
Remove all Polys with an OOI fraction of `>=fraction` in-place.

Added in 0.4.0.

Parameters **fraction** (*number*) – Minimum out of image fraction that a polygon has to have in order to be removed. A fraction of `1.0` removes only polygons that are `100%` outside of the image. A fraction of `0.0` removes all polygons.

Returns Reduced set of polygons, with those that had an out of image fraction greater or equal the given one removed. The object and its items may have been modified in-place.

Return type *imgaug.augmentables.polys.PolygonsOnImage*

shift(*self*, *x=0*, *y=0*, *top=None*, *right=None*, *bottom=None*, *left=None*)
Move the polygons along the x/y-axis.

The origin `(0, 0)` is at the top left of the image.

Parameters

- **x** (*number, optional*) – Value to be added to all x-coordinates. Positive values shift towards the right images.
- **y** (*number, optional*) – Value to be added to all y-coordinates. Positive values shift towards the bottom images.
- **top** (*None or int, optional*) – Deprecated since 0.4.0. Amount of pixels by which to shift all objects *from* the top (towards the bottom).
- **right** (*None or int, optional*) – Deprecated since 0.4.0. Amount of pixels by which to shift all objects *from* the right (towads the left).
- **bottom** (*None or int, optional*) – Deprecated since 0.4.0. Amount of pixels by which to shift all objects *from* the bottom (towards the top).
- **left** (*None or int, optional*) – Deprecated since 0.4.0. Amount of pixels by which to shift all objects *from* the left (towards the right).

Returns Shifted polygons.

Return type *imgaug.augmentables.polys.PolygonsOnImage*

shift_(*self*, *x=0*, *y=0*)
Move the polygons along the x/y-axis in-place.

The origin `(0, 0)` is at the top left of the image.

Added in 0.4.0.

Parameters

- **x** (*number, optional*) – Value to be added to all x-coordinates. Positive values shift towards the right images.
- **y** (*number, optional*) – Value to be added to all y-coordinates. Positive values shift towards the bottom images.

Returns Shifted polygons.

Return type *imgaug.augmentables.polys.PolygonsOnImage*

subdivide(*self*, *points_per_edge*)
Interpolate `N` points on each polygon.

Added in 0.4.0.

Parameters **points_per_edge** (*int*) – Number of points to interpolate on each edge.

Returns Subdivided polygons.

Return type *imgaug.augmentables.polys.PolygonsOnImage*

subdivide_(*self*, *points_per_edge*)
Interpolate `N` points on each polygon.

Added in 0.4.0.

Parameters **points_per_edge** (*int*) – Number of points to interpolate on each edge.

Returns Subdivided polygons.

Return type *imgaug.augmentables.polys.PolygonsOnImage*

to_keypoints_on_image(*self*)
Convert the polygons to one `KeypointsOnImage` instance.

Added in 0.4.0.

Returns A keypoints instance containing `N` coordinates for a total of `N` points in all exteriors of the polygons within this container. Order matches the order in `polygons`.

Return type *imgaug.augmentables.kps.KeypointsOnImage*

to_xy_array(*self*)
Convert all polygon coordinates to one array of shape `(N,2)`.

Added in 0.4.0.

Returns Array containing all xy-coordinates of all polygons within this instance.

Return type (N, 2) ndarray

imgaug.augmentables.polys.**recover_psois_**(*psois*, *psois_orig*, *recoverer*, *random_state*)
Apply a polygon recoverer to input polygons in-place.

Parameters

- **psois** (*list of imgaug.augmentables.polys.PolygonsOnImage or imgaug.augmentables.polys.PolygonsOnImage*) – The possibly broken polygons, e.g. after augmentation. The *recoverer* is applied to them.
- **psois_orig** (*list of imgaug.augmentables.polys.PolygonsOnImage or imgaug.augmentables.polys.PolygonsOnImage*) – Original polygons that were later changed to *psois*. They are an extra input to *recoverer*.
- **recoverer** (*imgaug.augmentables.polys._ConcavePolygonRecoverer*) – The polygon recoverer used to repair broken input polygons.
- **random_state** (*None or int or RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState*) – An RNG to use during the polygon recovery.

Returns List of repaired polygons. Note that this is *psois*, which was changed in-place.

Return type list of imgaug.augmentables.polys.PolygonsOnImage or *imgaug.augmentables.polys.PolygonsOnImage*

13.15 imgaug.augmentables.segmaps

Classes dealing with segmentation maps.

E.g. masks, semantic or instance segmentation maps.

imgaug.augmentables.segmaps.**SegmentationMapOnImage**(**args*, ***kwargs*)

Deprecated. Use `SegmentationMapsOnImage` instead. (Note the plural 'Maps' instead of old 'Map'.).

Object representing a segmentation map associated with an image.

class imgaug.augmentables.segmaps.**SegmentationMapsOnImage**(*arr*, *shape*, *nb_classes=None*)

Bases: *imgaug.augmentables.base.IAugmentable*

Object representing a segmentation map associated with an image.

Variables ***DEFAULT_SEGMENT_COLORS*** (*list of tuple of int*) – Standard RGB colors to use during drawing, ordered by class index.

Parameters

- **arr** (*(H,W) ndarray or (H,W,C) ndarray*) – Array representing the segmentation map(s). May have dtypes bool, int or uint.
- **shape** (*tuple of int*) – Shape of the image on which the segmentation map(s) is/are placed. **Not** the shape of the segmentation map(s) array, unless it is identical to the image shape (note the likely difference between the arrays in the number of channels). This is expected to be `(H, W)` or `(H, W, C)` with `C` usually being `3`. If there is no corresponding image, use `(H_arr, W_arr)` instead, where `H_arr` is the height of the segmentation map(s) array (analogous `W_arr`).
- **nb_classes** (*None or int, optional*) – Deprecated.

Methods

copy(self[, arr, shape])	Create a shallow copy of the segmentation map object.
deepcopy(self[, arr, shape])	Create a deep copy of the segmentation map object.
draw(self[, size, colors])	Render the segmentation map as an RGB image.
draw_on_image(self, image[, alpha, resize, ...])	Draw the segmentation map as an overlay over an image.
get_arr(self)	Return the seg.map array, with original dtype and shape ndim.
get_arr_int(self, *args, **kwargs)	**Deprecated**.
pad(self[, top, right, bottom, left, mode, cval])	Pad the segmentation maps at their top/right/bottom/left side.
pad_to_aspect_ratio(self, aspect_ratio[, ...])	Pad the segmentation maps until they match a target aspect ratio.
resize(self, sizes[, interpolation])	Resize the seg.map(s) array given a target size and interpolation.
scale(self, *args, **kwargs)	**Deprecated**.

DEFAULT_SEGMENT_COLORS = [(0, 0, 0), (230, 25, 75), (60, 180, 75), (255, 225, 25), (0,

copy(*self*, *arr=None*, *shape=None*)

Create a shallow copy of the segmentation map object.

Parameters

- **arr** (*None or (H,W) ndarray or (H,W,C) ndarray, optional*) – Optionally the *arr* attribute to use for the new segmentation map instance. Will be copied from the old instance if not provided. See `__init__()` for details.
- **shape** (*None or tuple of int, optional*) – Optionally the shape attribute to use for the the new segmentation map instance. Will be copied from the old instance if not provided. See `__init__()` for details.

Returns Shallow copy.

Return type *imgaug.augmentables.segmaps.SegmentationMapsOnImage*

deepcopy(*self*, *arr=None*, *shape=None*)
Create a deep copy of the segmentation map object.

Parameters

- **arr** (*None or (H,W) ndarray or (H,W,C) ndarray, optional*) – Optionally the *arr* attribute to use for the new segmentation map instance. Will be copied from the old instance if not provided. See `__init__()` for details.
- **shape** (*None or tuple of int, optional*) – Optionally the shape attribute to use for the the new segmentation map instance. Will be copied from the old instance if not provided. See `__init__()` for details.

Returns Deep copy.

Return type *imgaug.augmentables.segmaps.SegmentationMapsOnImage*

draw(*self*, *size=None*, *colors=None*)
Render the segmentation map as an RGB image.

Parameters

- **size** (*None or float or iterable of int or iterable of float, optional*) – Size of the rendered RGB image as `(height, width)`. See *imresize_single_image()* for details. If set to `None`, no resizing is performed and the size of the segmentation map array is used.
- **colors** (*None or list of tuple of int, optional*) – Colors to use. One for each class to draw. If `None`, then default colors will be used.

Returns Rendered segmentation map (dtype is `uint8`). One per `C` in the original input array `(H,W,C)`.

Return type list of (H,W,3) ndarray

draw_on_image(*self*, *image*, *alpha=0.75*, *resize='segmentation_map'*, *colors=None*, *draw_background=False*, *background_class_id=0*, *background_threshold=None*)
Draw the segmentation map as an overlay over an image.

Parameters

- **image** (*(H,W,3) ndarray*) – Image onto which to draw the segmentation map. Expected dtype is `uint8`.
- **alpha** (*float, optional*) – Alpha/opacity value to use for the mixing of image and segmentation map. Larger values mean that the segmentation map will be more visible and the image less visible.
- **resize** (*{'segmentation_map', 'image'}, optional*) – In case of size differences between the image and segmentation map, either the image or the segmentation map can be resized. This parameter controls which of the two will be resized to the other's size.

- **colors** (*None or list of tuple of int, optional*) – Colors to use. One for each class to draw. If `None`, then default colors will be used.
- **draw_background** (*bool, optional*) – If `True`, the background will be drawn like any other class. If `False`, the background will not be drawn, i.e. the respective background pixels will be identical with the image's RGB color at the corresponding spatial location and no color overlay will be applied.
- **background_class_id** (*int, optional*) – Class id to interpret as the background class. See *draw_background*.
- **background_threshold** (*None, optional*) – Deprecated. This parameter is ignored.

Returns Rendered overlays as `uint8` arrays. Always a **list** containing one RGB image per segmentation map array channel.

Return type list of (H,W,3) ndarray

get_arr(*self*)

Return the seg.map array, with original dtype and shape ndim.

Here, "original" denotes the dtype and number of shape dimensions that was used when the *SegmentationMapsOnImage* instance was created, i.e. upon the call of `SegmentationMapsOnImage.__init__()`. Internally, this class may use a different dtype and shape to simplify computations.

Note: The height and width may have changed compared to the original input due to e.g. pooling operations.

Returns Segmentation map array. Same dtype and number of dimensions as was originally used when the *SegmentationMapsOnImage* instance was created.

Return type ndarray

get_arr_int(*self*, **args*, ***kwargs*)

Deprecated. Use `SegmentationMapsOnImage.get_arr()` instead.

Return the seg.map array, with original dtype and shape ndim.

pad(*self*, *top=0*, *right=0*, *bottom=0*, *left=0*, *mode='constant'*, *cval=0*)

Pad the segmentation maps at their top/right/bottom/left side.

Parameters

- **top** (*int, optional*) – Amount of pixels to add at the top side of the segmentation map. Must be `0` or greater.
- **right** (*int, optional*) – Amount of pixels to add at the right side of the segmentation map. Must be `0` or greater.
- **bottom** (*int, optional*) – Amount of pixels to add at the bottom side of the segmentation map. Must be `0` or greater.
- **left** (*int, optional*) – Amount of pixels to add at the left side of the segmentation map. Must be `0` or greater.
- **mode** (*str, optional*) – Padding mode to use. See *pad()* for details.
- **cval** (*number, optional*) – Value to use for padding if *mode* is `constant`. See *pad()* for details.

Returns Padded segmentation map with height `H'=H+top+bottom` and width `W'=W+left+right`.

Return type *imgaug.augmentables.segmaps.SegmentationMapsOnImage*

pad_to_aspect_ratio(*self*, *aspect_ratio*, *mode='constant'*, *cval=0*, *return_pad_amounts=False*)
Pad the segmentation maps until they match a target aspect ratio.

Depending on which dimension is smaller (height or width), only the corresponding sides (left/right or top/bottom) will be padded. In each case, both of the sides will be padded equally.

Parameters

- **aspect_ratio** (*float*) – Target aspect ratio, given as width/height. E.g. `2.0` denotes the image having twice as much width as height.
- **mode** (*str, optional*) – Padding mode to use. See *pad()* for details.
- **cval** (*number, optional*) – Value to use for padding if *mode* is `constant`. See *pad()* for details.
- **return_pad_amounts** (*bool, optional*) – If `False`, then only the padded instance will be returned. If `True`, a tuple with two entries will be returned, where the first entry is the padded instance and the second entry are the amounts by which each array side was padded. These amounts are again a tuple of the form `(top, right, bottom, left)`, with each value being an integer.

Returns

- *imgaug.augmentables.segmaps.SegmentationMapsOnImage* – Padded segmentation map as *SegmentationMapsOnImage* instance.
- *tuple of int* – Amounts by which the instance's array was padded on each side, given as a tuple `(top, right, bottom, left)`. This tuple is only returned if *return_pad_amounts* was set to `True`.

resize(*self*, *sizes*, *interpolation='nearest'*)
Resize the seg.map(s) array given a target size and interpolation.

Parameters

- **sizes** (*float or iterable of int or iterable of float*) – New size of the array in `(height, width)`. See *imresize_single_image()* for details.
- **interpolation** (*None or str or int, optional*) – The interpolation to use during resize. Nearest neighbour interpolation (`"nearest"`) is almost always the best choice. See *imresize_single_image()* for details.

Returns Resized segmentation map object.

Return type *imgaug.augmentables.segmaps.SegmentationMapsOnImage*

scale(*self*, **args*, ***kwargs*)
Deprecated. Use `SegmentationMapsOnImage.resize()` instead. resize() has the exactly same interface.

Resize the seg.map(s) array given a target size and interpolation.

13.16 imgaug.augmentables.utils

Utility functions used in augmentable modules.

imgaug.augmentables.utils.**convert_cbaois_to_kpsois**(*cbaois*)
Convert coordinate-based augmentables to KeypointsOnImage instances.

Added in 0.4.0.

Parameters **cbaois** (*list of imgaug.augmentables.bbs.BoundingBoxesOnImage or list of imgaug.augmentables.bbs.PolygonsOnImage or list of imgaug.augmentables.bbs.LineStringsOnImage or imgaug.augmentables.bbs.BoundingBoxesOnImage or imgaug.augmentables.bbs.PolygonsOnImage or imgaug.augmentables.bbs.LineStringsOnImage*) – Coordinate-based augmentables to convert, e.g. bounding boxes.

Returns `KeypointsOnImage` instances containing the coordinates of input *cbaois*.

Return type list of imgaug.augmentables.kps.KeypointsOnImage or *imgaug.augmentables.kps.KeypointsOnImage*

imgaug.augmentables.utils.**copy_augmentables**(*augmentables*)

imgaug.augmentables.utils.**deepcopy_fast**(*obj*)

imgaug.augmentables.utils.**interpolate_point_pair**(*point_a*, *point_b*, *nb_steps*)
Interpolate `N` points on a line segment.

Parameters

- **point_a** (*iterable of number*) – Start point of the line segment, given as `(x,y)` coordinates.
- **point_b** (*iterable of number*) – End point of the line segment, given as `(x,y)` coordinates.
- **nb_steps** (*int*) – Number of points to interpolate between *point_a* and *point_b*.

Returns The interpolated points. Does not include *point_a*.

Return type list of tuple of number

imgaug.augmentables.utils.**interpolate_points**(*points*, *nb_steps*, *closed=True*)
Interpolate `N` on each line segment in a line string.

Parameters

- **points** (*iterable of iterable of number*) – Points on the line segments, each one given as `(x,y)` coordinates. They are assumed to form one connected line string.
- **nb_steps** (*int*) – Number of points to interpolate on each individual line string.
- **closed** (*bool, optional*) – If `True` the output contains the last point in *points*. Otherwise it does not (but it will contain the interpolated points leading to the last point).

Returns Coordinates of *points*, with additional *nb_steps* new points interpolated between each point pair. If *closed* is `False`, the last point in *points* is not returned.

Return type list of tuple of number

imgaug.augmentables.utils.**interpolate_points_by_max_distance**(*points*, *max_distance*, *closed=True*)
Interpolate points with distance `d` on a line string.

For a list of points `A, B, C`, if the distance between `A` and `B` is greater than *max_distance*, it will place at least one point between `A` and `B` at `A + max_distance * (B - A)`. Multiple points can be placed between the two points if they are far enough away from each other. The process is repeated for `B` and `C`.

Parameters

- **points** (*iterable of iterable of number*) – Points on the line segments, each one given as `(x,y)` coordinates. They are assumed to form one connected line string.
- **max_distance** (*number*) – Maximum distance between any two points in the result.
- **closed** (*bool, optional*) – If `True` the output contains the last point in *points*. Otherwise it does not (but it will contain the interpolated points leading to the last point).

Returns Coordinates of *points*, with interpolated points added to the iterable. If *closed* is `False`, the last point in *points* is not returned.

Return type list of tuple of number

imgaug.augmentables.utils.**invert_convert_cbaois_to_kpsois_**(*cbaois*, *kpsois*)

Invert the output of `convert_to_cbaois_to_kpsois()` in-place.

This function writes in-place into *cbaois*.

Added in 0.4.0.

Parameters

- **cbaois** (*list of imgaug.augmentables.bbs.BoundingBoxesOnImage or list of imgaug.augmentables.bbs.PolygonsOnImage or list of imgaug.augmentables.bbs.LineStringsOnImage or imgaug.augmentables.bbs.BoundingBoxesOnImage or imgaug.augmentables.bbs.PolygonsOnImage or imgaug.augmentables.bbs.LineStringsOnImage*) – Original coordinate-based augmentables before they were converted, i.e. the same inputs as provided to `convert_to_kpsois()`.
- **kpsois** (*list of imgaug.augmentables.kps.KeypointsOnImages or imgaug.augmentables.kps.KeypointsOnImages*) – Keypoints to convert back to the types of *cbaois*, i.e. the outputs of *convert_cbaois_to_kpsois()*.

Returns Parameter *cbaois*, with updated coordinates and shapes derived from *kpsois*. *cbaois* is modified in-place.

Return type list of imgaug.augmentables.bbs.BoundingBoxesOnImage or list of imgaug.augmentables.bbs.PolygonsOnImage or list of imgaug.augmentables.bbs.LineStringsOnImage or *imgaug.augmentables.bbs.BoundingBoxesOnImage* or imgaug.augmentables.bbs.PolygonsOnImage or imgaug.augmentables.bbs.LineStringsOnImage

imgaug.augmentables.utils.**normalize_shape**(*shape*)

Normalize a shape `tuple` or `array` to a shape `tuple`.

Parameters **shape** (*tuple of int or ndarray*) – The input to normalize. May optionally be an array.

Returns Shape `tuple`.

Return type tuple of int

imgaug.augmentables.utils.**project_coords**(*coords*, *from_shape*, *to_shape*)

Project coordinates from one image shape to another.

This performs a relative projection, e.g. a point at `60%` of the old image width will be at `60%` of the new image width after projection.

Parameters

- **coords** (*ndarray or list of tuple of number*) – Coordinates to project. Either an `(N,2)` numpy array or a `list` containing `(x,y)` coordinate `tuple` s.
- **from_shape** (*tuple of int or ndarray*) – Old image shape.

- **to_shape** (*tuple of int or ndarray*) – New image shape.

Returns Projected coordinates as `(N,2) float32` numpy array.

Return type ndarray

`imgaug.augmentables.utils.`**`project_coords_`**(*coords*, *from_shape*, *to_shape*)

Project coordinates from one image shape to another in-place.

This performs a relative projection, e.g. a point at `60%` of the old image width will be at `60%` of the new image width after projection.

Added in 0.4.0.

Parameters

- **coords** (*ndarray or list of tuple of number*) – Coordinates to project. Either an `(N,2)` numpy array or a `list` containing `(x,y)` coordinate `tuple` s.
- **from_shape** (*tuple of int or ndarray*) – Old image shape.
- **to_shape** (*tuple of int or ndarray*) – New image shape.

Returns Projected coordinates as `(N,2) float32` numpy array. This function may change the input data in-place.

Return type ndarray

13.17 imgaug.augmenters.arithmetic

Augmenters that perform simple arithmetic changes.

List of augmenters:

- *Add*
- *AddElementwise*
- *AdditiveGaussianNoise*
- *AdditiveLaplaceNoise*
- *AdditivePoissonNoise*
- *Multiply*
- *MultiplyElementwise*
- *Cutout*
- *Dropout*
- *CoarseDropout*
- *Dropout2d*
- *TotalDropout*
- *ReplaceElementwise*
- *ImpulseNoise*
- *SaltAndPepper*
- *CoarseSaltAndPepper*
- *Salt*

- *CoarseSalt*
- *Pepper*
- *CoarsePepper*
- *Invert*
- *Solarize*
- *ContrastNormalization*
- *JpegCompression*

class imgaug.augmenters.arithmetic.**Add**(*value=(-20, 20), per_channel=False, seed=None, name=None, random_state='deprecated', deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Add a value to all pixels in an image.

Supported dtypes:

See *add_scalar()*.

Parameters

- **value** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) –

  Value to add to all pixels.

  – If a number, exactly that value will always be used.

  – If a tuple `(a, b)`, then a value from the discrete interval `[a..b]` will be sampled per image.

  – If a list, then a random value will be sampled from that list per image.

  – If a `StochasticParameter`, then a value will be sampled per image from that parameter.

- **per_channel** (*bool or float or imgaug.parameters.StochasticParameter, optional*) – Whether to use (imagewise) the same sample(s) for all channels (`False`) or to sample value(s) for each channel (`True`). Setting this to `True` will therefore lead to different transformations per image *and* channel, otherwise only per image. If this value is a float `p`, then for `p` percent of all images *per_channel* will be treated as `True`. If it is a `StochasticParameter` it is expected to produce samples with values between `0.0` and `1.0`, where values `>0.5` will lead to per-channel behaviour (i.e. same as `True`).

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.

- **name** (*None or str, optional*) – See `__init__()`.

- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.

- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.Add(10)
```

Always adds a value of 10 to all channels of all pixels of all input images.

```
>>> aug = iaa.Add((-10, 10))
```

Adds a value from the discrete interval `[-10..10]` to all pixels of input images. The exact value is sampled per image.

```
>>> aug = iaa.Add((-10, 10), per_channel=True)
```

Adds a value from the discrete interval `[-10..10]` to all pixels of input images. The exact value is sampled per image *and* channel, i.e. to a red-channel it might add 5 while subtracting 7 from the blue channel of the same image.

```
>>> aug = iaa.Add((-10, 10), per_channel=0.5)
```

Identical to the previous example, but the *per_channel* feature is only active for 50 percent of all images.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.

Continued on next page

Table 51 – continued from previous page

find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.arithmetic.**AddElementwise**(*value=(-20, 20)*, *per_channel=False*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Add to the pixels of images values that are pixelwise randomly sampled.

While the Add Augmenter samples one value to add *per image* (and optionally per channel), this augmenter samples different values per image and *per pixel* (and optionally per channel), i.e. intensities of neighbouring pixels may be increased/decreased by different amounts.

Supported dtypes:

See *add_elementwise()*.

Parameters

- **value** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter, optional*) –

  Value to add to the pixels.

  – If an int, exactly that value will always be used.

  – If a tuple (a, b), then values from the discrete interval [a..b] will be sampled per image and pixel.

  – If a list of integers, a random value will be sampled from the list per image and pixel.

  - If a `StochasticParameter`, then values will be sampled per image and pixel from that parameter.

- **per_channel** (*bool or float or imgaug.parameters.StochasticParameter, optional*) – Whether to use (imagewise) the same sample(s) for all channels (`False`) or to sample value(s) for each channel (`True`). Setting this to `True` will therefore lead to different transformations per image *and* channel, otherwise only per image. If this value is a float `p`, then for `p` percent of all images *per_channel* will be treated as `True`. If it is a `StochasticParameter` it is expected to produce samples with values between `0.0` and `1.0`, where values `>0.5` will lead to per-channel behaviour (i.e. same as `True`).
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.AddElementwise(10)
```

Always adds a value of 10 to all channels of all pixels of all input images.

```
>>> aug = iaa.AddElementwise((-10, 10))
```

Samples per image and pixel a value from the discrete interval `[-10..10]` and adds that value to the respective pixel.

```
>>> aug = iaa.AddElementwise((-10, 10), per_channel=True)
```

Samples per image, pixel *and also channel* a value from the discrete interval `[-10..10]` and adds it to the respective pixel's channel value. Therefore, added values may differ between channels of the same pixel.

```
>>> aug = iaa.AddElementwise((-10, 10), per_channel=0.5)
```

Identical to the previous example, but the *per_channel* feature is only active for 50 percent of all images.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.

Continued on next page

Table 52 – continued from previous page

augment_batches(self, batches[, hooks, ...])	Augment multiple batches.
augment_bounding_boxes(self, ...[, parents, ...])	Augment a batch of bounding boxes.
augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.arithmetic.**AdditiveGaussianNoise**(*loc=0*, *scale=(0, 15)*, *per_channel=False*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.arithmetic.AddElementwise*

Add noise sampled from gaussian distributions elementwise to images.

This augmenter samples and adds noise elementwise, i.e. it can add different noise values to neighbouring pixels and is comparable to `AddElementwise`.

Supported dtypes:

See *AddElementwise*.

Parameters

- **loc** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) –

  Mean of the normal distribution from which the noise is sampled.

  – If a number, exactly that value will always be used.

  – If a tuple `(a, b)`, a random value from the interval `[a, b]` will be sampled per image.

  – If a list, then a random value will be sampled from that list per image.

  – If a `StochasticParameter`, a value will be sampled from the parameter per image.

- **scale** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Standard deviation of the normal distribution that generates the noise. Must be `>=0`. If `0` then *loc* will simply be added to all pixels.

  – If a number, exactly that value will always be used.

  – If a tuple `(a, b)`, a random value from the interval `[a, b]` will be sampled per image.

  – If a list, then a random value will be sampled from that list per image.

  – If a `StochasticParameter`, a value will be sampled from the parameter per image.

- **per_channel** (*bool or float or imgaug.parameters.StochasticParameter, optional*) – Whether to use (imagewise) the same sample(s) for all channels (`False`) or to sample value(s) for each channel (`True`). Setting this to `True` will therefore lead to different transformations per image *and* channel, otherwise only per image. If this value is a float `p`, then for `p` percent of all images *per_channel* will be treated as `True`. If it is a `StochasticParameter` it is expected to produce samples with values between `0.0` and `1.0`, where values `>0.5` will lead to per-channel behaviour (i.e. same as `True`).

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.

- **name** (*None or str, optional*) – See `__init__()`.

- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage

will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.

- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.AdditiveGaussianNoise(scale=0.1*255)
```

Adds gaussian noise from the distribution `N(0, 0.1*255)` to images. The samples are drawn per image and pixel.

```
>>> aug = iaa.AdditiveGaussianNoise(scale=(0, 0.1*255))
```

Adds gaussian noise from the distribution `N(0, s)` to images, where `s` is sampled per image from the interval `[0, 0.1*255]`.

```
>>> aug = iaa.AdditiveGaussianNoise(scale=0.1*255, per_channel=True)
```

Adds gaussian noise from the distribution `N(0, 0.1*255)` to images, where the noise value is different per image and pixel *and* channel (e.g. a different one for red, green and blue channels of the same pixel). This leads to "colorful" noise.

```
>>> aug = iaa.AdditiveGaussianNoise(scale=0.1*255, per_channel=0.5)
```

Identical to the previous example, but the *per_channel* feature is only active for 50 percent of all images.

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.

Continued on next page

Table 53 – continued from previous page

copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.arithmetic.**AdditiveLaplaceNoise**(*loc=0*, *scale=(0, 15)*, *per_channel=False*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.arithmetic.AddElementwise*

Add noise sampled from laplace distributions elementwise to images.

The laplace distribution is similar to the gaussian distribution, but puts more weight on the long tail. Hence, this noise will add more outliers (very high/low values). It is somewhere between gaussian noise and salt and pepper noise.

Values of around `255 * 0.05` for *scale* lead to visible noise (for `uint8`). Values of around `255 * 0.10` for *scale* lead to very visible noise (for `uint8`). It is recommended to usually set *per_channel* to `True`.

This augmenter samples and adds noise elementwise, i.e. it can add different noise values to neighbouring pixels and is comparable to `AddElementwise`.

Supported dtypes:

See *AddElementwise*.

Parameters

- **loc** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) –

  Mean of the laplace distribution that generates the noise.

  – If a number, exactly that value will always be used.

  – If a tuple `(a, b)`, a random value from the interval `[a, b]` will be sampled per image.

  – If a list, then a random value will be sampled from that list per image.

  – If a `StochasticParameter`, a value will be sampled from the parameter per image.

- **scale** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Standard deviation of the laplace distribution that generates the noise. Must be `>=0`. If `0` then only *loc* will be used. Recommended to be around `255*0.05`.

  – If a number, exactly that value will always be used.

  – If a tuple `(a, b)`, a random value from the interval `[a, b]` will be sampled per image.

  – If a list, then a random value will be sampled from that list per image.

  – If a `StochasticParameter`, a value will be sampled from the parameter per image.

- **per_channel** (*bool or float or imgaug.parameters.StochasticParameter, optional*) – Whether to use (imagewise) the same sample(s) for all channels (`False`) or to sample value(s) for each channel (`True`). Setting this to `True` will therefore lead to different transformations per image *and* channel, otherwise only per image. If this value is a float `p`, then for `p` percent of all images *per_channel* will be treated as `True`. If it is a `StochasticParameter` it is expected to produce samples with values between `0.0` and `1.0`, where values `>0.5` will lead to per-channel behaviour (i.e. same as `True`).

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.

- **name** (*None or str, optional*) – See `__init__()`.

- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.

- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.AdditiveLaplaceNoise(scale=0.1*255)
```

Adds laplace noise from the distribution `Laplace(0, 0.1*255)` to images. The samples are drawn per image and pixel.

```
>>> aug = iaa.AdditiveLaplaceNoise(scale=(0, 0.1*255))
```

Adds laplace noise from the distribution `Laplace(0, s)` to images, where `s` is sampled per image from the interval `[0, 0.1*255]`.

```
>>> aug = iaa.AdditiveLaplaceNoise(scale=0.1*255, per_channel=True)
```

Adds laplace noise from the distribution `Laplace(0, 0.1*255)` to images, where the noise value is different per image and pixel *and* channel (e.g. a different one for the red, green and blue channels of the same pixel). This leads to "colorful" noise.

```
>>> aug = iaa.AdditiveLaplaceNoise(scale=0.1*255, per_channel=0.5)
```

Identical to the previous example, but the *per_channel* feature is only active for 50 percent of all images.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.

Continued on next page

Table 54 – continued from previous page

get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.arithmetic.**AdditivePoissonNoise**(*lam=(0.0, 15.0)*, *per_channel=False*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.arithmetic.AddElementwise*

Add noise sampled from poisson distributions elementwise to images.

Poisson noise is comparable to gaussian noise, as e.g. generated via AdditiveGaussianNoise. As poisson distributions produce only positive numbers, the sign of the sampled values are here randomly flipped.

Values of around 10.0 for *lam* lead to visible noise (for uint8). Values of around 20.0 for *lam* lead to very visible noise (for uint8). It is recommended to usually set *per_channel* to True.

This augmenter samples and adds noise elementwise, i.e. it can add different noise values to neighbouring pixels and is comparable to AddElementwise.

Supported dtypes:

See *AddElementwise*.

Parameters

- **lam** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Lambda parameter of the poisson distribution. Must be >=0. Recommended values are around 0.0 to 10.0.
  - If a number, exactly that value will always be used.
  - If a tuple (a, b), a random value from the interval [a, b] will be sampled per image.
  - If a list, then a random value will be sampled from that list per image.
  - If a StochasticParameter, a value will be sampled from the parameter per image.
- **per_channel** (*bool or float or imgaug.parameters.StochasticParameter, optional*) – Whether to use (imagewise) the same sample(s) for all channels (False) or to sample

value(s) for each channel (`True`). Setting this to `True` will therefore lead to different transformations per image *and* channel, otherwise only per image. If this value is a float `p`, then for `p` percent of all images *per_channel* will be treated as `True`. If it is a `StochasticParameter` it is expected to produce samples with values between `0.0` and `1.0`, where values `>0.5` will lead to per-channel behaviour (i.e. same as `True`).

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.AdditivePoissonNoise(lam=5.0)
```

Adds poisson noise sampled from a poisson distribution with a `lambda` parameter of `5.0` to images. The samples are drawn per image and pixel.

```
>>> aug = iaa.AdditivePoissonNoise(lam=(0.0, 15.0))
```

Adds poisson noise sampled from `Poisson(x)` to images, where `x` is randomly sampled per image from the interval `[0.0, 15.0]`.

```
>>> aug = iaa.AdditivePoissonNoise(lam=5.0, per_channel=True)
```

Adds poisson noise sampled from `Poisson(5.0)` to images, where the values are different per image and pixel *and* channel (e.g. a different one for red, green and blue channels for the same pixel).

```
>>> aug = iaa.AdditivePoissonNoise(lam=(0.0, 15.0), per_channel=True)
```

Adds poisson noise sampled from `Poisson(x)` to images, with `x` being sampled from `uniform(0.0, 15.0)` per image and channel. This is the *recommended* configuration.

```
>>> aug = iaa.AdditivePoissonNoise(lam=(0.0, 15.0), per_channel=0.5)
```

Identical to the previous example, but the *per_channel* feature is only active for 50 percent of all images.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.

Continued on next page

Table 55 – continued from previous page

augment_batch_(self, batch[, parents, hooks])	Augment a single batch in-place.
augment_batches(self, batches[, hooks, ...])	Augment multiple batches.
augment_bounding_boxes(self, ...[, parents, ...])	Augment a batch of bounding boxes.
augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.arithmetic.**CoarseDropout**(*p=(0.02, 0.1)*, *size_px=None*, *size_percent=None*, *per_channel=False*, *min_size=3*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.arithmetic.MultiplyElementwise*

Set rectangular areas within images to zero.

In contrast to `Dropout`, these areas can have larger sizes. (E.g. you might end up with three large black rectangles in an image.) Note that the current implementation leads to correlated sizes, so if e.g. there is any thin and high rectangle that is dropped, there is a high likelihood that all other dropped areas are also thin and high.

This method is implemented by generating the dropout mask at a lower resolution (than the image has) and then upsampling the mask before dropping the pixels.

This augmenter is similar to Cutout. Usually, cutout is defined as an operation that drops exactly one rectangle from an image, while here `CoarseDropout` can drop multiple rectangles (with some correlation between the sizes of these rectangles).

Supported dtypes:

See *MultiplyElementwise*.

Parameters

- **p** (*float or tuple of float or imgaug.parameters.StochasticParameter, optional*) – The probability of any pixel being dropped (i.e. set to zero) in the lower-resolution dropout mask.
  - If a float, then that value will be used for all pixels. A value of `1.0` would mean, that all pixels will be dropped. A value of `0.0` would lead to no pixels being dropped.
  - If a tuple `(a, b)`, then a value `p` will be sampled from the interval `[a, b]` per image and be used as the dropout probability.
  - If a list, then a value will be sampled from that list per batch and used as the probability.
  - If a `StochasticParameter`, then this parameter will be used to determine per pixel whether it should be *kept* (sampled value of `>0.5`) or shouldn't be kept (sampled value of `<=0.5`). If you instead want to provide the probability as a stochastic parameter, you can usually do `imgaug.parameters.Binomial(1-p)` to convert parameter *p* to a 0/1 representation.
- **size_px** (*None or int or tuple of int or imgaug.parameters.StochasticParameter, optional*) – The size of the lower resolution image from which to sample the dropout mask in absolute pixel dimensions. Note that this means that *lower* values of this parameter lead to *larger* areas being dropped (as any pixel in the lower resolution image will correspond to a larger area at the original resolution).
  - If `None` then *size_percent* must be set.
  - If an integer, then that size will always be used for both height and width. E.g. a value of `3` would lead to a `3x3` mask, which is then upsampled to `HxW`, where `H` is the image size and `W` the image width.
  - If a tuple `(a, b)`, then two values `M`, `N` will be sampled from the discrete interval `[a..b]`. The dropout mask will then be generated at size `MxN` and upsampled to `HxW`.
  - If a `StochasticParameter`, then this parameter will be used to determine the sizes. It is expected to be discrete.

- **size_percent** (*None or float or tuple of float or imgaug.parameters.StochasticParameter, optional*) – The size of the lower resolution image from which to sample the dropout mask *in percent* of the input image. Note that this means that *lower* values of this parameter lead to *larger* areas being dropped (as any pixel in the lower resolution image will correspond to a larger area at the original resolution).
  - If `None` then *size_px* must be set.
  - If a float, then that value will always be used as the percentage of the height and width (relative to the original size). E.g. for value `p`, the mask will be sampled from `(p*H)x(p*W)` and later upsampled to `HxW`.
  - If a tuple `(a, b)`, then two values `m`, `n` will be sampled from the interval `(a, b)` and used as the size fractions, i.e the mask size will be `(m*H)x(n*W)`.
  - If a `StochasticParameter`, then this parameter will be used to sample the percentage values. It is expected to be continuous.
- **per_channel** (*bool or float or imgaug.parameters.StochasticParameter, optional*) – Whether to use (imagewise) the same sample(s) for all channels (`False`) or to sample value(s) for each channel (`True`). Setting this to `True` will therefore lead to different transformations per image *and* channel, otherwise only per image. If this value is a float `p`, then for `p` percent of all images *per_channel* will be treated as `True`. If it is a `StochasticParameter` it is expected to produce samples with values between `0.0` and `1.0`, where values `>0.5` will lead to per-channel behaviour (i.e. same as `True`).
- **min_size** (*int, optional*) – Minimum height and width of the low resolution mask. If *size_percent* or *size_px* leads to a lower value than this, *min_size* will be used instead. This should never have a value of less than `2`, otherwise one may end up with a `1x1` low resolution mask, leading easily to the whole image being dropped.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.CoarseDropout(0.02, size_percent=0.5)
```

Drops `2` percent of all pixels on a lower-resolution image that has `50` percent of the original image's size, leading to dropped areas that have roughly `2x2` pixels size.

```
>>> aug = iaa.CoarseDropout((0.0, 0.05), size_percent=(0.05, 0.5))
```

Generates a dropout mask at 5 to 50 percent of each input image's size. In that mask, 0 to 5 percent of all pixels are marked as being dropped. The mask is afterwards projected to the input image's size to apply the actual dropout operation.

```
>>> aug = iaa.CoarseDropout((0.0, 0.05), size_px=(2, 16))
```

Same as the previous example, but the lower resolution image has 2 to 16 pixels size. On images of e.g. 224x224` pixels in size this would lead to fairly large areas being dropped (height/width of ``224/2 to 224/16).

```
>>> aug = iaa.CoarseDropout(0.02, size_percent=0.5, per_channel=True)
```

Drops 2 percent of all pixels at 50 percent resolution (2x2 sizes) in a channel-wise fashion, i.e. it is unlikely for any pixel to have all channels set to zero (black pixels).

```
>>> aug = iaa.CoarseDropout(0.02, size_percent=0.5, per_channel=0.5)
```

Same as the previous example, but the *per_channel* feature is only active for 50 percent of all images.

Methods

__call__(self, *args, **kwargs)	Alias for *augment()*.
augment(self[, return_batch, hooks])	Augment a batch.
augment_batch(self, batch[, hooks])	**Deprecated**.
augment_batch_(self, batch[, parents, hooks])	Augment a single batch in-place.
augment_batches(self, batches[, hooks, ...])	Augment multiple batches.
augment_bounding_boxes(self, ...[, parents, ...])	Augment a batch of bounding boxes.
augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.

Continued on next page

Table 56 – continued from previous page

find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.arithmetic.**CoarsePepper**(*p=(0.02, 0.1)*, *size_px=None*, *size_percent=None*, *per_channel=False*, *min_size=3*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.arithmetic.ReplaceElementwise*

Replace rectangular areas in images with black-ish pixel noise.

Supported dtypes:

See *ReplaceElementwise*.

Parameters

- **p** (*float or tuple of float or list of float or imgaug.parameters.StochasticParameter, optional*) –

  Probability of changing a pixel to pepper noise.

  - If a float, then that value will always be used as the probability.
  - If a tuple `(a, b)`, then a probability will be sampled uniformly per image from the interval `[a, b]`.
  - If a list, then a random value will be sampled from that list per image.
  - If a `StochasticParameter`, then a lower-resolution mask will be sampled from that parameter per image. Any value `>0.5` in that mask will denote a spatial location that is to be replaced by pepper noise.

- **size_px** (*int or tuple of int or imgaug.parameters.StochasticParameter, optional*) – The size of the lower resolution image from which to sample the replacement mask in absolute pixel dimensions. Note that this means that *lower* values of this parameter lead to *larger* areas

being replaced (as any pixel in the lower resolution image will correspond to a larger area at the original resolution).

  - If `None` then *size_percent* must be set.
  - If an integer, then that size will always be used for both height and width. E.g. a value of `3` would lead to a `3x3` mask, which is then upsampled to `HxW`, where `H` is the image size and `W` the image width.
  - If a tuple `(a, b)`, then two values `M`, `N` will be sampled from the discrete interval `[a..b]`. The mask will then be generated at size `MxN` and upsampled to `HxW`.
  - If a `StochasticParameter`, then this parameter will be used to determine the sizes. It is expected to be discrete.

- **size_percent** (*float or tuple of float or imgaug.parameters.StochasticParameter, optional*) – The size of the lower resolution image from which to sample the replacement mask *in percent* of the input image. Note that this means that *lower* values of this parameter lead to *larger* areas being replaced (as any pixel in the lower resolution image will correspond to a larger area at the original resolution).
  - If `None` then *size_px* must be set.
  - If a float, then that value will always be used as the percentage of the height and width (relative to the original size). E.g. for value `p`, the mask will be sampled from `(p*H)x(p*W)` and later upsampled to `HxW`.
  - If a tuple `(a, b)`, then two values `m`, `n` will be sampled from the interval `(a, b)` and used as the size fractions, i.e the mask size will be `(m*H)x(n*W)`.
  - If a `StochasticParameter`, then this parameter will be used to sample the percentage values. It is expected to be continuous.
- **per_channel** (*bool or float or imgaug.parameters.StochasticParameter, optional*) – Whether to use (imagewise) the same sample(s) for all channels (`False`) or to sample value(s) for each channel (`True`). Setting this to `True` will therefore lead to different transformations per image *and* channel, otherwise only per image. If this value is a float `p`, then for `p` percent of all images *per_channel* will be treated as `True`. If it is a `StochasticParameter` it is expected to produce samples with values between `0.0` and `1.0`, where values `>0.5` will lead to per-channel behaviour (i.e. same as `True`).
- **min_size** (*int, optional*) – Minimum size of the low resolution mask, both width and height. If *size_percent* or *size_px* leads to a lower value than this, *min_size* will be used instead. This should never have a value of less than 2, otherwise one may end up with a `1x1` low resolution mask, leading easily to the whole image being replaced.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.CoarsePepper(0.05, size_percent=(0.01, 0.1))
```

Mark 5% of all pixels in a mask to be replaced by pepper noise. The mask has 1% to 10% the size of the input image. The mask is then upscaled to the input image size, leading to large rectangular areas being marked as to be replaced. These areas are then replaced in the input image by pepper noise.

Methods

__call__(self, *args, **kwargs)	Alias for *augment()*.
augment(self[, return_batch, hooks])	Augment a batch.
augment_batch(self, batch[, hooks])	**Deprecated**.
augment_batch_(self, batch[, parents, hooks])	Augment a single batch in-place.
augment_batches(self, batches[, hooks, ...])	Augment multiple batches.
augment_bounding_boxes(self, ...[, parents, ...])	Augment a batch of bounding boxes.
augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.

Continued on next page

Table 57 – continued from previous page

pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.arithmetic.**CoarseSalt**(*p=(0.02, 0.1), size_px=None, size_percent=None, per_channel=False, min_size=3, seed=None, name=None, random_state='deprecated', deterministic='deprecated'*)

Bases: *imgaug.augmenters.arithmetic.ReplaceElementwise*

Replace rectangular areas in images with white-ish pixel noise.

See also the similar `CoarseSaltAndPepper`.

Supported dtypes:

See *ReplaceElementwise*.

Parameters

- **p** (*float or tuple of float or list of float or imgaug.parameters.StochasticParameter, optional*) –

  Probability of changing a pixel to salt noise.

  - If a float, then that value will always be used as the probability.
  - If a tuple `(a, b)`, then a probability will be sampled uniformly per image from the interval `[a, b]`.
  - If a list, then a random value will be sampled from that list per image.
  - If a `StochasticParameter`, then a lower-resolution mask will be sampled from that parameter per image. Any value `>0.5` in that mask will denote a spatial location that is to be replaced by salt noise.

- **size_px** (*int or tuple of int or imgaug.parameters.StochasticParameter, optional*) – The size of the lower resolution image from which to sample the replacement mask in absolute pixel dimensions. Note that this means that *lower* values of this parameter lead to *larger* areas being replaced (as any pixel in the lower resolution image will correspond to a larger area at the original resolution).

  - If `None` then *size_percent* must be set.
  - If an integer, then that size will always be used for both height and width. E.g. a value of `3` would lead to a `3x3` mask, which is then upsampled to `HxW`, where `H` is the image size and `W` the image width.

  - If a tuple `(a, b)`, then two values `M`, `N` will be sampled from the discrete interval `[a..b]`. The mask will then be generated at size `MxN` and upsampled to `HxW`.
  - If a `StochasticParameter`, then this parameter will be used to determine the sizes. It is expected to be discrete.

- **size_percent** (*float or tuple of float or imgaug.parameters.StochasticParameter, optional*) – The size of the lower resolution image from which to sample the replacement mask *in percent* of the input image. Note that this means that *lower* values of this parameter lead to *larger* areas being replaced (as any pixel in the lower resolution image will correspond to a larger area at the original resolution).
  - If `None` then *size_px* must be set.
  - If a float, then that value will always be used as the percentage of the height and width (relative to the original size). E.g. for value `p`, the mask will be sampled from `(p*H)x(p*W)` and later upsampled to `HxW`.
  - If a tuple `(a, b)`, then two values `m`, `n` will be sampled from the interval `(a, b)` and used as the size fractions, i.e the mask size will be `(m*H)x(n*W)`.
  - If a `StochasticParameter`, then this parameter will be used to sample the percentage values. It is expected to be continuous.
- **per_channel** (*bool or float or imgaug.parameters.StochasticParameter, optional*) – Whether to use (imagewise) the same sample(s) for all channels (`False`) or to sample value(s) for each channel (`True`). Setting this to `True` will therefore lead to different transformations per image *and* channel, otherwise only per image. If this value is a float `p`, then for `p` percent of all images *per_channel* will be treated as `True`. If it is a `StochasticParameter` it is expected to produce samples with values between `0.0` and `1.0`, where values `>0.5` will lead to per-channel behaviour (i.e. same as `True`).
- **min_size** (*int, optional*) – Minimum height and width of the low resolution mask. If *size_percent* or *size_px* leads to a lower value than this, *min_size* will be used instead. This should never have a value of less than `2`, otherwise one may end up with a `1x1` low resolution mask, leading easily to the whole image being replaced.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.CoarseSalt(0.05, size_percent=(0.01, 0.1))
```

Mark `5%` of all pixels in a mask to be replaced by salt noise. The mask has `1%` to `10%` the size of the input image. The mask is then upscaled to the input image size, leading to large rectangular areas being marked as to be replaced. These areas are then replaced in the input image by salt noise.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.
`get_parameters`(self)	See *get_parameters()*.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`localize_random_state_`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`pool`(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
`remove_augmenters`(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
`remove_augmenters_`(self, func[, parents])	Remove in-place children of this augmenter that match a condition.

Continued on next page

Table 58 – continued from previous page

remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.arithmetic.**CoarseSaltAndPepper**(*p=(0.02, 0.1), size_px=None, size_percent=None, per_channel=False, min_size=3, seed=None, name=None, random_state='deprecated', deterministic='deprecated'*)

Bases: *imgaug.augmenters.arithmetic.ReplaceElementwise*

Replace rectangular areas in images with white/black-ish pixel noise.

This adds salt and pepper noise (noisy white-ish and black-ish pixels) to rectangular areas within the image. Note that this means that within these rectangular areas the color varies instead of each rectangle having only one color.

See also the similar `CoarseDropout`.

TODO replace dtype support with uint8 only, because replacement is geared towards that value range

Supported dtypes:

See *ReplaceElementwise*.

Parameters

- **p** (*float or tuple of float or list of float or imgaug.parameters.StochasticParameter, optional*) –

  Probability of changing a pixel to salt/pepper noise.

  - If a float, then that value will always be used as the probability.
  - If a tuple `(a, b)`, then a probability will be sampled uniformly per image from the interval `[a, b]`.
  - If a list, then a random value will be sampled from that list per image.
  - If a `StochasticParameter`, then a lower-resolution mask will be sampled from that parameter per image. Any value `>0.5` in that mask will denote a spatial location that is to be replaced by salt and pepper noise.

- **size_px** (*int or tuple of int or imgaug.parameters.StochasticParameter, optional*) – The size of the lower resolution image from which to sample the replacement mask in absolute pixel dimensions. Note that this means that *lower* values of this parameter lead to *larger* areas being replaced (as any pixel in the lower resolution image will correspond to a larger area at the original resolution).

  - If `None` then *size_percent* must be set.

  - If an integer, then that size will always be used for both height and width. E.g. a value of `3` would lead to a `3x3` mask, which is then upsampled to `HxW`, where `H` is the image size and `W` the image width.
  - If a tuple `(a, b)`, then two values `M`, `N` will be sampled from the discrete interval `[a..b]`. The mask will then be generated at size `MxN` and upsampled to `HxW`.
  - If a `StochasticParameter`, then this parameter will be used to determine the sizes. It is expected to be discrete.

- **size_percent** (*float or tuple of float or imgaug.parameters.StochasticParameter, optional*) – The size of the lower resolution image from which to sample the replacement mask *in percent* of the input image. Note that this means that *lower* values of this parameter lead to *larger* areas being replaced (as any pixel in the lower resolution image will correspond to a larger area at the original resolution).
  - If `None` then *size_px* must be set.
  - If a float, then that value will always be used as the percentage of the height and width (relative to the original size). E.g. for value `p`, the mask will be sampled from `(p*H)x(p*W)` and later upsampled to `HxW`.
  - If a tuple `(a, b)`, then two values `m`, `n` will be sampled from the interval `(a, b)` and used as the size fractions, i.e the mask size will be `(m*H)x(n*W)`.
  - If a `StochasticParameter`, then this parameter will be used to sample the percentage values. It is expected to be continuous.
- **per_channel** (*bool or float or imgaug.parameters.StochasticParameter, optional*) – Whether to use (imagewise) the same sample(s) for all channels (`False`) or to sample value(s) for each channel (`True`). Setting this to `True` will therefore lead to different transformations per image *and* channel, otherwise only per image. If this value is a float `p`, then for `p` percent of all images *per_channel* will be treated as `True`. If it is a `StochasticParameter` it is expected to produce samples with values between `0.0` and `1.0`, where values `>0.5` will lead to per-channel behaviour (i.e. same as `True`).
- **min_size** (*int, optional*) – Minimum height and width of the low resolution mask. If *size_percent* or *size_px* leads to a lower value than this, *min_size* will be used instead. This should never have a value of less than `2`, otherwise one may end up with a `1x1` low resolution mask, leading easily to the whole image being replaced.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.CoarseSaltAndPepper(0.05, size_percent=(0.01, 0.1))
```

Marks 5% of all pixels in a mask to be replaced by salt/pepper noise. The mask has 1% to 10% the size of the input image. The mask is then upscaled to the input image size, leading to large rectangular areas being marked as to be replaced. These areas are then replaced in the input image by salt/pepper noise.

```
>>> aug = iaa.CoarseSaltAndPepper(0.05, size_px=(4, 16))
```

Same as in the previous example, but the replacement mask before upscaling has a size between 4x4 and 16x16 pixels (the axis sizes are sampled independently, i.e. the mask may be rectangular).

```
>>> aug = iaa.CoarseSaltAndPepper(
>>>     0.05, size_percent=(0.01, 0.1), per_channel=True)
```

Same as in the first example, but mask and replacement are each sampled independently per image channel.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.

Continued on next page

Table 59 – continued from previous page

find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

imgaug.augmenters.arithmetic.**ContrastNormalization**(*alpha=1.0*, *per_channel=False*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Deprecated. Use `imgaug.contrast.LinearContrast` instead.

Change the contrast of images.

dtype support:

See `imgaug.augmenters.contrast.LinearContrast`.

Deprecated since 0.3.0.

Parameters

- **alpha** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Strength of the contrast normalization. Higher values than 1.0 lead to higher contrast, lower values decrease the contrast.
  - If a number, then that value will be used for all images.
  - If a tuple `(a, b)`, then a value will be sampled per image uniformly from the interval `[a, b]` and be used as the alpha value.
  - If a list, then a random value will be picked per image from that list.
  - If a `StochasticParameter`, then this parameter will be used to sample the alpha value per image.
- **per_channel** (*bool or float or imgaug.parameters.StochasticParameter, optional*) – Whether to use (imagewise) the same sample(s) for all channels (`False`) or to sample value(s) for each channel (`True`). Setting this to `True` will therefore lead to different transformations per image *and* channel, otherwise only per image. If this value is a float `p`, then for `p` percent of all images *per_channel* will be treated as `True`. If it is a

StochasticParameter it is expected to produce samples with values between 0.0 and 1.0, where values >0.5 will lead to per-channel behaviour (i.e. same as True).

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See __init__().
- **name** (*None or str, optional*) – See __init__().
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method to_deterministic() for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> iaa.ContrastNormalization((0.5, 1.5))
```

Decreases oder improves contrast per image by a random factor between 0.5 and 1.5. The factor 0.5 means that any difference from the center value (i.e. 128) will be halved, leading to less contrast.

```
>>> iaa.ContrastNormalization((0.5, 1.5), per_channel=0.5)
```

Same as before, but for 50 percent of all images the normalization is done independently per channel (i.e. factors can vary per channel for the same image). In the other 50 percent of all images, the factor is the same for all channels.

class imgaug.augmenters.arithmetic.**Cutout**(*nb_iterations=1*, *position='uniform'*, *size=0.2*, *squared=True*, *fill_mode='constant'*, *cval=128*, *fill_per_channel=False*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Fill one or more rectangular areas in an image using a fill mode.

See paper "Improved Regularization of Convolutional Neural Networks with Cutout" by DeVries and Taylor.

In contrast to the paper, this implementation also supports replacing image sub-areas with gaussian noise, random intensities or random RGB colors. It also supports non-squared areas. While the paper uses absolute pixel values for the size and position, this implementation uses relative values, which seems more appropriate for mixed-size datasets. The position parameter furthermore allows more flexibility, e.g. gaussian distributions around the center.

Note: This augmenter affects only image data. Other datatypes (e.g. segmentation map pixels or keypoints within the filled areas) are not affected.

Note: Gaussian fill mode will assume that float input images contain values in the interval [0.0, 1.0] and hence sample values from a gaussian within that interval, i.e. from N(0.5, std=0.5/3).

Added in 0.4.0.

Supported dtypes:

See *cutout_()*.

Parameters

- **nb_iterations** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter, optional*) –

  How many rectangular areas to fill.

  – If `int`: Exactly that many areas will be filled on all images.

  – If `tuple (a, b)`: A value from the interval `[a, b]` will be sampled per image.

  – If `list`: A random value will be sampled from that `list` per image.

  – If `StochasticParameter`: That parameter will be used to sample `(B,)` values per batch of `B` images.

- **position** (*{'uniform', 'normal', 'center', 'left-top', 'left-center', 'left-bottom', 'center-top', 'center-center', 'center-bottom', 'right-top', 'right-center', 'right-bottom'} or tuple of float or StochasticParameter or tuple of StochasticParameter, optional*) – Defines the position of each area to fill. Analogous to the definition in e.g. *CropToFixedSize*. Usually, `uniform` (anywhere in the image) or `normal` (anywhere in the image with preference around the center) are sane values.

- **size** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – The size of the rectangle to fill as a fraction of the corresponding image size, i.e. with value range `[0.0, 1.0]`. The size is sampled independently per image axis.

  – If `number`: Exactly that size is always used.

  – If `tuple (a, b)`: A value from the interval `[a, b]` will be sampled per area and axis.

  – If `list`: A random value will be sampled from that `list` per area and axis.

  – If `StochasticParameter`: That parameter will be used to sample `(N, 2)` values per batch, where `N` is the total number of areas to fill within the whole batch.

- **squared** (*bool or float or imgaug.parameters.StochasticParameter, optional*) – Whether to generate only squared areas cutout areas or allow rectangular ones. If this evaluates to a true-like value, the first value from *size* will be converted to absolute pixels and used for both axes.

  If this value is a float `p`, then for `p` percent of all areas to be filled *per_channel* will be treated as `True`. If it is a `StochasticParameter` it is expected to produce samples with values between `0.0` and `1.0`, where values `>0.5` will lead to per-channel behaviour (i.e. same as `True`).

- **fill_mode** (*str or list of str or imgaug.parameters.StochasticParameter, optional*) – Mode to use in order to fill areas. Corresponds to `mode` parameter in some other augmenters. Valid strings for the mode are:

  – `contant`: Fill each area with a single value.

  – `gaussian`: Fill each area with gaussian noise.

  Valid datatypes are:

  – If `str`: Exactly that mode will alaways be used.

  - If `list`: A random value will be sampled from that `list` per area.
  - If `StochasticParameter`: That parameter will be used to sample `(N,)` values per batch, where `N` is the total number of areas to fill within the whole batch.

- **cval** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – The value to use (i.e. the color) to fill areas if *fill_mode* is `constant.
  - If `number`: Exactly that value is used for all areas and channels.
  - If `tuple (a, b)`: A value from the interval `[a, b]` will be sampled per area (and channel if `per_channel=True`).
  - If `list`: A random value will be sampled from that `list` per area (and channel if `per_channel=True`).
  - If `StochasticParameter`: That parameter will be used to sample `(N, Cmax)` values per batch, where `N` is the total number of areas to fill within the whole batch and `Cmax` is the maximum number of channels in any image (usually `3`). If `per_channel=False`, only the first value of the second axis is used.
- **fill_per_channel** (*bool or float or imgaug.parameters.StochasticParameter, optional*) – Whether to fill each area in a channelwise fashion (`True`) or not (`False`). The behaviour per fill mode is:
  - `constant`: Whether to fill all channels with the same value (i.e, grayscale) or different values (i.e. usually RGB color).
  - `gaussian`: Whether to sample once from a gaussian and use the values for all channels (i.e. grayscale) or to sample channelwise (i.e. RGB colors)

  If this value is a float `p`, then for `p` percent of all areas to be filled *per_channel* will be treated as `True`. If it is a `StochasticParameter` it is expected to produce samples with values between `0.0` and `1.0`, where values `>0.5` will lead to per-channel behaviour (i.e. same as `True`).
- **name** (*None or str, optional*) – See `__init__()`.
- **deterministic** (*bool, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.bit_generator.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.Cutout(nb_iterations=2)
```

Fill per image two random areas, by default with grayish pixels.

```
>>> aug = iaa.Cutout(nb_iterations=(1, 5), size=0.2, squared=False)
```

Fill per image between one and five areas, each having `20%` of the corresponding size of the height and width (for non-square images this results in non-square areas to be filled).

```
>>> aug = iaa.Cutout(fill_mode="constant", cval=255)
```

Fill all areas with white pixels.

```
>>> aug = iaa.Cutout(fill_mode="constant", cval=(0, 255),
>>>                  fill_per_channel=0.5)
```

Fill 50% of all areas with a random intensity value between 0 and 256. Fill the other 50% of all areas with random colors.

```
>>> aug = iaa.Cutout(fill_mode="gaussian", fill_per_channel=True)
```

Fill areas with gaussian channelwise noise (i.e. usually RGB).

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`localize_random_state_`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.

Continued on next page

Table 60 – continued from previous page

pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.arithmetic.**Dropout**(*p=(0.0, 0.05)*, *per_channel=False*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.arithmetic.MultiplyElementwise*

Set a fraction of pixels in images to zero.

Supported dtypes:

See *MultiplyElementwise*.

Parameters

- **p** (*float or tuple of float or imgaug.parameters.StochasticParameter, optional*) –

  The probability of any pixel being dropped (i.e. to set it to zero).

  - If a float, then that value will be used for all images. A value of `1.0` would mean that all pixels will be dropped and `0.0` that no pixels will be dropped. A value of `0.05` corresponds to `5` percent of all pixels being dropped.
  - If a tuple `(a, b)`, then a value `p` will be sampled from the interval `[a, b]` per image and be used as the pixel's dropout probability.
  - If a list, then a value will be sampled from that list per batch and used as the probability.
  - If a `StochasticParameter`, then this parameter will be used to determine per pixel whether it should be *kept* (sampled value of `>0.5`) or shouldn't be kept (sampled value of `<=0.5`). If you instead want to provide the probability as a stochastic parameter, you can usually do `imgaug.parameters.Binomial(1-p)` to convert parameter *p* to a 0/1 representation.

- **per_channel** (*bool or float or imgaug.parameters.StochasticParameter, optional*) – Whether to use (imagewise) the same sample(s) for all channels (`False`) or to sample value(s) for each channel (`True`). Setting this to `True` will therefore lead to different transformations per image *and* channel, otherwise only per image. If this value is a float `p`, then for `p` percent of all images *per_channel* will be treated as `True`. If it is a `StochasticParameter` it is expected to produce samples with values between `0.0` and `1.0`, where values `>0.5` will lead to per-channel behaviour (i.e. same as `True`).

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.Dropout(0.02)
```

Drops 2 percent of all pixels.

```
>>> aug = iaa.Dropout((0.0, 0.05))
```

Drops in each image a random fraction of all pixels, where the fraction is uniformly sampled from the interval `[0.0, 0.05]`.

```
>>> aug = iaa.Dropout(0.02, per_channel=True)
```

Drops 2 percent of all pixels in a channelwise fashion, i.e. it is unlikely for any pixel to have all channels set to zero (black pixels).

```
>>> aug = iaa.Dropout(0.02, per_channel=0.5)
```

Identical to the previous example, but the *per_channel* feature is only active for 50 percent of all images.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.

Continued on next page

Table 61 – continued from previous page

augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.arithmetic.**Dropout2d**(*p=0.1*, *nb_keep_channels=1*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Drop random channels from images.

For image data, dropped channels will be filled with zeros.

Note: This augmenter may also set the arrays of heatmaps and segmentation maps to zero and remove all coordinate-based data (e.g. it removes all bounding boxes on images that were filled with zeros). It does so if and only if *all* channels of an image are dropped. If `nb_keep_channels >= 1` then that never happens.

Added in 0.4.0.

Supported dtypes:

- `uint8`: yes; fully tested
- `uint16`: yes; tested
- `uint32`: yes; tested
- `uint64`: yes; tested
- `int8`: yes; tested
- `int16`: yes; tested
- `int32`: yes; tested
- `int64`: yes; tested
- `float16`: yes; tested
- `float32`: yes; tested
- `float64`: yes; tested
- `float128`: yes; tested
- `bool`: yes; tested

Parameters

- **p** (*float or tuple of float or imgaug.parameters.StochasticParameter, optional*) –

  The probability of any channel to be dropped (i.e. set to zero).

  – If a `float`, then that value will be used for all channels. A value of `1.0` would mean, that all channels will be dropped. A value of `0.0` would lead to no channels being dropped.

  – If a tuple `(a, b)`, then a value `p` will be sampled from the interval `[a, b)` per batch and be used as the dropout probability.

  – If a list, then a value will be sampled from that list per batch and used as the probability.

  – If a `StochasticParameter`, then this parameter will be used to determine per channel whether it should be *kept* (sampled value of `>=0.5`) or shouldn't be kept (sampled value of `<0.5`). If you instead want to provide the probability as a stochastic parameter, you can usually do `imgaug.parameters.Binomial(1-p)` to convert parameter *p* to a 0/1 representation.

- **nb_keep_channels** (*int*) – Minimum number of channels to keep unaltered in all images. E.g. a value of `1` means that at least one channel in every image will not be dropped, even if `p=1.0`. Set to `0` to allow dropping all channels.

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.

- **name** (*None or str, optional*) – See `__init__()`.

- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.

- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.Dropout2d(p=0.5)
```

Create a dropout augmenter that drops on average half of all image channels. Dropped channels will be filled with zeros. At least one channel is kept unaltered in each image (default setting).

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.Dropout2d(p=0.5, nb_keep_channels=0)
```

Create a dropout augmenter that drops on average half of all image channels *and* may drop *all* channels in an image (i.e. images may contain nothing but zeros).

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.

Continued on next page

Table 62 – continued from previous page

find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.arithmetic.**ImpulseNoise**(*p=(0.0, 0.03)*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.arithmetic.SaltAndPepper*

Add impulse noise to images.

This is identical to SaltAndPepper, except that *per_channel* is always set to True.

Supported dtypes:

See *SaltAndPepper*.

Parameters

- **p** (*float or tuple of float or list of float or imgaug.parameters.StochasticParameter, optional*) –

  Probability of replacing a pixel to impulse noise.

  - If a float, then that value will always be used as the probability.
  - If a tuple (a, b), then a probability will be sampled uniformly per image from the interval [a, b].
  - If a list, then a random value will be sampled from that list per image.
  - If a StochasticParameter, then a image-sized mask will be sampled from that parameter per image. Any value >0.5 in that mask will be replaced with impulse noise noise.

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.ImpulseNoise(0.1)
```

Replace `10%` of all pixels with impulse noise.

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.

Continued on next page

Table 63 – continued from previous page

find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.arithmetic.**Invert**(*p=1*, *per_channel=False*, *min_value=None*, *max_value=None*, *threshold=None*, *invert_above_threshold=0.5*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Invert all values in images, e.g. turn `5` into `255-5=250`.

For the standard value range of 0-255 it converts `0` to `255`, `255` to `0` and `10` to `(255-10)=245`. Let `M` be the maximum value possible, `m` the minimum value possible, `v` a value. Then the distance of `v` to `m` is `d=abs(v-m)` and the new value is given by `v'=M-d`.

Supported dtypes:

See *invert_()*.

Parameters

- **p** (*float or imgaug.parameters.StochasticParameter, optional*) –

  The probability of an image to be inverted.

  – If a float, then that probability will be used for all images, i.e. *p* percent of all images will be inverted.

  – If a `StochasticParameter`, then that parameter will be queried per image and is expected to return values in the interval `[0.0, 1.0]`, where values `>0.5` mean that the image is supposed to be inverted. Recommended to be some form of `imgaug.parameters.Binomial`.

- **per_channel** (*bool or float or imgaug.parameters.StochasticParameter, optional*) – Whether to use (imagewise) the same sample(s) for all channels (`False`) or to sample value(s) for each channel (`True`). Setting this to `True` will therefore lead to different transformations per image *and* channel, otherwise only per image. If this value is a float `p`, then for `p` percent of all images *per_channel* will be treated as `True`. If it is a `StochasticParameter` it is expected to produce samples with values between `0.0` and `1.0`, where values `>0.5` will lead to per-channel behaviour (i.e. same as `True`).
- **min_value** (*None or number, optional*) – Minimum of the value range of input images, e.g. `0` for `uint8` images. If set to `None`, the value will be automatically derived from the image's dtype.
- **max_value** (*None or number, optional*) – Maximum of the value range of input images, e.g. `255` for `uint8` images. If set to `None`, the value will be automatically derived from the image's dtype.
- **threshold** (*None or number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – A threshold to use in order to invert only numbers above or below the threshold. If `None` no thresholding will be used.
  - If `None`: No thresholding will be used.
  - If `number`: The value will be used for all images.
  - If `tuple (a, b)`: A value will be uniformly sampled per image from the interval `[a, b)`.
  - If `list`: A random value will be picked from the list per image.
  - If `StochasticParameter`: Per batch of size `N`, the parameter will be queried once to return `(N,)` samples.
- **invert_above_threshold** (*bool or float or imgaug.parameters.StochasticParameter, optional*) – If `True`, only values `>=threshold` will be inverted. Otherwise, only values `<threshold` will be inverted. If a `number`, then expected to be in the interval `[0.0, 1.0]` and denoting an imagewise probability. If a `StochasticParameter` then `(N,)` values will be sampled from the parameter per batch of size `N` and interpreted as `True` if `>0.5`. If *threshold* is `None` this parameter has no effect.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.Invert(0.1)
```

Inverts the colors in 10 percent of all images.

```
>>> aug = iaa.Invert(0.1, per_channel=True)
```

Inverts the colors in 10 percent of all image channels. This may or may not lead to multiple channels in an image being inverted.

```
>>> aug = iaa.Invert(0.1, per_channel=0.5)
```

Identical to the previous example, but the *per_channel* feature is only active for 50 percent of all images.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.

Continued on next page

Table 64 – continued from previous page

localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

ALLOW_DTYPES_CUSTOM_MINMAX = [dtype('uint8'), dtype('uint16'), dtype('uint32'), dtype(

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.arithmetic.**JpegCompression**(*compression=(0, 100), seed=None, name=None, random_state='deprecated', deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Degrade the quality of images by JPEG-compressing them.

During JPEG compression, high frequency components (e.g. edges) are removed. With low compression (strength) only the highest frequency components are removed, while very high compression (strength) will lead to only the lowest frequency components "surviving". This lowers the image quality. For more details, see https://en.wikipedia.org/wiki/Compression_artifact.

Note that this augmenter still returns images as numpy arrays (i.e. saves the images with JPEG compression and then reloads them into arrays). It does not return the raw JPEG file content.

Supported dtypes:

See *compress_jpeg()*.

Parameters

- **compression** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Degree of compression used during JPEG compression within value range `[0, 100]`. Higher values denote stronger compression and will cause low-frequency components to disappear. Note that JPEG's compression strength is also often set as a *quality*, which is the inverse of this parameter. Common choices for the *quality* setting are around 80 to 95, depending on the image. This translates here to a *compression* parameter of around 20 to 5.
  - If a single number, then that value always will be used as the compression.
  - If a tuple `(a, b)`, then the compression will be a value sampled uniformly from the interval `[a, b]`.
  - If a list, then a random value will be sampled from that list per image and used as the compression.

  - If a `StochasticParameter`, then `N` samples will be drawn from that parameter per `N` input images, each representing the compression for the `n`-th image.

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.JpegCompression(compression=(70, 99))
```

Remove high frequency components in images via JPEG compression with a *compression strength* between `70` and `99` (randomly and uniformly sampled per image). This corresponds to a (very low) *quality* setting of `1` to `30`.

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.

Continued on next page

Table 65 – continued from previous page

copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.arithmetic.**Multiply**(*mul=(0.8, 1.2)*, *per_channel=False*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Multiply all pixels in an image with a random value sampled once per image.

This augmenter can be used to make images lighter or darker.

Supported dtypes:

See *multiply_scalar()*.

Parameters

- **mul** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) –

  The value with which to multiply the pixel values in each image.

  – If a number, then that value will always be used.

  - If a tuple `(a, b)`, then a value from the interval `[a, b]` will be sampled per image and used for all pixels.
  - If a list, then a random value will be sampled from that list per image.
  - If a `StochasticParameter`, then that parameter will be used to sample a new value per image.

- **per_channel** (*bool or float or imgaug.parameters.StochasticParameter, optional*) – Whether to use (imagewise) the same sample(s) for all channels (`False`) or to sample value(s) for each channel (`True`). Setting this to `True` will therefore lead to different transformations per image *and* channel, otherwise only per image. If this value is a float `p`, then for `p` percent of all images *per_channel* will be treated as `True`. If it is a `StochasticParameter` it is expected to produce samples with values between `0.0` and `1.0`, where values `>0.5` will lead to per-channel behaviour (i.e. same as `True`).
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.Multiply(2.0)
```

Multiplies all images by a factor of `2`, making the images significantly brighter.

```
>>> aug = iaa.Multiply((0.5, 1.5))
```

Multiplies images by a random value sampled uniformly from the interval `[0.5, 1.5]`, making some images darker and others brighter.

```
>>> aug = iaa.Multiply((0.5, 1.5), per_channel=True)
```

Identical to the previous example, but the sampled multipliers differ by image *and* channel, instead of only by image.

```
>>> aug = iaa.Multiply((0.5, 1.5), per_channel=0.5)
```

Identical to the previous example, but the *per_channel* feature is only active for 50 percent of all images.

Methods

__call__(self, *args, **kwargs)	Alias for *augment()*.
augment(self[, return_batch, hooks])	Augment a batch.
augment_batch(self, batch[, hooks])	**Deprecated**.
augment_batch_(self, batch[, parents, hooks])	Augment a single batch in-place.
augment_batches(self, batches[, hooks, ...])	Augment multiple batches.
augment_bounding_boxes(self, ...[, parents, ...])	Augment a batch of bounding boxes.
augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.

Continued on next page

Table 66 – continued from previous page

to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.arithmetic.**MultiplyElementwise**(*mul=(0.8, 1.2), per_channel=False, seed=None, name=None, random_state='deprecated', deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Multiply image pixels with values that are pixelwise randomly sampled.

While the `Multiply` Augmenter uses a constant multiplier *per image* (and optionally channel), this augmenter samples the multipliers to use per image and *per pixel* (and optionally per channel).

Supported dtypes:

See *multiply_elementwise()*.

Parameters

- **mul** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) –

  The value with which to multiply pixel values in the image.

  - If a number, then that value will always be used.
  - If a tuple `(a, b)`, then a value from the interval `[a, b]` will be sampled per image and pixel.
  - If a list, then a random value will be sampled from that list per image and pixel.
  - If a `StochasticParameter`, then that parameter will be used to sample a new value per image and pixel.

- **per_channel** (*bool or float or imgaug.parameters.StochasticParameter, optional*) – Whether to use (imagewise) the same sample(s) for all channels (`False`) or to sample value(s) for each channel (`True`). Setting this to `True` will therefore lead to different transformations per image *and* channel, otherwise only per image. If this value is a float `p`, then for `p` percent of all images *per_channel* will be treated as `True`. If it is a `StochasticParameter` it is expected to produce samples with values between `0.0` and `1.0`, where values `>0.5` will lead to per-channel behaviour (i.e. same as `True`).

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.

- **name** (*None or str, optional*) – See `__init__()`.

- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.

- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.MultiplyElementwise(2.0)
```

Multiply all images by a factor of `2.0`, making them significantly bighter.

```
>>> aug = iaa.MultiplyElementwise((0.5, 1.5))
```

Samples per image and pixel uniformly a value from the interval `[0.5, 1.5]` and multiplies the pixel with that value.

```
>>> aug = iaa.MultiplyElementwise((0.5, 1.5), per_channel=True)
```

Samples per image and pixel *and channel* uniformly a value from the interval `[0.5, 1.5]` and multiplies the pixel with that value. Therefore, used multipliers may differ between channels of the same pixel.

```
>>> aug = iaa.MultiplyElementwise((0.5, 1.5), per_channel=0.5)
```

Identical to the previous example, but the *per_channel* feature is only active for 50 percent of all images.

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).

Continued on next page

Table 67 – continued from previous page

deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.arithmetic.**Pepper**(*p=(0.0, 0.05)*, *per_channel=False*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.arithmetic.ReplaceElementwise*

Replace pixels in images with pepper noise, i.e. black-ish pixels.

This augmenter is similar to SaltAndPepper, but adds no salt noise to images.

This augmenter is similar to Dropout, but slower and the black pixels are not uniformly black.

Supported dtypes:

See *ReplaceElementwise*.

Parameters

- **p** (*float or tuple of float or list of float or imgaug.parameters.StochasticParameter, optional*) –

  Probability of replacing a pixel with pepper noise.

  – If a float, then that value will always be used as the probability.

  – If a tuple (a, b), then a probability will be sampled uniformly per image from the interval [a, b].

  - If a list, then a random value will be sampled from that list per image.
  - If a `StochasticParameter`, then a image-sized mask will be sampled from that parameter per image. Any value `>0.5` in that mask will be replaced with pepper noise.

- **per_channel** (*bool or float or imgaug.parameters.StochasticParameter, optional*) – Whether to use (imagewise) the same sample(s) for all channels (`False`) or to sample value(s) for each channel (`True`). Setting this to `True` will therefore lead to different transformations per image *and* channel, otherwise only per image. If this value is a float `p`, then for `p` percent of all images *per_channel* will be treated as `True`. If it is a `StochasticParameter` it is expected to produce samples with values between `0.0` and `1.0`, where values `>0.5` will lead to per-channel behaviour (i.e. same as `True`).
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.Pepper(0.05)
```

Replace `5%` of all pixels with pepper noise (black-ish colors).

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.

Continued on next page

Table 68 – continued from previous page

augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.arithmetic.**ReplaceElementwise**(*mask*, *replacement*, *per_channel=False*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Replace pixels in an image with new values.

Supported dtypes:

See *replace_elementwise_()*.

Parameters

- **mask** (*float or tuple of float or list of float or imgaug.parameters.StochasticParameter*) – Mask that indicates the pixels that are supposed to be replaced. The mask will be binarized

using a threshold of `0.5`. A value of `1` then indicates a pixel that is supposed to be replaced.

  - If this is a float, then that value will be used as the probability of being a `1` in the mask (sampled per image and pixel) and hence being replaced.
  - If a tuple `(a, b)`, then the probability will be uniformly sampled per image from the interval `[a, b]`.
  - If a list, then a random value will be sampled from that list per image and pixel.
  - If a `StochasticParameter`, then this parameter will be used to sample a mask per image.

- **replacement** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – The replacement to use at all locations that are marked as `1` in the mask.
  - If this is a number, then that value will always be used as the replacement.
  - If a tuple `(a, b)`, then the replacement will be sampled uniformly per image and pixel from the interval `[a, b]`.
  - If a list, then a random value will be sampled from that list per image and pixel.
  - If a `StochasticParameter`, then this parameter will be used sample replacement values per image and pixel.
- **per_channel** (*bool or float or imgaug.parameters.StochasticParameter, optional*) – Whether to use (imagewise) the same sample(s) for all channels (`False`) or to sample value(s) for each channel (`True`). Setting this to `True` will therefore lead to different transformations per image *and* channel, otherwise only per image. If this value is a float `p`, then for `p` percent of all images *per_channel* will be treated as `True`. If it is a `StochasticParameter` it is expected to produce samples with values between `0.0` and `1.0`, where values `>0.5` will lead to per-channel behaviour (i.e. same as `True`).
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = ReplaceElementwise(0.05, [0, 255])
```

Replaces `5` percent of all pixels in each image by either `0` or `255`.

```
>>> import imgaug.augmenters as iaa
>>> aug = ReplaceElementwise(0.1, [0, 255], per_channel=0.5)
```

For 50% of all images, replace 10% of all pixels with either the value 0 or the value 255 (same as in the previous example). For the other 50% of all images, replace *channelwise* 10% of all pixels with either the value 0 or the value 255. So, it will be very rare for each pixel to have all channels replaced by 255 or 0.

```
>>> import imgaug.augmenters as iaa
>>> import imgaug.parameters as iap
>>> aug = ReplaceElementwise(0.1, iap.Normal(128, 0.4*128), per_channel=0.5)
```

Replace 10% of all pixels by gaussian noise centered around 128. Both the replacement mask and the gaussian noise are sampled channelwise for 50% of all images.

```
>>> import imgaug.augmenters as iaa
>>> import imgaug.parameters as iap
>>> aug = ReplaceElementwise(
>>>     iap.FromLowerResolution(iap.Binomial(0.1), size_px=8),
>>>     iap.Normal(128, 0.4*128),
>>>     per_channel=0.5)
```

Replace 10% of all pixels by gaussian noise centered around 128. Sample the replacement mask at a lower resolution (8x8 pixels) and upscale it to the image size, resulting in coarse areas being replaced by gaussian noise.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.

Continued on next page

Table 69 – continued from previous page

find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.arithmetic.**Salt**(*p=(0.0, 0.03)*, *per_channel=False*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.arithmetic.ReplaceElementwise*

Replace pixels in images with salt noise, i.e. white-ish pixels.

This augmenter is similar to SaltAndPepper, but adds no pepper noise to images.

Supported dtypes:

See *ReplaceElementwise*.

Parameters

- **p** (*float or tuple of float or list of float or imgaug.parameters.StochasticParameter, optional*) –

  Probability of replacing a pixel with salt noise.

  – If a float, then that value will always be used as the probability.

  – If a tuple (a, b), then a probability will be sampled uniformly per image from the interval [a, b].

  – If a list, then a random value will be sampled from that list per image.

  – If a StochasticParameter, then a image-sized mask will be sampled from that parameter per image. Any value >0.5 in that mask will be replaced with salt noise.

- **per_channel** (*bool or float or imgaug.parameters.StochasticParameter, optional*) – Whether to use (imagewise) the same sample(s) for all channels (`False`) or to sample value(s) for each channel (`True`). Setting this to `True` will therefore lead to different transformations per image *and* channel, otherwise only per image. If this value is a float `p`, then for `p` percent of all images *per_channel* will be treated as `True`. If it is a `StochasticParameter` it is expected to produce samples with values between `0.0` and `1.0`, where values `>0.5` will lead to per-channel behaviour (i.e. same as `True`).
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.Salt(0.05)
```

Replace `5%` of all pixels with salt noise (white-ish colors).

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.

Continued on next page

Table 70 – continued from previous page

copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.arithmetic.**SaltAndPepper**(*p=(0.0, 0.03)*, *per_channel=False*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.arithmetic.ReplaceElementwise*

Replace pixels in images with salt/pepper noise (white/black-ish colors).

Supported dtypes:

See *ReplaceElementwise*.

Parameters

- **p** (*float or tuple of float or list of float or imgaug.parameters.StochasticParameter, optional*) –

  Probability of replacing a pixel to salt/pepper noise.

  – If a float, then that value will always be used as the probability.

  - If a tuple `(a, b)`, then a probability will be sampled uniformly per image from the interval `[a, b]`.
  - If a list, then a random value will be sampled from that list per image.
  - If a `StochasticParameter`, then a image-sized mask will be sampled from that parameter per image. Any value `>0.5` in that mask will be replaced with salt and pepper noise.

- **per_channel** (*bool or float or imgaug.parameters.StochasticParameter, optional*) – Whether to use (imagewise) the same sample(s) for all channels (`False`) or to sample value(s) for each channel (`True`). Setting this to `True` will therefore lead to different transformations per image *and* channel, otherwise only per image. If this value is a float `p`, then for `p` percent of all images *per_channel* will be treated as `True`. If it is a `StochasticParameter` it is expected to produce samples with values between `0.0` and `1.0`, where values `>0.5` will lead to per-channel behaviour (i.e. same as `True`).
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.SaltAndPepper(0.05)
```

Replace `5%` of all pixels with salt and pepper noise.

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.SaltAndPepper(0.05, per_channel=True)
```

Replace *channelwise* `5%` of all pixels with salt and pepper noise.

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.

Continued on next page

Table 71 – continued from previous page

augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.arithmetic.**Solarize**(*p=1*, *per_channel=False*, *min_value=None*, *max_value=None*, *threshold=(64, 192)*, *invert_above_threshold=True*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.arithmetic.Invert*

Invert all pixel values above a threshold.

This is the same as *Invert*, but sets a default threshold around 128 (+/- 64, decided per image) and default *invert_above_threshold* to True (i.e. only values above the threshold will be inverted).

See *Invert* for more details.

Added in 0.4.0.

Supported dtypes:

See *Invert*.

Parameters

- **p** (*float or imgaug.parameters.StochasticParameter*) – See *Invert*.
- **per_channel** (*bool or float or imgaug.parameters.StochasticParameter, optional*) – See *Invert*.
- **min_value** (*None or number, optional*) – See *Invert*.
- **max_value** (*None or number, optional*) – See *Invert*.
- **threshold** (*None or number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – See *Invert*.
- **invert_above_threshold** (*bool or float or imgaug.parameters.StochasticParameter, optional*) – See *Invert*.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See __init__().
- **name** (*None or str, optional*) – See __init__().
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method to_deterministic() for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.Solarize(0.5, threshold=(32, 128))
```

Invert the colors in 50 percent of all images for pixels with a value between 32 and 128 or more. The threshold is sampled once per image. The thresholding operation happens per channel.

Methods

__call__(self, *args, **kwargs)	Alias for *augment()*.
augment(self[, return_batch, hooks])	Augment a batch.
augment_batch(self, batch[, hooks])	**Deprecated**.

Continued on next page

Table 72 – continued from previous page

augment_batch_(self, batch[, parents, hooks])	Augment a single batch in-place.
augment_batches(self, batches[, hooks, ...])	Augment multiple batches.
augment_bounding_boxes(self, ...[, parents, ...])	Augment a batch of bounding boxes.
augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.arithmetic.**TotalDropout**(*p=1*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Drop all channels of a defined fraction of all images.

For image data, all components of dropped images will be filled with zeros.

Note: This augmenter also sets the arrays of heatmaps and segmentation maps to zero and removes all coordinate-based data (e.g. it removes all bounding boxes on images that were filled with zeros).

Added in 0.4.0.

Supported dtypes:

- `uint8`: yes; fully tested
- `uint16`: yes; tested
- `uint32`: yes; tested
- `uint64`: yes; tested
- `int8`: yes; tested
- `int16`: yes; tested
- `int32`: yes; tested
- `int64`: yes; tested
- `float16`: yes; tested
- `float32`: yes; tested
- `float64`: yes; tested
- `float128`: yes; tested
- `bool`: yes; tested

Parameters

- **p** (*float or tuple of float or imgaug.parameters.StochasticParameter, optional*) –

  The probability of an image to be filled with zeros.

  - If `float`: The value will be used for all images. A value of `1.0` would mean that all images will be set to zero. A value of `0.0` would lead to no images being set to zero.
  - If `tuple (a, b)`: A value `p` will be sampled from the interval `[a, b)` per batch and be used as the dropout probability.
  - If a list, then a value will be sampled from that list per batch and used as the probability.
  - If `StochasticParameter`: The parameter will be used to determine per image whether it should be *kept* (sampled value of `>=0.5`) or shouldn't be kept (sampled value of `<0.5`). If you instead want to provide the probability as a stochastic parameter, you can usually do `imgaug.parameters.Binomial(1-p)` to convert parameter *p* to a 0/1 representation.

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.

- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.TotalDropout(1.0)
```

Create an augmenter that sets *all* components of all images to zero.

```
>>> aug = iaa.TotalDropout(0.5)
```

Create an augmenter that sets *all* components of `50%` of all images to zero.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.

Continued on next page

Table 73 – continued from previous page

draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See get_parameters().
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See get_parameters().

imgaug.augmenters.arithmetic.**add_elementwise**(*image*, *values*)
Add an array of values to an image.

This method ensures that uint8 does not overflow during the addition.

Supported dtypes:

- uint8: yes; fully tested
- uint16: limited; tested (1)
- uint32: no
- uint64: no
- int8: limited; tested (1)
- int16: limited; tested (1)
- int32: no
- int64: no
- float16: limited; tested (1)
- float32: limited; tested (1)
- float64: no

- `float128`: no
- `bool`: limited; tested (1)
- (1) Non-uint8 dtypes can overflow. For floats, this can result in +/-inf.

Parameters

- **image** (*ndarray*) – Image array of shape `(H,W,[C])`.
- **values** (*ndarray*) – The values to add to the image. Expected to have the same height and width as *image* and either no channels or one channel or the same number of channels as *image*.

Returns Image with values added to it.

Return type ndarray

`imgaug.augmenters.arithmetic.`**`add_scalar`**(*image*, *value*)

Add a single scalar value or one scalar value per channel to an image.

This method ensures that `uint8` does not overflow during the addition.

Supported dtypes:

- `uint8`: yes; fully tested
- `uint16`: limited; tested (1)
- `uint32`: no
- `uint64`: no
- `int8`: limited; tested (1)
- `int16`: limited; tested (1)
- `int32`: no
- `int64`: no
- `float16`: limited; tested (1)
- `float32`: limited; tested (1)
- `float64`: no
- `float128`: no
- `bool`: limited; tested (1)
- (1) Non-uint8 dtypes can overflow. For floats, this can result in +/-inf.

Parameters

- **image** (*ndarray*) – Image array of shape `(H,W,[C])`. If *value* contains more than one value, the shape of the image is expected to be `(H,W,C)`.
- **value** (*number or ndarray*) – The value to add to the image. Either a single value or an array containing exactly one component per channel, i.e. `C` components.

Returns Image with value added to it.

Return type ndarray

imgaug.augmenters.arithmetic.**compress_jpeg**(*image*, *compression*)

Compress an image using jpeg compression.

Supported dtypes:

- uint8: yes; fully tested
- uint16: ?
- uint32: ?
- uint64: ?
- int8: ?
- int16: ?
- int32: ?
- int64: ?
- float16: ?
- float32: ?
- float64: ?
- float128: ?
- bool: ?

Parameters

- **image** (*ndarray*) – Image of dtype uint8 and shape (H,W,[C]). If C is provided, it must be 1 or 3.
- **compression** (*int*) – Strength of the compression in the interval [0, 100].

Returns Input image after applying jpeg compression to it and reloading the result into a new array. Same shape and dtype as the input.

Return type ndarray

imgaug.augmenters.arithmetic.**cutout**(*image*, *x1*, *y1*, *x2*, *y2*, *fill_mode='constant'*, *cval=0*, *fill_per_channel=False*, *seed=None*)

Fill a single area within an image using a fill mode.

This cutout method uses the top-left and bottom-right corner coordinates of the cutout region given as absolute pixel values.

Note: Gaussian fill mode will assume that float input images contain values in the interval [0.0, 1.0] and hence sample values from a gaussian within that interval, i.e. from N(0.5, std=0.5/3).

Supported dtypes:

See *cutout_()*.

Added in 0.4.0.

Parameters

- **image** (*ndarray*) – Image to modify.
- **x1** (*number*) – See *cutout_()*.
- **y1** (*number*) – See *cutout_()*.

- **x2** (*number*) – See *cutout_()*.
- **y2** (*number*) – See *cutout_()*.
- **fill_mode** (*{'constant', 'gaussian'}, optional*) – See *cutout_()*.
- **cval** (*number or tuple of number, optional*) – See *cutout_()*.
- **fill_per_channel** (*number or bool, optional*) – See *cutout_()*.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See *cutout_()*.

Returns Image with area filled in.

Return type ndarray

imgaug.augmenters.arithmetic.**cutout_**(*image*, *x1*, *y1*, *x2*, *y2*, *fill_mode='constant'*, *cval=0*, *fill_per_channel=False*, *seed=None*)

Fill a single area within an image using a fill mode (in-place).

This cutout method uses the top-left and bottom-right corner coordinates of the cutout region given as absolute pixel values.

Note: Gaussian fill mode will assume that float input images contain values in the interval `[0.0, 1.0]` and hence sample values from a gaussian within that interval, i.e. from `N(0.5, std=0.5/3)`.

Added in 0.4.0.

Supported dtypes:

minimum of (`_fill_rectangle_gaussian_()`, `_fill_rectangle_constant_()`

)

Parameters

- **image** (*ndarray*) – Image to modify. Might be modified in-place.
- **x1** (*number*) – X-coordinate of the top-left corner of the cutout region.
- **y1** (*number*) – Y-coordinate of the top-left corner of the cutout region.
- **x2** (*number*) – X-coordinate of the bottom-right corner of the cutout region.
- **y2** (*number*) – Y-coordinate of the bottom-right corner of the cutout region.
- **fill_mode** (*{'constant', 'gaussian'}, optional*) – Fill mode to use.
- **cval** (*number or tuple of number, optional*) – The constant value to use when filling with mode `constant`. May be an intensity value or color tuple.
- **fill_per_channel** (*number or bool, optional*) – Whether to fill in a channelwise fashion. If number then a value `>=0.5` will be interpreted as `True`.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – A random number generator to sample random values from. Usually an integer seed value or an `RNG` instance. See *imgaug.random.RNG* for details.

Returns Image with area filled in. The input image might have been modified in-place.

Return type ndarray

imgaug.augmenters.arithmetic.**invert**(*image*, *min_value=None*, *max_value=None*, *threshold=None*, *invert_above_threshold=True*)

Invert an array.

Supported dtypes:

See *invert_()*.

Parameters

- **image** (*ndarray*) – See *invert_()*.
- **min_value** (*None or number, optional*) – See *invert_()*.
- **max_value** (*None or number, optional*) – See *invert_()*.
- **threshold** (*None or number, optional*) – See *invert_()*.
- **invert_above_threshold** (*bool, optional*) – See *invert_()*.

Returns Inverted image.

Return type ndarray

imgaug.augmenters.arithmetic.**invert_**(*image*, *min_value=None*, *max_value=None*, *threshold=None*, *invert_above_threshold=True*)

Invert an array in-place.

Added in 0.4.0.

Supported dtypes:

if (min_value=None and max_value=None):

- `uint8`: yes; fully tested
- `uint16`: yes; tested
- `uint32`: yes; tested
- `uint64`: yes; tested
- `int8`: yes; tested
- `int16`: yes; tested
- `int32`: yes; tested
- `int64`: yes; tested
- `float16`: yes; tested
- `float32`: yes; tested
- `float64`: yes; tested
- `float128`: yes; tested
- `bool`: yes; tested

if (min_value!=None or max_value!=None):

- `uint8`: yes; fully tested
- `uint16`: yes; tested
- `uint32`: yes; tested
- `uint64`: no (1)
- `int8`: yes; tested

- int16: yes; tested
- int32: yes; tested
- int64: no (2)
- float16: yes; tested
- float32: yes; tested
- float64: no (2)
- float128: no (3)
- bool: no (4)

- (1) Not allowed due to numpy's clip converting from uint64 to float64.
- (2) Not allowed as int/float have to be increased in resolution when using min/max values.
- (3) Not tested.
- (4) Makes no sense when using min/max values.

Parameters

- **image** (*ndarray*) – Image array of shape (H,W,[C]). The array *might* be modified in-place.
- **min_value** (*None or number, optional*) – Minimum of the value range of input images, e.g. 0 for uint8 images. If set to None, the value will be automatically derived from the image's dtype.
- **max_value** (*None or number, optional*) – Maximum of the value range of input images, e.g. 255 for uint8 images. If set to None, the value will be automatically derived from the image's dtype.
- **threshold** (*None or number, optional*) – A threshold to use in order to invert only numbers above or below the threshold. If None no thresholding will be used.
- **invert_above_threshold** (*bool, optional*) – If True, only values >=threshold will be inverted. Otherwise, only values <threshold will be inverted. If *threshold* is None this parameter has no effect.

Returns Inverted image. This *can* be the same array as input in *image*, modified in-place.

Return type ndarray

imgaug.augmenters.arithmetic.**multiply_elementwise**(*image*, *multipliers*)

Multiply an image with an array of values.

This method ensures that uint8 does not overflow during the addition.

Supported dtypes:

- uint8: yes; fully tested
- uint16: limited; tested (1)
- uint32: no
- uint64: no
- int8: limited; tested (1)
- int16: limited; tested (1)

- int32: no
- int64: no
- float16: limited; tested (1)
- float32: limited; tested (1)
- float64: no
- float128: no
- bool: limited; tested (1)
- (1) Non-uint8 dtypes can overflow. For floats, this can result in +/-inf.

note:

```
Tests were only conducted for rather small multipliers, around
``-10.0`` to ``+10.0``.

In general, the multipliers sampled from `multipliers` must be in a
value range that corresponds to the input image's dtype. E.g. if the
input image has dtype ``uint16`` and the samples generated from
`multipliers` are ``float64``, this function will still force all
samples to be within the value range of ``float16``, as it has the
same number of bytes (two) as ``uint16``. This is done to make
overflows less likely to occur.
```

Parameters

- **image** (*ndarray*) – Image array of shape `(H,W,[C])`.
- **multipliers** (*ndarray*) – The multipliers with which to multiply the image. Expected to have the same height and width as *image* and either no channels or one channel or the same number of channels as *image*.

Returns Image, multiplied by *multipliers*.

Return type ndarray

imgaug.augmenters.arithmetic.**multiply_scalar**(*image*, *multiplier*)

Multiply an image by a single scalar or one scalar per channel.

This method ensures that uint8 does not overflow during the multiplication.

Supported dtypes:

- uint8: yes; fully tested
- uint16: limited; tested (1)
- uint32: no
- uint64: no
- int8: limited; tested (1)
- int16: limited; tested (1)
- int32: no
- int64: no
- float16: limited; tested (1)

- `float32`: limited; tested (1)
- `float64`: no
- `float128`: no
- `bool`: limited; tested (1)
- (1) Non-uint8 dtypes can overflow. For floats, this can result in +/-inf.

note:

```
Tests were only conducted for rather small multipliers, around
``-10.0`` to ``+10.0``.

In general, the multipliers sampled from `multiplier` must be in a
value range that corresponds to the input image's dtype. E.g. if the
input image has dtype ``uint16`` and the samples generated from
`multiplier` are ``float64``, this function will still force all
samples to be within the value range of ``float16``, as it has the
same number of bytes (two) as ``uint16``. This is done to make
overflows less likely to occur.
```

Parameters

- **image** (*ndarray*) – Image array of shape `(H,W,[C])`. If *value* contains more than one value, the shape of the image is expected to be `(H,W,C)`.
- **multiplier** (*number or ndarray*) – The multiplier to use. Either a single value or an array containing exactly one component per channel, i.e. `C` components.

Returns Image, multiplied by *multiplier*.

Return type ndarray

`imgaug.augmenters.arithmetic.`**`replace_elementwise_`**(*image*, *mask*, *replacements*)

Replace components in an image array with new values.

Supported dtypes:

- `uint8`: yes; fully tested
- `uint16`: yes; tested
- `uint32`: yes; tested
- `uint64`: no (1)
- `int8`: yes; tested
- `int16`: yes; tested
- `int32`: yes; tested
- `int64`: no (2)
- `float16`: yes; tested
- `float32`: yes; tested
- `float64`: yes; tested
- `float128`: no
- `bool`: yes; tested

- (1) `uint64` is currently not supported, because *clip_to_dtype_value_range_()* does not support it, which again is because numpy.clip() seems to not support it.
- (2) *int64* is disallowed due to being converted to *float64* by `numpy.clip()` since 1.17 (possibly also before?).

Parameters

- **image** (*ndarray*) – Image array of shape `(H,W,[C])`.
- **mask** (*ndarray*) – Mask of shape `(H,W,[C])` denoting which components to replace. If `C` is provided, it must be `1` or match the `C` of *image*. May contain floats in the interval `[0.0, 1.0]`.
- **replacements** (*iterable*) – Replacements to place in *image* at the locations defined by *mask*. This 1-dimensional iterable must contain exactly as many values as there are replaced components in *image*.

Returns Image with replaced components.

Return type ndarray

`imgaug.augmenters.arithmetic.`**`solarize`**(*image*, *threshold=128*)

Invert pixel values above a threshold.

Added in 0.4.0.

Supported dtypes:

See *solarize_()*.

Parameters

- **image** (*ndarray*) – See *solarize_()*.
- **threshold** (*None or number, optional*) – See *solarize_()*.

Returns Inverted image.

Return type ndarray

`imgaug.augmenters.arithmetic.`**`solarize_`**(*image*, *threshold=128*)

Invert pixel values above a threshold in-place.

This function is a wrapper around *invert()*.

This function performs the same transformation as `PIL.ImageOps.solarize()`.

Added in 0.4.0.

Supported dtypes:

See ~imgaug.augmenters.arithmetic.invert_(min_value=None and max_value=None).

Parameters

- **image** (*ndarray*) – See *invert_()*.
- **threshold** (*None or number, optional*) – See *invert_()*. Note: The default threshold is optimized for `uint8` images.

Returns Inverted image. This *can* be the same array as input in *image*, modified in-place.

Return type ndarray

13.18 imgaug.augmenters.artistic

Augmenters that apply artistic image filters.

List of augmenters:

- *Cartoon*

Added in 0.4.0.

class imgaug.augmenters.artistic.**Cartoon**(*blur_ksize=(1, 5)*, *segmentation_size=(0.8, 1.2)*, *saturation=(1.5, 2.5)*, *edge_prevalence=(0.9, 1.1)*, *from_colorspace='RGB'*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Convert the style of images to a more cartoonish one.

This augmenter was primarily designed for images with a size of `200` to `800` pixels. Smaller or larger images may cause issues.

Note that the quality of the results can currently not compete with learned style transfer, let alone human-made images. A lack of detected edges or also too many detected edges are probably the most significant drawbacks.

Added in 0.4.0.

Supported dtypes:

See *stylize_cartoon()*.

Parameters

- **blur_ksize** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Median filter kernel size. See *stylize_cartoon()* for details.
  - If `number`: That value will be used for all images.
  - If `tuple (a, b) of number`: A random value will be uniformly sampled per image from the interval `[a, b)`.
  - If `list`: A random value will be picked per image from the `list`.
  - If `StochasticParameter`: The parameter will be queried once per batch for `(N,)` values, where `N` is the number of images.
- **segmentation_size** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Mean-Shift segmentation size multiplier. See *stylize_cartoon()* for details.
  - If `number`: That value will be used for all images.
  - If `tuple (a, b) of number`: A random value will be uniformly sampled per image from the interval `[a, b)`.
  - If `list`: A random value will be picked per image from the `list`.
  - If `StochasticParameter`: The parameter will be queried once per batch for `(N,)` values, where `N` is the number of images.
- **saturation** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Saturation multiplier. See *stylize_cartoon()* for details.

  - If `number`: That value will be used for all images.
  - If `tuple (a, b) of number`: A random value will be uniformly sampled per image from the interval `[a, b)`.
  - If `list`: A random value will be picked per image from the `list`.
  - If `StochasticParameter`: The parameter will be queried once per batch for `(N,)` values, where `N` is the number of images.

- **edge_prevalence** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Multiplier for the prevalence of edges. See *stylize_cartoon()* for details.
  - If `number`: That value will be used for all images.
  - If `tuple (a, b) of number`: A random value will be uniformly sampled per image from the interval `[a, b)`.
  - If `list`: A random value will be picked per image from the `list`.
  - If `StochasticParameter`: The parameter will be queried once per batch for `(N,)` values, where `N` is the number of images.
- **from_colorspace** (*str, optional*) – The source colorspace. Use one of `imgaug.augmenters.color.CSPACE_*`. Defaults to `RGB`.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.Cartoon()
```

Create an example image, then apply a cartoon filter to it.

```
>>> aug = iaa.Cartoon(blur_ksize=3, segmentation_size=1.0,
>>>                   saturation=2.0, edge_prevalence=1.0)
```

Create a non-stochastic cartoon augmenter that produces decent-looking images.

Methods

__call__(self, *args, **kwargs)	Alias for *augment()*.
augment(self[, return_batch, hooks])	Augment a batch.
augment_batch(self, batch[, hooks])	**Deprecated**.
augment_batch_(self, batch[, parents, hooks])	Augment a single batch in-place.
augment_batches(self, batches[, hooks, ...])	Augment multiple batches.
augment_bounding_boxes(self, ...[, parents, ...])	Augment a batch of bounding boxes.
augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.

Continued on next page

Table 74 – continued from previous page

`to_deterministic`(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See *get_parameters()*.

imgaug.augmenters.artistic.**stylize_cartoon**(*image*, *blur_ksize=3*, *segmentation_size=1.0*, *saturation=2.0*, *edge_prevalence=1.0*, *suppress_edges=True*, *from_colorspace='RGB'*)

Convert the style of an image to a more cartoonish one.

This function was primarily designed for images with a size of `200` to `800` pixels. Smaller or larger images may cause issues.

Note that the quality of the results can currently not compete with learned style transfer, let alone human-made images. A lack of detected edges or also too many detected edges are probably the most significant drawbacks.

This method is loosely based on the one proposed in https://stackoverflow.com/a/11614479/3760780

Added in 0.4.0.

Supported dtypes:

- `uint8`: yes; fully tested
- `uint16`: no
- `uint32`: no
- `uint64`: no
- `int8`: no
- `int16`: no
- `int32`: no
- `int64`: no
- `float16`: no
- `float32`: no
- `float64`: no
- `float128`: no
- `bool`: no

Parameters

- **image** (*ndarray*) – A `(H,W,3) uint8` image array.
- **blur_ksize** (*int, optional*) – Kernel size of the median blur filter applied initially to the input image. Expected to be an odd value and `>=0`. If an even value, thn automatically increased to an odd one. If `<=1`, no blur will be applied.
- **segmentation_size** (*float, optional*) – Size multiplier to decrease/increase the base size of the initial mean-shift segmentation of the image. Expected to be `>=0`. Note that the base size is increased by roughly a factor of two for images with height and/or width `>=400`.
- **edge_prevalence** (*float, optional*) – Multiplier for the prevalance of edges. Higher values lead to more edges. Note that the default value of `1.0` is already fairly conservative, so there is limit effect from lowerin it further.

- **saturation** (*float, optional*) – Multiplier for the saturation. Set to `1.0` to not change the image's saturation.
- **suppress_edges** (*bool, optional*) – Whether to run edge suppression to remove blobs containing too many or too few edge pixels.
- **from_colorspace** (*str, optional*) – The source colorspace. Use one of `imgaug.augmenters.color.CSPACE_*`. Defaults to `RGB`.

Returns Image in cartoonish style.

Return type ndarray

13.19 imgaug.augmenters.base

Base classes and functions used by all/most augmenters.

This module is planned to contain *imgaug.augmenters.meta.Augmenter* in the future.

Added in 0.4.0.

exception imgaug.augmenters.base.**SuspiciousMultiImageShapeWarning**
Bases: `UserWarning`

Warning multi-image inputs that look like a single image.

exception imgaug.augmenters.base.**SuspiciousSingleImageShapeWarning**
Bases: `UserWarning`

Warning for single-image inputs that look like multiple images.

13.20 imgaug.augmenters.blend

Augmenters that blend two images with each other.

List of augmenters:

- *BlendAlpha*
- *BlendAlphaMask*
- *BlendAlphaElementwise*
- *BlendAlphaSimplexNoise*
- *BlendAlphaFrequencyNoise*
- *BlendAlphaSomeColors*
- *BlendAlphaHorizontalLinearGradient*
- *BlendAlphaVerticalLinearGradient*
- *BlendAlphaSegMapClassIds*
- *BlendAlphaBoundingBoxes*
- *BlendAlphaRegularGrid*
- *BlendAlphaCheckerboard*

imgaug.augmenters.blend.**Alpha**(*factor=0*, *first=None*, *second=None*, *per_channel=False*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Deprecated. Use `Alpha` instead. Alpha is deprecated. Use BlendAlpha instead. The order of parameters is the same. Parameter 'first' was renamed to 'foreground'. Parameter 'second' was renamed to 'background'.

See *BlendAlpha*.

Deprecated since 0.4.0.

imgaug.augmenters.blend.**AlphaElementwise**(*factor=0*, *first=None*, *second=None*, *per_channel=False*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Deprecated. Use `AlphaElementwise` instead. AlphaElementwise is deprecated. Use BlendAlphaElementwise instead. The order of parameters is the same. Parameter 'first' was renamed to 'foreground'. Parameter 'second' was renamed to 'background'.

See *BlendAlphaElementwise*.

Deprecated since 0.4.0.

class imgaug.augmenters.blend.**BlendAlpha**(*factor=(0.0, 1.0)*, *foreground=None*, *background=None*, *per_channel=False*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Alpha-blend two image sources using an alpha/opacity value.

The two image sources can be imagined as branches. If a source is not given, it is automatically the same as the input. Let `FG` be the foreground branch and `BG` be the background branch. Then the result images are defined as `factor * FG + (1-factor) * BG`, where `factor` is an overlay factor.

Note: It is not recommended to use `BlendAlpha` with augmenters that change the geometry of images (e.g. horizontal flips, affine transformations) if you *also* want to augment coordinates (e.g. keypoints, polygons, ...), as it is unclear which of the two coordinate results (foreground or background branch) should be used as the coordinates after augmentation.

Currently, if `factor >= 0.5` (per image), the results of the foreground branch are used as the new coordinates, otherwise the results of the background branch.

Added in 0.4.0. (Before that named *Alpha*.)

Supported dtypes:

See *blend_alpha()*.

Parameters

- **factor** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Opacity of the results of the foreground branch. Values close to `0.0` mean that the results from the background branch (see parameter *background*) make up most of the final image.
  - If float, then that value will be used for all images.
  - If tuple `(a, b)`, then a random value from the interval `[a, b]` will be sampled per image.
  - If a list, then a random value will be picked from that list per image.

  - If `StochasticParameter`, then that parameter will be used to sample a value per image.

- **foreground** (*None or imgaug.augmenters.meta.Augmenter or iterable of imgaug.augmenters.meta.Augmenter, optional*) – Augmenter(s) that make up the foreground branch. High alpha values will show this branch's results.
  - If `None`, then the input images will be reused as the output of the foreground branch.
  - If `Augmenter`, then that augmenter will be used as the branch.
  - If iterable of `Augmenter`, then that iterable will be converted into a `Sequential` and used as the augmenter.
- **background** (*None or imgaug.augmenters.meta.Augmenter or iterable of imgaug.augmenters.meta.Augmenter, optional*) – Augmenter(s) that make up the background branch. Low alpha values will show this branch's results.
  - If `None`, then the input images will be reused as the output of the background branch.
  - If `Augmenter`, then that augmenter will be used as the branch.
  - If iterable of `Augmenter`, then that iterable will be converted into a `Sequential` and used as the augmenter.
- **per_channel** (*bool or float or imgaug.parameters.StochasticParameter, optional*) – Whether to use the same factor for all channels (`False`) or to sample a new value for each channel (`True`). If this value is a float `p`, then for `p` percent of all images *per_channel* will be treated as True, otherwise as False.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.BlendAlpha(0.5, iaa.Grayscale(1.0))
```

Convert each image to pure grayscale and alpha-blend the result with the original image using an alpha of `50%`, thereby removing about `50%` of all color. This is equivalent to `iaa.Grayscale(0.5)`.

```
>>> aug = iaa.BlendAlpha((0.0, 1.0), iaa.Grayscale(1.0))
```

Same as in the previous example, but the alpha factor is sampled uniformly from the interval `[0.0, 1.0]` once per image, thereby removing a random fraction of all colors. This is equivalent to `iaa.Grayscale((0.0, 1.0))`.

```
>>> aug = iaa.BlendAlpha(
>>>     (0.0, 1.0),
>>>     iaa.Affine(rotate=(-20, 20)),
>>>     per_channel=0.5)
```

First, rotate each image by a random degree sampled uniformly from the interval `[-20, 20]`. Then, alpha-blend that new image with the original one using a random factor sampled uniformly from the interval `[0.0, 1.0]`. For `50%` of all images, the blending happens channel-wise and the factor is sampled independently per channel (`per_channel=0.5`). As a result, e.g. the red channel may look visibly rotated (factor near `1.0`), while the green and blue channels may not look rotated (factors near `0.0`).

```
>>> aug = iaa.BlendAlpha(
>>>     (0.0, 1.0),
>>>     foreground=iaa.Add(100),
>>>     background=iaa.Multiply(0.2))
```

Apply two branches of augmenters – `A` and `B` – *independently* to input images and alpha-blend the results of these branches using a factor `f`. Branch `A` increases image pixel intensities by `100` and `B` multiplies the pixel intensities by `0.2`. `f` is sampled uniformly from the interval `[0.0, 1.0]` per image. The resulting images contain a bit of `A` and a bit of `B`.

```
>>> aug = iaa.BlendAlpha([0.25, 0.75], iaa.MedianBlur(13))
```

Apply median blur to each image and alpha-blend the result with the original image using an alpha factor of either exactly `0.25` or exactly `0.75` (sampled once per image).

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).

Continued on next page

Table 75 – continued from previous page

deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	See *get_children_lists()*.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_children_lists(*self*)
See *get_children_lists()*.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.blend.**BlendAlphaBoundingBoxes**(*labels*, *foreground=None*, *background=None*, *nb_sample_labels=None*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.blend.BlendAlphaMask*

Blend images from two branches based on areas enclosed in bounding boxes.

This class generates masks that are `1.0` within bounding boxes of given labels. A mask pixel will be set to `1.0` if *at least* one bounding box covers the area and has one of the requested labels.

This class is a thin wrapper around *BlendAlphaMask* together with *BoundingBoxesMaskGen*.

Note: Avoid using augmenters as children that affect pixel locations (e.g. horizontal flips). See *BlendAlphaMask* for details.

Note: This class will produce an `AssertionError` if there are no bounding boxes in a batch.

Added in 0.4.0.

Supported dtypes:

See *BlendAlphaMask*.

Parameters

- **labels** (*None or str or list of str or imgaug.parameters.StochasticParameter*) – See *BoundingBoxesMaskGen*.
- **foreground** (*None or imgaug.augmenters.meta.Augmenter or iterable of imgaug.augmenters.meta.Augmenter, optional*) – Augmenter(s) that make up the foreground branch. High alpha values will show this branch's results.
  - If `None`, then the input images will be reused as the output of the foreground branch.
  - If `Augmenter`, then that augmenter will be used as the branch.
  - If iterable of `Augmenter`, then that iterable will be converted into a `Sequential` and used as the augmenter.
- **background** (*None or imgaug.augmenters.meta.Augmenter or iterable of imgaug.augmenters.meta.Augmenter, optional*) – Augmenter(s) that make up the background branch. Low alpha values will show this branch's results.
  - If `None`, then the input images will be reused as the output of the background branch.
  - If `Augmenter`, then that augmenter will be used as the branch.
  - If iterable of `Augmenter`, then that iterable will be converted into a `Sequential` and used as the augmenter.
- **nb_sample_labels** (*None or tuple of int or list of int or imgaug.parameters.StochasticParameter, optional*) – See *BoundingBoxesMaskGen*.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.BlendAlphaBoundingBoxes("person",
>>>                                   foreground=iaa.Grayscale(1.0))
```

Create an augmenter that removes color within bounding boxes having the label `person`.

```
>>> aug = iaa.BlendAlphaBoundingBoxes(["person", "car"],
>>>                                   foreground=iaa.AddToHue((-255, 255)))
```

Create an augmenter that randomizes the hue within bounding boxes that have the label `person` or `car`.

```
>>> aug = iaa.BlendAlphaBoundingBoxes(["person", "car"],
>>>                                   foreground=iaa.AddToHue((-255, 255)),
>>>                                   nb_sample_labels=1)
```

Create an augmenter that randomizes the hue within bounding boxes that have either the label `person` or `car`. Only one label is picked per image. Note that the sampling happens with replacement, so if `nb_sample_classes` would be `>1`, it could still lead to only one *unique* label being sampled.

```
>>> aug = iaa.BlendAlphaBoundingBoxes(None,
>>>                                   background=iaa.Multiply(0.0))
```

Create an augmenter that zeros all pixels (`Multiply(0.0)`) that are *not* (`background` branch) within bounding boxes of *any* (`None`) label. In other words, all pixels outside of bounding boxes become black. Note that we don't use `TotalDropout` here, because by default it will also remove all coordinate-based augmentables, which will break the blending of such inputs.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.

Continued on next page

Table 76 – continued from previous page

find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	See *get_children_lists()*.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.blend.**BlendAlphaCheckerboard**(*nb_rows*, *nb_cols*, *foreground=None*, *background=None*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.blend.BlendAlphaMask*

Blend images from two branches according to a checkerboard pattern.

This class generates for each image a mask following a checkboard layout of `H` rows and `W` columns. Each cell is then filled with either `1.0` or `0.0`. The cell at the top-left is always `1.0`. Its right and bottom neighbour cells are `0.0`. The 4-neighbours of any cell always have a value opposite to the cell's value (`0.0` vs. `1.0`).

This class is a thin wrapper around *BlendAlphaMask* together with *CheckerboardMaskGen*.

Note: Avoid using augmenters as children that affect pixel locations (e.g. horizontal flips). See *BlendAlphaMask* for details.

Added in 0.4.0.

Supported dtypes:

See *BlendAlphaMask*.

Parameters

- **nb_rows** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter*) – Number of rows of the checkerboard. See *CheckerboardMaskGen* for details.

- **nb_cols** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter*) – Number of columns of the checkerboard. Analogous to *nb_rows*. See `CheckerboardMaskGen` for details.
- **foreground** (*None or imgaug.augmenters.meta.Augmenter or iterable of imgaug.augmenters.meta.Augmenter, optional*) – Augmenter(s) that make up the foreground branch. High alpha values will show this branch's results.
  - If `None`, then the input images will be reused as the output of the foreground branch.
  - If `Augmenter`, then that augmenter will be used as the branch.
  - If iterable of `Augmenter`, then that iterable will be converted into a `Sequential` and used as the augmenter.
- **background** (*None or imgaug.augmenters.meta.Augmenter or iterable of imgaug.augmenters.meta.Augmenter, optional*) – Augmenter(s) that make up the background branch. Low alpha values will show this branch's results.
  - If `None`, then the input images will be reused as the output of the background branch.
  - If `Augmenter`, then that augmenter will be used as the branch.
  - If iterable of `Augmenter`, then that iterable will be converted into a `Sequential` and used as the augmenter.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.BlendAlphaCheckerboard(nb_rows=2, nb_cols=(1, 4),
>>>                                  foreground=iaa.AddToHue((-100, 100)))
```

Create an augmenter that places a `HxW` grid on each image, where `H` (rows) is always `2` and `W` is randomly and uniformly sampled from the interval `[1, 4]`. For half of the cells in the grid the hue is randomly modified, the other half of the cells is unaltered.

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.

Continued on next page

Table 77 – continued from previous page

augment_batch_(self, batch[, parents, hooks])	Augment a single batch in-place.
augment_batches(self, batches[, hooks, ...])	Augment multiple batches.
augment_bounding_boxes(self, ...[, parents, ...])	Augment a batch of bounding boxes.
augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	See *get_children_lists()*.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.blend.**BlendAlphaElementwise**(*factor=(0.0, 1.0)*, *foreground=None*, *background=None*, *per_channel=False*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.blend.BlendAlphaMask*

Alpha-blend two image sources using alpha/opacity values sampled per pixel.

This is the same as *BlendAlpha*, except that the opacity factor is sampled once per *pixel* instead of once per *image* (or a few times per image, if `BlendAlpha.per_channel` is set to `True`).

See *BlendAlpha* for more details.

This class is a wrapper around *BlendAlphaMask*.

Note: Avoid using augmenters as children that affect pixel locations (e.g. horizontal flips). See *BlendAlphaMask* for details.

Added in 0.4.0. (Before that named *AlphaElementwise*.)

Supported dtypes:

See *BlendAlphaMask*.

Parameters

- **factor** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Opacity of the results of the foreground branch. Values close to `0.0` mean that the results from the background branch (see parameter *background*) make up most of the final image.
  - If float, then that value will be used for all images.
  - If tuple `(a, b)`, then a random value from the interval `[a, b]` will be sampled per image.
  - If a list, then a random value will be picked from that list per image.
  - If `StochasticParameter`, then that parameter will be used to sample a value per image.
- **foreground** (*None or imgaug.augmenters.meta.Augmenter or iterable of imgaug.augmenters.meta.Augmenter, optional*) – Augmenter(s) that make up the foreground branch. High alpha values will show this branch's results.
  - If `None`, then the input images will be reused as the output of the foreground branch.
  - If `Augmenter`, then that augmenter will be used as the branch.
  - If iterable of `Augmenter`, then that iterable will be converted into a `Sequential` and used as the augmenter.
- **background** (*None or imgaug.augmenters.meta.Augmenter or iterable of imgaug.augmenters.meta.Augmenter, optional*) – Augmenter(s) that make up the background branch. Low alpha values will show this branch's results.
  - If `None`, then the input images will be reused as the output of the background branch.
  - If `Augmenter`, then that augmenter will be used as the branch.

  - If iterable of `Augmenter`, then that iterable will be converted into a `Sequential` and used as the augmenter.

- **per_channel** (*bool or float, optional*) – Whether to use the same factor for all channels (`False`) or to sample a new value for each channel (`True`). If this value is a float `p`, then for `p` percent of all images *per_channel* will be treated as True, otherwise as False.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.BlendAlphaElementwise(0.5, iaa.Grayscale(1.0))
```

Convert each image to pure grayscale and alpha-blend the result with the original image using an alpha of `50%` for all pixels, thereby removing about `50%` of all color. This is equivalent to `iaa.Grayscale(0.5)`. This is also equivalent to `iaa.BlendAlpha(0.5, iaa.Grayscale(1.0))`, as the opacity has a fixed value of `0.5` and is hence identical for all pixels.

```
>>> aug = iaa.BlendAlphaElementwise((0, 1.0), iaa.AddToHue(100))
```

Same as in the previous example, but here with hue-shift instead of grayscaling and additionally the alpha factor is sampled uniformly from the interval `[0.0, 1.0]` once per pixel, thereby shifting the hue by a random fraction for each pixel.

```
>>> aug = iaa.BlendAlphaElementwise(
>>>     (0.0, 1.0),
>>>     iaa.Affine(rotate=(-20, 20)),
>>>     per_channel=0.5)
```

First, rotate each image by a random degree sampled uniformly from the interval `[-20, 20]`. Then, alpha-blend that new image with the original one using a random factor sampled uniformly from the interval `[0.0, 1.0]` per pixel. For `50%` of all images, the blending happens channel-wise and the factor is sampled independently per pixel *and* channel (`per_channel=0.5`). As a result, e.g. the red channel may look visibly rotated (factor near `1.0`), while the green and blue channels may not look rotated (factors near `0.0`).

```
>>> aug = iaa.BlendAlphaElementwise(
>>>     (0.0, 1.0),
>>>     foreground=iaa.Add(100),
>>>     background=iaa.Multiply(0.2))
```

Apply two branches of augmenters – `A` and `B` – *independently* to input images and alpha-blend the results of these branches using a factor `f`. Branch `A` increases image pixel intensities by `100` and `B` multiplies the pixel

intensities by `0.2`. `f` is sampled uniformly from the interval `[0.0, 1.0]` per pixel. The resulting images contain a bit of `A` and a bit of `B`.

```
>>> aug = iaa.BlendAlphaElementwise([0.25, 0.75], iaa.MedianBlur(13))
```

Apply median blur to each image and alpha-blend the result with the original image using an alpha factor of either exactly `0.25` or exactly `0.75` (sampled once per pixel).

Attributes

factor

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	See *get_children_lists()*.
`get_parameters`(self)	See *get_parameters()*.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.

Continued on next page

Table 78 – continued from previous page

localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

factor

class imgaug.augmenters.blend.**BlendAlphaFrequencyNoise**(*exponent=(-4, 4), foreground=None, background=None, per_channel=False, size_px_max=(4, 16), upscale_method=None, iterations=(1, 3), aggregation_method=['avg', 'max'], sigmoid=0.5, sigmoid_thresh=None, seed=None, name=None, random_state='deprecated', deterministic='deprecated'*)

Bases: *imgaug.augmenters.blend.BlendAlphaElementwise*

Alpha-blend two image sources using frequency noise masks.

The alpha masks are sampled using frequency noise of varying scales, which can sometimes create large connected blobs of `1` s surrounded by `0` s and other times results in smaller patterns. If nearest neighbour upsampling is used, these blobs can be rectangular with sharp edges.

Added in 0.4.0. (Before that named *FrequencyNoiseAlpha*.)

Supported dtypes:

See *BlendAlphaElementwise*.

Parameters

- **exponent** (*number or tuple of number of list of number or imgaug.parameters.StochasticParameter, optional*) – Exponent to use when scaling in the frequency domain. Sane values are in the range `-4` (large blobs) to `4` (small patterns). To generate cloud-like structures, use roughly `-2`.
  - If number, then that number will be used as the exponent for all iterations.
  - If tuple of two numbers `(a, b)`, then a value will be sampled per iteration from the interval `[a, b]`.

  - If a list of numbers, then a value will be picked per iteration at random from that list.
  - If a `StochasticParameter`, then a value will be sampled from that parameter per iteration.
- **foreground** (*None or imgaug.augmenters.meta.Augmenter or iterable of imgaug.augmenters.meta.Augmenter, optional*) – Augmenter(s) that make up the foreground branch. High alpha values will show this branch's results.
  - If `None`, then the input images will be reused as the output of the foreground branch.
  - If `Augmenter`, then that augmenter will be used as the branch.
  - If iterable of `Augmenter`, then that iterable will be converted into a `Sequential` and used as the augmenter.
- **background** (*None or imgaug.augmenters.meta.Augmenter or iterable of imgaug.augmenters.meta.Augmenter, optional*) – Augmenter(s) that make up the background branch. Low alpha values will show this branch's results.
  - If `None`, then the input images will be reused as the output of the background branch.
  - If `Augmenter`, then that augmenter will be used as the branch.
  - If iterable of `Augmenter`, then that iterable will be converted into a `Sequential` and used as the augmenter.
- **per_channel** (*bool or float, optional*) – Whether to use the same factor for all channels (`False`) or to sample a new value for each channel (`True`). If this value is a float `p`, then for `p` percent of all images *per_channel* will be treated as `True`, otherwise as `False`.
- **size_px_max** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter, optional*) – The noise is generated in a low resolution environment. This parameter defines the maximum size of that environment (in pixels). The environment is initialized at the same size as the input image and then downscaled, so that no side exceeds *size_px_max* (aspect ratio is kept).
  - If `int`, then that number will be used as the size for all iterations.
  - If tuple of two `int` s `(a, b)`, then a value will be sampled per iteration from the discrete interval `[a..b]`.
  - If a list of `int` s, then a value will be picked per iteration at random from that list.
  - If a `StochasticParameter`, then a value will be sampled from that parameter per iteration.
- **upscale_method** (*None or imgaug.ALL or str or list of str or imgaug.parameters.StochasticParameter, optional*) – After generating the noise maps in low resolution environments, they have to be upscaled to the input image size. This parameter controls the upscaling method.
  - If `None`, then either `nearest` or `linear` or `cubic` is picked. Most weight is put on `linear`, followed by `cubic`.
  - If `imgaug.ALL`, then either `nearest` or `linear` or `area` or `cubic` is picked per iteration (all same probability).
  - If string, then that value will be used as the method (must be `nearest` or `linear` or `area` or `cubic`).
  - If list of string, then a random value will be picked from that list per iteration.

  - If `StochasticParameter`, then a random value will be sampled from that parameter per iteration.

- **iterations** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter, optional*) – How often to repeat the simplex noise generation process per image.
  - If `int`, then that number will be used as the iterations for all images.
  - If tuple of two `int` s `(a, b)`, then a value will be sampled per image from the discrete interval `[a..b]`.
  - If a list of `int` s, then a value will be picked per image at random from that list.
  - If a `StochasticParameter`, then a value will be sampled from that parameter per image.
- **aggregation_method** (*imgaug.ALL or str or list of str or imgaug.parameters.StochasticParameter, optional*) – The noise maps (from each iteration) are combined to one noise map using an aggregation process. This parameter defines the method used for that process. Valid methods are `min`, `max` or `avg`, where 'min' combines the noise maps by taking the (elementwise) minimum over all iteration's results, `max` the (elementwise) maximum and `avg` the (elementwise) average.
  - If `imgaug.ALL`, then a random value will be picked per image from the valid ones.
  - If a string, then that value will always be used as the method.
  - If a list of string, then a random value will be picked from that list per image.
  - If a `StochasticParameter`, then a random value will be sampled from that parameter per image.
- **sigmoid** (*bool or number, optional*) – Whether to apply a sigmoid function to the final noise maps, resulting in maps that have more extreme values (close to `0.0` or `1.0`).
  - If `bool`, then a sigmoid will always (`True`) or never (`False`) be applied.
  - If a number `p` with `0<=p<=1`, then a sigmoid will be applied to `p` percent of all final noise maps.
- **sigmoid_thresh** (*None or number or tuple of number or imgaug.parameters.StochasticParameter, optional*) – Threshold of the sigmoid, when applied. Thresholds above zero (e.g. `5.0`) will move the saddle point towards the right, leading to more values close to `0.0`.
  - If `None`, then `Normal(0, 5.0)` will be used.
  - If number, then that threshold will be used for all images.
  - If tuple of two numbers `(a, b)`, then a random value will be sampled per image from the range `[a, b]`.
  - If `StochasticParameter`, then a random value will be sampled from that parameter per image.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage

will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.

- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.BlendAlphaFrequencyNoise(foreground=iaa.EdgeDetect(1.0))
```

Detect per image all edges, mark them in a black and white image and then alpha-blend the result with the original image using frequency noise masks.

```
>>> aug = iaa.BlendAlphaFrequencyNoise(
>>>     foreground=iaa.EdgeDetect(1.0),
>>>     upscale_method="nearest")
```

Same as the first example, but using only linear upscaling to scale the frequency noise masks to the final image sizes, i.e. no nearest neighbour upsampling is used. This results in smooth edges.

```
>>> aug = iaa.BlendAlphaFrequencyNoise(
>>>     foreground=iaa.EdgeDetect(1.0),
>>>     upscale_method="linear")
```

Same as the first example, but using only linear upscaling to scale the frequency noise masks to the final image sizes, i.e. no nearest neighbour upsampling is used. This results in smooth edges.

```
>>> aug = iaa.BlendAlphaFrequencyNoise(
>>>     foreground=iaa.EdgeDetect(1.0),
>>>     upscale_method="linear",
>>>     exponent=-2,
>>>     sigmoid=False)
```

Same as in the previous example, but with the exponent set to a constant `-2` and the sigmoid deactivated, resulting in cloud-like patterns without sharp edges.

```
>>> aug = iaa.BlendAlphaFrequencyNoise(
>>>     foreground=iaa.EdgeDetect(1.0),
>>>     sigmoid_thresh=iap.Normal(10.0, 5.0))
```

Same as the first example, but using a threshold for the sigmoid function that is further to the right. This is more conservative, i.e. the generated noise masks will be mostly black (values around `0.0`), which means that most of the original images (parameter/branch *background*) will be kept, rather than using the results of the augmentation (parameter/branch *foreground*).

Attributes

factor

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.

Continued on next page

Table 79 – continued from previous page

augment(self[, return_batch, hooks])	Augment a batch.
augment_batch(self, batch[, hooks])	**Deprecated**.
augment_batch_(self, batch[, parents, hooks])	Augment a single batch in-place.
augment_batches(self, batches[, hooks, ...])	Augment multiple batches.
augment_bounding_boxes(self, ...[, parents, ...])	Augment a batch of bounding boxes.
augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	See *get_children_lists()*.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.blend.**BlendAlphaHorizontalLinearGradient**(*foreground=None*, *background=None*, *min_value=(0.0, 0.2)*, *max_value=(0.8, 1.0)*, *start_at=(0.0, 0.2)*, *end_at=(0.8, 1.0)*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.blend.BlendAlphaMask*

Blend images from two branches along a horizontal linear gradient.

This class generates a horizontal linear gradient mask (i.e. usually a mask with low values on the left and high values on the right) and alphas-blends between foreground and background branch using that mask.

This class is a thin wrapper around *BlendAlphaMask* together with *HorizontalLinearGradientMaskGen*.

Note: Avoid using augmenters as children that affect pixel locations (e.g. horizontal flips). See *BlendAlphaMask* for details.

Added in 0.4.0.

Supported dtypes:

See *BlendAlphaMask*.

Parameters

- **foreground** (*None or imgaug.augmenters.meta.Augmenter or iterable of imgaug.augmenters.meta.Augmenter, optional*) – Augmenter(s) that make up the foreground branch. High alpha values will show this branch's results.
  - If `None`, then the input images will be reused as the output of the foreground branch.
  - If `Augmenter`, then that augmenter will be used as the branch.
  - If iterable of `Augmenter`, then that iterable will be converted into a `Sequential` and used as the augmenter.
- **background** (*None or imgaug.augmenters.meta.Augmenter or iterable of imgaug.augmenters.meta.Augmenter, optional*) – Augmenter(s) that make up the background branch. Low alpha values will show this branch's results.
  - If `None`, then the input images will be reused as the output of the background branch.
  - If `Augmenter`, then that augmenter will be used as the branch.
  - If iterable of `Augmenter`, then that iterable will be converted into a `Sequential` and used as the augmenter.

- **min_value** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – See `HorizontalLinearGradientMaskGen`.
- **max_value** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – See `HorizontalLinearGradientMaskGen`.
- **start_at** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – See `HorizontalLinearGradientMaskGen`.
- **end_at** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – See `HorizontalLinearGradientMaskGen`.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.BlendAlphaHorizontalLinearGradient(iaa.AddToHue((-100, 100)))
```

Create an augmenter that randomizes the hue towards the right of the image.

```
>>> aug = iaa.BlendAlphaHorizontalLinearGradient(
>>>     iaa.TotalDropout(1.0),
>>>     min_value=0.2, max_value=0.8)
```

Create an augmenter that replaces pixels towards the right with darker and darker values. However it always keeps at least 20% (`1.0 - max_value`) of the original pixel value on the far right and always replaces at least 20% on the far left (`min_value=0.2`).

```
>>> aug = iaa.BlendAlphaHorizontalLinearGradient(
>>>     iaa.AveragePooling(11),
>>>     start_at=(0.0, 1.0), end_at=(0.0, 1.0))
```

Create an augmenter that blends with an average-pooled image according to a horizontal gradient that starts at a random x-coordinate and reaches its maximum at another random x-coordinate. Due to that randomness, the gradient may increase towards the left or right.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	See *get_children_lists()*.
`get_parameters`(self)	See *get_parameters()*.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`localize_random_state_`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`pool`(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
`remove_augmenters`(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
`remove_augmenters_`(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
`remove_augmenters_inplace`(self, func[, parents])	**Deprecated**.
`reseed`(self[, random_state, deterministic_too])	**Deprecated**.
`seed_`(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.

Continued on next page

Table 80 – continued from previous page

show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.blend.**BlendAlphaMask**(*mask_generator*, *foreground=None*, *background=None*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Alpha-blend two image sources using non-binary masks generated per image.

This augmenter queries for each image a mask generator to generate a `(H,W)` or `(H,W,C)` channelwise mask `[0.0, 1.0]`, where `H` is the image height and `W` the width. The mask will then be used to alpha-blend pixel- and possibly channel-wise between a foreground branch of augmenters and a background branch. (Both branches default to the identity operation if not provided.)

See also *BlendAlpha*.

Note: It is not recommended to use `BlendAlphaMask` with augmenters that change the geometry of images (e.g. horizontal flips, affine transformations) if you *also* want to augment coordinates (e.g. keypoints, polygons, ...), as it is unclear which of the two coordinate results (foreground or background branch) should be used as the final output coordinates after augmentation.

Currently, for keypoints the results of the foreground and background branch will be mixed. That means that for each coordinate the augmented result will be picked from the foreground or background branch based on the average alpha mask value at the corresponding spatial location.

For bounding boxes, line strings and polygons, either all objects (on an image) of the foreground or all of the background branch will be used, based on the average over the whole alpha mask.

Added in 0.4.0.

Supported dtypes:

See *blend_alpha()*.

Parameters

- **mask_generator** (*IBatchwiseMaskGenerator*) – A generator that will be queried per image to generate a mask.
- **foreground** (*None or imgaug.augmenters.meta.Augmenter or iterable of imgaug.augmenters.meta.Augmenter, optional*) – Augmenter(s) that make up the foreground branch. High alpha values will show this branch's results.
  - If `None`, then the input images will be reused as the output of the foreground branch (i.e. identity function).
  - If `Augmenter`, then that augmenter will be used as the branch.
  - If iterable of `Augmenter`, then that iterable will be converted into a `Sequential` and used as the augmenter.
- **background** (*None or imgaug.augmenters.meta.Augmenter or iterable of imgaug.augmenters.meta.Augmenter, optional*) – Augmenter(s) that make up the background branch. Low alpha values will show this branch's results.

  - If `None`, then the input images will be reused as the output of the background branch (i.e. identity function).
  - If `Augmenter`, then that augmenter will be used as the branch.
  - If iterable of `Augmenter`, then that iterable will be converted into a `Sequential` and used as the augmenter.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.BlendAlphaMask(
>>>     iaa.InvertMaskGen(0.5, iaa.VerticalLinearGradientMaskGen()),
>>>     iaa.Sequential([
>>>         iaa.Clouds(),
>>>         iaa.WithChannels([1, 2], iaa.Multiply(0.5))
>>>     ])
>>> )
```

Create an augmenter that sometimes adds clouds at the bottom and sometimes at the top of the image.

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.

Continued on next page

Table 81 – continued from previous page

augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	See *get_children_lists()*.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_children_lists(*self*)
See *get_children_lists()*.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.blend.**BlendAlphaRegularGrid**(*nb_rows, nb_cols, foreground=None, background=None, alpha=[0.0, 1.0], seed=None, name=None, random_state='deprecated', deterministic='deprecated'*)

Bases: *imgaug.augmenters.blend.BlendAlphaMask*

Blend images from two branches according to a regular grid.

This class generates for each image a mask that splits the image into a grid-like pattern of `H` rows and `W` columns. Each cell is then filled with an alpha value, sampled randomly per cell.

The difference to `AlphaBlendCheckerboard` is that this class samples random alpha values per grid cell, while in the checkerboard the alpha values follow a fixed pattern.

This class is a thin wrapper around *BlendAlphaMask* together with *RegularGridMaskGen*.

Note: Avoid using augmenters as children that affect pixel locations (e.g. horizontal flips). See *BlendAlphaMask* for details.

Added in 0.4.0.

Supported dtypes:

See *BlendAlphaMask*.

Parameters

- **nb_rows** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter*) – Number of rows of the checkerboard. See *CheckerboardMaskGen* for details.
- **nb_cols** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter*) – Number of columns of the checkerboard. Analogous to *nb_rows*. See *CheckerboardMaskGen* for details.
- **foreground** (*None or imgaug.augmenters.meta.Augmenter or iterable of imgaug.augmenters.meta.Augmenter, optional*) – Augmenter(s) that make up the foreground branch. High alpha values will show this branch's results.
  - If `None`, then the input images will be reused as the output of the foreground branch.
  - If `Augmenter`, then that augmenter will be used as the branch.
  - If iterable of `Augmenter`, then that iterable will be converted into a `Sequential` and used as the augmenter.
- **background** (*None or imgaug.augmenters.meta.Augmenter or iterable of imgaug.augmenters.meta.Augmenter, optional*) – Augmenter(s) that make up the background branch. Low alpha values will show this branch's results.
  - If `None`, then the input images will be reused as the output of the background branch.
  - If `Augmenter`, then that augmenter will be used as the branch.
  - If iterable of `Augmenter`, then that iterable will be converted into a `Sequential` and used as the augmenter.
- **alpha** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Alpha value of each cell.
  - If `number`: Exactly that value will be used for all images.
  - If `tuple (a, b)`: A random value will be uniformly sampled per image from the interval `[a, b]`.
  - If `list`: A random value will be picked per image from that list.
  - If `StochasticParameter`: That parameter will be queried once per batch for `(N,)` values – one per image.

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.BlendAlphaRegularGrid(nb_rows=(4, 6), nb_cols=(1, 4),
>>>                                 foreground=iaa.Multiply(0.0))
```

Create an augmenter that places a `HxW` grid on each image, where `H` (rows) is randomly and uniformly sampled from the interval `[4, 6]` and `W` is analogously sampled from the interval `[1, 4]`. Roughly half of the cells in the grid are filled with `0.0`, the remaining ones are unaltered. Which cells exactly are "dropped" is randomly decided per image. The resulting effect is similar to *CoarseDropout*.

```
>>> aug = iaa.BlendAlphaRegularGrid(nb_rows=2, nb_cols=2,
>>>                                 foreground=iaa.Multiply(0.0),
>>>                                 background=iaa.AveragePooling(8),
>>>                                 alpha=[0.0, 0.0, 1.0])
```

Create an augmenter that always placed `2x2` cells on each image and sets about `1/3` of them to zero (foreground branch) and the remaining `2/3` to a pixelated version (background branch).

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.

Continued on next page

Table 82 – continued from previous page

augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	See *get_children_lists()*.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.blend.**BlendAlphaSegMapClassIds**(*class_ids*, *foreground=None*, *background=None*, *nb_sample_classes=None*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.blend.BlendAlphaMask*

Blend images from two branches based on segmentation map ids.

This class generates masks that are `1.0` at pixel locations covered by specific classes in segmentation maps.

This class is a thin wrapper around *BlendAlphaMask* together with *SegMapClassIdsMaskGen*.

Note: Avoid using augmenters as children that affect pixel locations (e.g. horizontal flips). See *BlendAlphaMask* for details.

Note: Segmentation maps can have multiple channels. If that is the case then for each position `(x, y)` it is sufficient that any class id in any channel matches one of the desired class ids.

Note: This class will produce an `AssertionError` if there are no segmentation maps in a batch.

Added in 0.4.0.

Supported dtypes:

See *BlendAlphaMask*.

Parameters

- **class_ids** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter*) – See *SegMapClassIdsMaskGen*.
- **foreground** (*None or imgaug.augmenters.meta.Augmenter or iterable of imgaug.augmenters.meta.Augmenter, optional*) – Augmenter(s) that make up the foreground branch. High alpha values will show this branch's results.
  - If `None`, then the input images will be reused as the output of the foreground branch.
  - If `Augmenter`, then that augmenter will be used as the branch.
  - If iterable of `Augmenter`, then that iterable will be converted into a `Sequential` and used as the augmenter.
- **background** (*None or imgaug.augmenters.meta.Augmenter or iterable of imgaug.augmenters.meta.Augmenter, optional*) – Augmenter(s) that make up the background branch. Low alpha values will show this branch's results.
  - If `None`, then the input images will be reused as the output of the background branch.
  - If `Augmenter`, then that augmenter will be used as the branch.
  - If iterable of `Augmenter`, then that iterable will be converted into a `Sequential` and used as the augmenter.
- **nb_sample_classes** (*None or tuple of int or list of int or imgaug.parameters.StochasticParameter, optional*) – See *SegMapClassIdsMaskGen*.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.

- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.BlendAlphaSegMapClassIds(
>>>     [1, 3],
>>>     foreground=iaa.AddToHue((-100, 100)))
```

Create an augmenter that randomizes the hue wherever the segmentation maps contain the classes `1` or `3`.

```
>>> aug = iaa.BlendAlphaSegMapClassIds(
>>>     [1, 2, 3, 4],
>>>     nb_sample_classes=2,
>>>     foreground=iaa.GaussianBlur(3.0))
```

Create an augmenter that randomly picks `2` classes from the list `[1, 2, 3, 4]` and blurs the image content wherever these classes appear in the segmentation map. Note that as the sampling of class ids happens *with replacement*, it is not guaranteed to sample two *unique* class ids.

```
>>> aug = iaa.Sometimes(0.2,
>>>     iaa.BlendAlphaSegMapClassIds(
>>>         2,
>>>         background=iaa.TotalDropout(1.0)))
```

Create an augmenter that zeros for roughly every fifth image all image pixels that do *not* belong to class id `2` (note that the *background* branch was used, not the *foreground* branch). Example use case: Human body landmark detection where both the landmarks/keypoints and the body segmentation map are known. Train the model to detect landmarks and sometimes remove all non-body information to force the model to become more independent of the background.

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.

Continued on next page

Table 83 – continued from previous page

augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	See *get_children_lists()*.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.blend.**BlendAlphaSimplexNoise**(*foreground=None*, *background=None*, *per_channel=False*, *size_px_max=(2, 16)*, *upscale_method=None*, *iterations=(1, 3)*, *aggregation_method='max'*, *sigmoid=True*, *sigmoid_thresh=None*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.blend.BlendAlphaElementwise*

Alpha-blend two image sources using simplex noise alpha masks.

The alpha masks are sampled using a simplex noise method, roughly creating connected blobs of 1s surrounded

by 0s. If nearest neighbour upsampling is used, these blobs can be rectangular with sharp edges.

Added in 0.4.0. (Before that named *SimplexNoiseAlpha*.)

Supported dtypes:

See `BlendAlphaElementwise`.

Parameters

- **foreground** (*None or imgaug.augmenters.meta.Augmenter or iterable of imgaug.augmenters.meta.Augmenter, optional*) – Augmenter(s) that make up the foreground branch. High alpha values will show this branch's results.
  - If `None`, then the input images will be reused as the output of the foreground branch.
  - If `Augmenter`, then that augmenter will be used as the branch.
  - If iterable of `Augmenter`, then that iterable will be converted into a `Sequential` and used as the augmenter.
- **background** (*None or imgaug.augmenters.meta.Augmenter or iterable of imgaug.augmenters.meta.Augmenter, optional*) – Augmenter(s) that make up the background branch. Low alpha values will show this branch's results.
  - If `None`, then the input images will be reused as the output of the background branch.
  - If `Augmenter`, then that augmenter will be used as the branch.
  - If iterable of `Augmenter`, then that iterable will be converted into a `Sequential` and used as the augmenter.
- **per_channel** (*bool or float, optional*) – Whether to use the same factor for all channels (`False`) or to sample a new value for each channel (`True`). If this value is a float `p`, then for `p` percent of all images *per_channel* will be treated as `True`, otherwise as `False`.
- **size_px_max** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter, optional*) – The simplex noise is always generated in a low resolution environment. This parameter defines the maximum size of that environment (in pixels). The environment is initialized at the same size as the input image and then downscaled, so that no side exceeds *size_px_max* (aspect ratio is kept).
  - If int, then that number will be used as the size for all iterations.
  - If tuple of two `int` s `(a, b)`, then a value will be sampled per iteration from the discrete interval `[a..b]`.
  - If a list of `int` s, then a value will be picked per iteration at random from that list.
  - If a `StochasticParameter`, then a value will be sampled from that parameter per iteration.
- **upscale_method** (*None or imgaug.ALL or str or list of str or imgaug.parameters.StochasticParameter, optional*) – After generating the noise maps in low resolution environments, they have to be upscaled to the input image size. This parameter controls the upscaling method.
  - If `None`, then either `nearest` or `linear` or `cubic` is picked. Most weight is put on `linear`, followed by `cubic`.
  - If `imgaug.ALL`, then either `nearest` or `linear` or `area` or `cubic` is picked per iteration (all same probability).
  - If a string, then that value will be used as the method (must be `nearest` or `linear` or `area` or `cubic`).

  - If list of string, then a random value will be picked from that list per iteration.
  - If `StochasticParameter`, then a random value will be sampled from that parameter per iteration.

- **iterations** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter, optional*) –

  How often to repeat the simplex noise generation process per image.

  - If `int`, then that number will be used as the iterations for all images.
  - If tuple of two `int` s `(a, b)`, then a value will be sampled per image from the discrete interval `[a..b]`.
  - If a list of `int` s, then a value will be picked per image at random from that list.
  - If a `StochasticParameter`, then a value will be sampled from that parameter per image.

- **aggregation_method** (*imgaug.ALL or str or list of str or imgaug.parameters.StochasticParameter, optional*) – The noise maps (from each iteration) are combined to one noise map using an aggregation process. This parameter defines the method used for that process. Valid methods are `min`, `max` or `avg`, where `min` combines the noise maps by taking the (elementwise) minimum over all iteration's results, `max` the (elementwise) maximum and `avg` the (elementwise) average.

  - If `imgaug.ALL`, then a random value will be picked per image from the valid ones.
  - If a string, then that value will always be used as the method.
  - If a list of string, then a random value will be picked from that list per image.
  - If a `StochasticParameter`, then a random value will be sampled from that paramter per image.

- **sigmoid** (*bool or number, optional*) – Whether to apply a sigmoid function to the final noise maps, resulting in maps that have more extreme values (close to 0.0 or 1.0).

  - If `bool`, then a sigmoid will always (`True`) or never (`False`) be applied.
  - If a number `p` with `0<=p<=1`, then a sigmoid will be applied to `p` percent of all final noise maps.

- **sigmoid_thresh** (*None or number or tuple of number or imgaug.parameters.StochasticParameter, optional*) – Threshold of the sigmoid, when applied. Thresholds above zero (e.g. `5.0`) will move the saddle point towards the right, leading to more values close to 0.0.

  - If `None`, then `Normal(0, 5.0)` will be used.
  - If number, then that threshold will be used for all images.
  - If tuple of two numbers `(a, b)`, then a random value will be sampled per image from the interval `[a, b]`.
  - If `StochasticParameter`, then a random value will be sampled from that parameter per image.

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.

- **name** (*None or str, optional*) – See `__init__()`.

- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.BlendAlphaSimplexNoise(iaa.EdgeDetect(1.0))
```

Detect per image all edges, mark them in a black and white image and then alpha-blend the result with the original image using simplex noise masks.

```
>>> aug = iaa.BlendAlphaSimplexNoise(
>>>     iaa.EdgeDetect(1.0),
>>>     upscale_method="nearest")
```

Same as in the previous example, but using only nearest neighbour upscaling to scale the simplex noise masks to the final image sizes, i.e. no nearest linear upsampling is used. This leads to rectangles with sharp edges.

```
>>> aug = iaa.BlendAlphaSimplexNoise(
>>>     iaa.EdgeDetect(1.0),
>>>     upscale_method="linear")
```

Same as in the previous example, but using only linear upscaling to scale the simplex noise masks to the final image sizes, i.e. no nearest neighbour upsampling is used. This leads to rectangles with smooth edges.

```
>>> aug = iaa.BlendAlphaSimplexNoise(
>>>     iaa.EdgeDetect(1.0),
>>>     sigmoid_thresh=iap.Normal(10.0, 5.0))
```

Same as in the first example, but using a threshold for the sigmoid function that is further to the right. This is more conservative, i.e. the generated noise masks will be mostly black (values around `0.0`), which means that most of the original images (parameter/branch *background*) will be kept, rather than using the results of the augmentation (parameter/branch *foreground*).

Attributes

factor

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.

Continued on next page

Table 84 – continued from previous page

augment_bounding_boxes(self, ...[, parents, ...])	Augment a batch of bounding boxes.
augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	See *get_children_lists()*.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.blend.**BlendAlphaSomeColors**(*foreground=None, background=None, nb_bins=(5, 15), smoothness=(0.1, 0.3), alpha=[0.0, 1.0], rotation_deg=(0, 360), from_colorspace='RGB', seed=None, name=None, random_state='deprecated', deterministic='deprecated'*)

Bases: *imgaug.augmenters.blend.BlendAlphaMask*

Blend images from two branches using colorwise masks.

This class generates masks that "mark" a few colors and replace the pixels within these colors with the results of the foreground branch. The remaining pixels are replaced with the results of the background branch (usually the identity function). That allows to e.g. selectively grayscale a few colors, while keeping other colors unchanged.

This class is a thin wrapper around *BlendAlphaMask* together with *SomeColorsMaskGen*.

Note: The underlying mask generator will produce an `AssertionError` for batches that contain no images.

Note: Avoid using augmenters as children that affect pixel locations (e.g. horizontal flips). See *BlendAlphaMask* for details.

Added in 0.4.0.

Supported dtypes:

See *change_colorspaces_()*.

Parameters

- **foreground** (*None or imgaug.augmenters.meta.Augmenter or iterable of imgaug.augmenters.meta.Augmenter, optional*) – Augmenter(s) that make up the foreground branch. High alpha values will show this branch's results.
  - If `None`, then the input images will be reused as the output of the foreground branch.
  - If `Augmenter`, then that augmenter will be used as the branch.
  - If iterable of `Augmenter`, then that iterable will be converted into a `Sequential` and used as the augmenter.
- **background** (*None or imgaug.augmenters.meta.Augmenter or iterable of imgaug.augmenters.meta.Augmenter, optional*) – Augmenter(s) that make up the background branch. Low alpha values will show this branch's results.
  - If `None`, then the input images will be reused as the output of the background branch.
  - If `Augmenter`, then that augmenter will be used as the branch.
  - If iterable of `Augmenter`, then that iterable will be converted into a `Sequential` and used as the augmenter.
- **nb_bins** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter, optional*) – See *SomeColorsMaskGen*.
- **smoothness** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – See *SomeColorsMaskGen*.

- **alpha** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – See *SomeColorsMaskGen*.
- **rotation_deg** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – See *SomeColorsMaskGen*.
- **from_colorspace** (*str, optional*) – See *SomeColorsMaskGen*.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.BlendAlphaSomeColors(iaa.Grayscale(1.0))
```

Create an augmenter that turns randomly removes some colors in images by grayscaling them.

```
>>> aug = iaa.BlendAlphaSomeColors(iaa.TotalDropout(1.0))
```

Create an augmenter that removes some colors in images by replacing them with black pixels.

```
>>> aug = iaa.BlendAlphaSomeColors(
>>>     iaa.MultiplySaturation(0.5), iaa.MultiplySaturation(1.5))
```

Create an augmenter that desaturates some colors and increases the saturation of the remaining ones.

```
>>> aug = iaa.BlendAlphaSomeColors(
>>>     iaa.AveragePooling(7), alpha=[0.0, 1.0], smoothness=0.0)
```

Create an augmenter that applies average pooling to some colors. Each color tune is either selected (alpha of `1.0`) or not selected (`0.0`). There is no gradual change between similar colors.

```
>>> aug = iaa.BlendAlphaSomeColors(
>>>     iaa.AveragePooling(7), nb_bins=2, smoothness=0.0)
```

Create an augmenter that applies average pooling to some colors. Choose on average half of all colors in images for the blending operation.

```
>>> aug = iaa.BlendAlphaSomeColors(
>>>     iaa.AveragePooling(7), from_colorspace="BGR")
```

Create an augmenter that applies average pooling to some colors with input images being in BGR colorspace.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	See *get_children_lists()*.
`get_parameters`(self)	See *get_parameters()*.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`localize_random_state_`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`pool`(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
`remove_augmenters`(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
`remove_augmenters_`(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
`remove_augmenters_inplace`(self, func[, parents])	**Deprecated**.
`reseed`(self[, random_state, deterministic_too])	**Deprecated**.
`seed_`(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.

Continued on next page

Table 85 – continued from previous page

show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.blend.**BlendAlphaVerticalLinearGradient**(*foreground=None*, *background=None*, *min_value=(0.0, 0.2)*, *max_value=(0.8, 1.0)*, *start_at=(0.0, 0.2)*, *end_at=(0.8, 1.0)*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.blend.BlendAlphaMask*

Blend images from two branches along a vertical linear gradient.

This class generates a vertical linear gradient mask (i.e. usually a mask with low values on the left and high values on the right) and alphas-blends between foreground and background branch using that mask.

This class is a thin wrapper around *BlendAlphaMask* together with *VerticalLinearGradientMaskGen*.

Note: Avoid using augmenters as children that affect pixel locations (e.g. horizontal flips). See *BlendAlphaMask* for details.

Added in 0.4.0.

Supported dtypes:

See *BlendAlphaMask*.

Parameters

- **foreground** (*None or imgaug.augmenters.meta.Augmenter or iterable of imgaug.augmenters.meta.Augmenter, optional*) – Augmenter(s) that make up the foreground branch. High alpha values will show this branch's results.
  - If None, then the input images will be reused as the output of the foreground branch.
  - If Augmenter, then that augmenter will be used as the branch.
  - If iterable of Augmenter, then that iterable will be converted into a Sequential and used as the augmenter.
- **background** (*None or imgaug.augmenters.meta.Augmenter or iterable of imgaug.augmenters.meta.Augmenter, optional*) – Augmenter(s) that make up the background branch. Low alpha values will show this branch's results.

  - If `None`, then the input images will be reused as the output of the background branch.
  - If `Augmenter`, then that augmenter will be used as the branch.
  - If iterable of `Augmenter`, then that iterable will be converted into a `Sequential` and used as the augmenter.

- **min_value** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – See *VerticalLinearGradientMaskGen*.
- **max_value** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – See *VerticalLinearGradientMaskGen*.
- **start_at** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – See *VerticalLinearGradientMaskGen*.
- **end_at** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – See *VerticalLinearGradientMaskGen*.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.BlendAlphaVerticalLinearGradient(iaa.AddToHue((-100, 100)))
```

Create an augmenter that randomizes the hue towards the bottom of the image.

```
>>> aug = iaa.BlendAlphaVerticalLinearGradient(
>>>     iaa.TotalDropout(1.0),
>>>     min_value=0.2, max_value=0.8)
```

Create an augmenter that replaces pixels towards the bottom with darker and darker values. However it always keeps at least 20% (`1.0 - max_value`) of the original pixel value on the far bottom and always replaces at least 20% on the far top (`min_value=0.2`).

```
>>> aug = iaa.BlendAlphaVerticalLinearGradient(
>>>     iaa.AveragePooling(11),
>>>     start_at=(0.0, 1.0), end_at=(0.0, 1.0))
```

Create an augmenter that blends with an average-pooled image according to a vertical gradient that starts at a random y-coordinate and reaches its maximum at another random y-coordinate. Due to that randomness, the gradient may increase towards the bottom or top.

```
>>> aug = iaa.BlendAlphaVerticalLinearGradient(
>>>     iaa.Clouds(),
>>>     start_at=(0.15, 0.35), end_at=0.0)
```

Create an augmenter that draws clouds in roughly the top quarter of the image.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	See *get_children_lists()*.
`get_parameters`(self)	See *get_parameters()*.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`localize_random_state_`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.

Continued on next page

Table 86 – continued from previous page

pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.blend.**BoundingBoxesMaskGen**(*labels=None*, *nb_sample_labels=None*)

Bases: *imgaug.augmenters.blend.IBatchwiseMaskGenerator*

Generator that produces masks highlighting bounding boxes.

This class produces for each row (i.e. image + bounding boxes) in a batch a mask in which the inner areas of bounding box rectangles with given labels are marked (i.e. set to `1.0`). The labels may be provided as a fixed list of strings or a stochastic parameter from which labels will be sampled. If no labels are provided, all bounding boxes will be marked.

A pixel will be set to `1.0` if *at least* one bounding box at that location has one of the requested labels, even if there is *also* one bounding box at that location with a not requested label.

Note: This class will produce an `AssertionError` if there are no bounding boxes in a batch.

Added in 0.4.0.

Parameters

- **labels** (*None or str or list of str or imgaug.parameters.StochasticParameter*) – Labels of bounding boxes to select for.

  If *nb_sample_labels* is `None` then this is expected to be either also `None` (select all BBs) or a single `str` (select BBs with this one label) or a `list` of `str` s (always select BBs with these labels).

  If *nb_sample_labels* is set, then this parameter will be treated as a stochastic parameter with the following valid types:

  – If `None`: Ignore the sampling count and always use all bounding boxes.

  – If `str`: Exactly that label will be used for all images.

  – If `list` of `str`: `N` random values will be picked per image from that list and used as the labels.

  – If `StochasticParameter`: That parameter will be queried once per batch for `(sum(N),)` values.

  `N` denotes the number of labels to sample per segmentation map (derived from *nb_sample_labels*) and `sum(N)` denotes the sum of `N` s over all images.

- **nb_sample_labels** (*None or tuple of int or list of int or imgaug.parameters.StochasticParameter, optional*) – Number of labels to sample (with replacement) per image. As sampling happens with replacement, fewer *unique* labels may be sampled.
  - If `None`: *labels* is expected to also be `None` or a fixed value of labels to be used for all images.
  - If `int`: Exactly that many labels will be sampled for all images.
  - If `tuple (a, b)`: A random value will be uniformly sampled per image from the discrete interval `[a..b]`.
  - If `list`: A random value will be picked per image from that list.
  - If `StochasticParameter`: That parameter will be queried once per batch for `(B,)` values, where `B` is the number of images.

Methods

draw_masks(self, batch[, random_state])	See *draw_masks()*.
generate_mask(bbsoi, labels)	Generate a mask of the areas of bounding boxes with given labels.

draw_masks (*self*, *batch*, *random_state=None*)
See *draw_masks()*.

Added in 0.4.0.

classmethod generate_mask (*bbsoi*, *labels*)
Generate a mask of the areas of bounding boxes with given labels.

Added in 0.4.0.

Parameters

- **bbsoi** (*imgaug.augmentables.bbs.BoundingBoxesOnImage*) – The bounding boxes for which to generate the mask.
- **labels** (*None or iterable of str*) – Labels of the bounding boxes to set to `1.0`. For an `(x, y)` position, it is enough that *any* bounding box at the given location has one of the labels. If this is `None`, all bounding boxes will be marked.

Returns `float32` mask array with same height and width as `segmap.shape`. Values are in `[0.0, 1.0]`.

Return type ndarray

class imgaug.augmenters.blend.**CheckerboardMaskGen** (*nb_rows*, *nb_cols*)
Bases: *imgaug.augmenters.blend.IBatchwiseMaskGenerator*

Generate masks following a checkerboard-like pattern.

This mask generator splits each image into a regular grid of `H` rows and `W` columns. Each cell is then filled with either `1.0` or `0.0`. The cell at the top-left is always `1.0`. Its right and bottom neighbour cells are `0.0`. The 4-neighbours of any cell always have a value opposite to the cell's value (`0.0` vs. `1.0`).

Added in 0.4.0.

Parameters

- **nb_rows** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter, optional*) –

  Number of rows of the checkerboard.

  - If `int`: Exactly that value will be used for all images.
  - If `tuple (a, b)`: A random value will be uniformly sampled per image from the discrete interval `[a..b]`.
  - If `list`: A random value will be picked per image from that list.
  - If `StochasticParameter`: That parameter will be queried once per batch for `(N,)` values – one per image.

- **nb_cols** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter, optional*) – Number of columns of the checkerboard. Analogous to *nb_rows*.

Attributes

nb_cols	Get the number of columns of the checkerboard grid.
nb_rows	Get the number of rows of the checkerboard grid.

Methods

draw_masks(self, batch[, random_state])	See *draw_masks()*.
generate_mask(shape, nb_rows, nb_cols)	Generate a mask following a checkerboard pattern.

draw_masks (*self*, *batch*, *random_state=None*)

See *draw_masks()*.

Added in 0.4.0.

classmethod generate_mask (*shape*, *nb_rows*, *nb_cols*)

Generate a mask following a checkerboard pattern.

Added in 0.4.0.

Parameters

- **shape** (*tuple of int*) – Height and width of the output mask.
- **nb_rows** (*int*) – Number of rows of the checkerboard pattern.
- **nb_cols** (*int*) – Number of columns of the checkerboard pattern.

Returns `float32` mask array with same height and width as `segmap.shape`. Values are in `[0.0, 1.0]`.

Return type ndarray

nb_cols

Get the number of columns of the checkerboard grid.

Added in 0.4.0.

Returns The number of columns.

Return type int

nb_rows

Get the number of rows of the checkerboard grid.

Added in 0.4.0.

Returns The number of rows.

Return type int

imgaug.augmenters.blend.**FrequencyNoiseAlpha**(*exponent=(-4, 4), first=None, second=None, per_channel=False, size_px_max=(4, 16), upscale_method=None, iterations=(1, 3), aggregation_method=['avg', 'max'], sigmoid=0.5, sigmoid_thresh=None, seed=None, name=None, random_state='deprecated', deterministic='deprecated'*)

Deprecated. Use `BlendAlphaFrequencyNoise` instead. FrequencyNoiseAlpha is deprecated. Use BlendAlphaFrequencyNoise instead. The order of parameters is the same. Parameter 'first' was renamed to 'foreground'. Parameter 'second' was renamed to 'background'.

See *BlendAlphaFrequencyNoise*.

Deprecated since 0.4.0.

class imgaug.augmenters.blend.**HorizontalLinearGradientMaskGen**(*min_value=(0.0, 0.2), max_value=(0.8, 1.0), start_at=(0.0, 0.2), end_at=(0.8, 1.0)*)

Bases: `imgaug.augmenters.blend._LinearGradientMaskGen`

Generator that produces horizontal linear gradient masks.

This class receives batches and produces for each row (i.e. image) a horizontal linear gradient that matches the row's shape (i.e. image shape). The gradient increases linearly from a minimum value to a maximum value along the x-axis. The start and end points (i.e. where the minimum value starts to increase and where it reaches the maximum) may be defines as fractions of the width. E.g. for width `100` and `start=0.25`, `end=0.75`, the gradient would have its minimum in interval `[0px, 25px]` and its maximum in interval `[75px, 100px]`.

Note that this has nothing to do with a *derivative* along the x-axis.

Added in 0.4.0.

Parameters

- **min_value** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Minimum value that the mask will have up to the start point of the linear gradient. Note that *min_value* is allowed to be larger than *max_value*, in which case the gradient will start at the (higher) *min_value* and decrease towards the (lower) *max_value*.
  - If `number`: Exactly that value will be used for all images.
  - If `tuple (a, b)`: A random value will be uniformly sampled per image from the interval `[a, b]`.
  - If `list`: A random value will be picked per image from that list.
  - If `StochasticParameter`: That parameter will be queried once per batch for `(N,)` values – one per image.

- **max_value** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Maximum value that the mask will have at the end of the linear gradient.

  Datatypes are analogous to *min_value*.

- **start_at** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Position on the x-axis where the linear gradient starts, given as a fraction of the axis size. Interval is `[0.0, 1.0]`, where `0.0` is at the left of the image. If `end_at < start_at` the gradient will be inverted.

  Datatypes are analogous to *min_value*.

- **end_at** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Position on the x-axis where the linear gradient ends, given as a fraction of the axis size. Interval is `[0.0, 1.0]`, where `0.0` is at the right of the image.

  Datatypes are analogous to *min_value*.

Methods

`draw_masks`(self, batch[, random_state])	See *draw_masks()*.
generate_mask(shape, min_value, max_value, ...)	Generate a linear horizontal gradient mask.

classmethod generate_mask(*shape*, *min_value*, *max_value*, *start_at*, *end_at*)

Generate a linear horizontal gradient mask.

Added in 0.4.0.

Parameters

- **shape** (*tuple of int*) – Shape of the image. The mask will have the same height and width.
- **min_value** (*number*) – Minimum value of the gradient in interval `[0.0, 1.0]`.
- **max_value** (*number*) – Maximum value of the gradient in interval `[0.0, 1.0]`.
- **start_at** (*number*) – Position on the x-axis where the linear gradient starts, given as a fraction of the axis size. Interval is `[0.0, 1.0]`.
- **end_at** (*number*) – Position on the x-axis where the linear gradient ends, given as a fraction of the axis size. Interval is `[0.0, 1.0]`.

Returns `float32` mask array with same height and width as the image. Values are in `[0.0, 1.0]`.

Return type ndarray

class imgaug.augmenters.blend.**IBatchwiseMaskGenerator**

Bases: `object`

Interface for classes generating masks for batches.

Child classes are supposed to receive a batch and generate an iterable of masks, one per row (i.e. image), matching the row shape (i.e. image shape). This is used in *BlendAlphaMask*.

Added in 0.4.0.

Methods

draw_masks(self, batch[, random_state])	Generate a mask with given shape.

draw_masks (*self*, *batch*, *random_state=None*)
Generate a mask with given shape.

Parameters

- **batch** (*imgaug.augmentables.batches._BatchInAugmentation*) – Shape of the mask to sample.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – A seed or random number generator to use during the sampling process. If `None`, the global RNG will be used. See also `__init__()` for a similar parameter with more details.

Returns Masks, one per row in the batch. Each mask must be a `float32` array in interval `[0.0, 1.0]`. It must either have the same shape as the row (i.e. the image) or shape `(H, W)` if all channels are supposed to have the same mask.

Return type iterable of ndarray

class imgaug.augmenters.blend.**InvertMaskGen**(*p*, *child*)
Bases: *imgaug.augmenters.blend.IBatchwiseMaskGenerator*

Generator that inverts the outputs of other mask generators.

This class receives batches and calls for each row (i.e. image) a child mask generator to produce a mask. That mask is then inverted for `p%` of all rows, i.e. converted to `1.0 - mask`.

Added in 0.4.0.

Parameters

- **p** (*bool or float or imgaug.parameters.StochasticParameter, optional*) – Probability of inverting each mask produced by the other mask generator.
- **child** (*IBatchwiseMaskGenerator*) – The other mask generator to invert.

Methods

draw_masks(self, batch[, random_state])	See *draw_masks()*.

draw_masks (*self*, *batch*, *random_state=None*)
See *draw_masks()*.

Added in 0.4.0.

class imgaug.augmenters.blend.**RegularGridMaskGen**(*nb_rows*, *nb_cols*, *alpha=[0.0, 1.0]*)
Bases: *imgaug.augmenters.blend.IBatchwiseMaskGenerator*

Generate masks following a regular grid pattern.

This mask generator splits each image into a grid-like pattern of `H` rows and `W` columns. Each cell is then filled with an alpha value, sampled randomly per cell.

The difference to *CheckerboardMaskGen* is that this mask generator samples random alpha values per cell,

while in the checkerboard the alpha values follow a fixed pattern.

Added in 0.4.0.

Parameters

- **nb_rows** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter*) –

  Number of rows of the regular grid.

  - If `int`: Exactly that value will be used for all images.
  - If `tuple (a, b)`: A random value will be uniformly sampled per image from the discrete interval `[a..b]`.
  - If `list`: A random value will be picked per image from that list.
  - If `StochasticParameter`: That parameter will be queried once per batch for `(N,)` values – one per image.

- **nb_cols** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter*) – Number of columns of the checkerboard. Analogous to *nb_rows*.

- **alpha** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Alpha value of each cell.

  - If `number`: Exactly that value will be used for all images.
  - If `tuple (a, b)`: A random value will be uniformly sampled per image from the interval `[a, b]`.
  - If `list`: A random value will be picked per image from that list.
  - If `StochasticParameter`: That parameter will be queried once per batch for `(N,)` values – one per image.

Methods

draw_masks(self, batch[, random_state])	See *draw_masks()*.
generate_mask(shape, nb_rows, nb_cols, alphas)	Generate a mask following a checkerboard pattern.

draw_masks (*self*, *batch*, *random_state=None*)
See *draw_masks()*.

Added in 0.4.0.

classmethod generate_mask (*shape*, *nb_rows*, *nb_cols*, *alphas*)
Generate a mask following a checkerboard pattern.

Added in 0.4.0.

Parameters

- **shape** (*tuple of int*) – Height and width of the output mask.
- **nb_rows** (*int*) – Number of rows of the checkerboard pattern.
- **nb_cols** (*int*) – Number of columns of the checkerboard pattern.
- **alphas** (*ndarray*) – 1D or 2D array containing for each cell the alpha value, i.e. `nb_rows*nb_cols` values.

Returns `float32` mask array with same height and width as `segmap.shape`. Values are in `[0.0, 1.0]`.

Return type ndarray

class imgaug.augmenters.blend.**SegMapClassIdsMaskGen**(*class_ids*, *nb_sample_classes=None*)

Bases: *imgaug.augmenters.blend.IBatchwiseMaskGenerator*

Generator that produces masks highlighting segmentation map classes.

This class produces for each segmentation map in a batch a mask in which the locations of a set of provided classes are highlighted (i.e. `1.0`). The classes may be provided as a fixed list of class ids or a stochastic parameter from which class ids will be sampled.

The produced masks are initially of the same height and width as the segmentation map arrays and later upscaled to the image height and width.

Note: Segmentation maps can have multiple channels. If that is the case then for each position `(x, y)` it is sufficient that any class id in any channel matches one of the desired class ids.

Note: This class will produce an `AssertionError` if there are no segmentation maps in a batch.

Added in 0.4.0.

Parameters

- **class_ids** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter*) – Segmentation map classes to mark in the produced mask.

  If *nb_sample_classes* is `None` then this is expected to be either a single `int` (always mark this one class id) or a `list` of `int` s (always mark these class ids).

  If *nb_sample_classes* is set, then this parameter will be treated as a stochastic parameter with the following valid types:

  - If `int`: Exactly that class id will be used for all segmentation maps.
  - If `tuple (a, b)`: `N` random values will be uniformly sampled per segmentation map from the discrete interval `[a..b]` and used as the class ids.
  - If `list`: `N` random values will be picked per segmentation map from that list and used as the class ids.
  - If `StochasticParameter`: That parameter will be queried once per batch for `(sum(N),)` values.

  `N` denotes the number of classes to sample per segmentation map (derived from *nb_sample_classes*) and `sum(N)` denotes the sum of `N` s over all segmentation maps.

- **nb_sample_classes** (*None or tuple of int or list of int or imgaug.parameters.StochasticParameter, optional*) – Number of class ids to sample (with replacement) per segmentation map. As sampling happens with replacement, fewer *unique* class ids may be sampled.
  - If `None`: *class_ids* is expected to be a fixed value of class ids to be used for all segmentation maps.
  - If `int`: Exactly that many class ids will be sampled for all segmentation maps.

- If `tuple (a, b)`: A random value will be uniformly sampled per segmentation map from the discrete interval `[a..b]`.
- If `list` or `int`: A random value will be picked per segmentation map from that list.
- If `StochasticParameter`: That parameter will be queried once per batch for `(B,)` values, where `B` is the number of segmentation maps.

Methods

draw_masks(self, batch[, random_state])	See *draw_masks()*.
generate_mask(segmap, class_ids)	Generate a mask of where the segmentation map has the given classes.

draw_masks (*self*, *batch*, *random_state=None*)
See *draw_masks()*.

Added in 0.4.0.

classmethod generate_mask (*segmap*, *class_ids*)
Generate a mask of where the segmentation map has the given classes.

Added in 0.4.0.

Parameters

- **segmap** (*imgaug.augmentables.segmap.SegmentationMapsOnImage*) – The segmentation map for which to generate the mask.
- **class_ids** (*iterable of int*) – IDs of the classes to set to `1.0`. For an `(x, y)` position, it is enough that *any* channel at the given location to have one of these class ids to be marked as `1.0`.

Returns `float32` mask array with same height and width as `segmap.shape`. Values are in `[0.0, 1.0]`.

Return type ndarray

imgaug.augmenters.blend.**SimplexNoiseAlpha** (*first=None*, *second=None*, *per_channel=False*, *size_px_max=(2, 16)*, *upscale_method=None*, *iterations=(1, 3)*, *aggregation_method='max'*, *sigmoid=True*, *sigmoid_thresh=None*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Deprecated. Use `BlendAlphaSimplexNoise` instead. SimplexNoiseAlpha is deprecated. Use BlendAlphaSimplexNoise instead. The order of parameters is the same. Parameter 'first' was renamed to 'foreground'. Parameter 'second' was renamed to 'background'.

See *BlendAlphaSimplexNoise*.

Deprecated since 0.4.0.

class imgaug.augmenters.blend.**SomeColorsMaskGen** (*nb_bins=(5, 15)*, *smoothness=(0.1, 0.3)*, *alpha=[0.0, 1.0]*, *rotation_deg=(0, 360)*, *from_colorspace='RGB'*)

Bases: *imgaug.augmenters.blend.IBatchwiseMaskGenerator*

Generator that produces masks based on some similar colors in images.

This class receives batches for which to generate masks, iterates over the batch rows (i.e. images) and generates one mask per row. The mask contains high alpha values for some colors, while other colors get low mask values. Which colors are chosen is random. How wide or narrow the selection is (e.g. very specific blue tone or all blue-ish colors) is determined by the hyperparameters.

The color selection method performs roughly the following steps:

1. Split the full color range of the hue in `HSV` into `nb_bins` bins (i.e. `256/nb_bins` different possible hue tones).
2. Shift the bins by `rotation_deg` degrees. (This way, the `0th` bin does not always start at exactly `0deg` of hue.)
3. Sample `alpha` values for each bin.
4. Repeat the `nb_bins` bins until there are `256` bins.
5. Smoothen the alpha values of neighbouring bins using a gaussian kernel. The kernel's `sigma` is derived from `smoothness`.
6. Associate all hue values in the image with the corresponding bin's alpha value. This results in the alpha mask.

Note: This mask generator will produce an `AssertionError` for batches that contain no images.

Added in 0.4.0.

Supported dtypes:

See *change_colorspaces_()*.

Parameters

- **nb_bins** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter, optional*) – Number of bins. For `B` bins, each bin denotes roughly `360/B` degrees of colors in the hue channel. Lower values lead to a coarser selection of colors. Expected value range is `[2, 256]`.
  - If `int`: Exactly that value will be used for all images.
  - If `tuple (a, b)`: A random value will be uniformly sampled per image from the discrete interval `[a..b]`.
  - If `list`: A random value will be picked per image from that list.
  - If `StochasticParameter`: That parameter will be queried once per batch for `(N,)` values – one per image.
- **smoothness** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Strength of the 1D gaussian kernel applied to the sampled binwise alpha values. Larger values will lead to more similar grayscaling of neighbouring colors. Expected value range is `[0.0, 1.0]`.
  - If `number`: Exactly that value will be used for all images.
  - If `tuple (a, b)`: A random value will be uniformly sampled per image from the interval `[a, b]`.
  - If `list`: A random value will be picked per image from that list.
  - If `StochasticParameter`: That parameter will be queried once per batch for `(N,)` values – one per image.

- **alpha** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Parameter to sample binwise alpha blending factors from. Expected value range is `[0.0, 1.0]`. Note that the alpha values will be smoothed between neighbouring bins. Hence, it is usually a good idea to set this so that the probability distribution peaks are around `0.0` and `1.0`, e.g. via a list `[0.0, 1.0]` or a `Beta` distribution. It is not recommended to set this to a deterministic value, otherwise all bins and hence all pixels in the generated mask will have the same value.
  - If `number`: Exactly that value will be used for all bins.
  - If `tuple (a, b)`: A random value will be uniformly sampled per bin from the interval `[a, b]`.
  - If `list`: A random value will be picked per bin from that list.
  - If `StochasticParameter`: That parameter will be queried once per batch for `(N*B, )` values – one per image and bin.
- **rotation_deg** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Rotational shift of each bin as a fraction of `360` degrees. E.g. `0.0` will not shift any bins, while a value of `0.5` will shift by around `180` degrees. This shift is mainly used so that the `0th` bin does not always start at `0deg`. Expected value range is `[-360, 360]`. This parameter can usually be kept at the default value.
  - If `number`: Exactly that value will be used for all images.
  - If `tuple (a, b)`: A random value will be uniformly sampled per image from the interval `[a, b]`.
  - If `list`: A random value will be picked per image from that list.
  - If `StochasticParameter`: That parameter will be queried once per batch for `(N,)` values – one per image.
- **from_colorspace** (*str, optional*) – The source colorspace (of the input images). See *change_colorspace_()*.

Methods

draw_masks(self, batch[, random_state])	See *draw_masks()*.
generate_mask(image, binwise_alphas, sigma, ...)	Generate a colorwise alpha mask for a single image.

draw_masks (*self*, *batch*, *random_state=None*)
See *draw_masks()*.

classmethod generate_mask (*image*, *binwise_alphas*, *sigma*, *rotation_bins*, *from_colorspace*)
Generate a colorwise alpha mask for a single image.

Added in 0.4.0.

Parameters

- **image** (*ndarray*) – Image for which to generate the mask. Must have shape `(H,W,3)` in colorspace *from_colorspace*.
- **binwise_alphas** (*ndarray*) – Alpha values of shape `(B,)` with `B` in `[1, 256]` and values in interval `[0.0, 1.0]`. Will be upscaled to 256 bins by simple repetition. Each bin represents `1/256` th of the hue.

- **sigma** (*float*) – Sigma of the 1D gaussian kernel applied to the upscaled binwise alpha value array.
- **rotation_bins** (*int*) – By how much to rotate the 256 bin alpha array. The rotation is given in number of bins.
- **from_colorspace** (*str*) – Colorspace of the input image. One of `imgaug.augmenters.color.CSPACE_*`.

Returns `float32` mask array of shape `(H, W)` with values in `[0.0, 1.0]`

Return type ndarray

class imgaug.augmenters.blend.**StochasticParameterMaskGen**(*parameter*, *per_channel*)

Bases: *imgaug.augmenters.blend.IBatchwiseMaskGenerator*

Mask generator that queries stochastic parameters for mask values.

This class receives batches for which to generate masks, iterates over the batch rows (i.e. images) and generates one mask per row. For a row with shape `(H, W, C)` (= image shape), it generates either a `(H, W)` mask (if `per_channel` is false-like) or a `(H, W, C)` mask (if `per_channel` is true-like). The `per_channel` is sampled per batch for each row/image.

Added in 0.4.0.

Parameters

- **parameter** (*imgaug.parameters.StochasticParameter*) – Stochastic parameter to draw mask samples from. Expected to return values in interval `[0.0, 1.0]` (not all stochastic parameters do that) and must be able to handle sampling shapes `(H, W)` and `(H, W, C)` (all stochastic parameters should do that).
- **per_channel** (*bool or float or imgaug.parameters.StochasticParameter, optional*) – Whether to use the same mask for all channels (`False`) or to sample a new mask for each channel (`True`). If this value is a float `p`, then for `p` percent of all rows (i.e. images) *per_channel* will be treated as `True`, otherwise as `False`.

Methods

draw_masks(self, batch[, random_state])	See *draw_masks()*.

draw_masks(*self*, *batch*, *random_state=None*)

See *draw_masks()*.

class imgaug.augmenters.blend.**VerticalLinearGradientMaskGen**(*min_value=(0.0, 0.2)*, *max_value=(0.8, 1.0)*, *start_at=(0.0, 0.2)*, *end_at=(0.8, 1.0)*)

Bases: `imgaug.augmenters.blend._LinearGradientMaskGen`

Generator that produces vertical linear gradient masks.

See *HorizontalLinearGradientMaskGen* for details.

Added in 0.4.0.

Parameters

- **min_value** (*number or tuple of number or list of number or im-*

gaug.parameters.StochasticParameter, optional*) – Minimum value that the mask will have up to the start point of the linear gradient. Note that *min_value* is allowed to be larger than *max_value*, in which case the gradient will start at the (higher) *min_value* and decrease towards the (lower) *max_value*.

  - If `number`: Exactly that value will be used for all images.
  - If `tuple (a, b)`: A random value will be uniformly sampled per image from the interval `[a, b]`.
  - If `list`: A random value will be picked per image from that list.
  - If `StochasticParameter`: That parameter will be queried once per batch for `(N,)` values – one per image.

- **max_value** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Maximum value that the mask will have at the end of the linear gradient.

  Datatypes are analogous to *min_value*.

- **start_at** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Position on the y-axis where the linear gradient starts, given as a fraction of the axis size. Interval is `[0.0, 1.0]`, where `0.0` is at the top of the image. If `end_at < start_at` the gradient will be inverted.

  Datatypes are analogous to *min_value*.

- **end_at** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Position on the x-axis where the linear gradient ends, given as a fraction of the axis size. Interval is `[0.0, 1.0]`, where `1.0` is at the bottom of the image.

  Datatypes are analogous to *min_value*.

Methods

`draw_masks`(self, batch[, random_state])	See *draw_masks()*.
generate_mask(shape, min_value, max_value, ...)	Generate a linear horizontal gradient mask.

classmethod generate_mask(*shape*, *min_value*, *max_value*, *start_at*, *end_at*)

Generate a linear horizontal gradient mask.

Added in 0.4.0.

Parameters

- **shape** (*tuple of int*) – Shape of the image. The mask will have the same height and width.
- **min_value** (*number*) – Minimum value of the gradient in interval `[0.0, 1.0]`.
- **max_value** (*number*) – Maximum value of the gradient in interval `[0.0, 1.0]`.
- **start_at** (*number*) – Position on the x-axis where the linear gradient starts, given as a fraction of the axis size. Interval is `[0.0, 1.0]`.
- **end_at** (*number*) – Position on the x-axis where the linear gradient ends, given as a fraction of the axis size. Interval is `[0.0, 1.0]`.

Returns `float32` mask array with same height and width as the image. Values are in `[0.0, 1.0]`.

Return type ndarray

`imgaug.augmenters.blend.`**`blend_alpha`**(*image_fg*, *image_bg*, *alpha*, *eps=0.01*)

Blend two images using an alpha blending.

In alpha blending, the two images are naively mixed using a multiplier. Let `A` be the foreground image and `B` the background image and `a` is the alpha value. Each pixel intensity is then computed as `a * A_ij + (1-a) * B_ij`.

Supported dtypes:

- `uint8`: yes; fully tested
- `uint16`: yes; fully tested
- `uint32`: yes; fully tested
- `uint64`: yes; fully tested (1)
- `int8`: yes; fully tested
- `int16`: yes; fully tested
- `int32`: yes; fully tested
- `int64`: yes; fully tested (1)
- `float16`: yes; fully tested
- `float32`: yes; fully tested
- `float64`: yes; fully tested (1)
- `float128`: no (2)
- `bool`: yes; fully tested (2)

- (1) Tests show that these dtypes work, but a conversion to `float128` happens, which only has 96 bits of size instead of true 128 bits and hence not twice as much resolution. It is possible that these dtypes result in inaccuracies, though the tests did not indicate that.
- (2) Not available due to the input dtype having to be increased to an equivalent float dtype with two times the input resolution.
- (3) Mapped internally to `float16`.

Parameters

- **image_fg** (*(H,W,[C]) ndarray*) – Foreground image. Shape and dtype kind must match the one of the background image.
- **image_bg** (*(H,W,[C]) ndarray*) – Background image. Shape and dtype kind must match the one of the foreground image.
- **alpha** (*number or iterable of number or ndarray*) – The blending factor, between `0.0` and `1.0`. Can be interpreted as the opacity of the foreground image. Values around `1.0` result in only the foreground image being visible. Values around `0.0` result in only the background image being visible. Multiple alphas may be provided. In these cases, there must be exactly one alpha per channel in the foreground/background image. Alternatively, for `(H,W,C)` images, either one `(H,W)` array or an `(H,W,C)` array of alphas may be provided, denoting the elementwise alpha value.

- **eps** (*number, optional*) – Controls when an alpha is to be interpreted as exactly `1.0` or exactly `0.0`, resulting in only the foreground/background being visible and skipping the actual computation.

Returns **image_blend** – Blend of foreground and background image.

Return type (H,W,C) ndarray

13.21 imgaug.augmenters.blur

Augmenters that blur images.

List of augmenters:

- *GaussianBlur*
- *AverageBlur*
- *MedianBlur*
- *BilateralBlur*
- *MotionBlur*
- *MeanShiftBlur*

class imgaug.augmenters.blur.**AverageBlur**(*k=(1, 7)*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Blur an image by computing simple means over neighbourhoods.

The padding behaviour around the image borders is cv2's `BORDER_REFLECT_101`.

Supported dtypes:

- `uint8`: yes; fully tested
- `uint16`: yes; tested
- `uint32`: no (1)
- `uint64`: no (2)
- `int8`: yes; tested (3)
- `int16`: yes; tested
- `int32`: no (4)
- `int64`: no (5)
- `float16`: yes; tested (6)
- `float32`: yes; tested
- `float64`: yes; tested
- `float128`: no
- `bool`: yes; tested (7)

- (1) rejected by `cv2.blur()`
- (2) loss of resolution in `cv2.blur()` (result is `int32`)

- (3) `int8` is mapped internally to `int16`, `int8` itself leads to cv2 error "Unsupported combination of source format (=1), and buffer format (=4) in function 'getRowSumFilter'" in `cv2`
- (4) results too inaccurate
- (5) loss of resolution in `cv2.blur()` (result is `int32`)
- (6) `float16` is mapped internally to `float32`
- (7) `bool` is mapped internally to `float32`

Parameters

- **k** (*int or tuple of int or tuple of tuple of int or imgaug.parameters.StochasticParameter or tuple of StochasticParameter, optional*) –

  Kernel size to use.

  - If a single `int`, then that value will be used for the height and width of the kernel.
  - If a tuple of two `int` s `(a, b)`, then the kernel size will be sampled from the interval `[a..b]`.
  - If a tuple of two tuples of `int` s `((a, b), (c, d))`, then per image a random kernel height will be sampled from the interval `[a..b]` and a random kernel width will be sampled from the interval `[c..d]`.
  - If a `StochasticParameter`, then `N` samples will be drawn from that parameter per `N` input images, each representing the kernel size for the n-th image.
  - If a tuple `(a, b)`, where either `a` or `b` is a tuple, then `a` and `b` will be treated according to the rules above. This leads to different values for height and width of the kernel.

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.AverageBlur(k=5)
```

Blur all images using a kernel size of `5x5`.

```
>>> aug = iaa.AverageBlur(k=(2, 5))
```

Blur images using a varying kernel size, which is sampled (per image) uniformly from the interval `[2..5]`.

```
>>> aug = iaa.AverageBlur(k=((5, 7), (1, 3)))
```

Blur images using a varying kernel size, which's height is sampled (per image) uniformly from the interval `[5..7]` and which's width is sampled (per image) uniformly from `[1..3]`.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`localize_random_state_`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`pool`(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
`remove_augmenters`(self, func[, copy, ...])	Remove this augmenter or children that match a condition.

Continued on next page

Table 97 – continued from previous page

remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.blur.**BilateralBlur**(*d=(1, 9)*, *sigma_color=(10, 250)*, *sigma_space=(10, 250)*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Blur/Denoise an image using a bilateral filter.

Bilateral filters blur homogenous and textured areas, while trying to preserve edges.

See http://docs.opencv.org/2.4/modules/imgproc/doc/filtering.html#bilateralfilter for more information regarding the parameters.

Supported dtypes:

- uint8: yes; not tested
- uint16: ?
- uint32: ?
- uint64: ?
- int8: ?
- int16: ?
- int32: ?
- int64: ?
- float16: ?
- float32: ?
- float64: ?
- float128: ?
- bool: ?

Parameters

- **d** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter, optional*) – Diameter of each pixel neighborhood with value range [1 .. inf). High values for *d* lead to significantly worse performance. Values equal or less than 10 seem to be good. Use <5 for real-time applications.
  - If a single int, then that value will be used for the diameter.

  - If a tuple of two `int` s `(a, b)`, then the diameter will be a value sampled from the interval `[a..b]`.
  - If a list, then a random value will be sampled from that list per image.
  - If a `StochasticParameter`, then `N` samples will be drawn from that parameter per `N` input images, each representing the diameter for the n-th image. Expected to be discrete.

- **sigma_color** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Filter sigma in the color space with value range `[1, inf)`. A large value of the parameter means that farther colors within the pixel neighborhood (see *sigma_space*) will be mixed together, resulting in larger areas of semi-equal color.
  - If a single `int`, then that value will be used for the diameter.
  - If a tuple of two `int` s `(a, b)`, then the diameter will be a value sampled from the interval `[a, b]`.
  - If a list, then a random value will be sampled from that list per image.
  - If a `StochasticParameter`, then `N` samples will be drawn from that parameter per `N` input images, each representing the diameter for the n-th image. Expected to be discrete.
- **sigma_space** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Filter sigma in the coordinate space with value range `[1, inf)`. A large value of the parameter means that farther pixels will influence each other as long as their colors are close enough (see *sigma_color*).
  - If a single `int`, then that value will be used for the diameter.
  - If a tuple of two `int` s `(a, b)`, then the diameter will be a value sampled from the interval `[a, b]`.
  - If a list, then a random value will be sampled from that list per image.
  - If a `StochasticParameter`, then `N` samples will be drawn from that parameter per `N` input images, each representing the diameter for the n-th image. Expected to be discrete.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.BilateralBlur(
>>>     d=(3, 10), sigma_color=(10, 250), sigma_space=(10, 250))
```

Blur all images using a bilateral filter with a *max distance* sampled uniformly from the interval `[3, 10]` and wide ranges for *sigma_color* and *sigma_space*.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`localize_random_state_`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`pool`(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
`remove_augmenters`(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
`remove_augmenters_`(self, func[, parents])	Remove in-place children of this augmenter that match a condition.

Continued on next page

Table 98 – continued from previous page

remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.blur.**GaussianBlur**(*sigma=(0.0, 3.0)*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Augmenter to blur images using gaussian kernels.

Supported dtypes:

See ~imgaug.augmenters.blur.blur_gaussian_(backend="auto").

Parameters

- **sigma** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Standard deviation of the gaussian kernel. Values in the range 0.0 (no blur) to 3.0 (strong blur) are common.
  - If a single float, that value will always be used as the standard deviation.
  - If a tuple (a, b), then a random value from the interval [a, b] will be picked per image.
  - If a list, then a random value will be sampled per image from that list.
  - If a StochasticParameter, then N samples will be drawn from that parameter per N input images.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See __init__().
- **name** (*None or str, optional*) – See __init__().
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method to_deterministic() for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.GaussianBlur(sigma=1.5)
```

Blur all images using a gaussian kernel with a standard deviation of `1.5`.

```
>>> aug = iaa.GaussianBlur(sigma=(0.0, 3.0))
```

Blur images using a gaussian kernel with a random standard deviation sampled uniformly (per image) from the interval `[0.0, 3.0]`.

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.
`get_parameters`(self)	See *`get_parameters()`*.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`localize_random_state_`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`pool`(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
`remove_augmenters`(self, func[, copy, ...])	Remove this augmenter or children that match a condition.

Continued on next page

Table 99 – continued from previous page

remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.blur.**MeanShiftBlur**(*spatial_radius=(5.0, 40.0)*, *color_radius=(5.0, 40.0)*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Apply a pyramidic mean shift filter to each image.

See also *blur_mean_shift_()* for details.

This augmenter expects input images of shape `(H, W)` or `(H, W, 1)` or `(H, W, 3)`.

Note: This augmenter is quite slow.

Added in 0.4.0.

Supported dtypes:

See *blur_mean_shift_()*.

Parameters

- **spatial_radius** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) –

  Spatial radius for pixels that are assumed to be similar.

  – If `number`: Exactly that value will be used for all images.

  – If `tuple (a, b)`: A random value will be uniformly sampled per image from the interval `[a, b)`.

  – If `list`: A random value will be sampled from that `list` per image.

  – If `StochasticParameter`: The parameter will be queried once per batch for `(N,)` values with `N` denoting the number of images.

- **color_radius** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) –

  Color radius for pixels that are assumed to be similar.

  – If `number`: Exactly that value will be used for all images.

  – If `tuple (a, b)`: A random value will be uniformly sampled per image from the interval `[a, b)`.

  - If `list`: A random value will be sampled from that `list` per image.
  - If `StochasticParameter`: The parameter will be queried once per batch for `(N,)` values with `N` denoting the number of images.

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.MeanShiftBlur()
```

Create a mean shift blur augmenter.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.

Continued on next page

Table 100 – continued from previous page

copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.blur.**MedianBlur**(*k=(1, 7)*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Blur an image by computing median values over neighbourhoods.

Median blurring can be used to remove small dirt from images. At larger kernel sizes, its effects have some similarity with Superpixels.

Supported dtypes:

- uint8: yes; fully tested
- uint16: ?
- uint32: ?
- uint64: ?
- int8: ?
- int16: ?

- int32: ?
- int64: ?
- float16: ?
- float32: ?
- float64: ?
- float128: ?
- bool: ?

Parameters

- **k** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter, optional*) –

  Kernel size.

  - If a single int, then that value will be used for the height and width of the kernel. Must be an odd value.
  - If a tuple of two ints (a, b), then the kernel size will be an odd value sampled from the interval [a..b]. a and b must both be odd values.
  - If a list, then a random value will be sampled from that list per image.
  - If a StochasticParameter, then N samples will be drawn from that parameter per N input images, each representing the kernel size for the nth image. Expected to be discrete. If a sampled value is not odd, then that value will be increased by 1.

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See __init__().
- **name** (*None or str, optional*) – See __init__().
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method to_deterministic() for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.MedianBlur(k=5)
```

Blur all images using a kernel size of 5x5.

```
>>> aug = iaa.MedianBlur(k=(3, 7))
```

Blur images using varying kernel sizes, which are sampled uniformly from the interval [3..7]. Only odd values will be sampled, i.e. 3 or 5 or 7.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`localize_random_state_`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`pool`(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
`remove_augmenters`(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
`remove_augmenters_`(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
`remove_augmenters_inplace`(self, func[, parents])	**Deprecated**.
`reseed`(self[, random_state, deterministic_too])	**Deprecated**.
`seed_`(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.

Continued on next page

Table 101 – continued from previous page

show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.blur.**MotionBlur**(*k=(3, 7)*, *angle=(0, 360)*, *direction=(-1.0, 1.0)*, *order=1*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.convolutional.Convolve*

Blur images in a way that fakes camera or object movements.

Supported dtypes:

See *Convolve*.

Parameters

- **k** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter, optional*) –

  Kernel size to use.

  – If a single int, then that value will be used for the height and width of the kernel.

  – If a tuple of two int s (a, b), then the kernel size will be sampled from the interval [a..b].

  – If a list, then a random value will be sampled from that list per image.

  – If a StochasticParameter, then N samples will be drawn from that parameter per N input images, each representing the kernel size for the n-th image.

- **angle** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Angle of the motion blur in degrees (clockwise, relative to top center direction).

  – If a number, exactly that value will be used.

  – If a tuple (a, b), a random value from the interval [a, b] will be uniformly sampled per image.

  – If a list, then a random value will be sampled from that list per image.

  – If a StochasticParameter, a value will be sampled from the parameter per image.

- **direction** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Forward/backward direction of the motion blur. Lower values towards -1.0 will point the motion blur towards the back (with angle provided via *angle*). Higher values towards 1.0 will point the motion blur forward. A value of 0.0 leads to a uniformly (but still angled) motion blur.

  – If a number, exactly that value will be used.

  – If a tuple (a, b), a random value from the interval [a, b] will be uniformly sampled per image.

  – If a list, then a random value will be sampled from that list per image.

  – If a StochasticParameter, a value will be sampled from the parameter per image.

- **order** (*int or iterable of int or imgaug.ALL or imgaug.parameters.StochasticParameter, optional*) – Interpolation order to use when rotating the kernel according to *angle*. See `__init__()`. Recommended to be `0` or `1`, with `0` being faster, but less continuous/smooth as *angle* is changed, particularly around multiple of `45` degrees.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.MotionBlur(k=15)
```

Apply motion blur with a kernel size of `15x15` pixels to images.

```
>>> aug = iaa.MotionBlur(k=15, angle=[-45, 45])
```

Apply motion blur with a kernel size of `15x15` pixels and a blur angle of either `-45` or `45` degrees (randomly picked per image).

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.

Continued on next page

Table 102 – continued from previous page

augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

imgaug.augmenters.blur.**blur_gaussian_**(*image*, *sigma*, *ksize=None*, *backend='auto'*, *eps=0.001*)

Blur an image using gaussian blurring in-place.

This operation *may* change the input image in-place.

Supported dtypes:

if (backend="auto"):

- uint8: yes; fully tested (1)
- uint16: yes; tested (1)
- uint32: yes; tested (2)
- uint64: yes; tested (2)
- int8: yes; tested (1)
- int16: yes; tested (1)

- int32: yes; tested (1)
- int64: yes; tested (2)
- float16: yes; tested (1)
- float32: yes; tested (1)
- float64: yes; tested (1)
- float128: no
- bool: yes; tested (1)

- (1) Handled by cv2. See backend="cv2".
- (2) Handled by scipy. See backend="scipy".

if (backend="cv2"):

- uint8: yes; fully tested
- uint16: yes; tested
- uint32: no (2)
- uint64: no (3)
- int8: yes; tested (4)
- int16: yes; tested
- int32: yes; tested (5)
- int64: no (6)
- float16: yes; tested (7)
- float32: yes; tested
- float64: yes; tested
- float128: no (8)
- bool: yes; tested (1)

- (1) Mapped internally to float32. Otherwise causes TypeError: src data type = 0 is not supported.
- (2) Causes TypeError: src data type = 6 is not supported.
- (3) Causes cv2.error: OpenCV(3.4.5) (...)/filter.cpp:2957: error: (-213:The function/feature is not implemented) Unsupported combination of source format (=4), and buffer format (=5) in function 'getLinearRowFilter'.
- (4) Mapped internally to int16. Otherwise causes cv2.error: OpenCV(3.4.5) (...)/filter.cpp:2957: error: (-213:The function/feature is not implemented) Unsupported combination of source format (=1), and buffer format (=5) in function 'getLinearRowFilter'.
- (5) Mapped internally to float64. Otherwise causes cv2.error: OpenCV(3.4.5) (...)/filter.cpp:2957: error: (-213:The function/feature is not implemented) Unsupported combination of

source format (=4), and buffer format (=5) in function 'getLinearRowFilter'.

- (6) Causes cv2.error: OpenCV(3.4.5) (...)/filter.cpp:2957: error: (-213:The function/feature is not implemented) Unsupported combination of source format (=4), and buffer format (=5) in function 'getLinearRowFilter'.
- (7) Mapped internally to float32. Otherwise causes TypeError: src data type = 23 is not supported.
- (8) Causes TypeError: src data type = 13 is not supported.

if (backend="scipy"):

- uint8: yes; fully tested
- uint16: yes; tested
- uint32: yes; tested
- uint64: yes; tested
- int8: yes; tested
- int16: yes; tested
- int32: yes; tested
- int64: yes; tested
- float16: yes; tested (1)
- float32: yes; tested
- float64: yes; tested
- float128: no (2)
- bool: yes; tested (3)

- (1) Mapped internally to float32. Otherwise causes RuntimeError: array type dtype('float16') not supported.
- (2) Causes RuntimeError: array type dtype('float128') not supported.
- (3) Mapped internally to float32. Otherwise too inaccurate.

Parameters

- **image** (*numpy.ndarray*) – The image to blur. Expected to be of shape (H, W) or (H, W, C).
- **sigma** (*number*) – Standard deviation of the gaussian blur. Larger numbers result in more large-scale blurring, which is overall slower than small-scale blurring.
- **ksize** (*None or int, optional*) – Size in height/width of the gaussian kernel. This argument is only understood by the cv2 backend. If it is set to None, an appropriate value for *ksize* will automatically be derived from *sigma*. The value is chosen tighter for larger sigmas to avoid as much as possible very large kernel sizes and therey improve performance.
- **backend** (*{'auto', 'cv2', 'scipy'}, optional*) – Backend library to use. If auto, then the likely best library will be automatically picked per image. That is usually equivalent to cv2 (OpenCV) and it will fall back to scipy for datatypes not supported by OpenCV.

- **eps** (*number, optional*) – A threshold used to decide whether *sigma* can be considered zero.

Returns The blurred image. Same shape and dtype as the input. (Input image *might* have been altered in-place.)

Return type numpy.ndarray

imgaug.augmenters.blur.**blur_mean_shift_**(*image*, *spatial_window_radius*, *color_window_radius*)

Apply a pyramidic mean shift filter to the input image in-place.

This produces an output image that has similarity with one modified by a bilateral filter. That is different from mean shift *segmentation*, which averages the colors in segments found by mean shift clustering.

This function is a thin wrapper around `cv2.pyrMeanShiftFiltering`.

Note: This function does *not* change the image's colorspace to `RGB` before applying the mean shift filter. A non-`RGB` colorspace will hence influence the results.

Note: This function is quite slow.

Added in 0.4.0.

Supported dtypes:

- `uint8`: yes; fully tested
- `uint16`: no (1)
- `uint32`: no (1)
- `uint64`: no (1)
- `int8`: no (1)
- `int16`: no (1)
- `int32`: no (1)
- `int64`: no (1)
- `float16`: no (1)
- `float32`: no (1)
- `float64`: no (1)
- `float128`: no (1)
- `bool`: no (1)
- (1) Not supported by `cv2.pyrMeanShiftFiltering`.

Parameters

- **image** (*ndarray*) – `(H,W)` or `(H,W,1)` or `(H,W,3)` image to blur. Images with no or one channel will be temporarily tiled to have three channels.
- **spatial_window_radius** (*number*) – Spatial radius for pixels that are assumed to be similar.
- **color_window_radius** (*number*) – Color radius for pixels that are assumed to be similar.

Returns Blurred input image. Same shape and dtype as the input. (Input image *might* have been altered in-place.)

Return type ndarray

13.22 imgaug.augmenters.collections

Augmenters that are collections of other augmenters.

List of augmenters:

- *RandAugment*

Added in 0.4.0.

class imgaug.augmenters.collections.**RandAugment**(*n=2*, *m=(6, 12)*, *cval=128*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Sequential*

Apply RandAugment to inputs as described in the corresponding paper.

See paper:

```
Cubuk et al.

RandAugment: Practical automated data augmentation with a reduced
search space
```

Note: The paper contains essentially no hyperparameters for the individual augmentation techniques. The hyperparameters used here come mostly from the official code repository, which however seems to only contain code for CIFAR10 and SVHN, not for ImageNet. So some guesswork was involved and a few of the hyperparameters were also taken from https://github.com/ildoonet/pytorch-randaugment/blob/master/RandAugment/augmentations.py .

This implementation deviates from the code repository for all PIL enhance operations. In the repository these use a factor of `0.1 + M*1.8/M_max`, which would lead to a factor of `0.1` for the weakest `M` of `M=0`. For e.g. `Brightness` that would result in a basically black image. This definition is fine for AutoAugment (from where the code and hyperparameters are copied), which optimizes each transformation's `M` individually, but not for RandAugment, which uses a single fixed `M`. We hence redefine these hyperparameters to `1.0 + S * M * 0.9/M_max`, where `S` is randomly either `1` or `-1`.

We also note that it is not entirely clear which transformations were used in the ImageNet experiments. The paper lists some transformations in Figure 2, but names others in the text too (e.g. crops, flips, cutout). While Figure 2 lists the Identity function, this transformation seems to not appear in the repository (and in fact, the function `randaugment(N, M)` doesn't seem to exist in the repository either). So we also make a best guess here about what transformations might have been used.

Warning: This augmenter only works with image data, not e.g. bounding boxes. The used PIL-based affine transformations are not yet able to process non-image data. (This augmenter uses PIL-based affine transformations to ensure that outputs are as similar as possible to the paper's implementation.)

Added in 0.4.0.

Supported dtypes:

minimum of (*Fliplr*, *KeepSizeByResize*, *Crop*, *Sequential*, *SomeOf*, *Identity*, *Autocontrast*, *Equalize*, *Invert*, *Affine*, *Posterize*, *Solarize*, *EnhanceColor*, *EnhanceContrast*, *EnhanceBrightness*, *EnhanceSharpness*, *Cutout*, *FilterBlur*, *FilterSmooth*

)

Parameters **n** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter or None, optional*) – Parameter `N` in the paper, i.e. number of transformations to apply. The paper suggests `N=2` for ImageNet. See also parameter `n` in *SomeOf* for more details.

Note that horizontal flips (p=50%) and crops are always applied. This parameter only determines how many of the other transformations are applied per image.

m [int or tuple of int or list of int or imgaug.parameters.StochasticParameter or None, optional] Parameter `M` in the paper, i.e. magnitude/severity/strength of the applied transformations in interval `[0 .. 30]` with `M=0` being the weakest. The paper suggests for ImageNet `M=9` in case of ResNet-50 and `M=28` in case of EfficientNet-B7. This implementation uses a default value of `(6, 12)`, i.e. the value is uniformly sampled per image from the interval `[6 .. 12]`. This ensures greater diversity of transformations than using a single fixed value.

- If `int`: That value will always be used.
- If `tuple (a, b)`: A random value will be uniformly sampled per image from the discrete interval `[a .. b]`.
- If `list`: A random value will be picked from the list per image.
- If `StochasticParameter`: For `B` images in a batch, `B` values will be sampled per augmenter (provided the augmenter is dependent on the magnitude).

cval [number or tuple of number or list of number or imgaug.ALL or imgaug.parameters.StochasticParameter, optional] The constant value to use when filling in newly created pixels. See parameter *fillcolor* in *Affine* for details.

The paper's repository uses an RGB value of `125, 122, 113`. This implementation uses a single intensity value of `128`, which should work better for cases where input images don't have exactly `3` channels or come from a different dataset than used by the paper.

seed [None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional] See `__init__()`.

name [None or str, optional] See `__init__()`.

random_state [None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional] Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.

deterministic [bool, optional] Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.RandAugment(n=2, m=9)
```

Create a RandAugment augmenter similar to the suggested hyperparameters in the paper.

```
>>> aug = iaa.RandAugment(m=30)
```

Create a RandAugment augmenter with maximum magnitude/strength.

```
>>> aug = iaa.RandAugment(m=(0, 9))
```

Create a RandAugment augmenter that applies its transformations with a random magnitude between `0` (very weak) and `9` (recommended for ImageNet and ResNet-50). `m` is sampled per transformation.

```
>>> aug = iaa.RandAugment(n=(0, 3))
```

Create a RandAugment augmenter that applies `0` to `3` of its child transformations to images. Horizontal flips (p=50%) and crops are always applied.

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`add`(self, augmenter)	Add an augmenter to the list of child augmenters.
`append`(self, object, /)	Append object to the end of the list.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`clear`(self, /)	Remove all items from list.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`count`(self, value, /)	Return number of occurrences of value.
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`extend`(self, iterable, /)	Extend list by appending elements from the iterable.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.

Continued on next page

Table 103 – continued from previous page

find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	See *get_children_lists()*.
get_parameters(self)	See *get_parameters()*.
index(self, value[, start, stop])	Return first index of value.
insert(self, index, object, /)	Insert object before index.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
pop(self[, index])	Remove and return item at index (default last).
remove(self, value, /)	Remove first occurrence of value.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
reverse(self, /)	Reverse *IN PLACE*.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
sort(self, /, *[, key, reverse])	Stable sort *IN PLACE*.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See *get_parameters()*.

13.23 imgaug.augmenters.color

Augmenters that affect image colors or image colorspaces.

List of augmenters:

- *InColorspace* (deprecated)
- *WithColorspace*
- *WithBrightnessChannels*
- *MultiplyAndAddToBrightness*
- *MultiplyBrightness*
- *AddToBrightness*
- *WithHueAndSaturation*
- *MultiplyHueAndSaturation*
- *MultiplyHue*

- *MultiplySaturation*
- *RemoveSaturation*
- *AddToHueAndSaturation*
- *AddToHue*
- *AddToSaturation*
- *ChangeColorspace*
- *Grayscale*
- *ChangeColorTemperature*
- *KMeansColorQuantization*
- *UniformColorQuantization*
- *Posterize*

class imgaug.augmenters.color.**AddToBrightness**(*add=(-30, 30), to_colorspace=['YCrCb', 'HSV', 'HLS', 'Lab', 'Luv', 'YUV'], from_colorspace='RGB', seed=None, name=None, random_state='deprecated', deterministic='deprecated'*)

Bases: *imgaug.augmenters.color.MultiplyAndAddToBrightness*

Add to the brightness channels of input images.

This is a wrapper around *WithBrightnessChannels* and hence performs internally the same projection to random colorspaces.

Added in 0.4.0.

Supported dtypes:

See *MultiplyAndAddToBrightness*.

Parameters

- **add** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – See `Add`.
- **to_colorspace** (*imgaug.ALL or str or list of str or imgaug.parameters.StochasticParameter, optional*) – See *WithBrightnessChannels*.
- **from_colorspace** (*str, optional*) – See *WithBrightnessChannels*.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.AddToBrightness((-30, 30))
```

Convert each image to a colorspace with a brightness-related channel, extract that channel, add between `-30` and `30` and convert back to the original colorspace.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	See *get_children_lists()*.
`get_parameters`(self)	See *get_parameters()*.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`localize_random_state_`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`pool`(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.

Continued on next page

Table 104 – continued from previous page

remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.color.**AddToHue**(*value=(-255, 255), from_colorspace='RGB', seed=None, name=None, random_state='deprecated', deterministic='deprecated'*)

Bases: *imgaug.augmenters.color.AddToHueAndSaturation*

Add random values to the hue of images.

The augmenter first transforms images to HSV colorspace, then adds random values to the H channel and afterwards converts back to RGB.

If you want to change both the hue and the saturation, it is recommended to use `AddToHueAndSaturation` as otherwise the image will be converted twice to HSV and back to RGB.

This augmenter is a shortcut for `AddToHueAndSaturation(value_hue=...)`.

Supported dtypes:

See *AddToHueAndSaturation*.

Parameters

- **value** (*None or int or tuple of int or list of int or imgaug.parameters.StochasticParameter, optional*) – Value to add to the hue of all pixels. This is expected to be in the range `-255` to `+255` and will automatically be projected to an angular representation using `(hue/255) * (360/2)` (OpenCV's hue representation is in the range `[0, 180]` instead of `[0, 360]`).
  - If an integer, then that value will be used for all images.
  - If a tuple `(a, b)`, then a value from the discrete range `[a, b]` will be sampled per image.
  - If a list, then a random value will be sampled from that list per image.
  - If a StochasticParameter, then a value will be sampled from that parameter per image.
- **from_colorspace** (*str, optional*) – See *change_colorspace_()*.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage

will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.

- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.AddToHue((-50, 50))
```

Sample random values from the discrete uniform range `[-50..50]`, convert them to angular representation and add them to the hue, i.e. to the `H` channel in `HSV` colorspace.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.

Continued on next page

Table 105 – continued from previous page

get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.color.**AddToHueAndSaturation**(*value=None*, *value_hue=None*, *value_saturation=None*, *per_channel=False*, *from_colorspace='RGB'*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Increases or decreases hue and saturation by random values.

The augmenter first transforms images to HSV colorspace, then adds random values to the H and S channels and afterwards converts back to RGB.

This augmenter is faster than using WithHueAndSaturation in combination with Add.

TODO add float support

Supported dtypes:

See *change_colorspace_()*.

Parameters

- **value** (*None or int or tuple of int or list of int or imgaug.parameters.StochasticParameter, optional*) – Value to add to the hue *and* saturation of all pixels. It is expected to be in the range −255 to +255.
  - If this is None, *value_hue* and/or *value_saturation* may be set to values other than None.
  - If an integer, then that value will be used for all images.
  - If a tuple (a, b), then a value from the discrete range [a, b] will be sampled per image.
  - If a list, then a random value will be sampled from that list per image.
  - If a StochasticParameter, then a value will be sampled from that parameter per image.

- **value_hue** (*None or int or tuple of int or list of int or imgaug.parameters.StochasticParameter, optional*) – Value to add to the hue of all pixels. This is expected to be in the range `-255` to `+255` and will automatically be projected to an angular representation using `(hue/255) * (360/2)` (OpenCV's hue representation is in the range `[0, 180]` instead of `[0, 360]`). Only this or *value* may be set, not both.
  - If this and *value_saturation* are both `None`, *value* may be set to a non-`None` value.
  - If an integer, then that value will be used for all images.
  - If a tuple `(a, b)`, then a value from the discrete range `[a, b]` will be sampled per image.
  - If a list, then a random value will be sampled from that list per image.
  - If a StochasticParameter, then a value will be sampled from that parameter per image.
- **value_saturation** (*None or int or tuple of int or list of int or imgaug.parameters.StochasticParameter, optional*) – Value to add to the saturation of all pixels. It is expected to be in the range `-255` to `+255`. Only this or *value* may be set, not both.
  - If this and *value_hue* are both `None`, *value* may be set to a non-`None` value.
  - If an integer, then that value will be used for all images.
  - If a tuple `(a, b)`, then a value from the discrete range `[a, b]` will be sampled per image.
  - If a list, then a random value will be sampled from that list per image.
  - If a StochasticParameter, then a value will be sampled from that parameter per image.
- **per_channel** (*bool or float, optional*) – Whether to sample per image only one value from *value* and use it for both hue and saturation (`False`) or to sample independently one value for hue and one for saturation (`True`). If this value is a float `p`, then for `p` percent of all images *per_channel* will be treated as `True`, otherwise as `False`.

  This parameter has no effect is *value_hue* and/or *value_saturation* are used instead of *value*.
- **from_colorspace** (*str, optional*) – See *change_colorspace_()*.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.AddToHueAndSaturation((-50, 50), per_channel=True)
```

Add random values between `-50` and `50` to the hue and saturation (independently per channel and the same value for all pixels within that channel).

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`localize_random_state_`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`pool`(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
`remove_augmenters`(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
`remove_augmenters_`(self, func[, parents])	Remove in-place children of this augmenter that match a condition.

Continued on next page

Table 106 – continued from previous page

remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.color.**AddToSaturation**(*value=(-75, 75), from_colorspace='RGB', seed=None, name=None, random_state='deprecated', deterministic='deprecated'*)

Bases: *imgaug.augmenters.color.AddToHueAndSaturation*

Add random values to the saturation of images.

The augmenter first transforms images to HSV colorspace, then adds random values to the S channel and afterwards converts back to RGB.

If you want to change both the hue and the saturation, it is recommended to use `AddToHueAndSaturation` as otherwise the image will be converted twice to HSV and back to RGB.

This augmenter is a shortcut for `AddToHueAndSaturation(value_saturation=...)`.

Supported dtypes:

See *AddToHueAndSaturation*.

Parameters

- **value** (*None or int or tuple of int or list of int or imgaug.parameters.StochasticParameter, optional*) – Value to add to the saturation of all pixels. It is expected to be in the range `-255` to `+255`.
  - If an integer, then that value will be used for all images.
  - If a tuple `(a, b)`, then a value from the discrete range `[a, b]` will be sampled per image.
  - If a list, then a random value will be sampled from that list per image.
  - If a StochasticParameter, then a value will be sampled from that parameter per image.
- **from_colorspace** (*str, optional*) – See *change_colorspace_()*.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.

- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.AddToSaturation((-50, 50))
```

Sample random values from the discrete uniform range `[-50..50]`, and add them to the saturation, i.e. to the `S` channel in `HSV` colorspace.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.
`get_parameters`(self)	See *get_parameters()*.

Continued on next page

Table 107 – continued from previous page

localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.color.**ChangeColorTemperature**(*kelvin=(1000, 11000)*, *from_colorspace='RGB'*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Change the temperature to a provided Kelvin value.

Low Kelvin values around `1000` to `4000` will result in red, yellow or orange images. Kelvin values around `10000` to `40000` will result in progressively darker blue tones.

Color temperatures taken from http://www.vendian.org/mncharity/dir3/blackbody/UnstableURLs/bbr_color.html

Basic method to change color temperatures taken from https://stackoverflow.com/a/11888449

Added in 0.4.0.

Supported dtypes:

See *change_color_temperatures_()*.

Parameters kelvin (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Temperature in Kelvin. The temperatures of images will be modified to this value. Must be in the interval `[1000, 40000]`.

- If a number, exactly that value will always be used.
- If a `tuple (a, b)`, then a value from the interval `[a, b]` will be sampled per image.
- If a `list`, then a random value will be sampled from that

`list` per image. * If a `StochasticParameter`, then a value will be sampled per

image from that parameter.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.ChangeColorTemperature((1100, 10000))
```

Create an augmenter that changes the color temperature of images to a random value between 1100 and 10000 Kelvin.

Methods

__call__(self, *args, **kwargs)	Alias for *augment()*.
augment(self[, return_batch, hooks])	Augment a batch.
augment_batch(self, batch[, hooks])	**Deprecated**.
augment_batch_(self, batch[, parents, hooks])	Augment a single batch in-place.
augment_batches(self, batches[, hooks, ...])	Augment multiple batches.
augment_bounding_boxes(self, ...[, parents, ...])	Augment a batch of bounding boxes.
augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.

Continued on next page

Table 108 – continued from previous page

remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.color.**ChangeColorspace**(*to_colorspace*, *from_colorspace='RGB'*, *alpha=1.0*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Augmenter to change the colorspace of images.

Note: This augmenter is not tested. Some colorspaces might work, others might not.

..note:

```
This augmenter tries to project the colorspace value range on
0-255. It outputs dtype=uint8 images.
```

Supported dtypes:

See *change_colorspace_()*.

Parameters

- **to_colorspace** (*str or list of str or imgaug.parameters.StochasticParameter*) – The target colorspace. Allowed strings are: RGB, BGR, GRAY, CIE, YCrCb, HSV, HLS, Lab, Luv. These are also accessible via imgaug.augmenters.color.CSPACE_<NAME>, e.g. imgaug.augmenters.CSPACE_YCrCb.
  - If a string, it must be among the allowed colorspaces.
  - If a list, it is expected to be a list of strings, each one being an allowed colorspace. A random element from the list will be chosen per image.
  - If a StochasticParameter, it is expected to return string. A new sample will be drawn per image.
- **from_colorspace** (*str, optional*) – The source colorspace (of the input images). See *to_colorspace*. Only a single string is allowed.
- **alpha** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – The alpha value of the new colorspace

when overlayed over the old one. A value close to 1.0 means that mostly the new colorspace is visible. A value close to 0.0 means, that mostly the old image is visible.

- If an int or float, exactly that value will be used.
- If a tuple `(a, b)`, a random value from the range `a <= x <= b` will be sampled per image.
- If a list, then a random value will be sampled from that list per image.
- If a StochasticParameter, a value will be sampled from the parameter per image.

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.

Continued on next page

Table 109 – continued from previous page

draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See get_parameters().
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

BGR = 'BGR'

CIE = 'CIE'

COLORSPACES = {'BGR', 'CIE', 'GRAY', 'HLS', 'HSV', 'Lab', 'Luv', 'RGB', 'YCrCb'}

CV_VARS = {'BGR2CIE': <MagicMock id='140201070283520'>, 'BGR2GRAY': <MagicMock id='140

GRAY = 'GRAY'

HLS = 'HLS'

HSV = 'HSV'

Lab = 'Lab'

Luv = 'Luv'

RGB = 'RGB'

YCrCb = 'YCrCb'

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.color.**Grayscale**(*alpha=1*, *from_colorspace='RGB'*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.color.ChangeColorspace*

Augmenter to convert images to their grayscale versions.

Note: Number of output channels is still `3`, i.e. this augmenter just "removes" color.

TODO check dtype support

Supported dtypes:

See *change_colorspace_()*.

Parameters

- **alpha** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – The alpha value of the grayscale image when overlayed over the old image. A value close to 1.0 means, that mostly the new grayscale image is visible. A value close to 0.0 means, that mostly the old image is visible.
  - If a number, exactly that value will always be used.
  - If a tuple `(a, b)`, a random value from the range `a <= x <= b` will be sampled per image.
  - If a list, then a random value will be sampled from that list per image.
  - If a StochasticParameter, a value will be sampled from the parameter per image.
- **from_colorspace** (*str, optional*) – The source colorspace (of the input images). See *change_colorspace_()*.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.Grayscale(alpha=1.0)
```

Creates an augmenter that turns images to their grayscale versions.

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.Grayscale(alpha=(0.0, 1.0))
```

Creates an augmenter that turns images to their grayscale versions with an alpha value in the range `0 <= alpha <= 1`. An alpha value of 0.5 would mean, that the output image is 50 percent of the input image and 50 percent of the grayscale image (i.e. 50 percent of color removed).

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.
`get_parameters`(self)	See *get_parameters()*.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`localize_random_state_`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`pool`(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
`remove_augmenters`(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
`remove_augmenters_`(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
`remove_augmenters_inplace`(self, func[, parents])	**Deprecated**.
`reseed`(self[, random_state, deterministic_too])	**Deprecated**.
`seed_`(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.

Continued on next page

Table 110 – continued from previous page

show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

imgaug.augmenters.color.**InColorspace**(*to_colorspace*, *from_colorspace='RGB'*, *children=None*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Deprecated. Use `WithColorspace` instead.

Convert images to another colorspace.

class imgaug.augmenters.color.**KMeansColorQuantization**(*n_colors=(2, 16)*, *from_colorspace='RGB'*, *to_colorspace=['RGB', 'Lab']*, *max_size=128*, *interpolation='linear'*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: `imgaug.augmenters.color._AbstractColorQuantization`

Quantize colors using k-Means clustering.

This "collects" the colors from the input image, groups them into `k` clusters using k-Means clustering and replaces the colors in the input image using the cluster centroids.

This is slower than `UniformColorQuantization`, but adapts dynamically to the color range in the input image.

Note: This augmenter expects input images to be either grayscale or to have 3 or 4 channels and use colorspace *from_colorspace*. If images have 4 channels, it is assumed that the 4th channel is an alpha channel and it will not be quantized.

Supported dtypes:

if (image size <= max_size):

minimum of (*ChangeColorspace*, *quantize_kmeans()*

)

if (image size > max_size):

minimum of (*ChangeColorspace*, *quantize_kmeans()*, *imresize_single_image()*

)

Parameters

- **n_colors** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter, optional*) – Target number of colors in the generated output image. This corresponds to the number of clusters in k-Means, i.e. `k`. Sampled values below `2` will always be clipped to `2`.
  - If a number, exactly that value will always be used.
  - If a tuple `(a, b)`, then a value from the discrete interval `[a..b]` will be sampled per image.

  - If a list, then a random value will be sampled from that list per image.
  - If a `StochasticParameter`, then a value will be sampled per image from that parameter.

- **to_colorspace** (*None or str or list of str or imgaug.parameters.StochasticParameter*) – The colorspace in which to perform the quantization. See *change_colorspace_()* for valid values. This will be ignored for grayscale input images.
  - If `None` the colorspace of input images will not be changed.
  - If a string, it must be among the allowed colorspaces.
  - If a list, it is expected to be a list of strings, each one being an allowed colorspace. A random element from the list will be chosen per image.
  - If a StochasticParameter, it is expected to return string. A new sample will be drawn per image.
- **from_colorspace** (*str, optional*) – The colorspace of the input images. See *to_colorspace*. Only a single string is allowed.
- **max_size** (*int or None, optional*) – Maximum image size at which to perform the augmentation. If the width or height of an image exceeds this value, it will be downscaled before running the augmentation so that the longest side matches *max_size*. This is done to speed up the augmentation. The final output image has the same size as the input image. Use `None` to apply no downscaling.
- **interpolation** (*int or str, optional*) – Interpolation method to use during downscaling when *max_size* is exceeded. Valid methods are the same as in *imresize_single_image()*.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.KMeansColorQuantization()
```

Create an augmenter to apply k-Means color quantization to images using a random amount of colors, sampled uniformly from the interval `[2..16]`. It assumes the input image colorspace to be `RGB` and clusters colors randomly in `RGB` or `Lab` colorspace.

```
>>> aug = iaa.KMeansColorQuantization(n_colors=8)
```

Create an augmenter that quantizes images to (up to) eight colors.

```
>>> aug = iaa.KMeansColorQuantization(n_colors=(4, 16))
```

Create an augmenter that quantizes images to (up to) `n` colors, where `n` is randomly and uniformly sampled from the discrete interval `[4..16]`.

```
>>> aug = iaa.KMeansColorQuantization(
>>>     from_colorspace=iaa.CSPACE_BGR)
```

Create an augmenter that quantizes input images that are in `BGR` colorspace. The quantization happens in `RGB` or `Lab` colorspace, into which the images are temporarily converted.

```
>>> aug = iaa.KMeansColorQuantization(
>>>     to_colorspace=[iaa.CSPACE_RGB, iaa.CSPACE_HSV])
```

Create an augmenter that quantizes images by clustering colors randomly in either `RGB` or `HSV` colorspace. The assumed input colorspace of images is `RGB`.

Attributes

n_colors Alias for property `counts`.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.

Continued on next page

Table 111 – continued from previous page

find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

n_colors

Alias for property counts.

Added in 0.4.0.

class imgaug.augmenters.color.**MultiplyAndAddToBrightness**(*mul=(0.7, 1.3), add=(-30, 30), to_colorspace=['YCrCb', 'HSV', 'HLS', 'Lab', 'Luv', 'YUV'], from_colorspace='RGB', random_order=True, seed=None, name=None, random_state='deprecated', deterministic='deprecated'*)

Bases: *imgaug.augmenters.color.WithBrightnessChannels*

Multiply and add to the brightness channels of input images.

This is a wrapper around *WithBrightnessChannels* and hence performs internally the same projection to random colorspaces.

Added in 0.4.0.

Supported dtypes:

See *WithBrightnessChannels*.

Parameters

- **mul** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – See `Multiply`.
- **add** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – See `Add`.
- **to_colorspace** (*imgaug.ALL or str or list of str or imgaug.parameters.StochasticParameter, optional*) – See *`WithBrightnessChannels`*.
- **from_colorspace** (*str, optional*) – See *`WithBrightnessChannels`*.
- **random_order** (*bool, optional*) – Whether to apply the add and multiply operations in random order (`True`). If `False`, this augmenter will always first multiply and then add.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.MultiplyAndAddToBrightness(mul=(0.5, 1.5), add=(-30, 30))
```

Convert each image to a colorspace with a brightness-related channel, extract that channel, multiply it by a factor between `0.5` and `1.5`, add a value between `-30` and `30` and convert back to the original colorspace.

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.

Continued on next page

Table 112 – continued from previous page

augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	See *get_children_lists()*.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.color.**MultiplyBrightness**(*mul=(0.7, 1.3), to_colorspace=['YCrCb', 'HSV', 'HLS', 'Lab', 'Luv', 'YUV'], from_colorspace='RGB', seed=None, name=None, random_state='deprecated', deterministic='deprecated'*)

Bases: *imgaug.augmenters.color.MultiplyAndAddToBrightness*

Multiply the brightness channels of input images.

This is a wrapper around *WithBrightnessChannels* and hence performs internally the same projection to random colorspaces.

Added in 0.4.0.

Supported dtypes:

See *MultiplyAndAddToBrightness*.

Parameters

- **mul** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – See `Multiply`.
- **to_colorspace** (*imgaug.ALL or str or list of str or imgaug.parameters.StochasticParameter, optional*) – See *WithBrightnessChannels*.
- **from_colorspace** (*str, optional*) – See *WithBrightnessChannels*.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.MultiplyBrightness((0.5, 1.5))
```

Convert each image to a colorspace with a brightness-related channel, extract that channel, multiply it by a factor between `0.5` and `1.5`, and convert back to the original colorspace.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.

Continued on next page

Table 113 – continued from previous page

augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	See *get_children_lists()*.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.color.**MultiplyHue**(*mul=(-3.0, 3.0), from_colorspace='RGB', seed=None, name=None, random_state='deprecated', deterministic='deprecated'*)

Bases: *imgaug.augmenters.color.MultiplyHueAndSaturation*

Multiply the hue of images by random values.

The augmenter first transforms images to HSV colorspace, then multiplies the pixel values in the H channel and afterwards converts back to RGB.

This augmenter is a shortcut for `MultiplyHueAndSaturation(mul_hue=...)`.

Supported dtypes:

See *MultiplyHueAndSaturation*.

Parameters

- **mul** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Multiplier with which to multiply all hue values. This is expected to be in the range `-10.0` to `+10.0` and will automatically be projected to an angular representation using `(hue/255) * (360/2)` (OpenCV's hue representation is in the range `[0, 180]` instead of `[0, 360]`). Only this or *mul* may be set, not both.
  - If a number, then that multiplier will be used for all images.
  - If a tuple `(a, b)`, then a value from the continuous range `[a, b]` will be sampled per image.
  - If a list, then a random value will be sampled from that list per image.
  - If a StochasticParameter, then a value will be sampled from that parameter per image.
- **from_colorspace** (*str, optional*) – See *change_colorspace_()*.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.MultiplyHue((0.5, 1.5))
```

Multiply the hue channel of images using random values between `0.5` and `1.5`.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.

Continued on next page

Table 114 – continued from previous page

augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	See *get_children_lists()*.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.color.**MultiplyHueAndSaturation**(*mul=None*, *mul_hue=None*, *mul_saturation=None*, *per_channel=False*, *from_colorspace='RGB'*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.color.WithHueAndSaturation*

Multipy hue and saturation by random values.

The augmenter first transforms images to HSV colorspace, then multiplies the pixel values in the H and S channels and afterwards converts back to RGB.

This augmenter is a wrapper around `WithHueAndSaturation`.

Supported dtypes:

See *WithHueAndSaturation*.

Parameters

- **mul** (*None or number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Multiplier with which to multiply all hue *and* saturation values of all pixels. It is expected to be in the range `-10.0` to `+10.0`. Note that values of `0.0` or lower will remove all saturation.
  - If this is `None`, *mul_hue* and/or *mul_saturation* may be set to values other than `None`.
  - If a number, then that multiplier will be used for all images.
  - If a tuple `(a, b)`, then a value from the continuous range `[a, b]` will be sampled per image.
  - If a list, then a random value will be sampled from that list per image.
  - If a StochasticParameter, then a value will be sampled from that parameter per image.
- **mul_hue** (*None or number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Multiplier with which to multiply all hue values. This is expected to be in the range `-10.0` to `+10.0` and will automatically be projected to an angular representation using `(hue/255) * (360/2)` (OpenCV's hue representation is in the range `[0, 180]` instead of `[0, 360]`). Only this or *mul* may be set, not both.
  - If this and *mul_saturation* are both `None`, *mul* may be set to a non-`None` value.
  - If a number, then that multiplier will be used for all images.
  - If a tuple `(a, b)`, then a value from the continuous range `[a, b]` will be sampled per image.
  - If a list, then a random value will be sampled from that list per image.
  - If a StochasticParameter, then a value will be sampled from that parameter per image.
- **mul_saturation** (*None or number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Multiplier with which to multiply all saturation values. It is expected to be in the range `0.0` to `+10.0`. Only this or *mul* may be set, not both.
  - If this and *mul_hue* are both `None`, *mul* may be set to a non-`None` value.
  - If a number, then that value will be used for all images.

  - If a tuple `(a, b)`, then a value from the continuous range `[a, b]` will be sampled per image.
  - If a list, then a random value will be sampled from that list per image.
  - If a StochasticParameter, then a value will be sampled from that parameter per image.

- **per_channel** (*bool or float, optional*) – Whether to sample per image only one value from *mul* and use it for both hue and saturation (`False`) or to sample independently one value for hue and one for saturation (`True`). If this value is a float `p`, then for `p` percent of all images *per_channel* will be treated as `True`, otherwise as `False`.

  This parameter has no effect if *mul_hue* and/or *mul_saturation* are used instead of *mul*.

- **from_colorspace** (*str, optional*) – See *change_colorspace_()*.

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.

- **name** (*None or str, optional*) – See `__init__()`.

- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.

- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.MultiplyHueAndSaturation((0.5, 1.5), per_channel=True)
```

Multiply hue and saturation by random values between `0.5` and `1.5` (independently per channel and the same value for all pixels within that channel). The hue will be automatically projected to an angular representation.

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.MultiplyHueAndSaturation(mul_hue=(0.5, 1.5))
```

Multiply only the hue by random values between `0.5` and `1.5`.

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.MultiplyHueAndSaturation(mul_saturation=(0.5, 1.5))
```

Multiply only the saturation by random values between `0.5` and `1.5`.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.

Continued on next page

Table 115 – continued from previous page

augment_batches(self, batches[, hooks, ...])	Augment multiple batches.
augment_bounding_boxes(self, ...[, parents, ...])	Augment a batch of bounding boxes.
augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	See *get_children_lists()*.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.color.**MultiplySaturation**(*mul=(0.0, 3.0), from_colorspace='RGB', seed=None, name=None, random_state='deprecated', deterministic='deprecated'*)

Bases: *imgaug.augmenters.color.MultiplyHueAndSaturation*

Multiply the saturation of images by random values.

The augmenter first transforms images to HSV colorspace, then multiplies the pixel values in the H channel and afterwards converts back to RGB.

This augmenter is a shortcut for `MultiplyHueAndSaturation(mul_saturation=...)`.

Supported dtypes:

See *MultiplyHueAndSaturation*.

Parameters

- **mul** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Multiplier with which to multiply all saturation values. It is expected to be in the range `0.0` to `+10.0`.
  - If a number, then that value will be used for all images.
  - If a tuple `(a, b)`, then a value from the continuous range `[a, b]` will be sampled per image.
  - If a list, then a random value will be sampled from that list per image.
  - If a StochasticParameter, then a value will be sampled from that parameter per image.
- **from_colorspace** (*str, optional*) – See *change_colorspace_()*.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.MultiplySaturation((0.5, 1.5))
```

Multiply the saturation channel of images using random values between `0.5` and `1.5`.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	See *get_children_lists()*.
`get_parameters`(self)	See *get_parameters()*.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`localize_random_state_`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`pool`(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
`remove_augmenters`(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
`remove_augmenters_`(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
`remove_augmenters_inplace`(self, func[, parents])	**Deprecated**.
`reseed`(self[, random_state, deterministic_too])	**Deprecated**.
`seed_`(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
`show_grid`(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.

Continued on next page

Table 116 – continued from previous page

to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.color.**Posterize**(*nb_bits=(1, 8)*, *from_colorspace='RGB'*, *to_colorspace=None*, *max_size=None*, *interpolation='linear'*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.color.UniformColorQuantizationToNBits*

Alias for *UniformColorQuantizationToNBits*.

Added in 0.4.0.

Supported dtypes:

See *UniformColorQuantizationToNBits*.

Methods

__call__(self, *args, **kwargs)	Alias for *augment()*.
augment(self[, return_batch, hooks])	Augment a batch.
augment_batch(self, batch[, hooks])	**Deprecated**.
augment_batch_(self, batch[, parents, hooks])	Augment a single batch in-place.
augment_batches(self, batches[, hooks, ...])	Augment multiple batches.
augment_bounding_boxes(self, ...[, parents, ...])	Augment a batch of bounding boxes.
augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.

Continued on next page

Table 117 – continued from previous page

get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.color.**RemoveSaturation**(*mul=1*, *from_colorspace='RGB'*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.color.MultiplySaturation*

Decrease the saturation of images by varying degrees.

This creates images looking similar to *Grayscale*.

This augmenter is the same as `MultiplySaturation((0.0, 1.0))`.

Added in 0.4.0.

Supported dtypes:

See *MultiplySaturation*.

Parameters

- **mul** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – *Inverse* multiplier to use for the saturation values. High values denote stronger color removal. E.g. `1.0` will remove all saturation, `0.0` will remove nothing. Expected value range is `[0.0, 1.0]`.
  - If a number, then that value will be used for all images.
  - If a tuple `(a, b)`, then a value from the continuous range `[a, b]` will be sampled per image.
  - If a list, then a random value will be sampled from that list per image.
  - If a StochasticParameter, then a value will be sampled from that parameter per image.
- **from_colorspace** (*str, optional*) – See *change_colorspace_()*.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.

- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.RemoveSaturation((0.0, 1.0))
```

Create an augmenter that decreases saturation by varying degrees.

```
>>> aug = iaa.RemoveSaturation(1.0)
```

Create an augmenter that removes all saturation from input images. This is similar to *Grayscale*.

```
>>> aug = iaa.RemoveSaturation(from_colorspace=iaa.CSPACE_BGR)
```

Create an augmenter that decreases saturation of images in `BGR` colorspace by varying degrees.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.

Continued on next page

Table 118 – continued from previous page

copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	See *get_children_lists()*.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.color.**UniformColorQuantization**(*n_colors=(2, 16)*, *from_colorspace='RGB'*, *to_colorspace=None*, *max_size=None*, *interpolation='linear'*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: imgaug.augmenters.color._AbstractColorQuantization

Quantize colors into N bins with regular distance.

For uint8 images the equation is floor(v/q)*q + q/2 with q = 256/N, where v is a pixel intensity value and N is the target number of colors after quantization.

This augmenter is faster than KMeansColorQuantization, but the set of possible output colors is constant (i.e. independent of the input images). It may produce unsatisfying outputs for input images that are made up of very similar colors.

Note: This augmenter expects input images to be either grayscale or to have 3 or 4 channels and use colorspace *from_colorspace*. If images have 4 channels, it is assumed that the 4th channel is an alpha channel and it will

not be quantized.

Supported dtypes:

if (image size <= max_size):

minimum of (*ChangeColorspace*, *quantize_uniform_()*

)

if (image size > max_size):

minimum of (*ChangeColorspace*, *quantize_uniform_()*, *imresize_single_image()*

)

Parameters

- **n_colors** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter, optional*) –

  Target number of colors to use in the generated output image.

  - If a number, exactly that value will always be used.
  - If a tuple `(a, b)`, then a value from the discrete interval `[a..b]` will be sampled per image.
  - If a list, then a random value will be sampled from that list per image.
  - If a `StochasticParameter`, then a value will be sampled per image from that parameter.

- **to_colorspace** (*None or str or list of str or imgaug.parameters.StochasticParameter*) – The colorspace in which to perform the quantization. See *change_colorspace_()* for valid values. This will be ignored for grayscale input images.

  - If `None` the colorspace of input images will not be changed.
  - If a string, it must be among the allowed colorspaces.
  - If a list, it is expected to be a list of strings, each one being an allowed colorspace. A random element from the list will be chosen per image.
  - If a StochasticParameter, it is expected to return string. A new sample will be drawn per image.

- **from_colorspace** (*str, optional*) – The colorspace of the input images. See *to_colorspace*. Only a single string is allowed.

- **max_size** (*None or int, optional*) – Maximum image size at which to perform the augmentation. If the width or height of an image exceeds this value, it will be downscaled before running the augmentation so that the longest side matches *max_size*. This is done to speed up the augmentation. The final output image has the same size as the input image. Use `None` to apply no downscaling.

- **interpolation** (*int or str, optional*) – Interpolation method to use during downscaling when *max_size* is exceeded. Valid methods are the same as in *imresize_single_image()*.

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.

- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.UniformColorQuantization()
```

Create an augmenter to apply uniform color quantization to images using a random amount of colors, sampled uniformly from the discrete interval `[2..16]`.

```
>>> aug = iaa.UniformColorQuantization(n_colors=8)
```

Create an augmenter that quantizes images to (up to) eight colors.

```
>>> aug = iaa.UniformColorQuantization(n_colors=(4, 16))
```

Create an augmenter that quantizes images to (up to) `n` colors, where `n` is randomly and uniformly sampled from the discrete interval `[4..16]`.

```
>>> aug = iaa.UniformColorQuantization(
>>>     from_colorspace=iaa.CSPACE_BGR,
>>>     to_colorspace=[iaa.CSPACE_RGB, iaa.CSPACE_HSV])
```

Create an augmenter that uniformly quantizes images in either `RGB` or `HSV` colorspace (randomly picked per image). The input colorspace of all images has to be `BGR`.

Attributes

n_colors	Alias for property `counts`.

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.

Continued on next page

Table 119 – continued from previous page

augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

n_colors

Alias for property counts.

Added in 0.4.0.

class imgaug.augmenters.color.**UniformColorQuantizationToNBits**(*nb_bits=(1, 8)*, *from_colorspace='RGB'*, *to_colorspace=None*, *max_size=None*, *interpolation='linear'*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: imgaug.augmenters.color._AbstractColorQuantization

Quantize images by setting 8-B bits of each component to zero.

This augmenter sets the 8-B highest frequency (rightmost) bits of each array component to zero. For B bits this is equivalent to changing each component's intensity value v to v' = v & (2**(8-B) - 1), e.g. for B=3 this results in v' = c & ~(2**(3-1) - 1) = c & ~3 = c & ~0000 0011 = c & 1111 1100.

This augmenter behaves for B similarly to UniformColorQuantization(2**B), but quantizes each bin with interval (a, b) to a instead of to a + (b-a)/2.

This augmenter is comparable to PIL.ImageOps.posterize().

Note: This augmenter expects input images to be either grayscale or to have 3 or 4 channels and use colorspace *from_colorspace*. If images have 4 channels, it is assumed that the 4th channel is an alpha channel and it will not be quantized.

Added in 0.4.0.

Supported dtypes:

if (image size <= max_size):

minimum of (*ChangeColorspace*, *quantize_uniform()*

)

if (image size > max_size):

minimum of (*ChangeColorspace*, *quantize_uniform()*, *imresize_single_image()*

)

Parameters

- **nb_bits** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter, optional*) –

  Number of bits to keep in each image's array component.

  – If a number, exactly that value will always be used.

  – If a tuple (a, b), then a value from the discrete interval [a..b] will be sampled per image.

  – If a list, then a random value will be sampled from that list per image.

  – If a StochasticParameter, then a value will be sampled per image from that parameter.

- **to_colorspace** (*None or str or list of str or imgaug.parameters.StochasticParameter*) – The colorspace in which to perform the quantization. See `change_colorspace_()` for valid values. This will be ignored for grayscale input images.
  - If `None` the colorspace of input images will not be changed.
  - If a string, it must be among the allowed colorspaces.
  - If a list, it is expected to be a list of strings, each one being an allowed colorspace. A random element from the list will be chosen per image.
  - If a StochasticParameter, it is expected to return string. A new sample will be drawn per image.
- **from_colorspace** (*str, optional*) – The colorspace of the input images. See *to_colorspace*. Only a single string is allowed.
- **max_size** (*None or int, optional*) – Maximum image size at which to perform the augmentation. If the width or height of an image exceeds this value, it will be downscaled before running the augmentation so that the longest side matches *max_size*. This is done to speed up the augmentation. The final output image has the same size as the input image. Use `None` to apply no downscaling.
- **interpolation** (*int or str, optional*) – Interpolation method to use during downscaling when *max_size* is exceeded. Valid methods are the same as in `imresize_single_image()`.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.UniformColorQuantizationToNBits()
```

Create an augmenter to apply uniform color quantization to images using a random amount of bits to remove, sampled uniformly from the discrete interval `[1..8]`.

```
>>> aug = iaa.UniformColorQuantizationToNBits(nb_bits=(2, 8))
```

Create an augmenter that quantizes images by removing `8-B` rightmost bits from each component, where `B` is uniformly sampled from the discrete interval `[2..8]`.

```
>>> aug = iaa.UniformColorQuantizationToNBits(
>>>     from_colorspace=iaa.CSPACE_BGR,
>>>     to_colorspace=[iaa.CSPACE_RGB, iaa.CSPACE_HSV])
```

Create an augmenter that uniformly quantizes images in either RGB or HSV colorspace (randomly picked per image). The input colorspace of all images has to be BGR.

Methods

__call__(self, *args, **kwargs)	Alias for *augment()*.
augment(self[, return_batch, hooks])	Augment a batch.
augment_batch(self, batch[, hooks])	**Deprecated**.
augment_batch_(self, batch[, parents, hooks])	Augment a single batch in-place.
augment_batches(self, batches[, hooks, ...])	Augment multiple batches.
augment_bounding_boxes(self, ...[, parents, ...])	Augment a batch of bounding boxes.
augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.

Continued on next page

Table 120 – continued from previous page

remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.color.**WithBrightnessChannels**(*children=None, to_colorspace=['YCrCb', 'HSV', 'HLS', 'Lab', 'Luv', 'YUV'], from_colorspace='RGB', seed=None, name=None, random_state='deprecated', deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Augmenter to apply child augmenters to brightness-related image channels.

This augmenter first converts an image to a random colorspace containing a brightness-related channel (e.g. `V` in `HSV`), then extracts that channel and applies its child augmenters to this one channel. Afterwards, it reintegrates the augmented channel into the full image and converts back to the input colorspace.

Added in 0.4.0.

Supported dtypes:

See *change_colorspaces_()*.

Parameters

- **children** (*imgaug.augmenters.meta.Augmenter or list of imgaug.augmenters.meta.Augmenter or None, optional*) – One or more augmenters to apply to the brightness channels. They receive images with a single channel and have to modify these.
- **to_colorspace** (*imgaug.ALL or str or list of str or imgaug.parameters.StochasticParameter, optional*) – Colorspace in which to extract the brightness-related channels. Currently, `imgaug.augmenters.color.CSPACE_YCrCb`, `CSPACE_HSV`, `CSPACE_HLS`, `CSPACE_Lab`, `CSPACE_Luv`, `CSPACE_YUV`, `CSPACE_CIE` are supported.
  - If `imgaug.ALL`: Will pick imagewise a random colorspace from all supported colorspaces.
  - If `str`: Will always use this colorspace.
  - If `list` or `str`: Will pick imagewise a random colorspace from this list.
  - If *StochasticParameter*: A parameter that will be queried once per batch to generate all target colorspaces. Expected to return strings matching the `CSPACE_*` constants.
- **from_colorspace** (*str, optional*) – See *change_colorspace_()*.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.

- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.WithBrightnessChannels(iaa.Add((-50, 50)))
```

Add `-50` to `50` to the brightness-related channels of each image.

```
>>> aug = iaa.WithBrightnessChannels(
>>>     iaa.Add((-50, 50)), to_colorspace=[iaa.CSPACE_Lab, iaa.CSPACE_HSV])
```

Add `-50` to `50` to the brightness-related channels of each image, but pick those brightness-related channels only from `Lab` (`L`) and `HSV` (`V`) colorspaces.

```
>>> aug = iaa.WithBrightnessChannels(
>>>     iaa.Add((-50, 50)), from_colorspace=iaa.CSPACE_BGR)
```

Add `-50` to `50` to the brightness-related channels of each image, where the images are provided in `BGR` colorspace instead of the standard `RGB`.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.

Continued on next page

Table 121 – continued from previous page

copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	See *get_children_lists()*.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_children_lists(*self*)
See *get_children_lists()*.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.color.**WithColorspace**(*to_colorspace*, *from_colorspace='RGB'*, *children=None*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Apply child augmenters within a specific colorspace.

This augumenter takes a source colorspace A and a target colorspace B as well as children C. It changes images from A to B, then applies the child augmenters C and finally changes the colorspace back from B to A. See also ChangeColorspace() for more.

Supported dtypes:

See *change_colorspaces_()*.

Parameters

- **to_colorspace** (*str*) – See *change_colorspace_()*.
- **from_colorspace** (*str, optional*) – See *change_colorspace_()*.
- **children** (*imgaug.augmenters.meta.Augmenter or list of imgaug.augmenters.meta.Augmenter or None, optional*) – One or more augmenters to apply to converted images.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.WithColorspace(
>>>     to_colorspace=iaa.CSPACE_HSV,
>>>     from_colorspace=iaa.CSPACE_RGB,
>>>     children=iaa.WithChannels(
>>>         0,
>>>         iaa.Add((0, 50))
>>>     )
>>> )
```

Convert to `HSV` colorspace, add a value between `0` and `50` (uniformly sampled per image) to the Hue channel, then convert back to the input colorspace (`RGB`).

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.

Continued on next page

Table 122 – continued from previous page

augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	See *get_children_lists()*.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_children_lists(*self*)
See *get_children_lists()*.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.color.**WithHueAndSaturation**(*children=None*, *from_colorspace='RGB'*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Apply child augmenters to hue and saturation channels.

This augumenter takes an image in a source colorspace, converts it to HSV, extracts the H (hue) and S (saturation) channels, applies the provided child augmenters to these channels and finally converts back to the original colorspace.

The image array generated by this augmenter and provided to its children is in `int16` (**sic!** only augmenters that can handle `int16` arrays can be children!). The hue channel is mapped to the value range `[0, 255]`. Before converting back to the source colorspace, the saturation channel's values are clipped to `[0, 255]`. A modulo operation is applied to the hue channel's values, followed by a mapping from `[0, 255]` to `[0, 180]` (and finally the colorspace conversion).

Supported dtypes:

See *change_colorspaces_()*.

Parameters

- **from_colorspace** (*str, optional*) – See *change_colorspace_()*.
- **children** (*imgaug.augmenters.meta.Augmenter or list of imgaug.augmenters.meta.Augmenter or None, optional*) – One or more augmenters to apply to converted images. They receive `int16` images with two channels (hue, saturation) and have to modify these.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.WithHueAndSaturation(
>>>     iaa.WithChannels(0, iaa.Add((0, 50)))
>>> )
```

Create an augmenter that will add a random value between `0` and `50` (uniformly sampled per image) hue channel in HSV colorspace. It automatically accounts for the hue being in angular representation, i.e. if the angle goes beyond 360 degrees, it will start again at 0 degrees. The colorspace is finally converted back to `RGB` (default setting).

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.WithHueAndSaturation([
>>>     iaa.WithChannels(0, iaa.Add((-30, 10))),
>>>     iaa.WithChannels(1, [
>>>         iaa.Multiply((0.5, 1.5)),
>>>         iaa.LinearContrast((0.75, 1.25))
>>>     ])
>>> ])
```

Create an augmenter that adds a random value sampled uniformly from the range `[-30, 10]` to the hue and multiplies the saturation by a random factor sampled uniformly from `[0.5, 1.5]`. It also modifies the contrast of the saturation channel. After these steps, the `HSV` image is converted back to `RGB`.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	See *get_children_lists()*.
get_parameters(self)	See *get_parameters()*.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`localize_random_state_`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`pool`(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
`remove_augmenters`(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
`remove_augmenters_`(self, func[, parents])	Remove in-place children of this augmenter that match a condition.

Continued on next page

Table 123 – continued from previous page

remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_children_lists(*self*)
See *get_children_lists()*.

get_parameters(*self*)
See *get_parameters()*.

imgaug.augmenters.color.**change_color_temperature**(*image*, *kelvin*, *from_colorspace='RGB'*)

Change the temperature of an image to a given value in Kelvin.

Added in 0.4.0.

Supported dtypes:

See *change_color_temperatures_*.

Parameters

- **image** (*ndarray*) – The image which's color temperature is supposed to be changed. Expected to be of shape `(H,W,3)` array.
- **kelvin** (*number*) – The temperature in Kelvin. Expected value range is in the interval `(1000, 4000)`.
- **from_colorspace** (*str, optional*) – The source colorspace. See *change_colorspaces_()*. Defaults to `RGB`.

Returns Image with target color temperature.

Return type ndarray

imgaug.augmenters.color.**change_color_temperatures_**(*images*, *kelvins*, *from_colorspaces='RGB'*)

Change in-place the temperature of images to given values in Kelvin.

Added in 0.4.0.

Supported dtypes:

See *change_colorspace_*.

Parameters

- **images** (*ndarray or list of ndarray*) – The images which's color temperature is supposed to be changed. Either a list of `(H,W,3)` arrays or a single `(N,H,W,3)` array.
- **kelvins** (*iterable of number*) – Temperatures in Kelvin. One per image. Expected value range is in the interval `(1000, 4000)`.
- **from_colorspaces** (*str or list of str, optional*) – The source colorspace. See *change_colorspaces_()*. Defaults to `RGB`.

Returns Images with target color temperatures. The input array(s) might have been changed in-place.

Return type ndarray or list of ndarray

imgaug.augmenters.color.**change_colorspace_**(*image*, *to_colorspace*, *from_colorspace='RGB'*)

Change the colorspace of an image inplace.

Note: All outputs of this function are *uint8*. For some colorspaces this may not be optimal.

Note: Output grayscale images will still have three channels.

Supported dtypes:

- `uint8`: yes; fully tested
- `uint16`: no
- `uint32`: no
- `uint64`: no
- `int8`: no
- `int16`: no
- `int32`: no
- `int64`: no
- `float16`: no
- `float32`: no
- `float64`: no
- `float128`: no
- `bool`: no

Parameters

- **image** (*ndarray*) – The image to convert from one colorspace into another. Usually expected to have shape `(H,W,3)`.
- **to_colorspace** (*str*) – The target colorspace. See the `CSPACE` constants, e.g. `imgaug.augmenters.color.CSPACE_RGB`.
- **from_colorspace** (*str, optional*) – The source colorspace. Analogous to *to_colorspace*. Defaults to `RGB`.

Returns Image with target colorspace. *Can* be the same array instance as was originally provided (i.e. changed inplace). Grayscale images will still have three channels.

Return type ndarray

Examples

```
>>> import imgaug.augmenters as iaa
>>> import numpy as np
>>> # fake RGB image
```

(continues on next page)

(continued from previous page)

```
>>> image_rgb = np.arange(4*4*3).astype(np.uint8).reshape((4, 4, 3))
>>> image_bgr = iaa.change_colorspace_(np.copy(image_rgb), iaa.CSPACE_BGR)
```

imgaug.augmenters.color.**change_colorspaces_**(*images*, *to_colorspaces*, *from_colorspaces='RGB'*)

Change the colorspaces of a batch of images inplace.

Note: All outputs of this function are *uint8*. For some colorspaces this may not be optimal.

Note: Output grayscale images will still have three channels.

Supported dtypes:

See *change_colorspace_()*.

Parameters

- **images** (*ndarray or list of ndarray*) – The images to convert from one colorspace into another. Either a list of `(H,W,3)` arrays or a single `(N,H,W,3)` array.
- **to_colorspaces** (*str or iterable of str*) – The target colorspaces. Either a single string (all images will be converted to the same colorspace) or an iterable of strings (one per image). See the `CSPACE` constants, e.g. `imgaug.augmenters.color.CSPACE_RGB`.
- **from_colorspaces** (*str or list of str, optional*) – The source colorspace. Analogous to *to_colorspace*. Defaults to `RGB`.

Returns Images with target colorspaces. *Can* contain the same array instances as were originally provided (i.e. changed inplace). Grayscale images will still have three channels.

Return type ndarray or list of ndarray

Examples

```
>>> import imgaug.augmenters as iaa
>>> import numpy as np
>>> # fake RGB image
>>> image_rgb = np.arange(4*4*3).astype(np.uint8).reshape((4, 4, 3))
>>> images_rgb = [image_rgb, image_rgb, image_rgb]
>>> images_rgb_copy = [np.copy(image_rgb) for image_rgb in images_rgb]
>>> images_bgr = iaa.change_colorspaces_(images_rgb_copy, iaa.CSPACE_BGR)
```

Create three example `RGB` images and convert them to `BGR` colorspace.

```
>>> images_rgb_copy = [np.copy(image_rgb) for image_rgb in images_rgb]
>>> images_various = iaa.change_colorspaces_(
>>>     images_rgb_copy, [iaa.CSPACE_BGR, iaa.CSPACE_HSV, iaa.CSPACE_GRAY])
```

Chnage the colorspace of the first image to `BGR`, the one of the second image to `HSV` and the one of the third image to `grayscale` (note that in the latter case the image will still have shape `(H,W,3)`, not `(H,W,1)`).

imgaug.augmenters.color.**posterize**(*arr*, *nb_bits*)

Alias for *quantize_uniform_to_n_bits()*.

This function is an alias for *quantize_uniform_to_n_bits()* and was added for users familiar with the same function in PIL.

Added in 0.4.0.

Supported dtypes:

See *quantize_uniform_to_n_bits()*.

Parameters

- **arr** (*ndarray*) – See *quantize_uniform_to_n_bits()*.
- **nb_bits** (*int*) – See *quantize_uniform_to_n_bits()*.

Returns Array with quantized components.

Return type ndarray

imgaug.augmenters.color.**quantize_colors_kmeans**(*image*, *n_colors*, *n_max_iter=10*, *eps=1.0*)

Deprecated. Use `imgaug.augmenters.colors.quantize_kmeans` instead.

Outdated name of *quantize_kmeans()*.

Deprecated since 0.4.0.

imgaug.augmenters.color.**quantize_colors_uniform**(*image*, *n_colors*)

Deprecated. Use `imgaug.augmenters.colors.quantize_uniform` instead.

Outdated name for *quantize_uniform()*.

Deprecated since 0.4.0.

imgaug.augmenters.color.**quantize_kmeans**(*arr*, *nb_clusters*, *nb_max_iter=10*, *eps=1.0*)

Quantize an array into N bins using k-means clustering.

If the input is an image, this method returns in an image with a maximum of `N` colors. Similar colors are grouped to their mean. The k-means clustering happens across channels and not channelwise.

Code similar to https://docs.opencv.org/3.0-beta/doc/py_tutorials/py_ml/ py_kmeans/py_kmeans_opencv/py_kmeans_opencv.htm

> **Warning:** This function currently changes the RNG state of both OpenCV's internal RNG and imgaug's global RNG. This is necessary in order to ensure that the k-means clustering happens deterministically.

Added in 0.4.0. (Previously called `quantize_colors_kmeans()`.)

Supported dtypes:

- `uint8`: yes; fully tested
- `uint16`: no
- `uint32`: no
- `uint64`: no
- `int8`: no
- `int16`: no
- `int32`: no
- `int64`: no
- `float16`: no

- `float32`: no
- `float64`: no
- `float128`: no
- `bool`: no

Parameters

- **arr** (*ndarray*) – Array to quantize. Expected to be of shape `(H,W)` or `(H,W,C)` with `C` usually being `1` or `3`.
- **nb_clusters** (*int*) – Number of clusters to quantize into, i.e. `k` in k-means clustering. This corresponds to the maximum number of colors in an output image.
- **nb_max_iter** (*int, optional*) – Maximum number of iterations that the k-means clustering algorithm is run.
- **eps** (*float, optional*) – Minimum change of all clusters per k-means iteration. If all clusters change by less than this amount in an iteration, the clustering is stopped.

Returns Image with quantized colors.

Return type ndarray

Examples

```
>>> import imgaug.augmenters as iaa
>>> import numpy as np
>>> image = np.arange(4 * 4 * 3, dtype=np.uint8).reshape((4, 4, 3))
>>> image_quantized = iaa.quantize_kmeans(image, 6)
```

Generates a `4x4` image with `3` channels, containing consecutive values from `0` to `4*4*3`, leading to an equal number of colors. These colors are then quantized so that only `6` are remaining. Note that the six remaining colors do have to appear in the input image.

imgaug.augmenters.color.**quantize_uniform**(*arr*, *nb_bins*, *to_bin_centers=True*)

Quantize an array into N equally-sized bins.

See *quantize_uniform_()* for details.

Added in 0.4.0. (Previously called `quantize_colors_uniform()`.)

Supported dtypes:

See *quantize_uniform_()*.

Parameters

- **arr** (*ndarray*) – See *quantize_uniform_()*.
- **nb_bins** (*int*) – See *quantize_uniform_()*.
- **to_bin_centers** (*bool*) – See *quantize_uniform_()*.

Returns Array with quantized components.

Return type ndarray

imgaug.augmenters.color.**quantize_uniform_**(*arr*, *nb_bins*, *to_bin_centers=True*)

Quantize an array into N equally-sized bins in-place.

This can be used to quantize/posterize an image into N colors.

For `uint8` arrays the equation is `floor(v/q)*q + q/2` with `q = 256/N`, where `v` is a pixel intensity value and `N` is the target number of bins (roughly matches number of colors) after quantization.

Added in 0.4.0.

Supported dtypes:

- `uint8`: yes; fully tested
- `uint16`: no
- `uint32`: no
- `uint64`: no
- `int8`: no
- `int16`: no
- `int32`: no
- `int64`: no
- `float16`: no
- `float32`: no
- `float64`: no
- `float128`: no
- `bool`: no

Parameters

- **arr** (*ndarray*) – Array to quantize, usually an image. Expected to be of shape `(H,W)` or `(H,W,C)` with `C` usually being `1` or `3`. This array *may* be changed in-place.
- **nb_bins** (*int*) – Number of equally-sized bins to quantize into. This corresponds to the maximum number of colors in an output image.
- **to_bin_centers** (*bool*) – Whether to quantize each bin `(a, b)` to `a + (b-a)/2` (center of bin, `True`) or to `a` (lower boundary, `False`).

Returns Array with quantized components. This *may* be the input array with components changed in-place.

Return type ndarray

Examples

```
>>> import imgaug.augmenters as iaa
>>> import numpy as np
>>> image = np.arange(4 * 4 * 3, dtype=np.uint8).reshape((4, 4, 3))
>>> image_quantized = iaa.quantize_uniform_(np.copy(image), 6)
```

Generates a `4x4` image with `3` channels, containing consecutive values from `0` to `4*4*3`, leading to an equal number of colors. Each component is then quantized into one of `6` bins that regularly split up the value range of `[0..255]`, i.e. the resolution w.r.t. to the value range is reduced.

`imgaug.augmenters.color.`**`quantize_uniform_to_n_bits`**(*arr*, *nb_bits*)

Reduce each component in an array to a maximum number of bits.

See *`quantize_uniform_to_n_bits()`* for details.

Added in 0.4.0.

Supported dtypes:

See *quantize_uniform_to_n_bits_()*.

Parameters

- **arr** (*ndarray*) – See *quantize_uniform_to_n_bits()*.
- **nb_bits** (*int*) – See *quantize_uniform_to_n_bits()*.

Returns Array with quantized components.

Return type ndarray

imgaug.augmenters.color.**quantize_uniform_to_n_bits_**(*arr*, *nb_bits*)

Reduce each component in an array to a maximum number of bits in-place.

This operation sets the `8-B` highest frequency (rightmost) bits to zero. For `B` bits this is equivalent to changing each component's intensity value `v` to `v' = v & (2**(8-B) - 1)`, e.g. for `B=3` this results in `v' = c & ~(2**(3-1) - 1) = c & ~3 = c & ~0000 0011 = c & 1111 1100`.

This is identical to *quantize_uniform()* with `nb_bins=2**nb_bits` and `to_bin_centers=False`.

This function produces the same outputs as `PIL.ImageOps.posterize()`, but is significantly faster.

Added in 0.4.0.

Supported dtypes:

See *quantize_uniform_()*.

Parameters

- **arr** (*ndarray*) – Array to quantize, usually an image. Expected to be of shape `(H,W)` or `(H,W,C)` with `C` usually being `1` or `3`. This array *may* be changed in-place.
- **nb_bits** (*int*) – Number of bits to keep in each array component.

Returns Array with quantized components. This *may* be the input array with components changed in-place.

Return type ndarray

Examples

```
>>> import imgaug.augmenters as iaa
>>> import numpy as np
>>> image = np.arange(4 * 4 * 3, dtype=np.uint8).reshape((4, 4, 3))
>>> image_quantized = iaa.quantize_uniform_to_n_bits_(np.copy(image), 6)
```

Generates a `4x4` image with `3` channels, containing consecutive values from `0` to `4*4*3`, leading to an equal number of colors. These colors are then quantized so that each component's `8-6=2` rightmost bits are set to zero.

13.24 imgaug.augmenters.contrast

Augmenters that perform contrast changes.

List of augmenters:

- *GammaContrast*
- *SigmoidContrast*
- *LogContrast*
- *LinearContrast*
- *AllChannelsHistogramEqualization*
- *HistogramEqualization*
- *AllChannelsCLAHE*
- *CLAHE*

class imgaug.augmenters.contrast.**AllChannelsCLAHE**(*clip_limit=(0.1, 8)*, *tile_grid_size_px=(3, 12)*, *tile_grid_size_px_min=3*, *per_channel=False*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Apply CLAHE to all channels of images in their original colorspaces.

CLAHE (Contrast Limited Adaptive Histogram Equalization) performs histogram equilization within image patches, i.e. over local neighbourhoods.

In contrast to `imgaug.augmenters.contrast.CLAHE`, this augmenter operates directly on all channels of the input images. It does not perform any colorspace transformations and does not focus on specific channels (e.g. `L` in `Lab` colorspace).

Supported dtypes:

- `uint8`: yes; fully tested
- `uint16`: yes; tested
- `uint32`: no (1)
- `uint64`: no (2)
- `int8`: no (2)
- `int16`: no (2)
- `int32`: no (2)
- `int64`: no (2)
- `float16`: no (2)
- `float32`: no (2)
- `float64`: no (2)
- `float128`: no (1)
- `bool`: no (1)

- (1) rejected by cv2
- (2) results in error in cv2: `cv2.error: OpenCV(3.4.2) (...)/clahe.cpp:351: error: (-215:Assertion failed) src.type() == (((0) & ((1 <<`

```
3) - 1)) + (((1)-1) << 3)) || _src.type() == (((2) & ((1 <<
3) - 1)) + (((1)-1) << 3)) in function 'apply'
```

Parameters

- **clip_limit** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – See `imgaug.augmenters.contrast.CLAHE`.
- **tile_grid_size_px** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter or tuple of tuple of int or tuple of list of int or tuple of imgaug.parameters.StochasticParameter, optional*) – See `imgaug.augmenters.contrast.CLAHE`.
- **tile_grid_size_px_min** (*int, optional*) – See `imgaug.augmenters.contrast.CLAHE`.
- **per_channel** (*bool or float, optional*) – Whether to use the same value for all channels (`False`) or to sample a new value for each channel (`True`). If this value is a float `p`, then for `p` percent of all images *per_channel* will be treated as `True`, otherwise as `False`.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.AllChannelsCLAHE()
```

Create an augmenter that applies CLAHE to all channels of input images.

```
>>> aug = iaa.AllChannelsCLAHE(clip_limit=(1, 10))
```

Same as in the previous example, but the *clip_limit* used by CLAHE is uniformly sampled per image from the interval `[1, 10]`. Some images will therefore have stronger contrast than others (i.e. higher clip limit values).

```
>>> aug = iaa.AllChannelsCLAHE(clip_limit=(1, 10), per_channel=True)
```

Same as in the previous example, but the *clip_limit* is sampled per image *and* channel, leading to different levels of contrast for each channel.

Methods

__call__(self, *args, **kwargs)	Alias for *augment()*.
augment(self[, return_batch, hooks])	Augment a batch.
augment_batch(self, batch[, hooks])	**Deprecated**.
augment_batch_(self, batch[, parents, hooks])	Augment a single batch in-place.
augment_batches(self, batches[, hooks, ...])	Augment multiple batches.
augment_bounding_boxes(self, ...[, parents, ...])	Augment a batch of bounding boxes.
augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.

Continued on next page

Table 124 – continued from previous page

to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.contrast.**AllChannelsHistogramEqualization**(*seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Apply Histogram Eq. to all channels of images in their original colorspaces.

In contrast to imgaug.augmenters.contrast.HistogramEqualization, this augmenter operates directly on all channels of the input images. It does not perform any colorspace transformations and does not focus on specific channels (e.g. L in Lab colorspace).

Supported dtypes:

- uint8: yes; fully tested
- uint16: no (1)
- uint32: no (2)
- uint64: no (1)
- int8: no (1)
- int16: no (1)
- int32: no (1)
- int64: no (1)
- float16: no (2)
- float32: no (1)
- float64: no (1)
- float128: no (2)
- bool: no (1)

- (1) causes cv2 error: cv2.error: OpenCV(3.4.5) (...)/histogram.cpp:3345: error: (-215:Assertion failed) src.type() == CV_8UC1 in function 'equalizeHist'
- (2) rejected by cv2

Parameters

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See __init__().
- **name** (*None or str, optional*) – See __init__().

- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.AllChannelsHistogramEqualization()
```

Create an augmenter that applies histogram equalization to all channels of input images in the original colorspaces.

```
>>> aug = iaa.Alpha((0.0, 1.0), iaa.AllChannelsHistogramEqualization())
```

Same as in the previous example, but alpha-blends the contrast-enhanced augmented images with the original input images using random blend strengths. This leads to random strengths of the contrast adjustment.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.

Continued on next page

Table 125 – continued from previous page

draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See get_parameters().
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See get_parameters().

class imgaug.augmenters.contrast.**CLAHE**(*clip_limit=(0.1, 8)*, *tile_grid_size_px=(3, 12)*, *tile_grid_size_px_min=3*, *from_colorspace='RGB'*, *to_colorspace='Lab'*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: imgaug.augmenters.meta.Augmenter

Apply CLAHE to L/V/L channels in HLS/HSV/Lab colorspaces.

This augmenter applies CLAHE (Contrast Limited Adaptive Histogram Equalization) to images, a form of histogram equalization that normalizes within local image patches. The augmenter transforms input images to a target colorspace (e.g. Lab), extracts an intensity-related channel from the converted images (e.g. L for Lab), applies CLAHE to the channel and then converts the resulting image back to the original colorspace.

Grayscale images (images without channel axis or with only one channel axis) are automatically handled, *from_colorspace* does not have to be adjusted for them. For images with four channels (e.g. RGBA), the fourth channel is ignored in the colorspace conversion (e.g. from an RGBA image, only the RGB part is converted, normalized, converted back and concatenated with the input A channel). Images with unusual channel numbers (2, 5 or more than 5) are normalized channel-by-channel (same behaviour as AllChannelsCLAHE, though a warning will be raised).

If you want to apply CLAHE to each channel of the original input image's colorspace (without any colorspace conversion), use imgaug.augmenters.contrast.AllChannelsCLAHE instead.

Supported dtypes:

- `uint8`: yes; fully tested
- `uint16`: no (1)
- `uint32`: no (1)
- `uint64`: no (1)
- `int8`: no (1)
- `int16`: no (1)
- `int32`: no (1)
- `int64`: no (1)
- `float16`: no (1)
- `float32`: no (1)
- `float64`: no (1)
- `float128`: no (1)
- `bool`: no (1)

- (1) This augmenter uses *ChangeColorspace*, which is currently limited to `uint8`.

Parameters

- **clip_limit** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Clipping limit. Higher values result in stronger contrast. OpenCV uses a default of `40`, though values around `5` seem to already produce decent contrast.
  - If a number, then that value will be used for all images.
  - If a tuple `(a, b)`, then a value from the range `[a, b]` will be used per image.
  - If a list, then a random value will be sampled from that list per image.
  - If a `StochasticParameter`, then a value will be sampled per image from that parameter.
- **tile_grid_size_px** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter or tuple of tuple of int or tuple of list of int or tuple of imgaug.parameters.StochasticParameter, optional*) –

  Kernel size, i.e. size of each local neighbourhood in pixels.
  - If an `int`, then that value will be used for all images for both kernel height and width.
  - If a tuple `(a, b)`, then a value from the discrete interval `[a..b]` will be uniformly sampled per image.
  - If a list, then a random value will be sampled from that list per image and used for both kernel height and width.
  - If a `StochasticParameter`, then a value will be sampled per image from that parameter per image and used for both kernel height and width.
  - If a tuple of tuple of `int` given as `((a, b), (c, d))`, then two values will be sampled independently from the discrete ranges `[a..b]` and `[c..d]` per image and used as the kernel height and width.

  - If a tuple of lists of `int`, then two values will be sampled independently per image, one from the first list and one from the second, and used as the kernel height and width.
  - If a tuple of `StochasticParameter`, then two values will be sampled indepdently per image, one from the first parameter and one from the second, and used as the kernel height and width.

- **tile_grid_size_px_min** (*int, optional*) – Minimum kernel size in px, per axis. If the sampling results in a value lower than this minimum, it will be clipped to this value.
- **from_colorspace** (*{"RGB", "BGR", "HSV", "HLS", "Lab"}, optional*) – Colorspace of the input images. If any input image has only one or zero channels, this setting will be ignored and it will be assumed that the input is grayscale. If a fourth channel is present in an input image, it will be removed before the colorspace conversion and later re-added. See also *change_colorspace_()* for details.
- **to_colorspace** (*{"Lab", "HLS", "HSV"}, optional*) – Colorspace in which to perform CLAHE. For `Lab`, CLAHE will only be applied to the first channel (`L`), for `HLS` to the second (`L`) and for `HSV` to the third (`V`). To apply CLAHE to all channels of an input image (without colorspace conversion), see `imgaug.augmenters.contrast.AllChannelsCLAHE`.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.CLAHE()
```

Create a standard CLAHE augmenter.

```
>>> aug = iaa.CLAHE(clip_limit=(1, 10))
```

Create a CLAHE augmenter with a clip limit uniformly sampled from `[1..10]`, where `1` is rather low contrast and `10` is rather high contrast.

```
>>> aug = iaa.CLAHE(tile_grid_size_px=(3, 21))
```

Create a CLAHE augmenter with kernel sizes of `SxS`, where `S` is uniformly sampled from `[3..21]`. Sampling happens once per image.

```
>>> aug = iaa.CLAHE(
>>>     tile_grid_size_px=iap.Discretize(iap.Normal(loc=7, scale=2)),
>>>     tile_grid_size_px_min=3)
```

Create a CLAHE augmenter with kernel sizes of `SxS`, where `S` is sampled from `N(7, 2)`, but does not go below `3`.

```
>>> aug = iaa.CLAHE(tile_grid_size_px=((3, 21), [3, 5, 7]))
```

Create a CLAHE augmenter with kernel sizes of `HxW`, where `H` is uniformly sampled from `[3..21]` and `W` is randomly picked from the list `[3, 5, 7]`.

```
>>> aug = iaa.CLAHE(
>>>     from_colorspace=iaa.CSPACE_BGR,
>>>     to_colorspace=iaa.CSPACE_HSV)
```

Create a CLAHE augmenter that converts images from BGR colorspace to HSV colorspace and then applies the local histogram equalization to the `V` channel of the images (before converting back to `BGR`). Alternatively, `Lab` (default) or `HLS` can be used as the target colorspace. Grayscale images (no channels / one channel) are never converted and are instead directly normalized (i.e. *from_colorspace* does not have to be changed for them).

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.

Continued on next page

Table 126 – continued from previous page

get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See get_parameters().
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

BGR = 'BGR'

HLS = 'HLS'

HSV = 'HSV'

Lab = 'Lab'

RGB = 'RGB'

get_parameters(*self*)
See get_parameters().

class imgaug.augmenters.contrast.**GammaContrast**(*gamma=(0.7, 1.7)*, *per_channel=False*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: imgaug.augmenters.contrast._ContrastFuncWrapper

Adjust image contrast by scaling pixel values to 255*((v/255)**gamma).

Values in the range gamma=(0.5, 2.0) seem to be sensible.

Supported dtypes:

See adjust_contrast_gamma().

Parameters

- **gamma** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) –

  Exponent for the contrast adjustment. Higher values darken the image.

  – If a number, then that value will be used for all images.

  – If a tuple (a, b), then a value from the range [a, b] will be used per image.

  – If a list, then a random value will be sampled from that list per image.

  - If a `StochasticParameter`, then a value will be sampled per image from that parameter.

- **per_channel** (*bool or float, optional*) – Whether to use the same value for all channels (`False`) or to sample a new value for each channel (`True`). If this value is a float `p`, then for `p` percent of all images *per_channel* will be treated as `True`, otherwise as `False`.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.GammaContrast((0.5, 2.0))
```

Modify the contrast of images according to `255*((v/255)**gamma)`, where `v` is a pixel value and `gamma` is sampled uniformly from the interval `[0.5, 2.0]` (once per image).

```
>>> aug = iaa.GammaContrast((0.5, 2.0), per_channel=True)
```

Same as in the previous example, but `gamma` is sampled once per image *and* channel.

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.

Continued on next page

Table 127 – continued from previous page

augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.contrast.**HistogramEqualization**(*from_colorspace='RGB'*, *to_colorspace='Lab'*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Apply Histogram Eq. to L/V/L channels of images in HLS/HSV/Lab colorspaces.

This augmenter is similar to imgaug.augmenters.contrast.CLAHE.

The augmenter transforms input images to a target colorspace (e.g. Lab), extracts an intensity-related channel from the converted images (e.g. L for Lab), applies Histogram Equalization to the channel and then converts the resulting image back to the original colorspace.

Grayscale images (images without channel axis or with only one channel axis) are automatically handled, *from_colorspace* does not have to be adjusted for them. For images with four channels (e.g. RGBA), the fourth channel is ignored in the colorspace conversion (e.g. from an `RGBA` image, only the `RGB` part is converted, normalized, converted back and concatenated with the input `A` channel). Images with unusual channel numbers (2, 5 or more than 5) are normalized channel-by-channel (same behaviour as `AllChannelsHistogramEqualization`, though a warning will be raised).

If you want to apply HistogramEqualization to each channel of the original input image's colorspace (without any colorspace conversion), use `imgaug.augmenters.contrast.AllChannelsHistogramEqualization` instead.

Supported dtypes:

- `uint8`: yes; fully tested
- `uint16`: no (1)
- `uint32`: no (1)
- `uint64`: no (1)
- `int8`: no (1)
- `int16`: no (1)
- `int32`: no (1)
- `int64`: no (1)
- `float16`: no (1)
- `float32`: no (1)
- `float64`: no (1)
- `float128`: no (1)
- `bool`: no (1)

- (1) This augmenter uses *`AllChannelsHistogramEqualization`*, which only supports `uint8`.

Parameters

- **from_colorspace** (*{"RGB", "BGR", "HSV", "HLS", "Lab"}, optional*) – Colorspace of the input images. If any input image has only one or zero channels, this setting will be ignored and it will be assumed that the input is grayscale. If a fourth channel is present in an input image, it will be removed before the colorspace conversion and later re-added. See also *`change_colorspace_()`* for details.
- **to_colorspace** (*{"Lab", "HLS", "HSV"}, optional*) – Colorspace in which to perform Histogram Equalization. For `Lab`, the equalization will only be applied to the first channel (`L`), for `HLS` to the second (`L`) and for `HSV` to the third (`V`). To apply histogram equalization to all channels of an input image (without colorspace conversion), see `imgaug.augmenters.contrast.AllChannelsHistogramEqualization`.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.

- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.HistogramEqualization()
```

Create an augmenter that converts images to `HLS`/`HSV`/`Lab` colorspaces, extracts intensity-related channels (i.e. `L`/`V`/`L`), applies histogram equalization to these channels and converts back to the input colorspace.

```
>>> aug = iaa.Alpha((0.0, 1.0), iaa.HistogramEqualization())
```

Same as in the previous example, but alpha blends the result, leading to various strengths of contrast normalization.

```
>>> aug = iaa.HistogramEqualization(
>>>     from_colorspace=iaa.CSPACE_BGR,
>>>     to_colorspace=iaa.CSPACE_HSV)
```

Same as in the first example, but the colorspace of input images has to be `BGR` (instead of default `RGB`) and the histogram equalization is applied to the `V` channel in `HSV` colorspace.

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.

Continued on next page

Table 128 – continued from previous page

copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

BGR = 'BGR'

HLS = 'HLS'

HSV = 'HSV'

Lab = 'Lab'

RGB = 'RGB'

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.contrast.**LinearContrast**(*alpha=(0.6, 1.4)*, *per_channel=False*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: imgaug.augmenters.contrast._ContrastFuncWrapper

Adjust contrast by scaling each pixel to `127 + alpha*(v-127)`.

Supported dtypes:

See *adjust_contrast_linear()*.

Parameters

- **alpha** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Multiplier to linearly pronounce (`>1.0`), dampen (`0.0` to `1.0`) or invert (`<0.0`) the difference between each pixel value and the dtype's center value, e.g. `127` for `uint8`.
  - If a number, then that value will be used for all images.
  - If a tuple `(a, b)`, then a value from the interval `[a, b]` will be used per image.
  - If a list, then a random value will be sampled from that list per image.
  - If a `StochasticParameter`, then a value will be sampled per image from that parameter.
- **per_channel** (*bool or float, optional*) – Whether to use the same value for all channels (`False`) or to sample a new value for each channel (`True`). If this value is a float `p`, then for `p` percent of all images *per_channel* will be treated as `True`, otherwise as `False`.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.LinearContrast((0.4, 1.6))
```

Modify the contrast of images according to *127 + alpha*(v-127)*', where `v` is a pixel value and `alpha` is sampled uniformly from the interval `[0.4, 1.6]` (once per image).

```
>>> aug = iaa.LinearContrast((0.4, 1.6), per_channel=True)
```

Same as in the previous example, but `alpha` is sampled once per image *and* channel.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.

Continued on next page

Table 129 – continued from previous page

augment_batches(self, batches[, hooks, ...])	Augment multiple batches.
augment_bounding_boxes(self, ...[, parents, ...])	Augment a batch of bounding boxes.
augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.contrast.**LogContrast**(*gain=(0.4, 1.6), per_channel=False, seed=None, name=None, random_state='deprecated', deterministic='deprecated'*)

Bases: imgaug.augmenters.contrast._ContrastFuncWrapper

Adjust image contrast by scaling pixels to 255*gain*log_2(1+v/255).

This augmenter is fairly similar to imgaug.augmenters.arithmetic.Multiply.

Supported dtypes:

See *adjust_contrast_log()*.

Parameters

- **gain** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Multiplier for the logarithm result. Values around 1.0 lead to a contrast-adjusted images. Values above 1.0 quickly lead to partially broken images due to exceeding the datatype's value range.
  - If a number, then that value will be used for all images.
  - If a tuple (a, b), then a value from the interval [a, b] will uniformly sampled be used per image.
  - If a list, then a random value will be sampled from that list per image.
  - If a StochasticParameter, then a value will be sampled per image from that parameter.
- **per_channel** (*bool or float, optional*) – Whether to use the same value for all channels (False) or to sample a new value for each channel (True). If this value is a float p, then for p percent of all images *per_channel* will be treated as True, otherwise as False.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See __init__().
- **name** (*None or str, optional*) – See __init__().
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method to_deterministic() for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.LogContrast(gain=(0.6, 1.4))
```

Modify the contrast of images according to 255*gain*log_2(1+v/255), where v is a pixel value and gain is sampled uniformly from the interval [0.6, 1.4] (once per image).

```
>>> aug = iaa.LogContrast(gain=(0.6, 1.4), per_channel=True)
```

Same as in the previous example, but `gain` is sampled once per image *and* channel.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.
`get_parameters`(self)	See *get_parameters()*.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`localize_random_state_`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`pool`(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
`remove_augmenters`(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
`remove_augmenters_`(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
`remove_augmenters_inplace`(self, func[, parents])	**Deprecated**.

Continued on next page

Table 130 – continued from previous page

reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.contrast.**SigmoidContrast**(*gain=(5, 6)*, *cutoff=(0.3, 0.6)*, *per_channel=False*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: imgaug.augmenters.contrast._ContrastFuncWrapper

Adjust image contrast to 255*1/(1+exp(gain*(cutoff-I_ij/255))).

Values in the range gain=(5, 20) and cutoff=(0.25, 0.75) seem to be sensible.

A combination of gain=5.5 and cutof=0.45 is fairly close to the identity function.

Supported dtypes:

See *adjust_contrast_sigmoid()*.

Parameters

- **gain** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Multiplier for the sigmoid function's output. Higher values lead to quicker changes from dark to light pixels.
  - If a number, then that value will be used for all images.
  - If a tuple (a, b), then a value from the interval [a, b] will be sampled uniformly per image.
  - If a list, then a random value will be sampled from that list per image.
  - If a StochasticParameter, then a value will be sampled per image from that parameter.
- **cutoff** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Cutoff that shifts the sigmoid function in horizontal direction. Higher values mean that the switch from dark to light pixels happens later, i.e. the pixels will remain darker.
  - If a number, then that value will be used for all images.
  - If a tuple (a, b), then a value from the range [a, b] will be used per image.
  - If a list, then a random value will be sampled from that list per image.
  - If a StochasticParameter, then a value will be sampled per image from that parameter.
- **per_channel** (*bool or float, optional*) – Whether to use the same value for all channels (False) or to sample a new value for each channel (True). If this value is a float p, then for p percent of all images *per_channel* will be treated as True, otherwise as False.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See __init__().

- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.SigmoidContrast(gain=(3, 10), cutoff=(0.4, 0.6))
```

Modify the contrast of images according to `255*1/(1+exp(gain*(cutoff-v/255)))`, where `v` is a pixel value, `gain` is sampled uniformly from the interval `[3, 10]` (once per image) and `cutoff` is sampled uniformly from the interval `[0.4, 0.6]` (also once per image).

```
>>> aug = iaa.SigmoidContrast(
>>>     gain=(3, 10), cutoff=(0.4, 0.6), per_channel=True)
```

Same as in the previous example, but `gain` and `cutoff` are each sampled once per image *and* channel.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).

Continued on next page

Table 131 – continued from previous page

deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

imgaug.augmenters.contrast.**adjust_contrast_gamma**(*arr*, *gamma*)

Adjust image contrast by scaling pixel values to `255*((v/255)**gamma)`.

Supported dtypes:

- `uint8`: yes; fully tested (1) (2) (3)
- `uint16`: yes; tested (2) (3)
- `uint32`: yes; tested (2) (3)
- `uint64`: yes; tested (2) (3) (4)
- `int8`: limited; tested (2) (3) (5)
- `int16`: limited; tested (2) (3) (5)
- `int32`: limited; tested (2) (3) (5)
- `int64`: limited; tested (2) (3) (4) (5)
- `float16`: limited; tested (5)
- `float32`: limited; tested (5)
- `float64`: limited; tested (5)
- `float128`: no (6)
- `bool`: no (7)

- (1) Handled by `cv2`. Other dtypes are handled by `skimage`.
- (2) Normalization is done as `I_ij/max`, where `max` is the maximum value of the dtype, e.g. 255 for `uint8`. The normalization is reversed afterwards, e.g. `result*255` for `uint8`.
- (3) Integer-like values are not rounded after applying the contrast adjustment equation (before inverting the normalization to `[0.0, 1.0]` space), i.e. projection from continuous space to discrete happens according to floor function.
- (4) Note that scikit-image doc says that integers are converted to `float64` values before applying the contrast normalization method. This might lead to inaccuracies for large 64bit integer values. Tests showed no indication of that happening though.
- (5) Must not contain negative values. Values >=0 are fully supported.
- (6) Leads to error in scikit-image.
- (7) Does not make sense for contrast adjustments.

Parameters

- **arr** (*numpy.ndarray*) – Array for which to adjust the contrast. Dtype `uint8` is fastest.
- **gamma** (*number*) – Exponent for the contrast adjustment. Higher values darken the image.

Returns Array with adjusted contrast.

Return type numpy.ndarray

`imgaug.augmenters.contrast.`**`adjust_contrast_linear`**(*arr*, *alpha*)

Adjust contrast by scaling each pixel to `127 + alpha*(v-127)`.

Supported dtypes:

- `uint8`: yes; fully tested (1) (2)
- `uint16`: yes; tested (2)
- `uint32`: yes; tested (2)
- `uint64`: no (3)
- `int8`: yes; tested (2)
- `int16`: yes; tested (2)
- `int32`: yes; tested (2)
- `int64`: no (2)
- `float16`: yes; tested (2)
- `float32`: yes; tested (2)
- `float64`: yes; tested (2)
- `float128`: no (2)
- `bool`: no (4)

- (1) Handled by `cv2`. Other dtypes are handled by raw `numpy`.
- (2) Only tested for reasonable alphas with up to a value of around `100`.
- (3) Conversion to `float64` is done during augmentation, hence `uint64`, `int64`, and `float128` support cannot be guaranteed.

- (4) Does not make sense for contrast adjustments.

Parameters

- **arr** (*numpy.ndarray*) – Array for which to adjust the contrast. Dtype `uint8` is fastest.
- **alpha** (*number*) – Multiplier to linearly pronounce (`>1.0`), dampen (`0.0` to `1.0`) or invert (`<0.0`) the difference between each pixel value and the dtype's center value, e.g. `127` for `uint8`.

Returns Array with adjusted contrast.

Return type numpy.ndarray

`imgaug.augmenters.contrast.`**`adjust_contrast_log`**(*arr*, *gain*)

Adjust image contrast by scaling pixels to `255*gain*log_2(1+v/255)`.

Supported dtypes:

- `uint8`: yes; fully tested (1) (2) (3)
- `uint16`: yes; tested (2) (3)
- `uint32`: no; tested (2) (3) (8)
- `uint64`: no; tested (2) (3) (4) (8)
- `int8`: limited; tested (2) (3) (5)
- `int16`: limited; tested (2) (3) (5)
- `int32`: no; tested (2) (3) (5) (8)
- `int64`: no; tested (2) (3) (4) (5) (8)
- `float16`: limited; tested (5)
- `float32`: limited; tested (5)
- `float64`: limited; tested (5)
- `float128`: no (6)
- `bool`: no (7)

- (1) Handled by `cv2`. Other dtypes are handled by `skimage`.
- (2) Normalization is done as `I_ij/max`, where `max` is the maximum value of the dtype, e.g. 255 for `uint8`. The normalization is reversed afterwards, e.g. `result*255` for `uint8`.
- (3) Integer-like values are not rounded after applying the contrast adjustment equation (before inverting the normalization to `[0.0, 1.0]` space), i.e. projection from continuous space to discrete happens according to floor function.
- (4) Note that scikit-image doc says that integers are converted to `float64` values before applying the contrast normalization method. This might lead to inaccuracies for large 64bit integer values. Tests showed no indication of that happening though.
- (5) Must not contain negative values. Values >=0 are fully supported.
- (6) Leads to error in scikit-image.
- (7) Does not make sense for contrast adjustments.
- (8) No longer supported since numpy 1.17. Previously: 'yes' for `uint32`, `uint64`; 'limited' for `int32`, `int64`.

Parameters

- **arr** (*numpy.ndarray*) – Array for which to adjust the contrast. Dtype `uint8` is fastest.
- **gain** (*number*) – Multiplier for the logarithm result. Values around 1.0 lead to a contrast-adjusted images. Values above 1.0 quickly lead to partially broken images due to exceeding the datatype's value range.

Returns Array with adjusted contrast.

Return type numpy.ndarray

`imgaug.augmenters.contrast.`**`adjust_contrast_sigmoid`**(*arr*, *gain*, *cutoff*)

Adjust image contrast to `255*1/(1+exp(gain*(cutoff-I_ij/255)))`.

Supported dtypes:

- `uint8`: yes; fully tested (1) (2) (3)
- `uint16`: yes; tested (2) (3)
- `uint32`: yes; tested (2) (3)
- `uint64`: yes; tested (2) (3) (4)
- `int8`: limited; tested (2) (3) (5)
- `int16`: limited; tested (2) (3) (5)
- `int32`: limited; tested (2) (3) (5)
- `int64`: limited; tested (2) (3) (4) (5)
- `float16`: limited; tested (5)
- `float32`: limited; tested (5)
- `float64`: limited; tested (5)
- `float128`: no (6)
- `bool`: no (7)

- (1) Handled by `cv2`. Other dtypes are handled by `skimage`.
- (2) Normalization is done as `I_ij/max`, where `max` is the maximum value of the dtype, e.g. 255 for `uint8`. The normalization is reversed afterwards, e.g. `result*255` for `uint8`.
- (3) Integer-like values are not rounded after applying the contrast adjustment equation before inverting the normalization to `[0.0, 1.0]` space), i.e. projection from continuous space to discrete happens according to floor function.
- (4) Note that scikit-image doc says that integers are converted to `float64` values before applying the contrast normalization method. This might lead to inaccuracies for large 64bit integer values. Tests showed no indication of that happening though.
- (5) Must not contain negative values. Values >=0 are fully supported.
- (6) Leads to error in scikit-image.
- (7) Does not make sense for contrast adjustments.

Parameters

- **arr** (*numpy.ndarray*) – Array for which to adjust the contrast. Dtype `uint8` is fastest.

- **gain** (*number*) – Multiplier for the sigmoid function's output. Higher values lead to quicker changes from dark to light pixels.
- **cutoff** (*number*) – Cutoff that shifts the sigmoid function in horizontal direction. Higher values mean that the switch from dark to light pixels happens later, i.e. the pixels will remain darker.

Returns Array with adjusted contrast.

Return type numpy.ndarray

13.25 imgaug.augmenters.convolutional

Augmenters that are based on applying convolution kernels to images.

List of augmenters:

- *Convolve*
- *Sharpen*
- *Emboss*
- *EdgeDetect*
- *DirectedEdgeDetect*

For MotionBlur, see `blur.py`.

class imgaug.augmenters.convolutional.**Convolve**(*matrix=None*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Apply a convolution to input images.

Supported dtypes:

- `uint8`: yes; fully tested
- `uint16`: yes; tested
- `uint32`: no (1)
- `uint64`: no (2)
- `int8`: yes; tested (3)
- `int16`: yes; tested
- `int32`: no (2)
- `int64`: no (2)
- `float16`: yes; tested (4)
- `float32`: yes; tested
- `float64`: yes; tested
- `float128`: no (1)
- `bool`: yes; tested (4)

- (1) rejected by `cv2.filter2D()`.

- (2) causes error: cv2.error: OpenCV(3.4.2) (...)/filter.cpp:4487: error: (-213:The function/feature is not implemented) Unsupported combination of source format (=1), and destination format (=1) in function 'getLinearFilter'.
- (3) mapped internally to `int16`.
- (4) mapped internally to `float32`.

Parameters

- **matrix** (*None or (H, W) ndarray or imgaug.parameters.StochasticParameter or callable, optional*) –

  The weight matrix of the convolution kernel to apply.

  – If `None`, the input images will not be changed.

  – If a 2D numpy array, that array will always be used for all images and channels as the kernel.

  – If a callable, that method will be called for each image via `parameter(image, C, random_state)`. The function must either return a list of `C` matrices (i.e. one per channel) or a 2D numpy array (will be used for all channels) or a 3D `HxWxC` numpy array. If a list is returned, each entry may be `None`, which will result in no changes to the respective channel.

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> matrix = np.array([[0, -1, 0],
>>>                    [-1, 4, -1],
>>>                    [0, -1, 0]])
>>> aug = iaa.Convolve(matrix=matrix)
```

Convolves all input images with the kernel shown in the `matrix` variable.

```
>>> def gen_matrix(image, nb_channels, random_state):
>>>     matrix_A = np.array([[0, -1, 0],
>>>                          [-1, 4, -1],
>>>                          [0, -1, 0]])
>>>     matrix_B = np.array([[0, 1, 0],
>>>                          [1, -4, 1],
>>>                          [0, 1, 0]])
```

(continues on next page)

(continued from previous page)

```
>>>     if image.shape[0] % 2 == 0:
>>>         return [matrix_A] * nb_channels
>>>     else:
>>>         return [matrix_B] * nb_channels
>>> aug = iaa.Convolve(matrix=gen_matrix)
```

Convolves images that have an even height with matrix A and images having an odd height with matrix B.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`localize_random_state_`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`pool`(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.

Continued on next page

Table 132 – continued from previous page

remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.convolutional.**DirectedEdgeDetect**(*alpha=(0.0, 0.75)*, *direction=(0.0, 1.0)*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.convolutional.Convolve*

Detect edges from specified angles and alpha-blend with the input image.

This augmenter first detects edges along a certain angle. Usually, edges are detected in x- or y-direction, while here the edge detection kernel is rotated to match a specified angle. The result of applying the kernel is a black (non-edges) and white (edges) image. That image is alpha-blended with the input image.

Supported dtypes:

See *Convolve*.

Parameters

- **alpha** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Blending factor of the edge image. At `0.0`, only the original image is visible, at `1.0` only the edge image is visible.
  - If a number, exactly that value will always be used.
  - If a tuple `(a, b)`, a random value will be sampled from the interval `[a, b]` per image.
  - If a list, a random value will be sampled from that list per image.
  - If a `StochasticParameter`, a value will be sampled from that parameter per image.
- **direction** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Angle (in degrees) of edges to pronounce, where `0` represents `0` degrees and `1.0` represents 360 degrees (both clockwise, starting at the top). Default value is `(0.0, 1.0)`, i.e. pick a random angle per image.
  - If a number, exactly that value will always be used.
  - If a tuple `(a, b)`, a random value will be sampled from the interval `[a, b]` will be sampled per image.
  - If a list, then a random value will be sampled from that list per image.

– If a `StochasticParameter`, a value will be sampled from the parameter per image.

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.DirectedEdgeDetect(alpha=1.0, direction=0)
```

Turn input images into edge images in which edges are detected from the top side of the image (i.e. the top sides of horizontal edges are part of the edge image, while vertical edges are ignored).

```
>>> aug = iaa.DirectedEdgeDetect(alpha=1.0, direction=90/360)
```

Same as before, but edges are detected from the right. Horizontal edges are now ignored.

```
>>> aug = iaa.DirectedEdgeDetect(alpha=1.0, direction=(0.0, 1.0))
```

Same as before, but edges are detected from a random angle sampled uniformly from the interval `[0deg, 360deg]`.

```
>>> aug = iaa.DirectedEdgeDetect(alpha=(0.0, 0.3), direction=0)
```

Similar to the previous examples, but here the edge image is alpha-blended with the input image. The result is a mixture between the edge image and the input image. The blending factor is randomly sampled between `0%` and `30%`.

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.

Continued on next page

Table 133 – continued from previous page

augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.convolutional.**EdgeDetect**(*alpha=(0.0, 0.75)*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.convolutional.Convolve*

Generate a black & white edge image and alpha-blend it with the input image.

Supported dtypes:

See *Convolve*.

Parameters

- **alpha** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Blending factor of the edge image. At `0.0`, only the original image is visible, at `1.0` only the edge image is visible.
  - If a number, exactly that value will always be used.
  - If a tuple `(a, b)`, a random value will be sampled from the interval `[a, b]` per image.
  - If a list, a random value will be sampled from that list per image.
  - If a `StochasticParameter`, a value will be sampled from that parameter per image.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.EdgeDetect(alpha=(0.0, 1.0))
```

Detect edges in an image, mark them as black (non-edge) and white (edges) and alpha-blend the result with the original input image using a random blending factor between `0%` and `100%`.

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.

Continued on next page

Table 134 – continued from previous page

augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.convolutional.**Emboss**(*alpha=(0.0, 1.0)*, *strength=(0.25, 1.0)*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.convolutional.Convolve*

Emboss images and alpha-blend the result with the original input images.

The embossed version pronounces highlights and shadows, letting the image look as if it was recreated on a metal plate ("embossed").

Supported dtypes:

See *Convolve*.

Parameters

- **alpha** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Blending factor of the embossed image. At `0.0`, only the original image is visible, at `1.0` only its embossed version is visible.
  - If a number, exactly that value will always be used.
  - If a tuple `(a, b)`, a random value will be sampled from the interval `[a, b]` per image.
  - If a list, a random value will be sampled from that list per image.
  - If a `StochasticParameter`, a value will be sampled from that parameter per image.
- **strength** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Parameter that controls the strength of the embossing. Sane values are somewhere in the interval `[0.0, 2.0]` with `1.0` being the standard embossing effect. Default value is `1.0`.
  - If a number, exactly that value will always be used.
  - If a tuple `(a, b)`, a random value will be sampled from the interval `[a, b]` per image.
  - If a list, then a random value will be sampled from that list per image.
  - If a `StochasticParameter`, a value will be sampled from the parameter per image.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.Emboss(alpha=(0.0, 1.0), strength=(0.5, 1.5))
```

Emboss an image with a strength sampled uniformly from the interval `[0.5, 1.5]` and alpha-blend the result with the original input image using a random blending factor between `0%` and `100%`.

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.

Continued on next page

Table 135 – continued from previous page

augment_bounding_boxes(self, ...[, parents, ...])	Augment a batch of bounding boxes.
augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.convolutional.**Sharpen**(*alpha=(0.0, 0.2), lightness=(0.8, 1.2), seed=None, name=None, random_state='deprecated', deterministic='deprecated'*)

Bases: *imgaug.augmenters.convolutional.Convolve*

Sharpen images and alpha-blend the result with the original input images.

Supported dtypes:

See *Convolve*.

Parameters

- **alpha** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Blending factor of the sharpened image. At `0.0`, only the original image is visible, at `1.0` only its sharpened version is visible.
  - If a number, exactly that value will always be used.
  - If a tuple `(a, b)`, a random value will be sampled from the interval `[a, b]` per image.
  - If a list, a random value will be sampled from that list per image.
  - If a `StochasticParameter`, a value will be sampled from that parameter per image.
- **lightness** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Lightness/brightness of the sharped image. Sane values are somewhere in the interval `[0.5, 2.0]`. The value `0.0` results in an edge map. Values higher than `1.0` create bright images. Default value is `1.0`.
  - If a number, exactly that value will always be used.
  - If a tuple `(a, b)`, a random value will be sampled from the interval `[a, b]` per image.
  - If a list, a random value will be sampled from that list per image.
  - If a `StochasticParameter`, a value will be sampled from that parameter per image.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.Sharpen(alpha=(0.0, 1.0))
```

Sharpens input images and blends the sharpened image with the input image using a random blending factor between `0%` and `100%` (uniformly sampled).

```
>>> aug = iaa.Sharpen(alpha=(0.0, 1.0), lightness=(0.75, 2.0))
```

Sharpens input images with a variable *lightness* sampled uniformly from the interval `[0.75, 2.0]` and with a fully random blending factor (as in the above example).

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.
`get_parameters`(self)	See *get_parameters()*.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`localize_random_state_`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`pool`(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.

Continued on next page

Table 136 – continued from previous page

remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

13.26 imgaug.augmenters.debug

Augmenters that help with debugging.

List of augmenters:

- *SaveDebugImageEveryNBatches*

Added in 0.4.0.

class imgaug.augmenters.debug.**SaveDebugImageEveryNBatches**(*destination*, *interval*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: imgaug.augmenters.debug._SaveDebugImage

Visualize data in batches and save corresponding plots to a folder.

Added in 0.4.0.

Supported dtypes:

See *draw_debug_image()*.

Parameters

- **destination** (*str or _IImageDestination*) – Path to a folder. The saved images will follow a filename pattern of batch_<batch_id>.png. The latest image will additionally be saved to latest.png.
- **interval** (*int*) – Interval in batches. If set to N, every N th batch an image will be generated and saved, starting with the first observed batch. Note that the augmenter only counts batches that it sees. If it is executed conditionally or re-instantiated, it may not see all batches or the counter may be wrong in other ways.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See __init__().
- **name** (*None or str, optional*) – See __init__().
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage

will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.

- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> import tempfile
>>> folder_path = tempfile.mkdtemp()
>>> seq = iaa.Sequential([
>>>     iaa.Sequential([
>>>         iaa.Fliplr(0.5),
>>>         iaa.Crop(px=(0, 16))
>>>     ], random_order=True),
>>>     iaa.SaveDebugImageEveryNBatches(folder_path, 100)
>>> ])
```

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.

Continued on next page

Table 137 – continued from previous page

find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	Get the parameters of this augmenter.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)

Get the parameters of this augmenter.

Returns List of parameters of arbitrary types (usually child class of *StochasticParameter*, but not guaranteed to be).

Return type list

imgaug.augmenters.debug.**draw_debug_image**(*images*, *heatmaps=None*, *segmentation_maps=None*, *keypoints=None*, *bounding_boxes=None*, *polygons=None*, *line_strings=None*)

Generate a debug image grid of a single batch and various datatypes.

Added in 0.4.0.

Supported dtypes:

- uint8: yes; tested
- uint16: ?
- uint32: ?
- uint64: ?
- int8: ?
- int16: ?
- int32: ?
- int64: ?

- `float16`: ?
- `float32`: ?
- `float64`: ?
- `float128`: ?
- `bool`: ?

Parameters

- **images** (*ndarray or list of ndarray*) – Images in the batch. Must always be provided. Batches without images cannot be visualized.
- **heatmaps** (*None or list of imgaug.augmentables.heatmaps.HeatmapsOnImage, optional*) – Heatmaps on the provided images.
- **segmentation_maps** (*None or list of imgaug.augmentables.segmaps.SegmentationMapsOnImage, optional*) – Segmentation maps on the provided images.
- **keypoints** (*None or list of imgaug.augmentables.kps.KeypointsOnImage, optional*) – Keypoints on the provided images.
- **bounding_boxes** (*None or list of imgaug.augmentables.bbs.BoundingBoxesOnImage, optional*) – Bounding boxes on the provided images.
- **polygons** (*None or list of imgaug.augmentables.polys.PolygonsOnImage, optional*) – Polygons on the provided images.
- **line_strings** (*None or list of imgaug.augmentables.lines.LineStringsOnImage, optional*) – Line strings on the provided images.

Returns Visualized batch as RGB image.

Return type ndarray

Examples

```
>>> import numpy as np
>>> import imgaug.augmenters as iaa
>>> image = np.zeros((64, 64, 3), dtype=np.uint8)
>>> debug_image = iaa.draw_debug_image(images=[image, image])
```

Generate a debug image for two empty images.

```
>>> from imgaug.augmentables.kps import KeypointsOnImage
>>> kpsoi = KeypointsOnImage.from_xy_array([(10.5, 20.5), (30.5, 30.5)],
>>>                                        shape=image.shape)
>>> debug_image = iaa.draw_debug_image(images=[image, image],
>>>                                    keypoints=[kpsoi, kpsoi])
```

Generate a debug image for two empty images, each having two keypoints drawn on them.

```
>>> from imgaug.augmentables.batches import UnnormalizedBatch
>>> segmap_arr = np.zeros((32, 32, 1), dtype=np.int32)
>>> kp_tuples = [(10.5, 20.5), (30.5, 30.5)]
>>> batch = UnnormalizedBatch(images=[image, image],
>>>                           segmentation_maps=[segmap_arr, segmap_arr],
>>>                           keypoints=[kp_tuples, kp_tuples])
```

(continues on next page)

(continued from previous page)

```
>>> batch = batch.to_normalized_batch()
>>> debug_image = iaa.draw_debug_image(
>>>     images=batch.images_unaug,
>>>     segmentation_maps=batch.segmentation_maps_unaug,
>>>     keypoints=batch.keypoints_unaug)
```

Generate a debug image for two empty images, each having an empty segmentation map and two keypoints drawn on them. This example uses `UnnormalizedBatch` to show how to mostly evade going through imgaug classes.

13.27 imgaug.augmenters.edges

Augmenters that deal with edge detection.

List of augmenters:

- *Canny*

EdgeDetect and *DirectedEdgeDetect* are currently still in `convolutional.py`.

class imgaug.augmenters.edges.**Canny**(*alpha=(0.0, 1.0)*, *hysteresis_thresholds=((60, 140), (160, 240))*, *sobel_kernel_size=(3, 7)*, *colorizer=None*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Apply a canny edge detector to input images.

Supported dtypes:

- `uint8`: yes; fully tested
- `uint16`: no; not tested
- `uint32`: no; not tested
- `uint64`: no; not tested
- `int8`: no; not tested
- `int16`: no; not tested
- `int32`: no; not tested
- `int64`: no; not tested
- `float16`: no; not tested
- `float32`: no; not tested
- `float64`: no; not tested
- `float128`: no; not tested
- `bool`: no; not tested

Parameters

- **alpha** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Blending factor to use in alpha blending. A value close to 1.0 means that only the edge image is visible. A value close to 0.0 means that only the original image is visible. A value close to 0.5 means that the images

are merged according to *0.5*image + 0.5*edge_image*. If a sample from this parameter is 0, no action will be performed for the corresponding image.

  - If an int or float, exactly that value will be used.
  - If a tuple `(a, b)`, a random value from the range `a <= x <= b` will be sampled per image.
  - If a list, then a random value will be sampled from that list per image.
  - If a StochasticParameter, a value will be sampled from the parameter per image.

- **hysteresis_thresholds** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter or tuple of tuple of number or tuple of list of number or tuple of imgaug.parameters.StochasticParameter, optional*) – Min and max values to use in hysteresis thresholding. (This parameter seems to have not very much effect on the results.) Either a single parameter or a tuple of two parameters. If a single parameter is provided, the sampling happens once for all images with *(N,2)* samples being requested from the parameter, where each first value denotes the hysteresis minimum and each second the maximum. If a tuple of two parameters is provided, one sampling of *(N,)* values is independently performed per parameter (first parameter: hysteresis minimum, second: hysteresis maximum).
  - If this is a single number, both min and max value will always be exactly that value.
  - If this is a tuple of numbers `(a, b)`, two random values from the range `a <= x <= b` will be sampled per image.
  - If this is a list, two random values will be sampled from that list per image.
  - If this is a StochasticParameter, two random values will be sampled from that parameter per image.
  - If this is a tuple `(min, max)` with `min` and `max` both *not* being numbers, they will be treated according to the rules above (i.e. may be a number, tuple, list or StochasticParameter). A single value will be sampled per image and parameter.

- **sobel_kernel_size** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter, optional*) – Kernel size of the sobel operator initially applied to each image. This corresponds to `apertureSize` in `cv2.Canny()`. If a sample from this parameter is `<=1`, no action will be performed for the corresponding image. The maximum for this parameter is `7` (inclusive). Higher values are not accepted by OpenCV. If an even value `v` is sampled, it is automatically changed to `v-1`.
  - If this is a single integer, the kernel size always matches that value.
  - If this is a tuple of integers `(a, b)`, a random discrete value will be sampled from the range `a <= x <= b` per image.
  - If this is a list, a random value will be sampled from that list per image.
  - If this is a StochasticParameter, a random value will be sampled from that parameter per image.

- **colorizer** (*None or imgaug.augmenters.edges.IBinaryImageColorizer, optional*) – A strategy to convert binary edge images to color images. If this is `None`, an instance of `RandomColorBinaryImageColorizer` is created, which means that each edge image is converted into an `uint8` image, where edge and non-edge pixels each have a different color that was uniformly randomly sampled from the space of all `uint8` colors.

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.Canny()
```

Create an augmenter that generates random blends between images and their canny edge representations.

```
>>> aug = iaa.Canny(alpha=(0.0, 0.5))
```

Create a canny edge augmenter that generates edge images with a blending factor of max `50%`, i.e. the original (non-edge) image is always at least partially visible.

```
>>> aug = iaa.Canny(
>>>     alpha=(0.0, 0.5),
>>>     colorizer=iaa.RandomColorsBinaryImageColorizer(
>>>         color_true=255,
>>>         color_false=0
>>>     )
>>> )
```

Same as in the previous example, but the edge image always uses the color white for edges and black for the background.

```
>>> aug = iaa.Canny(alpha=(0.5, 1.0), sobel_kernel_size=[3, 7])
```

Create a canny edge augmenter that initially preprocesses images using a sobel filter with kernel size of either `3x3` or `13x13` and alpha-blends with result using a strength of `50%` (both images equally visible) to `100%` (only edge image visible).

```
>>> aug = iaa.Alpha(
>>>     (0.0, 1.0),
>>>     iaa.Canny(alpha=1),
>>>     iaa.MedianBlur(13)
>>> )
```

Create an augmenter that blends a canny edge image with a median-blurred version of the input image. The median blur uses a fixed kernel size of `13x13` pixels.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`localize_random_state_`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`pool`(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
`remove_augmenters`(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
`remove_augmenters_`(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
`remove_augmenters_inplace`(self, func[, parents])	**Deprecated**.
`reseed`(self[, random_state, deterministic_too])	**Deprecated**.
`seed_`(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.

Continued on next page

Table 138 – continued from previous page

show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.edges.**IBinaryImageColorizer**
Bases: object

Interface for classes that convert binary masks to color images.

Methods

colorize(self, image_binary, image_original, ...)	Convert a binary image to a colorized one.

colorize(*self*, *image_binary*, *image_original*, *nth_image*, *random_state*)
Convert a binary image to a colorized one.

Parameters

- **image_binary** (*ndarray*) – Boolean `(H,W)` image.
- **image_original** (*ndarray*) – Original `(H,W,C)` input image.
- **nth_image** (*int*) – Index of the image in the batch.
- **random_state** (*imgaug.random.RNG*) – Random state to use.

Returns Colorized form of *image_binary*.

Return type ndarray

class imgaug.augmenters.edges.**RandomColorsBinaryImageColorizer**(*color_true=(0, 255)*, *color_false=(0, 255)*)
Bases: *imgaug.augmenters.edges.IBinaryImageColorizer*

Colorizer using two randomly sampled foreground/background colors.

Parameters

- **color_true** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter, optional*) – Color of the foreground, i.e. all pixels in binary images that are `True`. This parameter will be queried once per image to generate `(3,)` samples denoting the color. (Note that even for grayscale images three values will be sampled and converted to grayscale according to `0.299*R + 0.587*G + 0.114*B`. This is the same equation that is also used by OpenCV.)
  - If an int, exactly that value will always be used, i.e. every color will be `(v, v, v)` for value `v`.
  - If a tuple `(a, b)`, three random values from the range `a <= x <= b` will be sampled per image.
  - If a list, then three random values will be sampled from that list per image.

– If a StochasticParameter, three values will be sampled from the parameter per image.

- **color_false** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter, optional*) – Analogous to *color_true*, but denotes the color for all pixels that are `False` in the binary input image.

Methods

colorize(self, image_binary, image_original, ...)	Convert a binary image to a colorized one.

colorize(*self*, *image_binary*, *image_original*, *nth_image*, *random_state*)
Convert a binary image to a colorized one.

Parameters

- **image_binary** (*ndarray*) – Boolean `(H,W)` image.
- **image_original** (*ndarray*) – Original `(H,W,C)` input image.
- **nth_image** (*int*) – Index of the image in the batch.
- **random_state** (*imgaug.random.RNG*) – Random state to use.

Returns Colorized form of *image_binary*.

Return type ndarray

13.28 imgaug.augmenters.flip

Augmenters that apply mirroring/flipping operations to images.

List of augmenters:

- *Fliplr*
- *Flipud*

class imgaug.augmenters.flip.**Fliplr**(*p=1*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)
Bases: *imgaug.augmenters.meta.Augmenter*

Flip/mirror input images horizontally.

Note: The default value for the probability is `0.0`. So, to flip *all* input images use `Fliplr(1.0)` and *not* just `Fliplr()`.

Supported dtypes:

See *fliplr()*.

Parameters

- **p** (*number or imgaug.parameters.StochasticParameter, optional*) – Probability of each image to get flipped.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.

- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.Fliplr(0.5)
```

Flip `50` percent of all images horizontally.

```
>>> aug = iaa.Fliplr(1.0)
```

Flip all images horizontally.

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.

Continued on next page

Table 141 – continued from previous page

`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`localize_random_state_`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`pool`(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
`remove_augmenters`(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
`remove_augmenters_`(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
`remove_augmenters_inplace`(self, func[, parents])	**Deprecated**.
`reseed`(self[, random_state, deterministic_too])	**Deprecated**.
`seed_`(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
`show_grid`(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
`to_deterministic`(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.flip.**Flipud**(*p=1*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Flip/mirror input images vertically.

Note: The default value for the probability is `0.0`. So, to flip *all* input images use `Flipud(1.0)` and *not* just `Flipud()`.

Supported dtypes:

See *flipud()*.

Parameters

- **p** (*number or imgaug.parameters.StochasticParameter, optional*) – Probability of each image to get flipped.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.

- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.Flipud(0.5)
```

Flip `50` percent of all images vertically.

```
>>> aug = iaa.Flipud(1.0)
```

Flip all images vertically.

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.

Continued on next page

Table 142 – continued from previous page

find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See *get_parameters()*.

imgaug.augmenters.flip.**HorizontalFlip**(**args*, ***kwargs*)
Alias for Fliplr.

imgaug.augmenters.flip.**VerticalFlip**(**args*, ***kwargs*)
Alias for Flipud.

imgaug.augmenters.flip.**fliplr**(*arr*)
Flip an image-like array horizontally.

Supported dtypes:

- uint8: yes; fully tested
- uint16: yes; fully tested
- uint32: yes; fully tested
- uint64: yes; fully tested
- int8: yes; fully tested
- int16: yes; fully tested
- int32: yes; fully tested
- int64: yes; fully tested
- float16: yes; fully tested
- float32: yes; fully tested
- float64: yes; fully tested

- `float128`: yes; fully tested
- `bool`: yes; fully tested

Parameters **arr** (*ndarray*) – A 2D/3D *(H, W, [C])* image array.

Returns Horizontally flipped array.

Return type ndarray

Examples

```
>>> import numpy as np
>>> import imgaug.augmenters.flip as flip
>>> arr = np.arange(16).reshape((4, 4))
>>> arr_flipped = flip.fliplr(arr)
```

Create a `4x4` array and flip it horizontally.

`imgaug.augmenters.flip.`**`flipud`**(*arr*)

Flip an image-like array vertically.

Supported dtypes:

- `uint8`: yes; fully tested
- `uint16`: yes; fully tested
- `uint32`: yes; fully tested
- `uint64`: yes; fully tested
- `int8`: yes; fully tested
- `int16`: yes; fully tested
- `int32`: yes; fully tested
- `int64`: yes; fully tested
- `float16`: yes; fully tested
- `float32`: yes; fully tested
- `float64`: yes; fully tested
- `float128`: yes; fully tested
- `bool`: yes; fully tested

Parameters **arr** (*ndarray*) – A 2D/3D *(H, W, [C])* image array.

Returns Vertically flipped array.

Return type ndarray

Examples

```
>>> import numpy as np
>>> import imgaug.augmenters.flip as flip
>>> arr = np.arange(16).reshape((4, 4))
>>> arr_flipped = flip.flipud(arr)
```

Create a `4x4` array and flip it vertically.

13.29 imgaug.augmenters.geometric

Augmenters that apply affine or similar transformations.

List of augmenters:

- *Affine*
- *ScaleX*
- *ScaleY*
- *TranslateX*
- *TranslateY*
- *Rotate*
- *ShearX*
- *ShearY*
- *AffineCv2*
- *PiecewiseAffine*
- *PerspectiveTransform*
- *ElasticTransformation*
- *Rot90*
- *WithPolarWarping*
- *Jigsaw*

class imgaug.augmenters.geometric.**Affine**(*scale=None*, *translate_percent=None*, *translate_px=None*, *rotate=None*, *shear=None*, *order=1*, *cval=0*, *mode='constant'*, *fit_output=False*, *backend='auto'*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Augmenter to apply affine transformations to images.

This is mostly a wrapper around the corresponding classes and functions in OpenCV and skimage.

Affine transformations involve:

- Translation ("move" image on the x-/y-axis)
- Rotation
- Scaling ("zoom" in/out)
- Shear (move one side of the image, turning a square into a trapezoid)

All such transformations can create "new" pixels in the image without a defined content, e.g. if the image is translated to the left, pixels are created on the right. A method has to be defined to deal with these pixel values. The parameters *cval* and *mode* of this class deal with this.

Some transformations involve interpolations between several pixels of the input image to generate output pixel values. The parameter *order* deals with the method of interpolation used for this.

Note: While this augmenter supports segmentation maps and heatmaps that have a different size than the corresponding image, it is strongly recommended to use the same aspect ratios. E.g. for an image of shape `(200, 100, 3)`, good segmap/heatmap array shapes also follow a `2:1` ratio and ideally are `(200, 100, C)`, `(100, 50, C)` or `(50, 25, C)`. Otherwise, transformations involving rotations or shearing will produce unaligned outputs. For performance reasons, there is no explicit validation of whether the aspect ratios are similar.

Supported dtypes:

if (backend="skimage", order in [0, 1]):

- `uint8`: yes; tested
- `uint16`: yes; tested
- `uint32`: yes; tested (1)
- `uint64`: no (2)
- `int8`: yes; tested
- `int16`: yes; tested
- `int32`: yes; tested (1)
- `int64`: no (2)
- `float16`: yes; tested
- `float32`: yes; tested
- `float64`: yes; tested
- `float128`: no (2)
- `bool`: yes; tested
- (1) scikit-image converts internally to float64, which might affect the accuracy of large integers. In tests this seemed to not be an issue.
- (2) results too inaccurate

if (backend="skimage", order in [3, 4]):

- `uint8`: yes; tested
- `uint16`: yes; tested
- `uint32`: yes; tested (1)
- `uint64`: no (2)
- `int8`: yes; tested
- `int16`: yes; tested
- `int32`: yes; tested (1)
- `int64`: no (2)
- `float16`: yes; tested

- `float32`: yes; tested
- `float64`: limited; tested (3)
- `float128`: no (2)
- `bool`: yes; tested
- (1) scikit-image converts internally to float64, which might affect the accuracy of large integers. In tests this seemed to not be an issue.
- (2) results too inaccurate
- (3) `NaN` around minimum and maximum of float64 value range

if (backend="skimage", order=5]):

- `uint8`: yes; tested
- `uint16`: yes; tested
- `uint32`: yes; tested (1)
- `uint64`: no (2)
- `int8`: yes; tested
- `int16`: yes; tested
- `int32`: yes; tested (1)
- `int64`: no (2)
- `float16`: yes; tested
- `float32`: yes; tested
- `float64`: limited; not tested (3)
- `float128`: no (2)
- `bool`: yes; tested
- (1) scikit-image converts internally to `float64`, which might affect the accuracy of large integers. In tests this seemed to not be an issue.
- (2) results too inaccurate
- (3) `NaN` around minimum and maximum of float64 value range

if (backend="cv2", order=0):

- `uint8`: yes; tested
- `uint16`: yes; tested
- `uint32`: no (1)
- `uint64`: no (2)
- `int8`: yes; tested
- `int16`: yes; tested
- `int32`: yes; tested
- `int64`: no (2)

- `float16`: yes; tested (3)
- `float32`: yes; tested
- `float64`: yes; tested
- `float128`: no (1)
- `bool`: yes; tested (3)

- (1) rejected by cv2
- (2) changed to `int32` by cv2
- (3) mapped internally to `float32`

if (backend="cv2", order=1):

- `uint8`: yes; fully tested
- `uint16`: yes; tested
- `uint32`: no (1)
- `uint64`: no (2)
- `int8`: yes; tested (3)
- `int16`: yes; tested
- `int32`: no (2)
- `int64`: no (2)
- `float16`: yes; tested (4)
- `float32`: yes; tested
- `float64`: yes; tested
- `float128`: no (1)
- `bool`: yes; tested (4)

- (1) rejected by cv2
- (2) causes cv2 error: `cv2.error: OpenCV(3.4.4) (...)imgwarp.cpp:1805: error: (-215:Assertion failed) ifunc != 0 in function 'remap'`
- (3) mapped internally to `int16`
- (4) mapped internally to `float32`

if (backend="cv2", order=3):

- `uint8`: yes; tested
- `uint16`: yes; tested
- `uint32`: no (1)
- `uint64`: no (2)
- `int8`: yes; tested (3)
- `int16`: yes; tested

- `int32`: no (2)
- `int64`: no (2)
- `float16`: yes; tested (4)
- `float32`: yes; tested
- `float64`: yes; tested
- `float128`: no (1)
- `bool`: yes; tested (4)

- (1) rejected by cv2
- (2) causes cv2 error: `cv2.error: OpenCV(3.4.4) (...)imgwarp.cpp:1805: error: (-215:Assertion failed) ifunc != 0 in function 'remap'`
- (3) mapped internally to `int16`
- (4) mapped internally to `float32`

Parameters

- **scale** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter or dict {"x": number/tuple/list/StochasticParameter, "y": number/tuple/list/StochasticParameter}, optional*) – Scaling factor to use, where `1.0` denotes "no change" and `0.5` is zoomed out to `50` percent of the original size.
  - If a single number, then that value will be used for all images.
  - If a tuple `(a, b)`, then a value will be uniformly sampled per image from the interval `[a, b]`. That value will be used identically for both x- and y-axis.
  - If a list, then a random value will be sampled from that list per image (again, used for both x- and y-axis).
  - If a `StochasticParameter`, then from that parameter a value will be sampled per image (again, used for both x- and y-axis).
  - If a dictionary, then it is expected to have the keys `x` and/or `y`. Each of these keys can have the same values as described above. Using a dictionary allows to set different values for the two axis and sampling will then happen *independently* per axis, resulting in samples that differ between the axes.
- **translate_percent** (*None or number or tuple of number or list of number or imgaug.parameters.StochasticParameter or dict {"x": number/tuple/list/StochasticParameter, "y": number/tuple/list/StochasticParameter}, optional*) – Translation as a fraction of the image height/width (x-translation, y-translation), where `0` denotes "no change" and `0.5` denotes "half of the axis size".
  - If `None` then equivalent to `0.0` unless *translate_px* has a value other than `None`.
  - If a single number, then that value will be used for all images.
  - If a tuple `(a, b)`, then a value will be uniformly sampled per image from the interval `[a, b]`. That sampled fraction value will be used identically for both x- and y-axis.
  - If a list, then a random value will be sampled from that list per image (again, used for both x- and y-axis).

- If a `StochasticParameter`, then from that parameter a value will be sampled per image (again, used for both x- and y-axis).
- If a dictionary, then it is expected to have the keys `x` and/or `y`. Each of these keys can have the same values as described above. Using a dictionary allows to set different values for the two axis and sampling will then happen *independently* per axis, resulting in samples that differ between the axes.

- **translate_px** (*None or int or tuple of int or list of int or imgaug.parameters.StochasticParameter or dict {"x": int/tuple/list/StochasticParameter, "y": int/tuple/list/StochasticParameter}, optional*) –

  Translation in pixels.

  - If `None` then equivalent to `0` unless *translate_percent* has a value other than `None`.
  - If a single int, then that value will be used for all images.
  - If a tuple `(a, b)`, then a value will be uniformly sampled per image from the discrete interval `[a..b]`. That number will be used identically for both x- and y-axis.
  - If a list, then a random value will be sampled from that list per image (again, used for both x- and y-axis).
  - If a `StochasticParameter`, then from that parameter a value will be sampled per image (again, used for both x- and y-axis).
  - If a dictionary, then it is expected to have the keys `x` and/or `y`. Each of these keys can have the same values as described above. Using a dictionary allows to set different values for the two axis and sampling will then happen *independently* per axis, resulting in samples that differ between the axes.

- **rotate** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Rotation in degrees (**NOT** radians), i.e. expected value range is around `[-360, 360]`. Rotation happens around the *center* of the image, not the top left corner as in some other frameworks.
  - If a number, then that value will be used for all images.
  - If a tuple `(a, b)`, then a value will be uniformly sampled per image from the interval `[a, b]` and used as the rotation value.
  - If a list, then a random value will be sampled from that list per image.
  - If a `StochasticParameter`, then this parameter will be used to sample the rotation value per image.

- **shear** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter or dict {"x": int/tuple/list/StochasticParameter, "y": int/tuple/list/StochasticParameter}, optional*) – Shear in degrees (**NOT** radians), i.e. expected value range is around `[-360, 360]`, with reasonable values being in the range of `[-45, 45]`.
  - If a number, then that value will be used for all images as the shear on the x-axis (no shear on the y-axis will be done).
  - If a tuple `(a, b)`, then two value will be uniformly sampled per image from the interval `[a, b]` and be used as the x- and y-shear value.
  - If a list, then two random values will be sampled from that list per image, denoting x- and y-shear.

  - If a `StochasticParameter`, then this parameter will be used to sample the x- and y-shear values per image.
  - If a dictionary, then similar to *translate_percent*, i.e. one `x` key and/or one `y` key are expected, denoting the shearing on the x- and y-axis respectively. The allowed datatypes are again `number`, `tuple (a, b)`, `list` or `StochasticParameter`.

- **order** (*int or iterable of int or imgaug.ALL or imgaug.parameters.StochasticParameter, optional*) –

  Interpolation order to use. Same meaning as in `skimage`:

  - `0`: `Nearest-neighbor`
  - `1`: `Bi-linear` (default)
  - `2`: `Bi-quadratic` (not recommended by skimage)
  - `3`: `Bi-cubic`
  - `4`: `Bi-quartic`
  - `5`: `Bi-quintic`

  Method `0` and `1` are fast, `3` is a bit slower, `4` and `5` are very slow. If the backend is `cv2`, the mapping to OpenCV's interpolation modes is as follows:

  - `0` -> `cv2.INTER_NEAREST`
  - `1` -> `cv2.INTER_LINEAR`
  - `2` -> `cv2.INTER_CUBIC`
  - `3` -> `cv2.INTER_CUBIC`
  - `4` -> `cv2.INTER_CUBIC`

  As datatypes this parameter accepts:

  - If a single `int`, then that order will be used for all images.
  - If a list, then a random value will be sampled from that list per image.
  - If `imgaug.ALL`, then equivalant to list `[0, 1, 3, 4, 5]` in case of `backend=skimage` and otherwise `[0, 1, 3]`.
  - If `StochasticParameter`, then that parameter is queried per image to sample the order value to use.

- **cval** (*number or tuple of number or list of number or imgaug.ALL or imgaug.parameters.StochasticParameter, optional*) – The constant value to use when filling in newly created pixels. (E.g. translating by 1px to the right will create a new 1px-wide column of pixels on the left of the image). The value is only used when *mode=constant*. The expected value range is `[0, 255]` for `uint8` images. It may be a float value.

  - If this is a single number, then that value will be used (e.g. 0 results in black pixels).
  - If a tuple `(a, b)`, then three values (for three image channels) will be uniformly sampled per image from the interval `[a, b]`.
  - If a list, then a random value will be sampled from that list per image.
  - If `imgaug.ALL` then equivalent to tuple ‘‘(0, 255)‘.
  - If a `StochasticParameter`, a new value will be sampled from the parameter per image.

- **mode** (*str or list of str or imgaug.ALL or imgaug.parameters.StochasticParameter, optional*) – Method to use when filling in newly created pixels. Same meaning as in `skimage` (and `numpy.pad()`):
  - `constant`: Pads with a constant value
  - `edge`: Pads with the edge values of array
  - `symmetric`: Pads with the reflection of the vector mirrored along the edge of the array.
  - `reflect`: Pads with the reflection of the vector mirrored on the first and last values of the vector along each axis.
  - `wrap`: Pads with the wrap of the vector along the axis. The first values are used to pad the end and the end values are used to pad the beginning.

  If `cv2` is chosen as the backend the mapping is as follows:
  - `constant` -> `cv2.BORDER_CONSTANT`
  - `edge` -> `cv2.BORDER_REPLICATE`
  - `symmetric` -> `cv2.BORDER_REFLECT`
  - `reflect` -> `cv2.BORDER_REFLECT_101`
  - `wrap` -> `cv2.BORDER_WRAP`

  The datatype of the parameter may be:
  - If a single string, then that mode will be used for all images.
  - If a list of strings, then a random mode will be picked from that list per image.
  - If `imgaug.ALL`, then a random mode from all possible modes will be picked.
  - If `StochasticParameter`, then the mode will be sampled from that parameter per image, i.e. it must return only the above mentioned strings.
- **fit_output** (*bool, optional*) – Whether to modify the affine transformation so that the whole output image is always contained in the image plane (`True`) or accept parts of the image being outside the image plane (`False`). This can be thought of as first applying the affine transformation and then applying a second transformation to "zoom in" on the new image so that it fits the image plane, This is useful to avoid corners of the image being outside of the image plane after applying rotations. It will however negate translation and scaling. Note also that activating this may lead to image sizes differing from the input image sizes. To avoid this, wrap `Affine` in *KeepSizeByResize*, e.g. `KeepSizeByResize(Affine(...))`.
- **backend** (*str, optional*) – Framework to use as a backend. Valid values are `auto`, `skimage` (scikit-image's warp) and `cv2` (OpenCV's warp). If `auto` is used, the augmenter will automatically try to use `cv2` whenever possible (order must be in `[0, 1, 3]`). It will silently fall back to skimage if order/dtype is not supported by cv2. cv2 is generally faster than skimage. It also supports RGB cvals, while skimage will resort to intensity cvals (i.e. 3x the same value as RGB). If `cv2` is chosen and order is `2` or `4`, it will automatically fall back to order `3`.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.

- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.Affine(scale=2.0)
```

Zoom in on all images by a factor of `2`.

```
>>> aug = iaa.Affine(translate_px=16)
```

Translate all images on the x- and y-axis by 16 pixels (towards the bottom right) and fill up any new pixels with zero (black values).

```
>>> aug = iaa.Affine(translate_percent=0.1)
```

Translate all images on the x- and y-axis by `10` percent of their width/height (towards the bottom right). The pixel values are computed per axis based on that axis' size. Fill up any new pixels with zero (black values).

```
>>> aug = iaa.Affine(rotate=35)
```

Rotate all images by `35` *degrees*. Fill up any new pixels with zero (black values).

```
>>> aug = iaa.Affine(shear=15)
```

Shear all images by `15` *degrees*. Fill up any new pixels with zero (black values).

```
>>> aug = iaa.Affine(translate_px=(-16, 16))
```

Translate all images on the x- and y-axis by a random value between `-16` and `16` pixels (to the bottom right) and fill up any new pixels with zero (black values). The translation value is sampled once per image and is the same for both axis.

```
>>> aug = iaa.Affine(translate_px={"x": (-16, 16), "y": (-4, 4)})
```

Translate all images on the x-axis by a random value between `-16` and `16` pixels (to the right) and on the y-axis by a random value between `-4` and `4` pixels to the bottom. The sampling happens independently per axis, so even if both intervals were identical, the sampled axis-wise values would likely be different. This also fills up any new pixels with zero (black values).

```
>>> aug = iaa.Affine(scale=2.0, order=[0, 1])
```

Same as in the above *scale* example, but uses (randomly) either nearest neighbour interpolation or linear interpolation. If *order* is not specified, `order=1` would be used by default.

```
>>> aug = iaa.Affine(translate_px=16, cval=(0, 255))
```

Same as in the *translate_px* example above, but newly created pixels are now filled with a random color (sampled once per image and the same for all newly created pixels within that image).

```
>>> aug = iaa.Affine(translate_px=16, mode=["constant", "edge"])
```

Similar to the previous example, but the newly created pixels are filled with black pixels in half of all images (mode `constant` with default *cval* being `0`) and in the other half of all images using `edge` mode, which repeats the color of the spatially closest pixel of the corresponding image edge.

```
>>> aug = iaa.Affine(shear={"y": (-45, 45)})
```

Shear images only on the y-axis. Set *shear* to `shear=(-45, 45)` to shear randomly on both axes, using for each image the same sample for both the x- and y-axis. Use `shear={"x": (-45, 45), "y": (-45, 45)}` to get independent samples per axis.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.

Continued on next page

Table 143 – continued from previous page

get_parameters(self)	See get_parameters().
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See get_parameters().

class imgaug.augmenters.geometric.**AffineCv2**(*scale=1.0*, *translate_percent=None*, *translate_px=None*, *rotate=0.0*, *shear=0.0*, *order=<MagicMock id='140201027995912'>*, *cval=0*, *mode=<MagicMock id='140201028024808'>*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: imgaug.augmenters.meta.Augmenter

Deprecated. Augmenter to apply affine transformations to images using cv2 (i.e. opencv) backend.

Warning: This augmenter is deprecated since 0.4.0. Use `Affine(..., backend='cv2')` instead.

Affine transformations involve:

- Translation ("move" image on the x-/y-axis)
- Rotation
- Scaling ("zoom" in/out)
- Shear (move one side of the image, turning a square into a trapezoid)

All such transformations can create "new" pixels in the image without a defined content, e.g. if the image is translated to the left, pixels are created on the right. A method has to be defined to deal with these pixel values. The parameters *cval* and *mode* of this class deal with this.

Some transformations involve interpolations between several pixels of the input image to generate output pixel values. The parameter *order* deals with the method of interpolation used for this.

Deprecated since 0.4.0.

Supported dtypes:

- uint8: yes; fully tested
- uint16: ?
- uint32: ?
- uint64: ?
- int8: ?
- int16: ?
- int32: ?
- int64: ?
- float16: ?
- float32: ?
- float64: ?
- float128: ?
- bool: ?

Parameters

- **scale** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter or dict {"x": number/tuple/list/StochasticParameter, "y": number/tuple/list/StochasticParameter}, optional*) – Scaling factor to use, where 1.0 denotes "no change" and 0.5 is zoomed out to 50 percent of the original size.
  - If a single number, then that value will be used for all images.
  - If a tuple (a, b), then a value will be uniformly sampled per image from the interval [a, b]. That value will be used identically for both x- and y-axis.
  - If a list, then a random value will be sampled from that list per image (again, used for both x- and y-axis).
  - If a StochasticParameter, then from that parameter a value will be sampled per image (again, used for both x- and y-axis).
  - If a dictionary, then it is expected to have the keys x and/or y. Each of these keys can have the same values as described above. Using a dictionary allows to set different values for the two axis and sampling will then happen *independently* per axis, resulting in samples that differ between the axes.
- **translate_percent** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter or dict {"x": number/tuple/list/StochasticParameter, "y": number/tuple/list/StochasticParameter}, optional*) – Translation as a fraction of the image height/width (x-translation, y-translation), where 0 denotes "no change" and 0.5 denotes "half of the axis size".
  - If None then equivalent to 0.0 unless *translate_px* has a value other than None.
  - If a single number, then that value will be used for all images.
  - If a tuple (a, b), then a value will be uniformly sampled per image from the interval [a, b]. That sampled fraction value will be used identically for both x- and y-axis.
  - If a list, then a random value will be sampled from that list per image (again, used for both x- and y-axis).

  - If a `StochasticParameter`, then from that parameter a value will be sampled per image (again, used for both x- and y-axis).
  - If a dictionary, then it is expected to have the keys `x` and/or `y`. Each of these keys can have the same values as described above. Using a dictionary allows to set different values for the two axis and sampling will then happen *independently* per axis, resulting in samples that differ between the axes.

- **translate_px** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter or dict {"x": int/tuple/list/StochasticParameter, "y": int/tuple/list/StochasticParameter}, optional*) –

  Translation in pixels.

  - If `None` then equivalent to `0` unless *translate_percent* has a value other than `None`.
  - If a single int, then that value will be used for all images.
  - If a tuple `(a, b)`, then a value will be uniformly sampled per image from the discrete interval `[a..b]`. That number will be used identically for both x- and y-axis.
  - If a list, then a random value will be sampled from that list per image (again, used for both x- and y-axis).
  - If a `StochasticParameter`, then from that parameter a value will be sampled per image (again, used for both x- and y-axis).
  - If a dictionary, then it is expected to have the keys `x` and/or `y`. Each of these keys can have the same values as described above. Using a dictionary allows to set different values for the two axis and sampling will then happen *independently* per axis, resulting in samples that differ between the axes.

- **rotate** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Rotation in degrees (**NOT** radians), i.e. expected value range is around `[-360, 360]`. Rotation happens around the *center* of the image, not the top left corner as in some other frameworks.
  - If a number, then that value will be used for all images.
  - If a tuple `(a, b)`, then a value will be uniformly sampled per image from the interval `[a, b]` and used as the rotation value.
  - If a list, then a random value will be sampled from that list per image.
  - If a `StochasticParameter`, then this parameter will be used to sample the rotation value per image.

- **shear** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Shear in degrees (**NOT** radians), i.e. expected value range is around `[-360, 360]`.
  - If a number, then that value will be used for all images.
  - If a tuple `(a, b)`, then a value will be uniformly sampled per image from the interval `[a, b]` and be used as the rotation value.
  - If a list, then a random value will be sampled from that list per image.
  - If a `StochasticParameter`, then this parameter will be used to sample the shear value per image.

- **order** (*int or list of int or str or list of str or imaug.ALL or imgaug.parameters.StochasticParameter, optional*) –

Interpolation order to use. Allowed are:

- `cv2.INTER_NEAREST` (nearest-neighbor interpolation)
- `cv2.INTER_LINEAR` (bilinear interpolation, used by default)
- **`cv2.INTER_CUBIC` (bicubic interpolation over `4x4` pixel** neighborhood)
- `cv2.INTER_LANCZOS4`
- string `nearest` (same as `cv2.INTER_NEAREST`)
- string `linear` (same as `cv2.INTER_LINEAR`)
- string `cubic` (same as `cv2.INTER_CUBIC`)
- string `lanczos4` (same as `cv2.INTER_LANCZOS`)

`INTER_NEAREST` (nearest neighbour interpolation) and `INTER_NEAREST` (linear interpolation) are the fastest.

- If a single `int`, then that order will be used for all images.
- If a string, then it must be one of: `nearest`, `linear`, `cubic`, `lanczos4`.
- If an iterable of `int`/`str`, then for each image a random value will be sampled from that iterable (i.e. list of allowed order values).
- If `imgaug.ALL`, then equivalant to list `[cv2.INTER_NEAREST, cv2.INTER_LINEAR, cv2.INTER_CUBIC, cv2.INTER_LANCZOS4]`.
- If `StochasticParameter`, then that parameter is queried per image to sample the order value to use.

- **cval** (*number or tuple of number or list of number or imaug.ALL or imgaug.parameters.StochasticParameter, optional*) – The constant value to use when filling in newly created pixels. (E.g. translating by 1px to the right will create a new 1px-wide column of pixels on the left of the image). The value is only used when *mode=constant*. The expected value range is `[0, 255]` for `uint8` images. It may be a float value.
  - If this is a single number, then that value will be used (e.g. 0 results in black pixels).
  - If a tuple `(a, b)`, then three values (for three image channels) will be uniformly sampled per image from the interval `[a, b]`.
  - If a list, then a random value will be sampled from that list per image.
  - If `imgaug.ALL` then equivalent to tuple ‘‘(0, 255)‘.
  - If a `StochasticParameter`, a new value will be sampled from the parameter per image.
- **mode** (*int or str or list of str or list of int or imgaug.ALL or imgaug.parameters.StochasticParameter,*) – optional Method to use when filling in newly created pixels. Same meaning as in OpenCV’s border mode. Let `abcdefgh` be an image’s content and `|` be an image boundary after which new pixels are filled in, then the valid modes and their behaviour are the following:
  - `cv2.BORDER_REPLICATE`: `aaaaaa|abcdefgh|hhhhhhh`
  - `cv2.BORDER_REFLECT`: `fedcba|abcdefgh|hgfedcb`
  - `cv2.BORDER_REFLECT_101`: `gfedcb|abcdefgh|gfedcba`
  - `cv2.BORDER_WRAP`: `cdefgh|abcdefgh|abcdefg`

- **cv2.BORDER_CONSTANT: iiiiii|abcdefgh|iiiiiii,** where `i` is the defined cval.
- `replicate`: Same as `cv2.BORDER_REPLICATE`.
- `reflect`: Same as `cv2.BORDER_REFLECT`.
- `reflect_101`: Same as `cv2.BORDER_REFLECT_101`.
- `wrap`: Same as `cv2.BORDER_WRAP`.
- `constant`: Same as `cv2.BORDER_CONSTANT`.

The datatype of the parameter may be:

- If a single `int`, then it must be one of the `cv2.BORDER_*` constants.
- If a single string, then it must be one of: `replicate`, `reflect`, `reflect_101`, `wrap`, `constant`.
- If a list of `int`/`str`, then per image a random mode will be picked from that list.
- If `imgaug.ALL`, then a random mode from all possible modes will be picked.
- If `StochasticParameter`, then the mode will be sampled from that parameter per image, i.e. it must return only the above mentioned strings.

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.AffineCv2(scale=2.0)
```

Zoom in on all images by a factor of `2`.

```
>>> aug = iaa.AffineCv2(translate_px=16)
```

Translate all images on the x- and y-axis by 16 pixels (towards the bottom right) and fill up any new pixels with zero (black values).

```
>>> aug = iaa.AffineCv2(translate_percent=0.1)
```

Translate all images on the x- and y-axis by `10` percent of their width/height (towards the bottom right). The pixel values are computed per axis based on that axis' size. Fill up any new pixels with zero (black values).

```
>>> aug = iaa.AffineCv2(rotate=35)
```

Rotate all images by `35` *degrees*. Fill up any new pixels with zero (black values).

```
>>> aug = iaa.AffineCv2(shear=15)
```

Shear all images by `15` *degrees*. Fill up any new pixels with zero (black values).

```
>>> aug = iaa.AffineCv2(translate_px=(-16, 16))
```

Translate all images on the x- and y-axis by a random value between `-16` and `16` pixels (to the bottom right) and fill up any new pixels with zero (black values). The translation value is sampled once per image and is the same for both axis.

```
>>> aug = iaa.AffineCv2(translate_px={"x": (-16, 16), "y": (-4, 4)})
```

Translate all images on the x-axis by a random value between `-16` and `16` pixels (to the right) and on the y-axis by a random value between `-4` and `4` pixels to the bottom. The sampling happens independently per axis, so even if both intervals were identical, the sampled axis-wise values would likely be different. This also fills up any new pixels with zero (black values).

```
>>> aug = iaa.AffineCv2(scale=2.0, order=[0, 1])
```

Same as in the above *scale* example, but uses (randomly) either nearest neighbour interpolation or linear interpolation. If *order* is not specified, `order=1` would be used by default.

```
>>> aug = iaa.AffineCv2(translate_px=16, cval=(0, 255))
```

Same as in the *translate_px* example above, but newly created pixels are now filled with a random color (sampled once per image and the same for all newly created pixels within that image).

```
>>> aug = iaa.AffineCv2(translate_px=16, mode=["constant", "replicate"])
```

Similar to the previous example, but the newly created pixels are filled with black pixels in half of all images (mode `constant` with default *cval* being `0`) and in the other half of all images using `replicate` mode, which repeats the color of the spatially closest pixel of the corresponding image edge.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.

Continued on next page

Table 144 – continued from previous page

augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.geometric.**ElasticTransformation**(*alpha=(0.0, 40.0)*, *sigma=(4.0, 8.0)*, *order=3*, *cval=0*, *mode='constant'*, *polygon_recoverer='auto'*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Transform images by moving pixels locally around using displacement fields.

The augmenter has the parameters `alpha` and `sigma`. `alpha` controls the strength of the displacement: higher values mean that pixels are moved further. `sigma` controls the smoothness of the displacement: higher values lead to smoother patterns – as if the image was below water – while low values will cause indivdual pixels to be moved very differently from their neighbours, leading to noisy and pixelated images.

A relation of 10:1 seems to be good for `alpha` and `sigma`, e.g. `alpha=10` and `sigma=1` or `alpha=50`, `sigma=5`. For `128x128` a setting of `alpha=(0, 70.0)`, `sigma=(4.0, 6.0)` may be a good choice and will lead to a water-like effect.

Code here was initially inspired by https://gist.github.com/chsasank/4d8f68caf01f041a6453e67fb30f8f5a

For a detailed explanation, see

```
Simard, Steinkraus and Platt
Best Practices for Convolutional Neural Networks applied to Visual
Document Analysis
in Proc. of the International Conference on Document Analysis and
Recognition, 2003
```

Note: For coordinate-based inputs (keypoints, bounding boxes, polygons, ...), this augmenter still has to perform an image-based augmentation, which will make it significantly slower for such inputs than other augmenters. See *Performance*.

Supported dtypes:

- `uint8`: yes; fully tested (1)
- `uint16`: yes; tested (1)
- `uint32`: yes; tested (2)
- `uint64`: limited; tested (3)
- `int8`: yes; tested (1) (4) (5)
- `int16`: yes; tested (4) (6)
- `int32`: yes; tested (4) (6)
- `int64`: limited; tested (3)
- `float16`: yes; tested (1)
- `float32`: yes; tested (1)
- `float64`: yes; tested (1)
- `float128`: no
- `bool`: yes; tested (1) (7)

- (1) Always handled by `cv2`.
- (2) Always handled by `scipy`.
- (3) Only supported for `order != 0`. Will fail for `order=0`.
- (4) Mapped internally to `float64` when `order=1`.
- (5) Mapped internally to `int16` when `order>=2`.

- (6) Handled by `cv2` when `order=0` or `order=1`, otherwise by `scipy`.
- (7) Mapped internally to `float32`.

Parameters

- **alpha** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Strength of the distortion field. Higher values mean that pixels are moved further with respect to the distortion field's direction. Set this to around 10 times the value of *sigma* for visible effects.
  - If number, then that value will be used for all images.
  - If tuple `(a, b)`, then a random value will be uniformly sampled per image from the interval `[a, b]`.
  - If a list, then for each image a random value will be sampled from that list.
  - If `StochasticParameter`, then that parameter will be used to sample a value per image.
- **sigma** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Standard deviation of the gaussian kernel used to smooth the distortion fields. Higher values (for `128x128` images around 5.0) lead to more water-like effects, while lower values (for `128x128` images around `1.0` and lower) lead to more noisy, pixelated images. Set this to around 1/10th of *alpha* for visible effects.
  - If number, then that value will be used for all images.
  - If tuple `(a, b)`, then a random value will be uniformly sampled per image from the interval `[a, b]`.
  - If a list, then for each image a random value will be sampled from that list.
  - If `StochasticParameter`, then that parameter will be used to sample a value per image.
- **order** (*int or list of int or imaug.ALL or imgaug.parameters.StochasticParameter, optional*) – Interpolation order to use. Same meaning as in `scipy.ndimage.map_coordinates()` and may take any integer value in the range `0` to `5`, where orders close to `0` are faster.
  - If a single int, then that order will be used for all images.
  - If a tuple `(a, b)`, then a random value will be uniformly sampled per image from the interval `[a, b]`.
  - If a list, then for each image a random value will be sampled from that list.
  - If `imgaug.ALL`, then equivalant to list `[0, 1, 2, 3, 4, 5]`.
  - If `StochasticParameter`, then that parameter is queried per image to sample the order value to use.
- **cval** (*number or tuple of number or list of number or imgaug.ALL or imgaug.parameters.StochasticParameter, optional*) – The constant intensity value used to fill in new pixels. This value is only used if *mode* is set to `constant`. For standard `uint8` images (value range `0` to `255`), this value may also should also be in the range `0` to `255`. It may be a `float` value, even for images with integer dtypes.
  - If this is a single number, then that value will be used (e.g. `0` results in black pixels).

  - If a tuple `(a, b)`, then a random value will be uniformly sampled per image from the interval `[a, b]`.
  - If a list, then a random value will be picked from that list per image.
  - If `imgaug.ALL`, a value from the discrete range `[0..255]` will be sampled per image.
  - If a `StochasticParameter`, a new value will be sampled from the parameter per image.

- **mode** (*str or list of str or imgaug.ALL or imgaug.parameters.StochasticParameter, optional*) – Parameter that defines the handling of newly created pixels. May take the same values as in `scipy.ndimage.map_coordinates()`, i.e. `constant`, `nearest`, `reflect` or `wrap`.
  - If a single string, then that mode will be used for all images.
  - If a list of strings, then per image a random mode will be picked from that list.
  - If `imgaug.ALL`, then a random mode from all possible modes will be picked.
  - If `StochasticParameter`, then the mode will be sampled from that parameter per image, i.e. it must return only the above mentioned strings.
- **polygon_recoverer** (*'auto' or None or imgaug.augmentables.polygons._ConcavePolygonRecoverer, optional*) – The class to use to repair invalid polygons. If `"auto"`, a new instance of :class'imgaug.augmentables.polygons._ConcavePolygonRecoverer' will be created. If `None`, no polygon recoverer will be used. If an object, then that object will be used and must provide a `recover_from()` method, similar to `_ConcavePolygonRecoverer`.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.ElasticTransformation(alpha=50.0, sigma=5.0)
```

Apply elastic transformations with a strength/alpha of `50.0` and smoothness of `5.0` to all images.

```
>>> aug = iaa.ElasticTransformation(alpha=(0.0, 70.0), sigma=5.0)
```

Apply elastic transformations with a strength/alpha that comes from the interval `[0.0, 70.0]` (randomly picked per image) and with a smoothness of `5.0`.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`localize_random_state_`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`pool`(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
`remove_augmenters`(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
`remove_augmenters_`(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
`remove_augmenters_inplace`(self, func[, parents])	**Deprecated**.
`reseed`(self[, random_state, deterministic_too])	**Deprecated**.
`seed_`(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.

Continued on next page

Table 145 – continued from previous page

show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

KEYPOINT_AUG_ALPHA_THRESH = 0.05

KEYPOINT_AUG_SIGMA_THRESH = 1.0

NB_NEIGHBOURING_KEYPOINTS = 3

NEIGHBOURING_KEYPOINTS_DISTANCE = 1.0

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.geometric.**Jigsaw**(*nb_rows=(3, 10)*, *nb_cols=(3, 10)*, *max_steps=1*, *allow_pad=True*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Move cells within images similar to jigsaw patterns.

Note: This augmenter will by default pad images until their height is a multiple of *nb_rows*. Analogous for *nb_cols*.

Note: This augmenter will resize heatmaps and segmentation maps to the image size, then apply similar padding as for the corresponding images and resize back to the original map size. That also means that images may change in shape (due to padding), but heatmaps/segmaps will not change. For heatmaps/segmaps, this deviates from pad augmenters that will change images and heatmaps/segmaps in corresponding ways and then keep the heatmaps/segmaps at the new size.

Warning: This augmenter currently only supports augmentation of images, heatmaps, segmentation maps and keypoints. Other augmentables, i.e. bounding boxes, polygons and line strings, will result in errors.

Added in 0.4.0.

Supported dtypes:

See *apply_jigsaw()*.

Parameters

- **nb_rows** (*int or list of int or tuple of int or imgaug.parameters.StochasticParameter, optional*) –

  How many rows the jigsaw pattern should have.

  – If a single `int`, then that value will be used for all images.

  – If a tuple `(a, b)`, then a random value will be uniformly sampled per image from the discrete interval `[a..b]`.

  – If a list, then for each image a random value will be sampled from that list.

  - If `StochasticParameter`, then that parameter is queried per image to sample the value to use.

- **nb_cols** (*int or list of int or tuple of int or imgaug.parameters.StochasticParameter, optional*) –

  How many cols the jigsaw pattern should have.

  - If a single `int`, then that value will be used for all images.
  - If a tuple `(a, b)`, then a random value will be uniformly sampled per image from the discrete interval `[a..b]`.
  - If a list, then for each image a random value will be sampled from that list.
  - If `StochasticParameter`, then that parameter is queried per image to sample the value to use.

- **max_steps** (*int or list of int or tuple of int or imgaug.parameters.StochasticParameter, optional*) –

  How many steps each jigsaw cell may be moved.

  - If a single `int`, then that value will be used for all images.
  - If a tuple `(a, b)`, then a random value will be uniformly sampled per image from the discrete interval `[a..b]`.
  - If a list, then for each image a random value will be sampled from that list.
  - If `StochasticParameter`, then that parameter is queried per image to sample the value to use.

- **allow_pad** (*bool, optional*) – Whether to allow automatically padding images until they are evenly divisible by `nb_rows` and `nb_cols`.

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.

- **name** (*None or str, optional*) – See `__init__()`.

- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.

- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.Jigsaw(nb_rows=10, nb_cols=10)
```

Create a jigsaw augmenter that splits images into `10x10` cells and shifts them around by `0` to `2` steps (default setting).

```
>>> aug = iaa.Jigsaw(nb_rows=(1, 4), nb_cols=(1, 4))
```

Create a jigsaw augmenter that splits each image into `1` to `4` cells along each axis.

```
>>> aug = iaa.Jigsaw(nb_rows=10, nb_cols=10, max_steps=(1, 5))
```

Create a jigsaw augmenter that moves the cells in each image by a random amount between `1` and `5` times (decided per image). Some images will be barely changed, some will be fairly distorted.

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.
`get_parameters`(self)	Get the parameters of this augmenter.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`localize_random_state_`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`pool`(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
`remove_augmenters`(self, func[, copy, ...])	Remove this augmenter or children that match a condition.

Continued on next page

Table 146 – continued from previous page

remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
Get the parameters of this augmenter.

Returns List of parameters of arbitrary types (usually child class of *StochasticParameter*, but not guaranteed to be).

Return type list

class imgaug.augmenters.geometric.**PerspectiveTransform**(*scale=(0.0, 0.06), cval=0, mode='constant', keep_size=True, fit_output=False, polygon_recoverer='auto', seed=None, name=None, random_state='deprecated', deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Apply random four point perspective transformations to images.

Each of the four points is placed on the image using a random distance from its respective corner. The distance is sampled from a normal distribution. As a result, most transformations don't change the image very much, while some "focus" on polygons far inside the image.

The results of this augmenter have some similarity with Crop.

Code partially from http://www.pyimagesearch.com/2014/08/25/4-point-opencv-getperspective-transform-example/

Supported dtypes:

if (keep_size=False):

- uint8: yes; fully tested
- uint16: yes; tested
- uint32: no (1)
- uint64: no (2)
- int8: yes; tested (3)
- int16: yes; tested
- int32: no (2)
- int64: no (2)
- float16: yes; tested (4)

- `float32`: yes; tested
- `float64`: yes; tested
- `float128`: no (1)
- `bool`: yes; tested (4)
- (1) rejected by opencv
- (2) leads to opencv error: cv2.error: `OpenCV(3.4.4) (...)imgwarp.cpp:1805: error: (-215:Assertion failed) ifunc != 0 in function 'remap'`.
- (3) mapped internally to `int16`.
- (4) mapped intenally to `float32`.

if (keep_size=True):

minimum of (`imgaug.augmenters.geometric.PerspectiveTransform(keep_size=False)`, *imresize_many_images()*

)

Parameters

- **scale** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Standard deviation of the normal distributions. These are used to sample the random distances of the subimage's corners from the full image's corners. The sampled values reflect percentage values (with respect to image height/width). Recommended values are in the range `0.0` to `0.1`.
  - If a single number, then that value will always be used as the scale.
  - If a tuple `(a, b)` of numbers, then a random value will be uniformly sampled per image from the interval `(a, b)`.
  - If a list of values, a random value will be picked from the list per image.
  - If a `StochasticParameter`, then that parameter will be queried to draw one value per image.
- **keep_size** (*bool, optional*) – Whether to resize image's back to their original size after applying the perspective transform. If set to `False`, the resulting images may end up having different shapes and will always be a list, never an array.
- **cval** (*number or tuple of number or list of number or imaug.ALL or imgaug.parameters.StochasticParameter, optional*) – The constant value used to fill up pixels in the result image that didn't exist in the input image (e.g. when translating to the left, some new pixels are created at the right). Such a fill-up with a constant value only happens, when *mode* is `constant`. The expected value range is `[0, 255]` for `uint8` images. It may be a float value.
  - If this is a single int or float, then that value will be used (e.g. 0 results in black pixels).
  - If a tuple `(a, b)`, then a random value is uniformly sampled per image from the interval `[a, b]`.
  - If a list, then a random value will be sampled from that list per image.
  - If `imgaug.ALL`, then equivalent to tuple `(0, 255)`.

  - If a `StochasticParameter`, a new value will be sampled from the parameter per image.

- **mode** (*int or str or list of str or list of int or imgaug.ALL or imgaug.parameters.StochasticParameter, optional*) – Parameter that defines the handling of newly created pixels. Same meaning as in OpenCV's border mode. Let `abcdefgh` be an image's content and `|` be an image boundary, then:

  - `cv2.BORDER_REPLICATE`: `aaaaaa|abcdefgh|hhhhhhh`
  - `cv2.BORDER_CONSTANT`: `iiiiii|abcdefgh|iiiiiii`, where `i` is the defined cval.
  - `replicate`: Same as `cv2.BORDER_REPLICATE`.
  - `constant`: Same as `cv2.BORDER_CONSTANT`.

  The datatype of the parameter may be:

  - If a single `int`, then it must be one of `cv2.BORDER_*`.
  - If a single string, then it must be one of: `replicate`, `reflect`, `reflect_101`, `wrap`, `constant`.
  - If a list of ints/strings, then per image a random mode will be picked from that list.
  - If `imgaug.ALL`, then a random mode from all possible modes will be picked per image.
  - If `StochasticParameter`, then the mode will be sampled from that parameter per image, i.e. it must return only the above mentioned strings.

- **fit_output** (*bool, optional*) – If `True`, the image plane size and position will be adjusted to still capture the whole image after perspective transformation. (Followed by image resizing if *keep_size* is set to `True`.) Otherwise, parts of the transformed image may be outside of the image plane. This setting should not be set to `True` when using large *scale* values as it could lead to very large images.

  Added in 0.4.0.

- **polygon_recoverer** (*'auto' or None or imgaug.augmentables.polygons._ConcavePolygonRecoverer, optional*) – The class to use to repair invalid polygons. If `"auto"`, a new instance of :class'imgaug.augmentables.polygons._ConcavePolygonRecoverer' will be created. If `None`, no polygon recoverer will be used. If an object, then that object will be used and must provide a `recover_from()` method, similar to `_ConcavePolygonRecoverer`.

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.

- **name** (*None or str, optional*) – See `__init__()`.

- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.

- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.PerspectiveTransform(scale=(0.01, 0.15))
```

Apply perspective transformations using a random scale between `0.01` and `0.15` per image, where the scale is roughly a measure of how far the perspective transformation's corner points may be distanced from the image's corner points. Higher scale values lead to stronger "zoom-in" effects (and thereby stronger distortions).

```
>>> aug = iaa.PerspectiveTransform(scale=(0.01, 0.15), keep_size=False)
```

Same as in the previous example, but images are not resized back to the input image size after augmentation. This will lead to smaller output images.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.

Continued on next page

Table 147 – continued from previous page

localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.geometric.**PiecewiseAffine**(*scale=(0.0, 0.04)*, *nb_rows=(2, 4)*, *nb_cols=(2, 4)*, *order=1*, *cval=0*, *mode='constant'*, *absolute_scale=False*, *polygon_recoverer=None*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Apply affine transformations that differ between local neighbourhoods.

This augmenter places a regular grid of points on an image and randomly moves the neighbourhood of these point around via affine transformations. This leads to local distortions.

This is mostly a wrapper around scikit-image's PiecewiseAffine. See also Affine for a similar technique.

Note: This augmenter is very slow. See *Performance*. Try to use ElasticTransformation instead, which is at least 10x faster.

Note: For coordinate-based inputs (keypoints, bounding boxes, polygons, ...), this augmenter still has to perform an image-based augmentation, which will make it significantly slower for such inputs than other augmenters. See *Performance*.

Supported dtypes:

- uint8: yes; fully tested
- uint16: yes; tested (1)
- uint32: yes; tested (1) (2)
- uint64: no (3)

- `int8`: yes; tested (1)
- `int16`: yes; tested (1)
- `int32`: yes; tested (1) (2)
- `int64`: no (3)
- `float16`: yes; tested (1)
- `float32`: yes; tested (1)
- `float64`: yes; tested (1)
- `float128`: no (3)
- `bool`: yes; tested (1) (4)

- (1) Only tested with *order* set to `0`.
- (2) scikit-image converts internally to `float64`, which might introduce inaccuracies. Tests showed that these inaccuracies seemed to not be an issue.
- (3) Results too inaccurate.
- (4) Mapped internally to `float64`.

Parameters

- **scale** (*float or tuple of float or imgaug.parameters.StochasticParameter, optional*) – Each point on the regular grid is moved around via a normal distribution. This scale factor is equivalent to the normal distribution's sigma. Note that the jitter (how far each point is moved in which direction) is multiplied by the height/width of the image if `absolute_scale=False` (default), so this scale can be the same for different sized images. Recommended values are in the range `0.01` to `0.05` (weak to strong augmentations).
  - If a single `float`, then that value will always be used as the scale.
  - If a tuple `(a, b)` of `float` s, then a random value will be uniformly sampled per image from the interval `[a, b]`.
  - If a list, then a random value will be picked from that list per image.
  - If a `StochasticParameter`, then that parameter will be queried to draw one value per image.
- **nb_rows** (*int or tuple of int or imgaug.parameters.StochasticParameter, optional*) – Number of rows of points that the regular grid should have. Must be at least `2`. For large images, you might want to pick a higher value than `4`. You might have to then adjust scale to lower values.
  - If a single `int`, then that value will always be used as the number of rows.
  - If a tuple `(a, b)`, then a value from the discrete interval `[a..b]` will be uniformly sampled per image.
  - If a list, then a random value will be picked from that list per image.
  - If a StochasticParameter, then that parameter will be queried to draw one value per image.
- **nb_cols** (*int or tuple of int or imgaug.parameters.StochasticParameter, optional*) – Number of columns. Analogous to *nb_rows*.

- **order** (*int or list of int or imgaug.ALL or imgaug.parameters.StochasticParameter, optional*) – See `__init__()`.
- **cval** (*int or float or tuple of float or imgaug.ALL or imgaug.parameters.StochasticParameter, optional*) – See `__init__()`.
- **mode** (*str or list of str or imgaug.ALL or imgaug.parameters.StochasticParameter, optional*) – See `__init__()`.
- **absolute_scale** (*bool, optional*) – Take *scale* as an absolute value rather than a relative value.
- **polygon_recoverer** (*'auto' or None or imgaug.augmentables.polygons._ConcavePolygonRecoverer, optional*) – The class to use to repair invalid polygons. If `"auto"`, a new instance of :class`imgaug.augmentables.polygons._ConcavePolygonRecoverer` will be created. If `None`, no polygon recoverer will be used. If an object, then that object will be used and must provide a `recover_from()` method, similar to `_ConcavePolygonRecoverer`.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.PiecewiseAffine(scale=(0.01, 0.05))
```

Place a regular grid of points on each image and then randomly move each point around by `1` to `5` percent (with respect to the image height/width). Pixels between these points will be moved accordingly.

```
>>> aug = iaa.PiecewiseAffine(scale=(0.01, 0.05), nb_rows=8, nb_cols=8)
```

Same as the previous example, but uses a denser grid of `8x8` points (default is `4x4`). This can be useful for large images.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.

Continued on next page

Table 148 – continued from previous page

augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.geometric.**Rot90**(*k=1*, *keep_size=True*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)
Bases: *imgaug.augmenters.meta.Augmenter*

Rotate images clockwise by multiples of 90 degrees.

This could also be achieved using `Affine`, but `Rot90` is significantly more efficient.

Supported dtypes:

if (keep_size=False):

- `uint8`: yes; fully tested
- `uint16`: yes; tested
- `uint32`: yes; tested
- `uint64`: yes; tested
- `int8`: yes; tested
- `int16`: yes; tested
- `int32`: yes; tested
- `int64`: yes; tested
- `float16`: yes; tested
- `float32`: yes; tested
- `float64`: yes; tested
- `float128`: yes; tested
- `bool`: yes; tested

if (keep_size=True):

minimum of (`imgaug.augmenters.geometric.Rot90(keep_size=False)`, *imresize_many_images()*

)

Parameters

- **k** (*int or list of int or tuple of int or imaug.ALL or imgaug.parameters.StochasticParameter, optional*) –

  How often to rotate clockwise by 90 degrees.

  - If a single `int`, then that value will be used for all images.
  - If a tuple `(a, b)`, then a random value will be uniformly sampled per image from the discrete interval `[a..b]`.
  - If a list, then for each image a random value will be sampled from that list.
  - If `imgaug.ALL`, then equivalant to list `[0, 1, 2, 3]`.
  - If `StochasticParameter`, then that parameter is queried per image to sample the value to use.

- **keep_size** (*bool, optional*) – After rotation by an odd-valued *k* (e.g. 1 or 3), the resulting image may have a different height/width than the original image. If this parameter is set to `True`, then the rotated image will be resized to the input image's size. Note that this might also cause the augmented image to look distorted.

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.

- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.Rot90(1)
```

Rotate all images by 90 degrees. Resize these images afterwards to keep the size that they had before augmentation. This may cause the images to look distorted.

```
>>> aug = iaa.Rot90([1, 3])
```

Rotate all images by 90 or 270 degrees. Resize these images afterwards to keep the size that they had before augmentation. This may cause the images to look distorted.

```
>>> aug = iaa.Rot90((1, 3))
```

Rotate all images by 90, 180 or 270 degrees. Resize these images afterwards to keep the size that they had before augmentation. This may cause the images to look distorted.

```
>>> aug = iaa.Rot90((1, 3), keep_size=False)
```

Rotate all images by 90, 180 or 270 degrees. Does not resize to the original image size afterwards, i.e. each image's size may change.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.

Continued on next page

Table 149 – continued from previous page

augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.geometric.**Rotate**(*rotate=(-30, 30)*, *order=1*, *cval=0*, *mode='constant'*, *fit_output=False*, *backend='auto'*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.geometric.Affine*

Apply affine rotation on the y-axis to input data.

This is a wrapper around *Affine*. It is the same as `Affine(rotate=<value>)`.

Added in 0.4.0.

Supported dtypes:

See *Affine*.

Parameters

- **rotate** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – See *Affine*.
- **order** (*int or iterable of int or imgaug.ALL or imgaug.parameters.StochasticParameter, optional*) – See *Affine*.
- **cval** (*number or tuple of number or list of number or imgaug.ALL or imgaug.parameters.StochasticParameter, optional*) – See *Affine*.
- **mode** (*str or list of str or imgaug.ALL or imgaug.parameters.StochasticParameter, optional*) – See *Affine*.
- **fit_output** (*bool, optional*) – See *Affine*.
- **backend** (*str, optional*) – See *Affine*.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.Rotate((-45, 45))
```

Create an augmenter that rotates images by a random value between `-45` and `45` degress.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.

Continued on next page

Table 150 – continued from previous page

augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.geometric.**ScaleX**(*scale=(0.5, 1.5)*, *order=1*, *cval=0*, *mode='constant'*, *fit_output=False*, *backend='auto'*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.geometric.Affine*

Apply affine scaling on the x-axis to input data.

This is a wrapper around *Affine*.

Added in 0.4.0.

Supported dtypes:

See *Affine*.

Parameters

- **scale** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Analogous to `scale` in *Affine*, except that this scale value only affects the x-axis. No dictionary input is allowed.
- **order** (*int or iterable of int or imgaug.ALL or imgaug.parameters.StochasticParameter, optional*) – See *Affine*.
- **cval** (*number or tuple of number or list of number or imgaug.ALL or imgaug.parameters.StochasticParameter, optional*) – See *Affine*.
- **mode** (*str or list of str or imgaug.ALL or imgaug.parameters.StochasticParameter, optional*) – See *Affine*.
- **fit_output** (*bool, optional*) – See *Affine*.
- **backend** (*str, optional*) – See *Affine*.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.ScaleX((0.5, 1.5))
```

Create an augmenter that scales images along the width to sizes between `50%` and `150%`. This does not change the image shape (i.e. height and width), only the pixels within the image are remapped and potentially new ones are filled in.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.

Continued on next page

Table 151 – continued from previous page

augment_bounding_boxes(self, ...[, parents, ...])	Augment a batch of bounding boxes.
augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.geometric.**ScaleY**(*scale=(0.5, 1.5)*, *order=1*, *cval=0*, *mode='constant'*, *fit_output=False*, *backend='auto'*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.geometric.Affine*

Apply affine scaling on the y-axis to input data.

This is a wrapper around *Affine*.

Added in 0.4.0.

Supported dtypes:

See *Affine*.

Parameters

- **scale** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Analogous to `scale` in *Affine*, except that this scale value only affects the y-axis. No dictionary input is allowed.
- **order** (*int or iterable of int or imgaug.ALL or imgaug.parameters.StochasticParameter, optional*) – See *Affine*.
- **cval** (*number or tuple of number or list of number or imgaug.ALL or imgaug.parameters.StochasticParameter, optional*) – See *Affine*.
- **mode** (*str or list of str or imgaug.ALL or imgaug.parameters.StochasticParameter, optional*) – See *Affine*.
- **fit_output** (*bool, optional*) – See *Affine*.
- **backend** (*str, optional*) – See *Affine*.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.ScaleY((0.5, 1.5))
```

Create an augmenter that scales images along the height to sizes between `50%` and `150%`. This does not change the image shape (i.e. height and width), only the pixels within the image are remapped and potentially new ones are filled in.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.
`get_parameters`(self)	See *get_parameters()*.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`localize_random_state_`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`pool`(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
`remove_augmenters`(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
`remove_augmenters_`(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
`remove_augmenters_inplace`(self, func[, parents])	**Deprecated**.
`reseed`(self[, random_state, deterministic_too])	**Deprecated**.
`seed_`(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.

Continued on next page

Table 152 – continued from previous page

show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.geometric.**ShearX**(*shear=(-30, 30)*, *order=1*, *cval=0*, *mode='constant'*, *fit_output=False*, *backend='auto'*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.geometric.Affine*

Apply affine shear on the x-axis to input data.

This is a wrapper around *Affine*.

Added in 0.4.0.

Supported dtypes:

See *Affine*.

Parameters

- **shear** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Analogous to shear in *Affine*, except that this shear value only affects the x-axis. No dictionary input is allowed.
- **order** (*int or iterable of int or imgaug.ALL or imgaug.parameters.StochasticParameter, optional*) – See *Affine*.
- **cval** (*number or tuple of number or list of number or imgaug.ALL or imgaug.parameters.StochasticParameter, optional*) – See *Affine*.
- **mode** (*str or list of str or imgaug.ALL or imgaug.parameters.StochasticParameter, optional*) – See *Affine*.
- **fit_output** (*bool, optional*) – See *Affine*.
- **backend** (*str, optional*) – See *Affine*.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See __init__().
- **name** (*None or str, optional*) – See __init__().
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method to_deterministic() for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.ShearX((-20, 20))
```

Create an augmenter that shears images along the x-axis by random amounts between `-20` and `20` degrees.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.
`get_parameters`(self)	See *get_parameters()*.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`localize_random_state_`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`pool`(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.

Continued on next page

Table 153 – continued from previous page

remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.geometric.**ShearY**(*shear=(-30, 30)*, *order=1*, *cval=0*, *mode='constant'*, *fit_output=False*, *backend='auto'*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.geometric.Affine*

Apply affine shear on the y-axis to input data.

This is a wrapper around *Affine*.

Added in 0.4.0.

Supported dtypes:

See *Affine*.

Parameters

- **shear** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Analogous to shear in *Affine*, except that this shear value only affects the y-axis. No dictionary input is allowed.
- **order** (*int or iterable of int or imgaug.ALL or imgaug.parameters.StochasticParameter, optional*) – See *Affine*.
- **cval** (*number or tuple of number or list of number or imgaug.ALL or imgaug.parameters.StochasticParameter, optional*) – See *Affine*.
- **mode** (*str or list of str or imgaug.ALL or imgaug.parameters.StochasticParameter, optional*) – See *Affine*.
- **fit_output** (*bool, optional*) – See *Affine*.
- **backend** (*str, optional*) – See *Affine*.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See __init__().
- **name** (*None or str, optional*) – See __init__().
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.

- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.ShearY((-20, 20))
```

Create an augmenter that shears images along the y-axis by random amounts between `-20` and `20` degrees.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.
`get_parameters`(self)	See *get_parameters()*.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.

Continued on next page

Table 154 – continued from previous page

localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.geometric.**TranslateX**(*percent=None*, *px=None*, *order=1*, *cval=0*, *mode='constant'*, *fit_output=False*, *backend='auto'*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.geometric.Affine*

Apply affine translation on the x-axis to input data.

This is a wrapper around *Affine*.

Added in 0.4.0.

Supported dtypes:

See *Affine*.

Parameters

- **percent** (*None or number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Analogous to translate_percent in *Affine*, except that this translation value only affects the x-axis. No dictionary input is allowed.
- **px** (*None or int or tuple of int or list of int or imgaug.parameters.StochasticParameter or dict {"x": int/tuple/list/StochasticParameter, "y": int/tuple/list/StochasticParameter}, optional*) – Analogous to translate_px in *Affine*, except that this translation value only affects the x-axis. No dictionary input is allowed.
- **order** (*int or iterable of int or imgaug.ALL or imgaug.parameters.StochasticParameter, optional*) – See *Affine*.
- **cval** (*number or tuple of number or list of number or imgaug.ALL or imgaug.parameters.StochasticParameter, optional*) – See *Affine*.
- **mode** (*str or list of str or imgaug.ALL or imgaug.parameters.StochasticParameter, optional*) – See *Affine*.
- **fit_output** (*bool, optional*) – See *Affine*.
- **backend** (*str, optional*) – See *Affine*.

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.TranslateX(px=(-20, 20))
```

Create an augmenter that translates images along the x-axis by `-20` to `20` pixels.

```
>>> aug = iaa.TranslateX(percent=(-0.1, 0.1))
```

Create an augmenter that translates images along the x-axis by `-10%` to `10%` (relative to the x-axis size).

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.

Continued on next page

Table 155 – continued from previous page

copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.geometric.**TranslateY**(*percent=None*, *px=None*, *order=1*, *cval=0*, *mode='constant'*, *fit_output=False*, *backend='auto'*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.geometric.Affine*

Apply affine translation on the y-axis to input data.

This is a wrapper around *Affine*.

Added in 0.4.0.

Supported dtypes:

See *Affine*.

Parameters

- **percent** (*None or number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Analogous to translate_percent in *Affine*, except that this translation value only affects the y-axis. No dictionary input is allowed.
- **px** (*None or int or tuple of int or list of int or imgaug.parameters.StochasticParameter or*

dict {"x": int/tuple/list/StochasticParameter, "y": int/tuple/list/StochasticParameter}, optional) – Analogous to `translate_px` in *Affine*, except that this translation value only affects the y-axis. No dictionary input is allowed.

- **order** (*int or iterable of int or imgaug.ALL or imgaug.parameters.StochasticParameter, optional*) – See *Affine*.
- **cval** (*number or tuple of number or list of number or imgaug.ALL or imgaug.parameters.StochasticParameter, optional*) – See *Affine*.
- **mode** (*str or list of str or imgaug.ALL or imgaug.parameters.StochasticParameter, optional*) – See *Affine*.
- **fit_output** (*bool, optional*) – See *Affine*.
- **backend** (*str, optional*) – See *Affine*.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.TranslateY(px=(-20, 20))
```

Create an augmenter that translates images along the y-axis by `-20` to `20` pixels.

```
>>> aug = iaa.TranslateY(percent=(-0.1, 0.1))
```

Create an augmenter that translates images along the y-axis by `-10%` to `10%` (relative to the y-axis size).

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.

Continued on next page

Table 156 – continued from previous page

augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.geometric.**WithPolarWarping**(*children*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Augmenter that applies other augmenters in a polar-transformed space.

This augmenter first transforms an image into a polar representation, then applies its child augmenter, then

transforms back to cartesian space. The polar representation is still in the image's input dtype (i.e. `uint8` stays `uint8`) and can be visualized. It can be thought of as an "unrolled" version of the image, where previously circular lines appear straight. Hence, applying child augmenters in that space can lead to circular effects. E.g. replacing rectangular pixel areas in the polar representation with black pixels will lead to curved black areas in the cartesian result.

This augmenter can create new pixels in the image. It will fill these with black pixels. For segmentation maps it will fill with class id `0`. For heatmaps it will fill with `0.0`.

This augmenter is limited to arrays with a height and/or width of `32767` or less.

Warning: When augmenting coordinates in polar representation, it is possible that these are shifted outside of the polar image, but are inside the image plane after transforming back to cartesian representation, usually on newly created pixels (i.e. black backgrounds). These coordinates are currently not removed. It is recommended to not use very strong child transformations when also augmenting coordinate-based augmentables.

Warning: For bounding boxes, this augmenter suffers from the same problem as affine rotations applied to bounding boxes, i.e. the resulting bounding boxes can have unintuitive (seemingly wrong) appearance. This is due to coordinates being "rotated" that are inside the bounding box, but do not fall on the object and actually are background. It is recommended to use this augmenter with caution when augmenting bounding boxes.

Warning: For polygons, this augmenter should not be combined with augmenters that perform automatic polygon recovery for invalid polygons, as the polygons will frequently appear broken in polar representation and their "fixed" version will be very broken in cartesian representation. Augmenters that perform such polygon recovery are currently `PerspectiveTransform`, `PiecewiseAffine` and `ElasticTransformation`.

Added in 0.4.0.

Supported dtypes:

- `uint8`: yes; fully tested
- `uint16`: yes; tested
- `uint32`: no (1)
- `uint64`: no (2)
- `int8`: yes; tested
- `int16`: yes; tested
- `int32`: yes; tested
- `int64`: no (2)
- `float16`: yes; tested (3)
- `float32`: yes; tested
- `float64`: yes; tested
- `float128`: no (1)
- `bool`: yes; tested (4)

- (1) OpenCV produces error `TypeError: Expected cv::UMat for argument 'src'`
- (2) OpenCV produces array of nothing but zeros.
- (3) Mapepd to `float32`.
- (4) Mapped to `uint8`.

Parameters

- **children** (*imgaug.augmenters.meta.Augmenter or list of imgaug.augmenters.meta.Augmenter or None, optional*) – One or more augmenters to apply to images after they were transformed to polar representation.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.WithPolarWarping(iaa.CropAndPad(percent=(-0.1, 0.1)))
```

Apply cropping and padding in polar representation, then warp back to cartesian representation.

```
>>> aug = iaa.WithPolarWarping(
>>>     iaa.Affine(
>>>         translate_percent={"x": (-0.2, 0.2), "y": (-0.2, 0.2)},
>>>         rotate=(-35, 35),
>>>         scale=(0.8, 1.2),
>>>         shear={"x": (-15, 15), "y": (-15, 15)}
>>>     )
>>> )
```

Apply affine transformations in polar representation.

```
>>> aug = iaa.WithPolarWarping(iaa.AveragePooling((2, 8)))
```

Apply average pooling in polar representation. This leads to circular bins.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	See *get_children_lists()*.
get_parameters(self)	See *get_parameters()*.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`localize_random_state_`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`pool`(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
`remove_augmenters`(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
`remove_augmenters_`(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
`remove_augmenters_inplace`(self, func[, parents])	**Deprecated**.
`reseed`(self[, random_state, deterministic_too])	**Deprecated**.
`seed_`(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
`show_grid`(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.

Continued on next page

Table 157 – continued from previous page

to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.
warpPolarCoords(src, dsize, center, ...)	

get_children_lists(*self*)
See *get_children_lists()*.

get_parameters(*self*)
See *get_parameters()*.

classmethod warpPolarCoords(*src*, *dsize*, *center*, *maxRadius*, *flags*)

imgaug.augmenters.geometric.**apply_jigsaw**(*arr*, *destinations*)
Move cells of an image similar to a jigsaw puzzle.

This function will split the image into `rows x cols` cells and move each cell to the target index given in *destinations*.

Added in 0.4.0.

Supported dtypes:

- `uint8`: yes; fully tested
- `uint16`: yes; fully tested
- `uint32`: yes; fully tested
- `uint64`: yes; fully tested
- `int8`: yes; fully tested
- `int16`: yes; fully tested
- `int32`: yes; fully tested
- `int64`: yes; fully tested
- `float16`: yes; fully tested
- `float32`: yes; fully tested
- `float64`: yes; fully tested
- `float128`: yes; fully tested
- `bool`: yes; fully tested

Parameters

- **arr** (*ndarray*) – Array with at least two dimensions denoting height and width.
- **destinations** (*ndarray*) – 2-dimensional array containing for each cell the id of the destination cell. The order is expected to a flattened c-order, i.e. row by row. The height of the image must be evenly divisible by the number of rows in this array. Analogous for the width and columns.

Returns Modified image with cells moved according to *destinations*.

Return type ndarray

imgaug.augmenters.geometric.**apply_jigsaw_to_coords**(*coords*, *destinations*, *image_shape*)
Move coordinates on an image similar to a jigsaw puzzle.

This is the same as *apply_jigsaw()*, but moves coordinates within the cells.

Added in 0.4.0.

Parameters

- **coords** (*ndarray*) – `(N, 2)` array denoting xy-coordinates.
- **destinations** (*ndarray*) – See *apply_jigsaw()*.
- **image_shape** (*tuple of int*) – `(height, width, ...)` shape of the image on which the coordinates are placed. Only height and width are required.

Returns Moved coordinates.

Return type ndarray

`imgaug.augmenters.geometric.`**generate_jigsaw_destinations**(*nb_rows*, *nb_cols*, *max_steps*, *seed*, *connectivity=4*)

Generate a destination pattern for *apply_jigsaw()*.

Added in 0.4.0.

Parameters

- **nb_rows** (*int*) – Number of rows to split the image into.
- **nb_cols** (*int*) – Number of columns to split the image into.
- **max_steps** (*int*) – Maximum number of cells that each cell may be moved.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState*) – Seed value or alternatively RNG to use. If `None` the global RNG will be used.
- **connectivity** (*int, optional*) – Whether a diagonal move of a cell counts as one step (`connectivity=8`) or two steps (`connectivity=4`).

Returns 2-dimensional array containing for each cell the id of the target cell.

Return type ndarray

13.30 imgaug.augmenters.imgcorruptlike

Augmenters that wrap methods from `imagecorruptions` package.

See **'https://github.com/bethgelab/imagecorruptions'_** for the package.

The package is derived from **'https://github.com/hendrycks/robustness'_**. The corresponding paper is:

```
Hendrycks, Dan and Dietterich, Thomas G.
Benchmarking Neural Network Robustness to Common Corruptions and
Surface Variations
```

with the newer version being:

```
Hendrycks, Dan and Dietterich, Thomas G.
Benchmarking Neural Network Robustness to Common Corruptions and
Perturbations
```

List of augmenters:

- *GaussianNoise*
- *ShotNoise*
- *ImpulseNoise*
- *SpeckleNoise*
- *GaussianBlur*
- *GlassBlur*
- *DefocusBlur*
- *MotionBlur*
- *ZoomBlur*
- *Fog*
- *Frost*
- *Snow*
- *Spatter*
- *Contrast*
- *Brightness*
- *Saturate*
- *JpegCompression*
- *Pixelate*
- *ElasticTransform*

Note: The functions provided here have identical outputs to the ones in `imagecorruptions` when called using the `corrupt()` function of that package. E.g. the outputs are always `uint8` and not `float32` or `float64`.

Example usage:

```
>>> # Skip the doctests in this file as the imagecorruptions package is
>>> # not available in all python versions that are otherwise supported
>>> # by imgaug.
>>>
>>> import imgaug as ia
>>> import imgaug.augmenters as iaa
>>> import numpy as np
>>> image = np.zeros((64, 64, 3), dtype=np.uint8)
>>> names, funcs = iaa.imgcorruptlike.get_corruption_names("validation")
>>> for name, func in zip(names, funcs):
>>>     image_aug = func(image, severity=5, seed=1)
>>>     image_aug = ia.draw_text(image_aug, x=20, y=20, text=name)
>>>     ia.imshow(image_aug)

Use e.g. ``iaa.imgcorruptlike.GaussianNoise(severity=2)(images=...)`` to
create and apply a specific augmenter.
```

Added in 0.4.0.

class imgaug.augmenters.imgcorruptlike.**Brightness**(*severity=(1, 5), seed=None, name=None, random_state='deprecated', deterministic='deprecated'*)

Bases: imgaug.augmenters.imgcorruptlike._ImgcorruptAugmenterBase

Wrapper around imagecorruptions.corruptions.brightness.

Note: This augmenter only affects images. Other data is not changed.

Added in 0.4.0.

Supported dtypes:

See *apply_brightness()*.

Parameters

- **severity** (*int, optional*) – Strength of the corruption, with valid values being 1 <= severity <= 5.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See __init__().
- **name** (*None or str, optional*) – See __init__().
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method to_deterministic() for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> # doctest: +SKIP
>>> import imgaug.augmenters as iaa
>>> aug = iaa.imgcorruptlike.Brightness(severity=2)
```

Create an augmenter around imagecorruptions.corruptions.brightness. Apply it to images using e.g. aug(images=[image1, image2, ...]).

Methods

__call__(self, *args, **kwargs)	Alias for *augment()*.
augment(self[, return_batch, hooks])	Augment a batch.
augment_batch(self, batch[, hooks])	**Deprecated**.
augment_batch_(self, batch[, parents, hooks])	Augment a single batch in-place.
augment_batches(self, batches[, hooks, ...])	Augment multiple batches.

Continued on next page

Table 158 – continued from previous page

augment_bounding_boxes(self, ...[, parents, ...])	Augment a batch of bounding boxes.
augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.imgcorruptlike.**Contrast**(*severity=(1, 5), seed=None, name=None, random_state='deprecated', deterministic='deprecated'*)

Bases: imgaug.augmenters.imgcorruptlike._ImgcorruptAugmenterBase

Wrapper around imagecorruptions.corruptions.contrast.

Note: This augmenter only affects images. Other data is not changed.

Added in 0.4.0.

Supported dtypes:

See *apply_contrast()*.

Parameters

- **severity** (*int, optional*) – Strength of the corruption, with valid values being 1 <= severity <= 5.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See __init__().
- **name** (*None or str, optional*) – See __init__().
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method to_deterministic() for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> # doctest: +SKIP
>>> import imgaug.augmenters as iaa
>>> aug = iaa.imgcorruptlike.Contrast(severity=2)
```

Create an augmenter around imagecorruptions.corruptions.contrast. Apply it to images using e.g. aug(images=[image1, image2, ...]).

Methods

__call__(self, *args, **kwargs)	Alias for *augment()*.
augment(self[, return_batch, hooks])	Augment a batch.
augment_batch(self, batch[, hooks])	**Deprecated**.
augment_batch_(self, batch[, parents, hooks])	Augment a single batch in-place.
augment_batches(self, batches[, hooks, ...])	Augment multiple batches.

Continued on next page

Table 159 – continued from previous page

augment_bounding_boxes(self, ...[, parents, ...])	Augment a batch of bounding boxes.
augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.imgcorruptlike.**DefocusBlur**(*severity=(1, 5), seed=None, name=None, random_state='deprecated', deterministic='deprecated'*)

Bases: imgaug.augmenters.imgcorruptlike._ImgcorruptAugmenterBase

Wrapper around imagecorruptions.corruptions.defocus_blur.

Note: This augmenter only affects images. Other data is not changed.

Added in 0.4.0.

Supported dtypes:

See *apply_defocus_blur()*.

Parameters

- **severity** (*int, optional*) – Strength of the corruption, with valid values being 1 <= severity <= 5.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See __init__().
- **name** (*None or str, optional*) – See __init__().
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method to_deterministic() for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> # doctest: +SKIP
>>> import imgaug.augmenters as iaa
>>> aug = iaa.imgcorruptlike.DefocusBlur(severity=2)
```

Create an augmenter around imagecorruptions.corruptions.defocus_blur. Apply it to images using e.g. aug(images=[image1, image2, ...]).

Methods

__call__(self, *args, **kwargs)	Alias for *augment()*.
augment(self[, return_batch, hooks])	Augment a batch.
augment_batch(self, batch[, hooks])	**Deprecated**.
augment_batch_(self, batch[, parents, hooks])	Augment a single batch in-place.
augment_batches(self, batches[, hooks, ...])	Augment multiple batches.

Continued on next page

Table 160 – continued from previous page

augment_bounding_boxes(self, ...[, parents, ...])	Augment a batch of bounding boxes.
augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.imgcorruptlike.**ElasticTransform**(*severity=(1, 5), seed=None, name=None, random_state='deprecated', deterministic='deprecated'*)

Bases: imgaug.augmenters.imgcorruptlike._ImgcorruptAugmenterBase

Wrapper around imagecorruptions.corruptions.elastic_transform.

Warning: This augmenter can currently only transform image-data. Batches containing heatmaps, segmentation maps and coordinate-based augmentables will be rejected with an error. Use *ElasticTransformation* if you have to transform such inputs.

Added in 0.4.0.

Supported dtypes:

See *apply_elastic_transform()*.

Parameters

- **severity** (*int, optional*) – Strength of the corruption, with valid values being 1 <= severity <= 5.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See __init__().
- **name** (*None or str, optional*) – See __init__().
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method to_deterministic() for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> # doctest: +SKIP
>>> import imgaug.augmenters as iaa
>>> aug = iaa.imgcorruptlike.ElasticTransform(severity=2)
```

Create an augmenter around imagecorruptions.corruptions.elastic_transform. Apply it to images using e.g. aug(images=[image1, image2, ...]).

Methods

__call__(self, *args, **kwargs)	Alias for *augment()*.
augment(self[, return_batch, hooks])	Augment a batch.

Continued on next page

Table 161 – continued from previous page

augment_batch(self, batch[, hooks])	**Deprecated**.
augment_batch_(self, batch[, parents, hooks])	Augment a single batch in-place.
augment_batches(self, batches[, hooks, ...])	Augment multiple batches.
augment_bounding_boxes(self, ...[, parents, ...])	Augment a batch of bounding boxes.
augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.imgcorruptlike.**Fog**(*severity=(1, 5), seed=None, name=None, random_state='deprecated', deterministic='deprecated'*)

Bases: imgaug.augmenters.imgcorruptlike._ImgcorruptAugmenterBase

Wrapper around imagecorruptions.corruptions.fog.

Note: This augmenter only affects images. Other data is not changed.

Added in 0.4.0.

Supported dtypes:

See *apply_fog()*.

Parameters

- **severity** (*int, optional*) – Strength of the corruption, with valid values being 1 <= severity <= 5.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See __init__().
- **name** (*None or str, optional*) – See __init__().
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method to_deterministic() for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> # doctest: +SKIP
>>> import imgaug.augmenters as iaa
>>> aug = iaa.imgcorruptlike.Fog(severity=2)
```

Create an augmenter around imagecorruptions.corruptions.fog. Apply it to images using e.g. aug(images=[image1, image2, ...]).

Methods

__call__(self, *args, **kwargs)	Alias for *augment()*.
augment(self[, return_batch, hooks])	Augment a batch.
augment_batch(self, batch[, hooks])	**Deprecated**.
augment_batch_(self, batch[, parents, hooks])	Augment a single batch in-place.
augment_batches(self, batches[, hooks, ...])	Augment multiple batches.
augment_bounding_boxes(self, ...[, parents, ...])	Augment a batch of bounding boxes.

Continued on next page

Table 162 – continued from previous page

augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.imgcorruptlike.**Frost**(*severity=(1, 5)*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: imgaug.augmenters.imgcorruptlike._ImgcorruptAugmenterBase

Wrapper around imagecorruptions.corruptions.frost.

Note: This augmenter only affects images. Other data is not changed.

Added in 0.4.0.

Supported dtypes:

See *apply_frost()*.

Parameters

- **severity** (*int, optional*) – Strength of the corruption, with valid values being `1 <= severity <= 5`.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> # doctest: +SKIP
>>> import imgaug.augmenters as iaa
>>> aug = iaa.imgcorruptlike.Frost(severity=2)
```

Create an augmenter around `imagecorruptions.corruptions.frost`. Apply it to images using e.g. `aug(images=[image1, image2, ...])`.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.

Continued on next page

Table 163 – continued from previous page

augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.imgcorruptlike.**GaussianBlur**(*severity=(1, 5)*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: imgaug.augmenters.imgcorruptlike._ImgcorruptAugmenterBase

Wrapper around imagecorruptions.corruptions.gaussian_blur.

Note: This augmenter only affects images. Other data is not changed.

Added in 0.4.0.

Supported dtypes:

See *apply_gaussian_blur()*.

Parameters

- **severity** (*int, optional*) – Strength of the corruption, with valid values being `1 <= severity <= 5`.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> # doctest: +SKIP
>>> import imgaug.augmenters as iaa
>>> aug = iaa.imgcorruptlike.GaussianBlur(severity=2)
```

Create an augmenter around `imagecorruptions.corruptions.gaussian_blur`. Apply it to images using e.g. `aug(images=[image1, image2, ...])`.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.

Continued on next page

Table 164 – continued from previous page

augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.imgcorruptlike.**GaussianNoise**(*severity=(1, 5)*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: imgaug.augmenters.imgcorruptlike._ImgcorruptAugmenterBase

Wrapper around imagecorruptions.corruptions.gaussian_noise.

Note: This augmenter only affects images. Other data is not changed.

Added in 0.4.0.

Supported dtypes:

See *apply_gaussian_noise()*.

Parameters

- **severity** (*int, optional*) – Strength of the corruption, with valid values being `1 <= severity <= 5`.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> # doctest: +SKIP
>>> import imgaug.augmenters as iaa
>>> aug = iaa.imgcorruptlike.GaussianNoise(severity=2)
```

Create an augmenter around `imagecorruptions.corruptions.gaussian_noise`. Apply it to images using e.g. `aug(images=[image1, image2, ...])`.

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.

Continued on next page

Table 165 – continued from previous page

copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.imgcorruptlike.**GlassBlur**(*severity=(1, 5)*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: imgaug.augmenters.imgcorruptlike._ImgcorruptAugmenterBase

Wrapper around imagecorruptions.corruptions.glass_blur.

Note: This augmenter only affects images. Other data is not changed.

Added in 0.4.0.

Supported dtypes:

See *apply_glass_blur()*.

Parameters

- **severity** (*int, optional*) – Strength of the corruption, with valid values being 1 <= severity <= 5.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or*

numpy.random.RandomState, optional) – See `__init__()`.

- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> # doctest: +SKIP
>>> import imgaug.augmenters as iaa
>>> aug = iaa.imgcorruptlike.GlassBlur(severity=2)
```

Create an augmenter around `imagecorruptions.corruptions.glass_blur`. Apply it to images using e.g. `aug(images=[image1, image2, ...])`.

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.

Continued on next page

Table 166 – continued from previous page

find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.imgcorruptlike.**ImpulseNoise**(*severity=(1, 5)*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: imgaug.augmenters.imgcorruptlike._ImgcorruptAugmenterBase

Wrapper around imagecorruptions.corruptions.impulse_noise.

Note: This augmenter only affects images. Other data is not changed.

Added in 0.4.0.

Supported dtypes:

See *apply_impulse_noise()*.

Parameters

- **severity** (*int, optional*) – Strength of the corruption, with valid values being 1 <= severity <= 5.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See __init__().
- **name** (*None or str, optional*) – See __init__().
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or*

numpy.random.RandomState, optional) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.

- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> # doctest: +SKIP
>>> import imgaug.augmenters as iaa
>>> aug = iaa.imgcorruptlike.ImpulseNoise(severity=2)
```

Create an augmenter around `imagecorruptions.corruptions.impulse_noise`. Apply it to images using e.g. `aug(images=[image1, image2, ...])`.

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.

Continued on next page

Table 167 – continued from previous page

get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.imgcorruptlike.**JpegCompression**(*severity=(1, 5)*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: imgaug.augmenters.imgcorruptlike._ImgcorruptAugmenterBase

Wrapper around imagecorruptions.corruptions.jpeg_compression.

Note: This augmenter only affects images. Other data is not changed.

Added in 0.4.0.

Supported dtypes:

See *apply_jpeg_compression()*.

Parameters

- **severity** (*int, optional*) – Strength of the corruption, with valid values being 1 <= severity <= 5.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See __init__().
- **name** (*None or str, optional*) – See __init__().
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.

- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> # doctest: +SKIP
>>> import imgaug.augmenters as iaa
>>> aug = iaa.imgcorruptlike.JpegCompression(severity=2)
```

Create an augmenter around `imagecorruptions.corruptions.jpeg_compression`. Apply it to images using e.g. `aug(images=[image1, image2, ...])`.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.
`get_parameters`(self)	See *get_parameters()*.

Continued on next page

Table 168 – continued from previous page

localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.imgcorruptlike.**MotionBlur**(*severity=(1, 5), seed=None, name=None, random_state='deprecated', deterministic='deprecated'*)

Bases: imgaug.augmenters.imgcorruptlike._ImgcorruptAugmenterBase

Wrapper around imagecorruptions.corruptions.motion_blur.

Note: This augmenter only affects images. Other data is not changed.

Added in 0.4.0.

Supported dtypes:

See *apply_motion_blur()*.

Parameters

- **severity** (*int, optional*) – Strength of the corruption, with valid values being 1 <= severity <= 5.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See __init__().
- **name** (*None or str, optional*) – See __init__().
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method to_deterministic() for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> # doctest: +SKIP
>>> import imgaug.augmenters as iaa
>>> aug = iaa.imgcorruptlike.MotionBlur(severity=2)
```

Create an augmenter around `imagecorruptions.corruptions.motion_blur`. Apply it to images using e.g. `aug(images=[image1, image2, ...])`.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.
`get_parameters`(self)	See *get_parameters()*.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`localize_random_state_`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`pool`(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.

Continued on next page

Table 169 – continued from previous page

remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.imgcorruptlike.**Pixelate**(*severity=(1, 5), seed=None, name=None, random_state='deprecated', deterministic='deprecated'*)

Bases: imgaug.augmenters.imgcorruptlike._ImgcorruptAugmenterBase

Wrapper around imagecorruptions.corruptions.pixelate.

Note: This augmenter only affects images. Other data is not changed.

Added in 0.4.0.

Supported dtypes:

See *apply_pixelate()*.

Parameters

- **severity** (*int, optional*) – Strength of the corruption, with valid values being 1 <= severity <= 5.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See __init__().
- **name** (*None or str, optional*) – See __init__().
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method to_deterministic() for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> # doctest: +SKIP
>>> import imgaug.augmenters as iaa
>>> aug = iaa.imgcorruptlike.Pixelate(severity=2)
```

Create an augmenter around `imagecorruptions.corruptions.pixelate`. Apply it to images using e.g. `aug(images=[image1, image2, ...])`.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.
`get_parameters`(self)	See *get_parameters()*.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`localize_random_state_`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`pool`(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
`remove_augmenters`(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
`remove_augmenters_`(self, func[, parents])	Remove in-place children of this augmenter that match a condition.

Continued on next page

Table 170 – continued from previous page

remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.imgcorruptlike.**Saturate**(*severity=(1, 5), seed=None, name=None, random_state='deprecated', deterministic='deprecated'*)

Bases: imgaug.augmenters.imgcorruptlike._ImgcorruptAugmenterBase

Wrapper around imagecorruptions.corruptions.saturate.

Note: This augmenter only affects images. Other data is not changed.

Added in 0.4.0.

Supported dtypes:

See *apply_saturate()*.

Parameters

- **severity** (*int, optional*) – Strength of the corruption, with valid values being 1 <= severity <= 5.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See __init__().
- **name** (*None or str, optional*) – See __init__().
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method to_deterministic() for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> # doctest: +SKIP
>>> import imgaug.augmenters as iaa
>>> aug = iaa.imgcorruptlike.Saturate(severity=2)
```

Create an augmenter around imagecorruptions.corruptions.saturate. Apply it to images using e.g. aug(images=[image1, image2, ...]).

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.
`get_parameters`(self)	See *get_parameters()*.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`localize_random_state_`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`pool`(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
`remove_augmenters`(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
`remove_augmenters_`(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
`remove_augmenters_inplace`(self, func[, parents])	**Deprecated**.
`reseed`(self[, random_state, deterministic_too])	**Deprecated**.
`seed_`(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.

Continued on next page

Table 171 – continued from previous page

show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.imgcorruptlike.**ShotNoise**(*severity=(1, 5), seed=None, name=None, random_state='deprecated', deterministic='deprecated'*)

Bases: imgaug.augmenters.imgcorruptlike._ImgcorruptAugmenterBase

Wrapper around imagecorruptions.shot_noise.

Note: This augmenter only affects images. Other data is not changed.

Added in 0.4.0.

Supported dtypes:

See *apply_shot_noise()*.

Parameters

- **severity** (*int, optional*) – Strength of the corruption, with valid values being 1 <= severity <= 5.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See __init__().
- **name** (*None or str, optional*) – See __init__().
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method to_deterministic() for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> # doctest: +SKIP
>>> import imgaug.augmenters as iaa
>>> aug = iaa.imgcorruptlike.ShotNoise(severity=2)
```

Create an augmenter around imagecorruptions.corruptions.shot_noise. Apply it to images using e.g. aug(images=[image1, image2, ...]).

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.
`get_parameters`(self)	See *get_parameters()*.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`localize_random_state_`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`pool`(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
`remove_augmenters`(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
`remove_augmenters_`(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
`remove_augmenters_inplace`(self, func[, parents])	**Deprecated**.
`reseed`(self[, random_state, deterministic_too])	**Deprecated**.
`seed_`(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
`show_grid`(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.

Continued on next page

Table 172 – continued from previous page

to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.imgcorruptlike.**Snow**(*severity=(1, 5)*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: imgaug.augmenters.imgcorruptlike._ImgcorruptAugmenterBase

Wrapper around imagecorruptions.corruptions.snow.

Note: This augmenter only affects images. Other data is not changed.

Added in 0.4.0.

Supported dtypes:

See *apply_snow()*.

Parameters

- **severity** (*int, optional*) – Strength of the corruption, with valid values being 1 <= severity <= 5.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See __init__().
- **name** (*None or str, optional*) – See __init__().
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method to_deterministic() for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> # doctest: +SKIP
>>> import imgaug.augmenters as iaa
>>> aug = iaa.imgcorruptlike.Snow(severity=2)
```

Create an augmenter around imagecorruptions.corruptions.snow. Apply it to images using e.g. aug(images=[image1, image2, ...]).

Methods

__call__(self, *args, **kwargs)	Alias for *augment()*.
augment(self[, return_batch, hooks])	Augment a batch.
augment_batch(self, batch[, hooks])	**Deprecated**.

Continued on next page

Table 173 – continued from previous page

augment_batch_(self, batch[, parents, hooks])	Augment a single batch in-place.
augment_batches(self, batches[, hooks, ...])	Augment multiple batches.
augment_bounding_boxes(self, ...[, parents, ...])	Augment a batch of bounding boxes.
augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.imgcorruptlike.**Spatter**(*severity=(1, 5), seed=None, name=None, random_state='deprecated', deterministic='deprecated'*)

Bases: imgaug.augmenters.imgcorruptlike._ImgcorruptAugmenterBase

Wrapper around imagecorruptions.corruptions.spatter.

Note: This augmenter only affects images. Other data is not changed.

Added in 0.4.0.

Supported dtypes:

See *apply_spatter()*.

Parameters

- **severity** (*int, optional*) – Strength of the corruption, with valid values being 1 <= severity <= 5.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See __init__().
- **name** (*None or str, optional*) – See __init__().
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method to_deterministic() for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> # doctest: +SKIP
>>> import imgaug.augmenters as iaa
>>> aug = iaa.imgcorruptlike.Spatter(severity=2)
```

Create an augmenter around imagecorruptions.corruptions.spatter. Apply it to images using e.g. aug(images=[image1, image2, ...]).

Methods

__call__(self, *args, **kwargs)	Alias for *augment()*.
augment(self[, return_batch, hooks])	Augment a batch.
augment_batch(self, batch[, hooks])	**Deprecated**.
augment_batch_(self, batch[, parents, hooks])	Augment a single batch in-place.
augment_batches(self, batches[, hooks, ...])	Augment multiple batches.

Continued on next page

Table 174 – continued from previous page

augment_bounding_boxes(self, ...[, parents, ...])	Augment a batch of bounding boxes.
augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.imgcorruptlike.**SpeckleNoise**(*severity=(1, 5), seed=None, name=None, random_state='deprecated', deterministic='deprecated'*)

Bases: imgaug.augmenters.imgcorruptlike._ImgcorruptAugmenterBase

Wrapper around imagecorruptions.corruptions.speckle_noise.

Note: This augmenter only affects images. Other data is not changed.

Added in 0.4.0.

Supported dtypes:

See *apply_speckle_noise()*.

Parameters

- **severity** (*int, optional*) – Strength of the corruption, with valid values being `1 <= severity <= 5`.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> # doctest: +SKIP
>>> import imgaug.augmenters as iaa
>>> aug = iaa.imgcorruptlike.SpeckleNoise(severity=2)
```

Create an augmenter around imagecorruptions.corruptions.speckle_noise. Apply it to images using e.g. `aug(images=[image1, image2, ...])`.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.

Continued on next page

Table 175 – continued from previous page

augment_bounding_boxes(self, ...[, parents, ...])	Augment a batch of bounding boxes.
augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.imgcorruptlike.**ZoomBlur**(*severity=(1, 5), seed=None, name=None, random_state='deprecated', deterministic='deprecated'*)

Bases: imgaug.augmenters.imgcorruptlike._ImgcorruptAugmenterBase

Wrapper around imagecorruptions.corruptions.zoom_blur.

Note: This augmenter only affects images. Other data is not changed.

Added in 0.4.0.

Supported dtypes:

See *apply_zoom_blur()*.

Parameters

- **severity** (*int, optional*) – Strength of the corruption, with valid values being 1 <= severity <= 5.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See __init__().
- **name** (*None or str, optional*) – See __init__().
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method to_deterministic() for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> # doctest: +SKIP
>>> import imgaug.augmenters as iaa
>>> aug = iaa.imgcorruptlike.ZoomBlur(severity=2)
```

Create an augmenter around imagecorruptions.corruptions.zoom_blur. Apply it to images using e.g. aug(images=[image1, image2, ...]).

Methods

__call__(self, *args, **kwargs)	Alias for *augment()*.
augment(self[, return_batch, hooks])	Augment a batch.
augment_batch(self, batch[, hooks])	**Deprecated**.
augment_batch_(self, batch[, parents, hooks])	Augment a single batch in-place.
augment_batches(self, batches[, hooks, ...])	Augment multiple batches.

Continued on next page

Table 176 – continued from previous page

augment_bounding_boxes(self, ...[, parents, ...])	Augment a batch of bounding boxes.
augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

imgaug.augmenters.imgcorruptlike.**apply_brightness**(*x*, *severity=1*, *seed=None*)

Apply brightness from imagecorruptions.

Added in 0.4.0.

Supported dtypes:

See `_call_imgcorrupt_func()`.

Parameters

- **x** (*ndarray*) – Image array. Expected to have shape `(H,W)`, `(H,W,1)` or `(H,W,3)` with dtype `uint8` and a minimum height/width of `32`.
- **severity** (*int, optional*) – Strength of the corruption, with valid values being `1 <= severity <= 5`.
- **seed** (*None or int, optional*) – Seed for the random number generation to use.

Returns Corrupted image.

Return type ndarray

`imgaug.augmenters.imgcorruptlike.`**`apply_contrast`**(*x*, *severity=1*, *seed=None*)

Apply `contrast` from `imagecorruptions`.

Added in 0.4.0.

Supported dtypes:

See `_call_imgcorrupt_func()`.

Parameters

- **x** (*ndarray*) – Image array. Expected to have shape `(H,W)`, `(H,W,1)` or `(H,W,3)` with dtype `uint8` and a minimum height/width of `32`.
- **severity** (*int, optional*) – Strength of the corruption, with valid values being `1 <= severity <= 5`.
- **seed** (*None or int, optional*) – Seed for the random number generation to use.

Returns Corrupted image.

Return type ndarray

`imgaug.augmenters.imgcorruptlike.`**`apply_defocus_blur`**(*x*, *severity=1*, *seed=None*)

Apply `defocus_blur` from `imagecorruptions`.

Added in 0.4.0.

Supported dtypes:

See `_call_imgcorrupt_func()`.

Parameters

- **x** (*ndarray*) – Image array. Expected to have shape `(H,W)`, `(H,W,1)` or `(H,W,3)` with dtype `uint8` and a minimum height/width of `32`.
- **severity** (*int, optional*) – Strength of the corruption, with valid values being `1 <= severity <= 5`.
- **seed** (*None or int, optional*) – Seed for the random number generation to use.

Returns Corrupted image.

Return type ndarray

`imgaug.augmenters.imgcorruptlike.`**`apply_elastic_transform`**(*image*, *severity=1*, *seed=None*)

Apply `elastic_transform` from `imagecorruptions`.

Added in 0.4.0.

Supported dtypes:

See `_call_imgcorrupt_func()`.

Parameters

- **image** (*ndarray*) – Image array. Expected to have shape `(H,W)`, `(H,W,1)` or `(H,W,3)` with dtype `uint8` and a minimum height/width of `32`.
- **severity** (*int, optional*) – Strength of the corruption, with valid values being `1 <= severity <= 5`.
- **seed** (*None or int, optional*) – Seed for the random number generation to use.

Returns Corrupted image.

Return type ndarray

`imgaug.augmenters.imgcorruptlike.`**`apply_fog`**(*x*, *severity=1*, *seed=None*)

Apply `fog` from `imagecorruptions`.

Added in 0.4.0.

Supported dtypes:

See `_call_imgcorrupt_func()`.

Parameters

- **x** (*ndarray*) – Image array. Expected to have shape `(H,W)`, `(H,W,1)` or `(H,W,3)` with dtype `uint8` and a minimum height/width of `32`.
- **severity** (*int, optional*) – Strength of the corruption, with valid values being `1 <= severity <= 5`.
- **seed** (*None or int, optional*) – Seed for the random number generation to use.

Returns Corrupted image.

Return type ndarray

`imgaug.augmenters.imgcorruptlike.`**`apply_frost`**(*x*, *severity=1*, *seed=None*)

Apply `frost` from `imagecorruptions`.

Added in 0.4.0.

Supported dtypes:

See `_call_imgcorrupt_func()`.

Parameters

- **x** (*ndarray*) – Image array. Expected to have shape `(H,W)`, `(H,W,1)` or `(H,W,3)` with dtype `uint8` and a minimum height/width of `32`.
- **severity** (*int, optional*) – Strength of the corruption, with valid values being `1 <= severity <= 5`.
- **seed** (*None or int, optional*) – Seed for the random number generation to use.

Returns Corrupted image.

Return type ndarray

`imgaug.augmenters.imgcorruptlike.`**`apply_gaussian_blur`**(*x*, *severity=1*, *seed=None*)

Apply `gaussian_blur` from `imagecorruptions`.

Added in 0.4.0.

Supported dtypes:

See `_call_imgcorrupt_func()`.

Parameters

- **x** (*ndarray*) – Image array. Expected to have shape `(H,W)`, `(H,W,1)` or `(H,W,3)` with dtype `uint8` and a minimum height/width of `32`.
- **severity** (*int, optional*) – Strength of the corruption, with valid values being `1 <= severity <= 5`.
- **seed** (*None or int, optional*) – Seed for the random number generation to use.

Returns Corrupted image.

Return type ndarray

`imgaug.augmenters.imgcorruptlike.`**`apply_gaussian_noise`**(*x*, *severity=1*, *seed=None*)

Apply `gaussian_noise` from `imagecorruptions`.

Added in 0.4.0.

Supported dtypes:

See `_call_imgcorrupt_func()`.

Parameters

- **x** (*ndarray*) – Image array. Expected to have shape `(H,W)`, `(H,W,1)` or `(H,W,3)` with dtype `uint8` and a minimum height/width of `32`.
- **severity** (*int, optional*) – Strength of the corruption, with valid values being `1 <= severity <= 5`.
- **seed** (*None or int, optional*) – Seed for the random number generation to use.

Returns Corrupted image.

Return type ndarray

`imgaug.augmenters.imgcorruptlike.`**`apply_glass_blur`**(*x*, *severity=1*, *seed=None*)

Apply `glass_blur` from `imagecorruptions`.

Added in 0.4.0.

Supported dtypes:

See `_call_imgcorrupt_func()`.

Parameters

- **x** (*ndarray*) – Image array. Expected to have shape `(H,W)`, `(H,W,1)` or `(H,W,3)` with dtype `uint8` and a minimum height/width of `32`.
- **severity** (*int, optional*) – Strength of the corruption, with valid values being `1 <= severity <= 5`.
- **seed** (*None or int, optional*) – Seed for the random number generation to use.

Returns Corrupted image.

Return type ndarray

`imgaug.augmenters.imgcorruptlike.`**`apply_impulse_noise`**(*x*, *severity=1*, *seed=None*)

Apply `impulse_noise` from `imagecorruptions`.

Added in 0.4.0.

Supported dtypes:

See `_call_imgcorrupt_func()`.

Parameters

- **x** (*ndarray*) – Image array. Expected to have shape `(H,W)`, `(H,W,1)` or `(H,W,3)` with dtype `uint8` and a minimum height/width of `32`.
- **severity** (*int, optional*) – Strength of the corruption, with valid values being `1 <= severity <= 5`.
- **seed** (*None or int, optional*) – Seed for the random number generation to use.

Returns Corrupted image.

Return type ndarray

`imgaug.augmenters.imgcorruptlike.`**`apply_jpeg_compression`**(*x*, *severity=1*, *seed=None*)

Apply `jpeg_compression` from `imagecorruptions`.

Added in 0.4.0.

Supported dtypes:

See `_call_imgcorrupt_func()`.

Parameters

- **x** (*ndarray*) – Image array. Expected to have shape `(H,W)`, `(H,W,1)` or `(H,W,3)` with dtype `uint8` and a minimum height/width of `32`.
- **severity** (*int, optional*) – Strength of the corruption, with valid values being `1 <= severity <= 5`.
- **seed** (*None or int, optional*) – Seed for the random number generation to use.

Returns Corrupted image.

Return type ndarray

`imgaug.augmenters.imgcorruptlike.`**`apply_motion_blur`**(*x*, *severity=1*, *seed=None*)

Apply `motion_blur` from `imagecorruptions`.

Added in 0.4.0.

Supported dtypes:

See `_call_imgcorrupt_func()`.

Parameters

- **x** (*ndarray*) – Image array. Expected to have shape `(H,W)`, `(H,W,1)` or `(H,W,3)` with dtype `uint8` and a minimum height/width of `32`.
- **severity** (*int, optional*) – Strength of the corruption, with valid values being `1 <= severity <= 5`.
- **seed** (*None or int, optional*) – Seed for the random number generation to use.

Returns Corrupted image.

Return type ndarray

`imgaug.augmenters.imgcorruptlike.`**`apply_pixelate`**(*x*, *severity=1*, *seed=None*)

Apply `pixelate` from `imagecorruptions`.

Added in 0.4.0.

Supported dtypes:

See `_call_imgcorrupt_func()`.

Parameters

- **x** (*ndarray*) – Image array. Expected to have shape `(H,W)`, `(H,W,1)` or `(H,W,3)` with dtype `uint8` and a minimum height/width of `32`.
- **severity** (*int, optional*) – Strength of the corruption, with valid values being `1 <= severity <= 5`.
- **seed** (*None or int, optional*) – Seed for the random number generation to use.

Returns Corrupted image.

Return type ndarray

`imgaug.augmenters.imgcorruptlike.`**`apply_saturate`**(*x*, *severity=1*, *seed=None*)

Apply `saturate` from `imagecorruptions`.

Added in 0.4.0.

Supported dtypes:

See `_call_imgcorrupt_func()`.

Parameters

- **x** (*ndarray*) – Image array. Expected to have shape `(H,W)`, `(H,W,1)` or `(H,W,3)` with dtype `uint8` and a minimum height/width of `32`.
- **severity** (*int, optional*) – Strength of the corruption, with valid values being `1 <= severity <= 5`.
- **seed** (*None or int, optional*) – Seed for the random number generation to use.

Returns Corrupted image.

Return type ndarray

`imgaug.augmenters.imgcorruptlike.`**`apply_shot_noise`**(*x*, *severity=1*, *seed=None*)

Apply `shot_noise` from `imagecorruptions`.

Added in 0.4.0.

Supported dtypes:

See `_call_imgcorrupt_func()`.

Parameters

- **x** (*ndarray*) – Image array. Expected to have shape `(H,W)`, `(H,W,1)` or `(H,W,3)` with dtype `uint8` and a minimum height/width of `32`.
- **severity** (*int, optional*) – Strength of the corruption, with valid values being `1 <= severity <= 5`.
- **seed** (*None or int, optional*) – Seed for the random number generation to use.

Returns Corrupted image.

Return type ndarray

`imgaug.augmenters.imgcorruptlike.`**`apply_snow`**(*x*, *severity=1*, *seed=None*)

Apply `snow` from `imagecorruptions`.

Added in 0.4.0.

Supported dtypes:

See `_call_imgcorrupt_func()`.

Parameters

- **x** (*ndarray*) – Image array. Expected to have shape `(H,W)`, `(H,W,1)` or `(H,W,3)` with dtype `uint8` and a minimum height/width of `32`.
- **severity** (*int, optional*) – Strength of the corruption, with valid values being `1 <= severity <= 5`.
- **seed** (*None or int, optional*) – Seed for the random number generation to use.

Returns Corrupted image.

Return type ndarray

`imgaug.augmenters.imgcorruptlike.`**`apply_spatter`**(*x*, *severity=1*, *seed=None*)
Apply `spatter` from `imagecorruptions`.

Added in 0.4.0.

Supported dtypes:

See `_call_imgcorrupt_func()`.

Parameters

- **x** (*ndarray*) – Image array. Expected to have shape `(H,W)`, `(H,W,1)` or `(H,W,3)` with dtype `uint8` and a minimum height/width of `32`.
- **severity** (*int, optional*) – Strength of the corruption, with valid values being `1 <= severity <= 5`.
- **seed** (*None or int, optional*) – Seed for the random number generation to use.

Returns Corrupted image.

Return type ndarray

`imgaug.augmenters.imgcorruptlike.`**`apply_speckle_noise`**(*x*, *severity=1*, *seed=None*)
Apply `speckle_noise` from `imagecorruptions`.

Added in 0.4.0.

Supported dtypes:

See `_call_imgcorrupt_func()`.

Parameters

- **x** (*ndarray*) – Image array. Expected to have shape `(H,W)`, `(H,W,1)` or `(H,W,3)` with dtype `uint8` and a minimum height/width of `32`.
- **severity** (*int, optional*) – Strength of the corruption, with valid values being `1 <= severity <= 5`.
- **seed** (*None or int, optional*) – Seed for the random number generation to use.

Returns Corrupted image.

Return type ndarray

`imgaug.augmenters.imgcorruptlike.`**`apply_zoom_blur`**(*x*, *severity=1*, *seed=None*)
Apply `zoom_blur` from `imagecorruptions`.

Added in 0.4.0.

Supported dtypes:

See `_call_imgcorrupt_func()`.

Parameters

- **x** (*ndarray*) – Image array. Expected to have shape `(H,W)`, `(H,W,1)` or `(H,W,3)` with dtype `uint8` and a minimum height/width of `32`.
- **severity** (*int, optional*) – Strength of the corruption, with valid values being `1 <= severity <= 5`.
- **seed** (*None or int, optional*) – Seed for the random number generation to use.

Returns Corrupted image.

Return type ndarray

`imgaug.augmenters.imgcorruptlike.`**`get_corruption_names`**(*subset='common'*)

Get a named subset of image corruption functions.

Note: This function returns the augmentation names (as strings) *and* the corresponding augmentation functions, while `get_corruption_names()` in `imagecorruptions` only returns the augmentation names.

Added in 0.4.0.

Parameters **subset** (*{'common', 'validation', 'all'}, optional.*) – Name of the subset of image corruption functions.

Returns

- *list of str* – Names of the corruption methods, e.g. "gaussian_noise".
- *list of callable* – Function corresponding to the name. Is one of the `apply_*()` functions in this module. Apply e.g. via `func(image, severity=2, seed=123)`.

13.31 imgaug.augmenters.meta

Augmenters that don't apply augmentations themselves, but are needed for meta usage.

List of augmenters:

- *Augmenter* (base class for all augmenters)
- *Sequential*
- *SomeOf*
- *OneOf*
- *Sometimes*
- *WithChannels*
- *Identity*
- *Noop*
- *Lambda*
- *AssertLambda*
- *AssertShape*

- *ChannelShuffle*

Note: *WithColorspace* is in `color.py`.

class imgaug.augmenters.meta.**AssertLambda**(*func_images=None*, *func_heatmaps=None*, *func_segmentation_maps=None*, *func_keypoints=None*, *func_bounding_boxes=None*, *func_polygons=None*, *func_line_strings=None*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Lambda*

Assert conditions based on lambda-function to be the case for input data.

This augmenter applies a lambda function to each image or other input. The lambda function must return `True` or `False`. If `False` is returned, an assertion error is produced.

This is useful to ensure that generic assumption about the input data are actually the case and error out early otherwise.

Supported dtypes:

- `uint8`: yes; fully tested
- `uint16`: yes; tested
- `uint32`: yes; tested
- `uint64`: yes; tested
- `int8`: yes; tested
- `int16`: yes; tested
- `int32`: yes; tested
- `int64`: yes; tested
- `float16`: yes; tested
- `float32`: yes; tested
- `float64`: yes; tested
- `float128`: yes; tested
- `bool`: yes; tested

Parameters

- **func_images** (*None or callable, optional*) – The function to call for each batch of images. It must follow the form:

```
function(images, random_state, parents, hooks)
```

and return either `True` (valid input) or `False` (invalid input). It essentially re-uses the interface of `_augment_images()`.

- **func_heatmaps** (*None or callable, optional*) – The function to call for each batch of heatmaps. It must follow the form:

```
function(heatmaps, random_state, parents, hooks)
```

and return either `True` (valid input) or `False` (invalid input). It essentially re-uses the interface of `_augment_heatmaps()`.

- **func_segmentation_maps** (*None or callable, optional*) – The function to call for each batch of segmentation maps. It must follow the form:

```
function(segmaps, random_state, parents, hooks)
```

and return either `True` (valid input) or `False` (invalid input). It essentially re-uses the interface of `_augment_segmentation_maps()`.

- **func_keypoints** (*None or callable, optional*) – The function to call for each batch of keypoints. It must follow the form:

```
function(keypoints_on_images, random_state, parents, hooks)
```

and return either `True` (valid input) or `False` (invalid input). It essentially re-uses the interface of `_augment_keypoints()`.

- **func_bounding_boxes** (*None or callable, optional*) – The function to call for each batch of bounding boxes. It must follow the form:

```
function(bounding_boxes_on_images, random_state, parents, hooks)
```

and return either `True` (valid input) or `False` (invalid input). It essentially re-uses the interface of `_augment_bounding_boxes()`.

Added in 0.4.0.

- **func_polygons** (*None or callable, optional*) – The function to call for each batch of polygons. It must follow the form:

```
function(polygons_on_images, random_state, parents, hooks)
```

and return either `True` (valid input) or `False` (invalid input). It essentially re-uses the interface of `_augment_polygons()`.

- **func_line_strings** (*None or callable, optional*) – The function to call for each batch of line strings. It must follow the form:

```
function(line_strings_on_images, random_state, parents, hooks)
```

and return either `True` (valid input) or `False` (invalid input). It essentially re-uses the interface of `_augment_line_strings()`.

Added in 0.4.0.

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.

- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.
`get_parameters`(self)	See *get_parameters()*.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`localize_random_state_`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`pool`(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
`remove_augmenters`(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
`remove_augmenters_`(self, func[, parents])	Remove in-place children of this augmenter that match a condition.

Continued on next page

Table 177 – continued from previous page

remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.meta.**AssertShape**(*shape*, *check_images=True*, *check_heatmaps=True*, *check_segmentation_maps=True*, *check_keypoints=True*, *check_bounding_boxes=True*, *check_polygons=True*, *check_line_strings=True*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Lambda*

Assert that inputs have a specified shape.

Supported dtypes:

- uint8: yes; fully tested
- uint16: yes; tested
- uint32: yes; tested
- uint64: yes; tested
- int8: yes; tested
- int16: yes; tested
- int32: yes; tested
- int64: yes; tested
- float16: yes; tested
- float32: yes; tested
- float64: yes; tested
- float128: yes; tested
- bool: yes; tested

Parameters

- **shape** (*tuple*) – The expected shape, given as a tuple. The number of entries in the tuple must match the number of dimensions, i.e. it must contain four entries for (N, H, W, C). If only a single entity is augmented, e.g. via *augment_image()*, then N is 1 in the input to this augmenter. Images that don't have a channel axis will automatically have one assigned, i.e. C is at least 1. For each component of the tuple one of the following datatypes may be used:
  - If a component is None, any value for that dimensions is accepted.

  - If a component is `int`, exactly that value (and no other one) will be accepted for that dimension.
  - If a component is a `tuple` of two `int` s with values `a` and `b`, only a value within the interval `[a, b)` will be accepted for that dimension.
  - If an entry is a `list` of `int` s, only a value from that `list` will be accepted for that dimension.

- **check_images** (*bool, optional*) – Whether to validate input images via the given shape.
- **check_heatmaps** (*bool, optional*) – Whether to validate input heatmaps via the given shape. The number of heatmaps will be verified as `N`. For each *`HeatmapsOnImage`* instance its array's height and width will be verified as `H` and `W`, but not the channel count.
- **check_segmentation_maps** (*bool, optional*) – Whether to validate input segmentation maps via the given shape. The number of segmentation maps will be verified as `N`. For each *`SegmentationMapOnImage`* instance its array's height and width will be verified as `H` and `W`, but not the channel count.
- **check_keypoints** (*bool, optional*) – Whether to validate input keypoints via the given shape. This will check (a) the number of keypoints and (b) for each *`KeypointsOnImage`* instance the `.shape` attribute, i.e. the shape of the corresponding image.
- **check_bounding_boxes** (*bool, optional*) – Whether to validate input bounding boxes via the given shape. This will check (a) the number of bounding boxes and (b) for each *`BoundingBoxesOnImage`* instance the `.shape` attribute, i.e. the shape of the corresponding image.

  Added in 0.4.0.
- **check_polygons** (*bool, optional*) – Whether to validate input polygons via the given shape. This will check (a) the number of polygons and (b) for each *`PolygonsOnImage`* instance the `.shape` attribute, i.e. the shape of the corresponding image.
- **check_line_strings** (*bool, optional*) – Whether to validate input line strings via the given shape. This will check (a) the number of line strings and (b) for each *`LineStringsOnImage`* instance the `.shape` attribute, i.e. the shape of the corresponding image.

  Added in 0.4.0.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> seq = iaa.Sequential([
>>>     iaa.AssertShape((None, 32, 32, 3)),
>>>     iaa.Fliplr(0.5)
>>> ])
```

Verify first for each image batch if it contains a variable number of `32x32` images with `3` channels each. Only if that check succeeds, the horizontal flip will be executed. Otherwise an assertion error will be raised.

```
>>> seq = iaa.Sequential([
>>>     iaa.AssertShape((None, (32, 64), 32, [1, 3])),
>>>     iaa.Fliplr(0.5)
>>> ])
```

Similar to the above example, but now the height may be in the interval `[32, 64)` and the number of channels may be either `1` or `3`.

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.

Continued on next page

Table 178 – continued from previous page

find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.meta.**Augmenter**(*seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: object

Base class for Augmenter objects. All augmenters derive from this class.

Parameters

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Seed to use for this augmenter's random number generator (RNG) or alternatively an RNG itself. Setting this parameter allows to control/influence the random number sampling of this specific augmenter without affecting other augmenters. Usually, there is no need to set this parameter.
  - If None: The global RNG is used (shared by all augmenters).
  - If int: The value will be used as a seed for a new *RNG* instance.
  - If *RNG*: The RNG instance will be used without changes.
  - If Generator: A new *RNG* instance will be created, containing that generator.
  - If BitGenerator: Will be wrapped in a Generator. Then similar behaviour to Generator parameters.
  - If SeedSequence: Will be wrapped in a new bit generator and Generator. Then similar behaviour to Generator parameters.
  - If RandomState: Similar behaviour to Generator. Outdated in numpy 1.17+.

  If a new bit generator has to be created, it will be an instance of numpy.random.SFC64.

  Added in 0.4.0.

- **name** (*None or str, optional*) – Name given to the Augmenter instance. This name is used when converting the instance to a string, e.g. for `print` statements. It is also used for `find`, `remove` or similar operations on augmenters with children. If `None`, `UnnamedX` will be used as the name, where `X` is the Augmenter's class name.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
augment(self[, return_batch, hooks])	Augment a batch.
augment_batch(self, batch[, hooks])	**Deprecated**.
*augment_batch_*(self, batch[, parents, hooks])	Augment a single batch in-place.
augment_batches(self, batches[, hooks, ...])	Augment multiple batches.
augment_bounding_boxes(self, ...[, parents, ...])	Augment a batch of bounding boxes.
augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
*copy_random_state_*(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	Get the parameters of this augmenter.

Continued on next page

Table 179 – continued from previous page

localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
*localize_random_state_*(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
*remove_augmenters_*(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
*seed_*(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

augment (*self*, *return_batch=False*, *hooks=None*, ***kwargs*)

Augment a batch.

This method is a wrapper around *UnnormalizedBatch* and *augment_batch()*. Hence, it supports the same datatypes as *UnnormalizedBatch*.

If *return_batch* was set to `False` (the default), the method will return a tuple of augmentables. It will return the same types of augmentables (but in augmented form) as input into the method. This behaviour is partly specific to the python version:

- In **python 3.6+** (if `return_batch=False`):
  - Any number of augmentables may be provided as input.
  - None of the provided named arguments *has to be image* or *images* (but of coarse you *may* provide them).
  - The return order matches the order of the named arguments, e.g. `x_aug, y_aug, z_aug = augment(X=x, Y=y, Z=z)`.
- In **python <3.6** (if `return_batch=False`):
  - One or two augmentables may be used as input, not more than that.
  - One of the input arguments has to be *image* or *images*.
  - The augmented images are *always* returned first, independent of the input argument order, e.g. `a_aug, b_aug = augment(b=b, images=a)`. This also means that the output of the function can only be one of the following three cases: a batch, list/array of images, tuple of images and something (like images + segmentation maps).

If *return_batch* was set to `True`, an instance of *UnnormalizedBatch* will be returned. The output is the same for all python version and any number or combination of augmentables may be provided.

So, to keep code downward compatible for python <3.6, use one of the following three options:

- Use `batch = augment(images=X, ..., return_batch=True)`.
- Call `images = augment(images=X)`.
- Call `images, other = augment(images=X, <something_else>=Y)`.

All augmentables must be provided as named arguments. E.g. `augment(<array>)` will crash, but `augment(images=<array>)` will work.

Parameters

- **image** (*None or (H,W,C) ndarray or (H,W) ndarray, optional*) – The image to augment. Only this or *images* can be set, not both. If *return_batch* is `False` and the python version is below 3.6, either this or *images* **must** be provided.
- **images** (*None or (N,H,W,C) ndarray or (N,H,W) ndarray or iterable of (H,W,C) ndarray or iterable of (H,W) ndarray, optional*) – The images to augment. Only this or *image* can be set, not both. If *return_batch* is `False` and the python version is below 3.6, either this or *image* **must** be provided.
- **heatmaps** (*None or (N,H,W,C) ndarray or imgaug.augmentables.heatmaps.HeatmapsOnImage or iterable of (H,W,C) ndarray or iterable of imgaug.augmentables.heatmaps.HeatmapsOnImage, optional*) – The heatmaps to augment. If anything else than *HeatmapsOnImage*, then the number of heatmaps must match the number of images provided via parameter *images*. The number is contained either in `N` or the first iterable's size.
- **segmentation_maps** (*None or (N,H,W) ndarray or imgaug.augmentables.segmaps.SegmentationMapsOnImage or iterable of (H,W) ndarray or iterable of imgaug.augmentables.segmaps.SegmentationMapsOnImage, optional*) – The segmentation maps to augment. If anything else than *SegmentationMapsOnImage*, then the number of segmaps must match the number of images provided via parameter *images*. The number is contained either in `N` or the first iterable's size.
- **keypoints** (*None or list of (N,K,2) ndarray or tuple of number or imgaug.augmentables.kps.Keypoint or iterable of (K,2) ndarray or iterable of tuple of number or iterable of imgaug.augmentables.kps.Keypoint or iterable of imgaug.augmentables.kps.KeypointOnImage or iterable of iterable of tuple of number or iterable of iterable of imgaug.augmentables.kps.Keypoint, optional*) – The keypoints to augment. If a tuple (or iterable(s) of tuple), then iterpreted as `(x,y)` coordinates and must hence contain two numbers. A single tuple represents a single coordinate on one image, an iterable of tuples the coordinates on one image and an iterable of iterable of tuples the coordinates on several images. Analogous if *Keypoint* instances are used instead of tuples. If an ndarray, then `N` denotes the number of images and `K` the number of keypoints on each image. If anything else than *KeypointsOnImage* is provided, then the number of keypoint groups must match the number of images provided via parameter *images*. The number is contained e.g. in `N` or in case of "iterable of iterable of tuples" in the first iterable's size.
- **bounding_boxes** (*None or (N,B,4) ndarray or tuple of number or imgaug.augmentables.bbs.BoundingBox or imgaug.augmentables.bbs.BoundingBoxesOnImage or iterable of (B,4) ndarray or iterable of tuple of number or iterable of imgaug.augmentables.bbs.BoundingBox or iterable of imgaug.augmentables.bbs.BoundingBoxesOnImage or iterable of iterable of tuple of number or iterable of iterable imgaug.augmentables.bbs.BoundingBox, optional*) – The bounding boxes to augment. This is analogous to the *keypoints* parameter. However, each tuple – and also the last index in case of arrays – has size `4`, denoting the bounding box coordinates `x1`, `y1`, `x2` and `y2`.
- **polygons** (*None or (N,#polys,#points,2) ndarray or imgaug.augmentables.polys.Polygon or imgaug.augmentables.polys.PolygonsOnImage or iterable of (#polys,#points,2) ndarray or iterable of tuple of number or iterable of imgaug.augmentables.kps.Keypoint or iterable of imgaug.augmentables.polys.Polygon or iterable of imgaug.augmentables.polys.PolygonsOnImage or iterable of iterable of (#points,2)*

ndarray or iterable of iterable of tuple of number or iterable of iterable of imgaug.augmentables.kps.Keypoint or iterable of iterable of imgaug.augmentables.polys.Polygon or iterable of iterable of iterable of tuple of number or iterable of iterable of iterable of tuple of imgaug.augmentables.kps.Keypoint, optional) – The polygons to augment. This is similar to the *keypoints* parameter. However, each polygon may be made up of several ‘‘(x,y) ‘‘coordinates (three or more are required for valid polygons). The following datatypes will be interpreted as a single polygon on a single image:

- `imgaug.augmentables.polys.Polygon`
- `iterable of tuple of number`
- `iterable of imgaug.augmentables.kps.Keypoint`

The following datatypes will be interpreted as multiple polygons on a single image:

- `imgaug.augmentables.polys.PolygonsOnImage`
- `iterable of imgaug.augmentables.polys.Polygon`
- `iterable of iterable of tuple of number`
- `iterable of iterable of imgaug.augmentables.kps.Keypoint`
- `iterable of iterable of imgaug.augmentables.polys.Polygon`

The following datatypes will be interpreted as multiple polygons on multiple images:

- `(N,#polys,#points,2) ndarray`
- `iterable of (#polys,#points,2) ndarray`
- `iterable of iterable of (#points,2) ndarray`
- `iterable of iterable of iterable of tuple of number`
- `iterable of iterable of iterable of tuple of imgaug.augmentables.kps.Keypoint`

- **line_strings** (*None or (N,#lines,#points,2) ndarray or imgaug.augmentables.lines.LineString or imgaug.augmentables.lines.LineStringOnImage or iterable of (#polys,#points,2) ndarray or iterable of tuple of number or iterable of imgaug.augmentables.kps.Keypoint or iterable of imgaug.augmentables.lines.LineString or iterable of imgaug.augmentables.lines.LineStringOnImage or iterable of iterable of (#points,2) ndarray or iterable of iterable of tuple of number or iterable of iterable of imgaug.augmentables.kps.Keypoint or iterable of iterable of imgaug.augmentables.lines.LineString or iterable of iterable of iterable of tuple of number or iterable of iterable of iterable of tuple of imgaug.augmentables.kps.Keypoint, optional*) – The line strings to augment. See *polygons*, which behaves similarly.
- **return_batch** (*bool, optional*) – Whether to return an instance of `UnnormalizedBatch`. If the python version is below 3.6 and more than two augmentables were provided (e.g. images, keypoints and polygons), then this must be set to `True`. Otherwise an error will be raised.
- **hooks** (*None or imgaug.imgaug.HooksImages, optional*) – Hooks object to dynamically interfere with the augmentation process.

Returns If *return_batch* was set to `True`, a instance of `UnnormalizedBatch` will be returned. If *return_batch* was set to `False`, a tuple of augmentables will be returned, e.g. `(augmented images, augmented keypoints)`. The datatypes match the input datatypes of the corresponding named arguments. In python <3.6, augmented images are

always the first entry in the returned tuple. In python 3.6+ the order matches the order of the named arguments.

Return type tuple or *imgaug.augmentables.batches.UnnormalizedBatch*

Examples

```
>>> import numpy as np
>>> import imgaug as ia
>>> import imgaug.augmenters as iaa
>>> aug = iaa.Affine(rotate=(-25, 25))
>>> image = np.zeros((64, 64, 3), dtype=np.uint8)
>>> keypoints = [(10, 20), (30, 32)]  # (x,y) coordinates
>>> images_aug, keypoints_aug = aug.augment(
>>>     image=image, keypoints=keypoints)
```

Create a single image and a set of two keypoints on it, then augment both by applying a random rotation between `-25` deg and `+25` deg. The sampled rotation value is automatically aligned between image and keypoints. Note that in python <3.6, augmented images will always be returned first, independent of the order of the named input arguments. So `keypoints_aug, images_aug = aug.augment(keypoints=keypoints, image=image)` would **not** be correct (but in python 3.6+ it would be).

```
>>> import numpy as np
>>> import imgaug as ia
>>> import imgaug.augmenters as iaa
>>> from imgaug.augmentables.bbs import BoundingBox
>>> aug = iaa.Affine(rotate=(-25, 25))
>>> images = [np.zeros((64, 64, 3), dtype=np.uint8),
>>>           np.zeros((32, 32, 3), dtype=np.uint8)]
>>> keypoints = [[(10, 20), (30, 32)],   # KPs on first image
>>>              [(22, 10), (12, 14)]]   # KPs on second image
>>> bbs = [
>>>           [BoundingBox(x1=5, y1=5, x2=50, y2=45)],
>>>           [BoundingBox(x1=4, y1=6, x2=10, y2=15),
>>>            BoundingBox(x1=8, y1=9, x2=16, y2=30)]
>>>       ]  # one BB on first image, two BBs on second image
>>> batch_aug = aug.augment(
>>>     images=images, keypoints=keypoints, bounding_boxes=bbs,
>>>     return_batch=True)
```

Create two images of size `64x64` and `32x32`, two sets of keypoints (each containing two keypoints) and two sets of bounding boxes (the first containing one bounding box, the second two bounding boxes). These augmentables are then augmented by applying random rotations between `-25` deg and `+25` deg to them. The rotation values are sampled by image and aligned between all augmentables on the same image. The method finally returns an instance of *UnnormalizedBatch* from which the augmented data can be retrieved via `batch_aug.images_aug`, `batch_aug.keypoints_aug`, and `batch_aug.bounding_boxes_aug`. In python 3.6+, *return_batch* can be kept at `False` and the augmented data can be retrieved as `images_aug, keypoints_aug, bbs_aug = augment(...)`.

augment_batch(*self*, *batch*, *hooks=None*)

Deprecated. Use `augment_batch_()` instead. *augment_batch()* was renamed to *augment_batch_()* as it changes all **_unaug* attributes of batches in-place. Note that *augment_batch_()* has now a *parents* parameter. Calls of the style *augment_batch(batch, hooks)* must be changed to *augment_batch(batch, hooks=hooks)*.

Augment a single batch.

Deprecated since 0.4.0.

augment_batch_ (*self*, *batch*, *parents=None*, *hooks=None*)
Augment a single batch in-place.

Added in 0.4.0.

Parameters

- **batch** (*imgaug.augmentables.batches.Batch or imgaug.augmentables.batches.UnnormalizedBatch or imgaug.augmentables.batch._BatchInAugmentation*) – A single batch to augment.

  If *imgaug.augmentables.batches.UnnormalizedBatch* or *imgaug.augmentables.batches.Batch*, then the `*_aug` attributes may be modified in-place, while the `*_unaug` attributes will not be modified. If `imgaug.augmentables.batches._BatchInAugmentation`, then all attributes may be modified in-place.

- **parents** (*None or list of imgaug.augmenters.Augmenter, optional*) – Parent augmenters that have previously been called before the call to this function. Usually you can leave this parameter as `None`. It is set automatically for child augmenters.

- **hooks** (*None or imgaug.HooksImages, optional*) – HooksImages object to dynamically interfere with the augmentation process.

Returns Augmented batch.

Return type *imgaug.augmentables.batches.Batch* or *imgaug.augmentables.batches.UnnormalizedBatch*

augment_batches (*self*, *batches*, *hooks=None*, *background=False*)
Augment multiple batches.

In contrast to other `augment_*` method, this one **yields** batches instead of returning a full list. This is more suited for most training loops.

This method also also supports augmentation on multiple cpu cores, activated via the *background* flag. If the *background* flag is activated, an instance of *Pool* will be spawned using all available logical CPU cores and an `output_buffer_size` of `C*10`, where `C` is the number of logical CPU cores. I.e. a maximum of `C*10` batches will be somewhere in the augmentation pipeline (or waiting to be retrieved by downstream functions) before this method temporarily stops the loading of new batches from *batches*.

Parameters

- **batches** (*imgaug.augmentables.batches.Batch or imgaug.augmentables.batches.UnnormalizedBatch or iterable of imgaug.augmentables.batches.Batch or iterable of imgaug.augmentables.batches.UnnormalizedBatch*) – A single batch or a list of batches to augment.

- **hooks** (*None or imgaug.HooksImages, optional*) – HooksImages object to dynamically interfere with the augmentation process.

- **background** (*bool, optional*) – Whether to augment the batches in background processes. If `True`, hooks can currently not be used as that would require pickling functions. Note that multicore augmentation distributes the batches onto different CPU cores. It does *not* split the data *within* batches. It is therefore *not* sensible to use `background=True` to augment a single batch. Only use it for multiple batches. Note also that multicore augmentation needs some time to start. It is therefore not recommended to use it for very few batches.

Yields *imgaug.augmentables.batches.Batch or imgaug.augmentables.batches.UnnormalizedBatch or iterable of imgaug.augmentables.batches.Batch or iterable of imgaug.augmentables.batches.UnnormalizedBatch* – Augmented batches.

augment_bounding_boxes(*self*, *bounding_boxes_on_images*, *parents=None*, *hooks=None*)

Augment a batch of bounding boxes.

This is the corresponding function to *Augmenter.augment_images()*, just for bounding boxes. Usually you will want to call *Augmenter.augment_images()* with a list of images, e.g. `augment_images([A, B, C])` and then `augment_bounding_boxes()` with the corresponding list of bounding boxes on these images, e.g. `augment_bounding_boxes([Abb, Bbb, Cbb])`, where `Abb` are the bounding boxes on image `A`.

Make sure to first convert the augmenter(s) to deterministic states before augmenting images and their corresponding bounding boxes, e.g. by

```
>>> import imgaug.augmenters as iaa
>>> from imgaug.augmentables.bbs import BoundingBox
>>> from imgaug.augmentables.bbs import BoundingBoxesOnImage
>>> A = B = C = np.ones((10, 10), dtype=np.uint8)
>>> Abb = Bbb = Cbb = BoundingBoxesOnImage([
>>>     BoundingBox(1, 1, 9, 9)], (10, 10))
>>> seq = iaa.Fliplr(0.5)
>>> seq_det = seq.to_deterministic()
>>> imgs_aug = seq_det.augment_images([A, B, C])
>>> bbs_aug = seq_det.augment_bounding_boxes([Abb, Bbb, Cbb])
```

Otherwise, different random values will be sampled for the image and bounding box augmentations, resulting in different augmentations (e.g. images might be rotated by `30deg` and bounding boxes by `-10deg`). Also make sure to call *Augmenter.to_deterministic()* again for each new batch, otherwise you would augment all batches in the same way.

Note that there is also *Augmenter.augment()*, which automatically handles the random state alignment.

Parameters

- **bounding_boxes_on_images** (*imgaug.augmentables.bbs.BoundingBoxesOnImage or list of imgaug.augmentables.bbs.BoundingBoxesOnImage*) – The bounding boxes to augment. Either a single instance of *BoundingBoxesOnImage* or a list of such instances, with each one of them containing the bounding boxes of a single image.
- **parents** (*None or list of imgaug.augmenters.meta.Augmenter, optional*) – Parent augmenters that have previously been called before the call to this function. Usually you can leave this parameter as `None`. It is set automatically for child augmenters.
- **hooks** (*None or imgaug.imgaug.HooksKeypoints, optional*) – *HooksKeypoints* object to dynamically interfere with the augmentation process.

Returns Augmented bounding boxes.

Return type *imgaug.augmentables.bbs.BoundingBoxesOnImage* or list of imgaug.augmentables.bbs.BoundingBoxesOnImage

augment_heatmaps(*self*, *heatmaps*, *parents=None*, *hooks=None*)

Augment a batch of heatmaps.

Parameters

- **heatmaps** (*imgaug.augmentables.heatmaps.HeatmapsOnImage or list of imgaug.augmentables.heatmaps.HeatmapsOnImage*) – Heatmap(s) to augment. Either

a single heatmap or a list of heatmaps.

- **parents** (*None or list of imgaug.augmenters.meta.Augmenter, optional*) – Parent augmenters that have previously been called before the call to this function. Usually you can leave this parameter as `None`. It is set automatically for child augmenters.
- **hooks** (*None or imaug.imgaug.HooksHeatmaps, optional*) – `HooksHeatmaps` object to dynamically interfere with the augmentation process.

Returns Corresponding augmented heatmap(s).

Return type *imgaug.augmentables.heatmaps.HeatmapsOnImage* or list of imgaug.augmentables.heatmaps.HeatmapsOnImage

augment_image(*self*, *image*, *hooks=None*)

Augment a single image.

Parameters

- **image** (*(H,W,C) ndarray or (H,W) ndarray*) – The image to augment. Channel-axis is optional, but expected to be the last axis if present. In most cases, this array should be of dtype `uint8`, which is supported by all augmenters. Support for other dtypes varies by augmenter – see the respective augmenter-specific documentation for more details.
- **hooks** (*None or imgaug.HooksImages, optional*) – HooksImages object to dynamically interfere with the augmentation process.

Returns The corresponding augmented image.

Return type ndarray

augment_images(*self*, *images*, *parents=None*, *hooks=None*)

Augment a batch of images.

Parameters

- **images** (*(N,H,W,C) ndarray or (N,H,W) ndarray or list of (H,W,C) ndarray or list of (H,W) ndarray*) – Images to augment. The input can be a list of numpy arrays or a single array. Each array is expected to have shape `(H, W, C)` or `(H, W)`, where `H` is the height, `W` is the width and `C` are the channels. The number of channels may differ between images. If a list is provided, the height, width and channels may differ between images within the provided batch. In most cases, the image array(s) should be of dtype `uint8`, which is supported by all augmenters. Support for other dtypes varies by augmenter – see the respective augmenter-specific documentation for more details.
- **parents** (*None or list of imgaug.augmenters.Augmenter, optional*) – Parent augmenters that have previously been called before the call to this function. Usually you can leave this parameter as `None`. It is set automatically for child augmenters.
- **hooks** (*None or imgaug.imgaug.HooksImages, optional*) – `HooksImages` object to dynamically interfere with the augmentation process.

Returns Corresponding augmented images. If the input was an `ndarray`, the output is also an `ndarray`, unless the used augmentations have led to different output image sizes (as can happen in e.g. cropping).

Return type ndarray or list

Examples

```
>>> import imgaug.augmenters as iaa
>>> import numpy as np
>>> aug = iaa.GaussianBlur((0.0, 3.0))
>>> # create empty example images
>>> images = np.zeros((2, 64, 64, 3), dtype=np.uint8)
>>> images_aug = aug.augment_images(images)
```

Create 2 empty (i.e. black) example numpy images and apply gaussian blurring to them.

augment_keypoints(*self*, *keypoints_on_images*, *parents=None*, *hooks=None*)

Augment a batch of keypoints/landmarks.

This is the corresponding function to *Augmenter.augment_images()*, just for keypoints/landmarks (i.e. points on images). Usually you will want to call *Augmenter.augment_images()* with a list of images, e.g. `augment_images([A, B, C])` and then `augment_keypoints()` with the corresponding list of keypoints on these images, e.g. `augment_keypoints([Ak, Bk, Ck])`, where `Ak` are the keypoints on image `A`.

Make sure to first convert the augmenter(s) to deterministic states before augmenting images and their corresponding keypoints, e.g. by

```
>>> import imgaug.augmenters as iaa
>>> from imgaug.augmentables.kps import Keypoint
>>> from imgaug.augmentables.kps import KeypointsOnImage
>>> A = B = C = np.zeros((10, 10), dtype=np.uint8)
>>> Ak = Bk = Ck = KeypointsOnImage([Keypoint(2, 2)], (10, 10))
>>> seq = iaa.Fliplr(0.5)
>>> seq_det = seq.to_deterministic()
>>> imgs_aug = seq_det.augment_images([A, B, C])
>>> kps_aug = seq_det.augment_keypoints([Ak, Bk, Ck])
```

Otherwise, different random values will be sampled for the image and keypoint augmentations, resulting in different augmentations (e.g. images might be rotated by `30deg` and keypoints by `-10deg`). Also make sure to call *Augmenter.to_deterministic()* again for each new batch, otherwise you would augment all batches in the same way.

Note that there is also *Augmenter.augment()*, which automatically handles the random state alignment.

Parameters

- **keypoints_on_images** (*imgaug.augmentables.kps.KeypointsOnImage or list of imgaug.augmentables.kps.KeypointsOnImage*) – The keypoints/landmarks to augment. Either a single instance of *KeypointsOnImage* or a list of such instances. Each instance must contain the keypoints of a single image.
- **parents** (*None or list of imgaug.augmenters.meta.Augmenter, optional*) – Parent augmenters that have previously been called before the call to this function. Usually you can leave this parameter as `None`. It is set automatically for child augmenters.
- **hooks** (*None or imgaug.imgaug.HooksKeypoints, optional*) – *HooksKeypoints* object to dynamically interfere with the augmentation process.

Returns Augmented keypoints.

Return type *imgaug.augmentables.kps.KeypointsOnImage* or list of imgaug.augmentables.kps.KeypointsOnImage

augment_line_strings(*self*, *line_strings_on_images*, *parents=None*, *hooks=None*)
Augment a batch of line strings.

This is the corresponding function to `Augmenter.augment_images()`, just for line strings. Usually you will want to call *Augmenter.augment_images()* with a list of images, e.g. `augment_images([A, B, C])` and then `augment_line_strings()` with the corresponding list of line strings on these images, e.g. `augment_line_strings([A_line, B_line, C_line])`, where `A_line` are the line strings on image `A`.

Make sure to first convert the augmenter(s) to deterministic states before augmenting images and their corresponding line strings, e.g. by

```
>>> import imgaug.augmenters as iaa
>>> from imgaug.augmentables.lines import LineString
>>> from imgaug.augmentables.lines import LineStringsOnImage
>>> A = B = C = np.ones((10, 10), dtype=np.uint8)
>>> A_line = B_line = C_line = LineStringsOnImage(
>>>     [LineString([(0, 0), (1, 0), (1, 1), (0, 1)])],
>>>     shape=(10, 10))
>>> seq = iaa.Fliplr(0.5)
>>> seq_det = seq.to_deterministic()
>>> imgs_aug = seq_det.augment_images([A, B, C])
>>> lines_aug = seq_det.augment_line_strings([A_line, B_line, C_line])
```

Otherwise, different random values will be sampled for the image and line string augmentations, resulting in different augmentations (e.g. images might be rotated by `30deg` and line strings by `-10deg`). Also make sure to call `to_deterministic()` again for each new batch, otherwise you would augment all batches in the same way.

Note that there is also *Augmenter.augment()*, which automatically handles the random state alignment.

Parameters

- **line_strings_on_images** (*imgaug.augmentables.lines.LineStringsOnImage or list of imgaug.augmentables.lines.LineStringsOnImage*) – The line strings to augment. Either a single instance of *LineStringsOnImage* or a list of such instances, with each one of them containing the line strings of a single image.
- **parents** (*None or list of imgaug.augmenters.meta.Augmenter, optional*) – Parent augmenters that have previously been called before the call to this function. Usually you can leave this parameter as None. It is set automatically for child augmenters.
- **hooks** (*None or imgaug.imgaug.HooksKeypoints, optional*) – *HooksKeypoints* object to dynamically interfere with the augmentation process.

Returns Augmented line strings.

Return type *imgaug.augmentables.lines.LineStringsOnImage* or list of imgaug.augmentables.lines.LineStringsOnImage

augment_polygons(*self*, *polygons_on_images*, *parents=None*, *hooks=None*)
Augment a batch of polygons.

This is the corresponding function to *Augmenter.augment_images()*, just for polygons. Usually you will want to call `Augmenter.augment_images()` with a list of images, e.g. `augment_images([A, B, C])` and then `augment_polygons()` with the corresponding list of polygons on these images, e.g. `augment_polygons([A_poly, B_poly, C_poly])`, where `A_poly` are the polygons on image `A`.

Make sure to first convert the augmenter(s) to deterministic states before augmenting images and their corresponding polygons, e.g. by

```
>>> import imgaug.augmenters as iaa
>>> from imgaug.augmentables.polys import Polygon, PolygonsOnImage
>>> A = B = C = np.ones((10, 10), dtype=np.uint8)
>>> Apoly = Bpoly = Cpoly = PolygonsOnImage(
>>>     [Polygon([(0, 0), (1, 0), (1, 1), (0, 1)])],
>>>     shape=(10, 10))
>>> seq = iaa.Fliplr(0.5)
>>> seq_det = seq.to_deterministic()
>>> imgs_aug = seq_det.augment_images([A, B, C])
>>> polys_aug = seq_det.augment_polygons([Apoly, Bpoly, Cpoly])
```

Otherwise, different random values will be sampled for the image and polygon augmentations, resulting in different augmentations (e.g. images might be rotated by `30deg` and polygons by `-10deg`). Also make sure to call `to_deterministic()` again for each new batch, otherwise you would augment all batches in the same way.

Note that there is also *Augmenter.augment()*, which automatically handles the random state alignment.

Parameters

- **polygons_on_images** (*imgaug.augmentables.polys.PolygonsOnImage or list of imgaug.augmentables.polys.PolygonsOnImage*) – The polygons to augment. Either a single instance of *PolygonsOnImage* or a list of such instances, with each one of them containing the polygons of a single image.
- **parents** (*None or list of imgaug.augmenters.meta.Augmenter, optional*) – Parent augmenters that have previously been called before the call to this function. Usually you can leave this parameter as `None`. It is set automatically for child augmenters.
- **hooks** (*None or imgaug.imgaug.HooksKeypoints, optional*) – *HooksKeypoints* object to dynamically interfere with the augmentation process.

Returns Augmented polygons.

Return type *imgaug.augmentables.polys.PolygonsOnImage* or list of imgaug.augmentables.polys.PolygonsOnImage

augment_segmentation_maps(*self*, *segmaps*, *parents=None*, *hooks=None*)

Augment a batch of segmentation maps.

Parameters

- **segmaps** (*imgaug.augmentables.segmaps.SegmentationMapsOnImage or list of imgaug.augmentables.segmaps.SegmentationMapsOnImage*) – Segmentation map(s) to augment. Either a single segmentation map or a list of segmentation maps.
- **parents** (*None or list of imgaug.augmenters.meta.Augmenter, optional*) – Parent augmenters that have previously been called before the call to this function. Usually you can leave this parameter as `None`. It is set automatically for child augmenters.
- **hooks** (*None or imgaug.HooksHeatmaps, optional*) – *HooksHeatmaps* object to dynamically interfere with the augmentation process.

Returns Corresponding augmented segmentation map(s).

Return type *imgaug.augmentables.segmaps.SegmentationMapsOnImage* or list of imgaug.augmentables.segmaps.SegmentationMapsOnImage

copy(*self*)
Create a shallow copy of this Augmenter instance.

Returns Shallow copy of this Augmenter instance.

Return type *imgaug.augmenters.meta.Augmenter*

copy_random_state(*self*, *source*, *recursive=True*, *matching='position'*, *matching_tolerant=True*, *copy_determinism=False*)
Copy the RNGs from a source augmenter sequence.

Parameters

- **source** (*imgaug.augmenters.meta.Augmenter*) – See `copy_random_state_()`.
- **recursive** (*bool, optional*) – See `copy_random_state_()`.
- **matching** (*{'position', 'name'}, optional*) – See `copy_random_state_()`.
- **matching_tolerant** (*bool, optional*) – See `copy_random_state_()`.
- **copy_determinism** (*bool, optional*) – See `copy_random_state_()`.

Returns Copy of the augmenter itself (with copied RNGs).

Return type *imgaug.augmenters.meta.Augmenter*

copy_random_state_(*self*, *source*, *recursive=True*, *matching='position'*, *matching_tolerant=True*, *copy_determinism=False*)
Copy the RNGs from a source augmenter sequence (in-place).

Note: The source augmenters are not allowed to use the global RNG. Call `localize_random_state_()` once on the source to localize all random states.

Parameters

- **source** (*imgaug.augmenters.meta.Augmenter*) – The source augmenter(s) from where to copy the RNG(s). The source may have children (e.g. the source can be a `Sequential`).
- **recursive** (*bool, optional*) – Whether to copy the RNGs of the source augmenter *and* all of its children (`True`) or just the source augmenter (`False`).
- **matching** (*{'position', 'name'}, optional*) – Defines the matching mode to use during recursive copy. This is used to associate source augmenters with target augmenters. If `position` then the target and source sequences of augmenters are turned into flattened lists and are associated based on their list indices. If `name` then the target and source augmenters are matched based on their names (i.e. `augmenter.name`).
- **matching_tolerant** (*bool, optional*) – Whether to use tolerant matching between source and target augmenters. If set to `False`: Name matching will raise an exception for any target augmenter which's name does not appear among the source augmenters. Position matching will raise an exception if source and target augmenter have an unequal number of children.
- **copy_determinism** (*bool, optional*) – Whether to copy the `deterministic` attributes from source to target augmenters too.

Returns The augmenter itself.

Return type *imgaug.augmenters.meta.Augmenter*

deepcopy(*self*)
Create a deep copy of this Augmenter instance.

Returns Deep copy of this Augmenter instance.

Return type *imgaug.augmenters.meta.Augmenter*

draw_grid(*self*, *images*, *rows*, *cols*)
Augment images and draw the results as a single grid-like image.

This method applies this augmenter to the provided images and returns a grid image of the results. Each cell in the grid contains a single augmented version of an input image.

If multiple input images are provided, the row count is multiplied by the number of images and each image gets its own row. E.g. for `images = [A, B]`, `rows=2`, `cols=3`:

```
A A A
B B B
A A A
B B B
```

for `images = [A]`, `rows=2`, `cols=3`:

```
A A A
A A A
```

Parameters

- **images** (*(N,H,W,3) ndarray or (H,W,3) ndarray or (H,W) ndarray or list of (H,W,3) ndarray or list of (H,W) ndarray*) – List of images to augment and draw in the grid. If a list, then each element is expected to have shape `(H, W)` or `(H, W, 3)`. If a single array, then it is expected to have shape `(N, H, W, 3)` or `(H, W, 3)` or `(H, W)`.
- **rows** (*int*) – Number of rows in the grid. If `N` input images are given, this value will automatically be multiplied by `N` to create rows for each image.
- **cols** (*int*) – Number of columns in the grid.

Returns The generated grid image with augmented versions of the input images. Here, `Hg` and `Wg` reference the output size of the grid, and *not* the sizes of the input images.

Return type (Hg, Wg, 3) ndarray

find_augmenters(*self*, *func*, *parents=None*, *flat=True*)
Find augmenters that match a condition.

This function will compare this augmenter and all of its children with a condition. The condition is a lambda function.

Parameters

- **func** (*callable*) – A function that receives a *Augmenter* instance and a list of parent *Augmenter* instances and must return `True`, if that augmenter is valid match or `False` otherwise.
- **parents** (*None or list of imgaug.augmenters.meta.Augmenter, optional*) – List of parent augmenters. Intended for nested calls and can usually be left as `None`.
- **flat** (*bool, optional*) – Whether to return the result as a flat list (`True`) or a nested list (`False`). In the latter case, the nesting matches each augmenters position among the children.

Returns Nested list if *flat* was set to `False`. Flat list if *flat* was set to `True`.

Return type list of imgaug.augmenters.meta.Augmenter

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.Sequential([
>>>     iaa.Fliplr(0.5, name="fliplr"),
>>>     iaa.Flipud(0.5, name="flipud")
>>> ])
>>> print(aug.find_augmenters(lambda a, parents: a.name == "fliplr"))
```

Return the first child augmenter (`Fliplr` instance).

find_augmenters_by_name (*self*, *name*, *regex=False*, *flat=True*)
Find augmenter(s) by name.

Parameters

- **name** (*str*) – Name of the augmenter(s) to search for.
- **regex** (*bool, optional*) – Whether *name* parameter is a regular expression.
- **flat** (*bool, optional*) – See *find_augmenters()*.

Returns **augmenters** – Nested list if *flat* was set to `False`. Flat list if *flat* was set to `True`.

Return type list of imgaug.augmenters.meta.Augmenter

find_augmenters_by_names (*self*, *names*, *regex=False*, *flat=True*)
Find augmenter(s) by names.

Parameters

- **names** (*list of str*) – Names of the augmenter(s) to search for.
- **regex** (*bool, optional*) – Whether *names* is a list of regular expressions. If it is, an augmenter is considered a match if *at least* one of these expressions is a match.
- **flat** (*boolean, optional*) – See *find_augmenters()*.

Returns **augmenters** – Nested list if *flat* was set to `False`. Flat list if *flat* was set to `True`.

Return type list of imgaug.augmenters.meta.Augmenter

get_all_children (*self*, *flat=False*)
Get all children of this augmenter as a list.

If the augmenter has no children, the returned list is empty.

Parameters **flat** (*bool*) – If set to `True`, the returned list will be flat.

Returns The children as a nested or flat list.

Return type list of imgaug.augmenters.meta.Augmenter

get_children_lists (*self*)
Get a list of lists of children of this augmenter.

For most augmenters, the result will be a single empty list. For augmenters with children it will often be a list with one sublist containing all children. In some cases the augmenter will contain multiple distinct lists of children, e.g. an if-list and an else-list. This will lead to a result consisting of a single list with multiple sublists, each representing the respective sublist of children.

E.g. for an if/else-augmenter that executes the children `A1`, `A2` if a condition is met and otherwise executes the children `B1`, `B2`, `B3` the result will be `[[A1, A2], [B1, B2, B3]]`.

IMPORTANT: While the topmost list may be newly created, each of the sublist must be editable inplace resulting in a changed children list of the augmenter. E.g. if an Augmenter `IfElse(condition, [A1, A2], [B1, B2, B3])` returns `[[A1, A2], [B1, B2, B3]]` for a call to *get_children_lists()* and `A2` is removed inplace from `[A1, A2]`, then the children lists of `IfElse(...)` must also change to `[A1]`, `[B1, B2, B3]`. This is used in `remove_augmenters_()`.

Returns One or more lists of child augmenter. Can also be a single empty list.

Return type list of list of imgaug.augmenters.meta.Augmenter

get_parameters(*self*)
Get the parameters of this augmenter.

Returns List of parameters of arbitrary types (usually child class of *StochasticParameter*, but not guaranteed to be).

Return type list

localize_random_state(*self*, *recursive=True*)
Assign augmenter-specific RNGs to this augmenter and its children.

See *Augmenter.localize_random_state_()* for more details.

Parameters **recursive** (*bool, optional*) – See *localize_random_state_()*.

Returns Copy of the augmenter and its children, with localized RNGs.

Return type *imgaug.augmenters.meta.Augmenter*

localize_random_state_(*self*, *recursive=True*)
Assign augmenter-specific RNGs to this augmenter and its children.

This method iterates over this augmenter and all of its children and replaces any pointer to the global RNG with a new local (i.e. augmenter-specific) RNG.

A random number generator (RNG) is used for the sampling of random values. The global random number generator exists exactly once throughout the library and is shared by many augmenters. A local RNG (usually) exists within exactly one augmenter and is only used by that augmenter.

Usually there is no need to change global into local RNGs. The only noteworthy exceptions are

- Whenever you want to use determinism (so that the global RNG is not accidentally reverted).
- Whenever you want to copy RNGs from one augmenter to another. (Copying the global RNG would usually not be useful. Copying the global RNG from augmenter A to B, then executing A and then B would result in B's (global) RNG's state having already changed because of A's sampling. So the samples of A and B would differ.)

The case of determinism is handled automatically by *to_deterministic()*. Only when you copy RNGs (via *copy_random_state()*), you need to call this function first.

Parameters **recursive** (*bool, optional*) – Whether to localize the RNGs of the augmenter's children too.

Returns Returns itself (with localized RNGs).

Return type *imgaug.augmenters.meta.Augmenter*

pool(*self*, *processes=None*, *maxtasksperchild=None*, *seed=None*)
Create a pool used for multicore augmentation.

Parameters

- **processes** (*None or int, optional*) – Same as in `__init__()`. The number of background workers. If `None`, the number of the machine's CPU cores will be used (this counts hyperthreads as CPU cores). If this is set to a negative value `p`, then `P - abs(p)` will be used, where `P` is the number of CPU cores. E.g. `-1` would use all cores except one (this is useful to e.g. reserve one core to feed batches to the GPU).
- **maxtasksperchild** (*None or int, optional*) – Same as for `__init__()`. The number of tasks done per worker process before the process is killed and restarted. If `None`, worker processes will not be automatically restarted.
- **seed** (*None or int, optional*) – Same as for `__init__()`. The seed to use for child processes. If `None`, a random seed will be used.

Returns Pool for multicore augmentation.

Return type *imgaug.multicore.Pool*

Examples

```
>>> import numpy as np
>>> import imgaug as ia
>>> import imgaug.augmenters as iaa
>>> from imgaug.augmentables.batches import Batch
>>>
>>> aug = iaa.Add(1)
>>> images = np.zeros((16, 128, 128, 3), dtype=np.uint8)
>>> batches = [Batch(images=np.copy(images)) for _ in range(100)]
>>> with aug.pool(processes=-1, seed=2) as pool:
>>>     batches_aug = pool.map_batches(batches, chunksize=8)
>>> print(np.sum(batches_aug[0].images_aug[0]))
49152
```

Create `100` batches of empty images. Each batch contains `16` images of size `128x128`. The batches are then augmented on all CPU cores except one (`processes=-1`). After augmentation, the sum of pixel values from the first augmented image is printed.

```
>>> import numpy as np
>>> import imgaug as ia
>>> import imgaug.augmenters as iaa
>>> from imgaug.augmentables.batches import Batch
>>>
>>> aug = iaa.Add(1)
>>> images = np.zeros((16, 128, 128, 3), dtype=np.uint8)
>>> def generate_batches():
>>>     for _ in range(100):
>>>         yield Batch(images=np.copy(images))
>>>
>>> with aug.pool(processes=-1, seed=2) as pool:
>>>     batches_aug = pool.imap_batches(generate_batches(), chunksize=8)
>>>     batch_aug = next(batches_aug)
>>>     print(np.sum(batch_aug.images_aug[0]))
49152
```

Same as above. This time, a generator is used to generate batches of images. Again, the first augmented image's sum of pixels is printed.

remove_augmenters(*self*, *func*, *copy=True*, *identity_if_topmost=True*, *noop_if_topmost=None*)
Remove this augmenter or children that match a condition.

Parameters

- **func** (*callable*) – Condition to match per augmenter. The function must expect the augmenter itself and a list of parent augmenters and returns `True` if that augmenter is supposed to be removed, or `False` otherwise. E.g. `lambda a, parents: a.name == "fliplr" and len(parents) == 1` removes an augmenter with name `fliplr` if it is the direct child of the augmenter upon which `remove_augmenters()` was initially called.
- **copy** (*bool, optional*) – Whether to copy this augmenter and all if its children before removing. If `False`, removal is performed in-place.
- **identity_if_topmost** (*bool, optional*) – If `True` and the condition (lambda function) leads to the removal of the topmost augmenter (the one this function is called on initially), then that topmost augmenter will be replaced by an instance of *Noop* (i.e. an augmenter that doesn't change its inputs). If `False`, `None` will be returned in these cases. This can only be `False` if copy is set to `True`.
- **noop_if_topmost** (*bool, optional*) – Deprecated since 0.4.0.

Returns This augmenter after the removal was performed. `None` is returned if the condition was matched for the topmost augmenter, *copy* was set to `True` and *noop_if_topmost* was set to `False`.

Return type *imgaug.augmenters.meta.Augmenter* or None

Examples

```
>>> import imgaug.augmenters as iaa
>>> seq = iaa.Sequential([
>>>     iaa.Fliplr(0.5, name="fliplr"),
>>>     iaa.Flipud(0.5, name="flipud"),
>>> ])
>>> seq = seq.remove_augmenters(lambda a, parents: a.name == "fliplr")
```

This removes the augmenter `Fliplr` from the `Sequential` object's children.

remove_augmenters_(*self*, *func*, *parents=None*)
Remove in-place children of this augmenter that match a condition.

This is functionally identical to `remove_augmenters()` with `copy=False`, except that it does not affect the topmost augmenter (the one on which this function is initially called on).

Added in 0.4.0.

Parameters

- **func** (*callable*) – See *remove_augmenters()*.
- **parents** (*None or list of imgaug.augmenters.meta.Augmenter, optional*) – List of parent *Augmenter* instances that lead to this augmenter. If `None`, an empty list will be used. This parameter can usually be left empty and will be set automatically for children.

Examples

```
>>> import imgaug.augmenters as iaa
>>> seq = iaa.Sequential([
>>>     iaa.Fliplr(0.5, name="fliplr"),
>>>    iaa.Flipud(0.5, name="flipud"),
>>> ])
>>> seq.remove_augmenters_(lambda a, parents: a.name == "fliplr")
```

This removes the augmenter `Fliplr` from the `Sequential` object's children.

remove_augmenters_inplace(*self*, *func*, *parents=None*)

Deprecated. Use `remove_augmenters_` instead.

Old name for `remove_augmenters_()`.

Deprecated since 0.4.0.

reseed(*self*, *random_state=None*, *deterministic_too=False*)

Deprecated. Use `imgaug.augmenters.meta.Augmenter.seed_` instead.

Old name of *seed_()*.

Deprecated since 0.4.0.

seed_(*self*, *entropy=None*, *deterministic_too=False*)

Seed this augmenter and all of its children.

This method assigns a new random number generator to the augmenter and all of its children (if it has any). The new random number generator is *derived* from the provided seed or RNG – or from the global random number generator if `None` was provided. Note that as child RNGs are *derived*, they do not all use the same seed.

If this augmenter or any child augmenter had a random number generator that pointed to the global random state, it will automatically be replaced with a local random state. This is similar to what *localize_random_state()* does.

This method is useful when augmentations are run in the background (i.e. on multiple cores). It should be called before sending this *Augmenter* instance to a background worker or once within each worker with different seeds (i.e., if `N` workers are used, the function should be called `N` times). Otherwise, all background workers will use the same seeds and therefore apply the same augmentations. Note that *Augmenter.augment_batches()* and *Augmenter.pool()* already do this automatically.

Added in 0.4.0.

Parameters

- **entropy** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – A seed or random number generator that is used to derive new random number generators for this augmenter and its children. If an `int` is provided, it will be interpreted as a seed. If `None` is provided, the global random number generator will be used.
- **deterministic_too** (*bool, optional*) – Whether to also change the seed of an augmenter `A`, if `A` is deterministic. This is the case both when this augmenter object is `A` or one of its children is `A`.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.Sequential([
>>>     iaa.Crop(px=(0, 10)),
>>>     iaa.Crop(px=(0, 10))
>>> ])
>>> aug.seed_(1)
```

Seed an augmentation sequence containing two crop operations. Even though the same seed was used, the two operations will still sample different pixel amounts to crop as the child-specific seed is merely derived from the provided seed.

show_grid(*self*, *images*, *rows*, *cols*)

Augment images and plot the results as a single grid-like image.

This calls *draw_grid()* and simply shows the results. See that method for details.

Parameters

- **images** (*(N,H,W,3) ndarray or (H,W,3) ndarray or (H,W) ndarray or list of (H,W,3) ndarray or list of (H,W) ndarray*) – List of images to augment and draw in the grid. If a list, then each element is expected to have shape `(H, W)` or `(H, W, 3)`. If a single array, then it is expected to have shape `(N, H, W, 3)` or `(H, W, 3)` or `(H, W)`.
- **rows** (*int*) – Number of rows in the grid. If `N` input images are given, this value will automatically be multiplied by `N` to create rows for each image.
- **cols** (*int*) – Number of columns in the grid.

to_deterministic(*self*, *n=None*)

Convert this augmenter from a stochastic to a deterministic one.

A stochastic augmenter samples pseudo-random values for each parameter, image and batch. A deterministic augmenter also samples new values for each parameter and image, but not batch. Instead, for consecutive batches it will sample the same values (provided the number of images and their sizes don't change). From a technical perspective this means that a deterministic augmenter starts each batch's augmentation with a random number generator in the same state (i.e. same seed), instead of advancing that state from batch to batch.

Using determinism is useful to (a) get the same augmentations for two or more image batches (e.g. for stereo cameras), (b) to augment images and corresponding data on them (e.g. segmentation maps or bounding boxes) in the same way.

Parameters **n** (*None or int, optional*) – Number of deterministic augmenters to return. If `None` then only one *Augmenter* instance will be returned. If `1` or higher, a list containing `n` *Augmenter* instances will be returned.

Returns A single Augmenter object if *n* was None, otherwise a list of Augmenter objects (even if *n* was `1`).

Return type *imgaug.augmenters.meta.Augmenter* or list of imgaug.augmenters.meta.Augmenter

class imgaug.augmenters.meta.**ChannelShuffle**(*p=1.0*, *channels=None*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Randomize the order of channels in input images.

Supported dtypes:

- `uint8`: yes; fully tested
- `uint16`: yes; tested
- `uint32`: yes; tested
- `uint64`: yes; tested
- `int8`: yes; tested
- `int16`: yes; tested
- `int32`: yes; tested
- `int64`: yes; tested
- `float16`: yes; tested
- `float32`: yes; tested
- `float64`: yes; tested
- `float128`: yes; tested
- `bool`: yes; tested

Parameters

- **p** (*float or imgaug.parameters.StochasticParameter, optional*) – Probability of shuffling channels in any given image. May be a fixed probability as a `float`, or a *`StochasticParameter`* that returns `0` s and `1` s.
- **channels** (*None or imgaug.ALL or list of int, optional*) – Which channels are allowed to be shuffled with each other. If this is `None` or `imgaug.ALL`, then all channels may be shuffled. If it is a `list` of `int` s, then only the channels with indices in that list may be shuffled. (Values start at `0`. All channel indices in the list must exist in each image.)
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.ChannelShuffle(0.35)
```

Shuffle all channels of `35%` of all images.

```
>>> aug = iaa.ChannelShuffle(0.35, channels=[0, 1])
```

Shuffle only channels `0` and `1` of `35%` of all images. As the new channel orders `0, 1` and `1, 0` are both valid outcomes of the shuffling, it means that for `0.35 * 0.5 = 0.175` or `17.5%` of all images the order of channels `0` and `1` is inverted.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`localize_random_state_`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`pool`(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
`remove_augmenters`(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
`remove_augmenters_`(self, func[, parents])	Remove in-place children of this augmenter that match a condition.

Continued on next page

Table 180 – continued from previous page

remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.meta.**ClipCBAsToImagePlanes**(*seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Clip coordinate-based augmentables to areas within the image plane.

This augmenter inspects all coordinate-based augmentables (e.g. bounding boxes, line strings) within a given batch and from each of them parts that are outside of the image plane. Parts within the image plane will be retained. This may e.g. shrink down bounding boxes. For keypoints, it removes any single points outside of the image plane. Any augmentable that is completely outside of the image plane will be removed.

Added in 0.4.0.

Supported dtypes:

- uint8: yes; fully tested
- uint16: yes; fully tested
- uint32: yes; fully tested
- uint64: yes; fully tested
- int8: yes; fully tested
- int16: yes; fully tested
- int32: yes; fully tested
- int64: yes; fully tested
- float16: yes; fully tested
- float32: yes; fully tested
- float64: yes; fully tested
- float128: yes; fully tested
- bool: yes; fully tested

Parameters

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See __init__().
- **name** (*None or str, optional*) – See __init__().

- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.Sequential([
>>>     iaa.Affine(translate_px={"x": (-100, 100)}),
>>>     iaa.ClipCBAsToImagePlanes()
>>> ])
```

Translate input data on the x-axis by `-100` to `100` pixels, then cut all coordinate-based augmentables (e.g. bounding boxes) down to areas that are within the image planes of their corresponding images.

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.

Continued on next page

Table 181 – continued from previous page

find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See get_parameters().
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See get_parameters().

class imgaug.augmenters.meta.**Identity**(*seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: imgaug.augmenters.meta.Augmenter

Augmenter that does not change the input data.

This augmenter is useful e.g. during validation/testing as it allows to re-use the training code without actually performing any augmentation.

Added in 0.4.0.

Supported dtypes:

- uint8: yes; fully tested
- uint16: yes; tested
- uint32: yes; tested
- uint64: yes; tested
- int8: yes; tested
- int16: yes; tested
- int32: yes; tested
- int64: yes; tested
- float16: yes; tested

- `float32`: yes; tested
- `float64`: yes; tested
- `float128`: yes; tested
- `bool`: yes; tested

Parameters

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.Identity()
```

Create an augmenter that does not change inputs.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.

Continued on next page

Table 182 – continued from previous page

copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.meta.**Lambda**(*func_images=None*, *func_heatmaps=None*, *func_segmentation_maps=None*, *func_keypoints=None*, *func_bounding_boxes='keypoints'*, *func_polygons='keypoints'*, *func_line_strings='keypoints'*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Augmenter that calls a lambda function for each input batch.

This is useful to add missing functions to a list of augmenters.

Supported dtypes:

- uint8: yes; fully tested
- uint16: yes; tested

- `uint32`: yes; tested
- `uint64`: yes; tested
- `int8`: yes; tested
- `int16`: yes; tested
- `int32`: yes; tested
- `int64`: yes; tested
- `float16`: yes; tested
- `float32`: yes; tested
- `float64`: yes; tested
- `float128`: yes; tested
- `bool`: yes; tested

Parameters

- **func_images** (*None or callable, optional*) – The function to call for each batch of images. It must follow the form:

```
function(images, random_state, parents, hooks)
```

  and return the changed images (may be transformed in-place). This is essentially the interface of `_augment_images()`. If this is `None` instead of a function, the images will not be altered.

- **func_heatmaps** (*None or callable, optional*) – The function to call for each batch of heatmaps. It must follow the form:

```
function(heatmaps, random_state, parents, hooks)
```

  and return the changed heatmaps (may be transformed in-place). This is essentially the interface of `_augment_heatmaps()`. If this is `None` instead of a function, the heatmaps will not be altered.

- **func_segmentation_maps** (*None or callable, optional*) – The function to call for each batch of segmentation maps. It must follow the form:

```
function(segmaps, random_state, parents, hooks)
```

  and return the changed segmaps (may be transformed in-place). This is essentially the interface of `_augment_segmentation_maps()`. If this is `None` instead of a function, the segmentatio maps will not be altered.

- **func_keypoints** (*None or callable, optional*) – The function to call for each batch of keypoints. It must follow the form:

```
function(keypoints_on_images, random_state, parents, hooks)
```

  and return the changed keypoints (may be transformed in-place). This is essentially the interface of `_augment_keypoints()`. If this is `None` instead of a function, the keypoints will not be altered.

- **func_bounding_boxes** (*"keypoints" or None or callable, optional*) – The function to call for each batch of bounding boxes. It must follow the form:

```
function(bounding_boxes_on_images, random_state, parents, hooks)
```

and return the changed bounding boxes (may be transformed in-place). This is essentially the interface of `_augment_bounding_boxes()`. If this is `None` instead of a function, the bounding boxes will not be altered. If this is the string `"keypoints"` instead of a function, the bounding boxes will automatically be augmented by transforming their corner vertices to keypoints and calling *func_keypoints*.

Added in 0.4.0.

- **func_polygons** (*"keypoints" or None or callable, optional*) – The function to call for each batch of polygons. It must follow the form:

```
function(polygons_on_images, random_state, parents, hooks)
```

and return the changed polygons (may be transformed in-place). This is essentially the interface of `_augment_polygons()`. If this is `None` instead of a function, the polygons will not be altered. If this is the string `"keypoints"` instead of a function, the polygons will automatically be augmented by transforming their corner vertices to keypoints and calling *func_keypoints*.

- **func_line_strings** (*"keypoints" or None or callable, optional*) – The function to call for each batch of line strings. It must follow the form:

```
function(line_strings_on_images, random_state, parents, hooks)
```

and return the changed line strings (may be transformed in-place). This is essentially the interface of `_augment_line_strings()`. If this is `None` instead of a function, the line strings will not be altered. If this is the string `"keypoints"` instead of a function, the line strings will automatically be augmented by transforming their corner vertices to keypoints and calling *func_keypoints*.

Added in 0.4.0.

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>>
>>> def func_images(images, random_state, parents, hooks):
>>>     images[:, ::2, :, :] = 0
>>>     return images
```

(continues on next page)

(continued from previous page)

```
>>>
>>> aug = iaa.Lambda(
>>>     func_images=func_images
>>> )
```

Replace every second row in input images with black pixels. Leave other data (e.g. heatmaps, keypoints) unchanged.

```
>>> def func_images(images, random_state, parents, hooks):
>>>     images[:, ::2, :, :] = 0
>>>     return images
>>>
>>> def func_heatmaps(heatmaps, random_state, parents, hooks):
>>>     for heatmaps_i in heatmaps:
>>>         heatmaps.arr_0to1[::2, :, :] = 0
>>>     return heatmaps
>>>
>>> def func_keypoints(keypoints_on_images, random_state, parents, hooks):
>>>     return keypoints_on_images
>>>
>>> aug = iaa.Lambda(
>>>     func_images=func_images,
>>>     func_heatmaps=func_heatmaps,
>>>     func_keypoints=func_keypoints
>>> )
```

Replace every second row in images with black pixels, set every second row in heatmaps to zero and leave other data (e.g. keypoints) unchanged.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.

Continued on next page

Table 183 – continued from previous page

copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.meta.**Noop**(*seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Identity*

Alias for augmenter *Identity*.

It is recommended to now use *Identity*. *Noop* might be deprecated in the future.

Supported dtypes:

See *Identity*.

Parameters

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See __init__().
- **name** (*None or str, optional*) – See __init__().

- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.Noop()
```

Create an augmenter that does not change inputs.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.

Continued on next page

Table 184 – continued from previous page

get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.meta.**OneOf**(*children*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.SomeOf*

Augmenter that always executes exactly one of its children.

Supported dtypes:

See *imgaug.augmenters.meta.SomeOf*.

Parameters

- **children** (*imgaug.augmenters.meta.Augmenter or list of imgaug.augmenters.meta.Augmenter*) – The choices of augmenters to apply.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See __init__().
- **name** (*None or str, optional*) – See __init__().
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method to_deterministic() for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> images = [np.ones((10, 10), dtype=np.uint8)]  # dummy example images
>>> seq = iaa.OneOf([
>>>     iaa.Fliplr(1.0),
>>>     iaa.Flipud(1.0)
>>> ])
>>> images_aug = seq.augment_images(images)
```

Flip each image either horizontally or vertically.

```
>>> images = [np.ones((10, 10), dtype=np.uint8)]  # dummy example images
>>> seq = iaa.OneOf([
>>>     iaa.Fliplr(1.0),
>>>     iaa.Sequential([
>>>         iaa.GaussianBlur(1.0),
>>>         iaa.Dropout(0.05),
>>>         iaa.AdditiveGaussianNoise(0.1*255)
>>>     ]),
>>>     iaa.Noop()
>>> ])
>>> images_aug = seq.augment_images(images)
```

Either flip each image horizontally, or add blur+dropout+noise or do nothing.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`add`(self, augmenter)	Add an augmenter to the list of child augmenters.
`append`(self, object, /)	Append object to the end of the list.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`clear`(self, /)	Remove all items from list.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.

Continued on next page

Table 185 – continued from previous page

copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
count(self, value, /)	Return number of occurrences of value.
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
extend(self, iterable, /)	Extend list by appending elements from the iterable.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	See *get_children_lists()*.
get_parameters(self)	See *get_parameters()*.
index(self, value[, start, stop])	Return first index of value.
insert(self, index, object, /)	Insert object before index.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
pop(self[, index])	Remove and return item at index (default last).
remove(self, value, /)	Remove first occurrence of value.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
reverse(self, /)	Reverse *IN PLACE*.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
sort(self, /, *[, key, reverse])	Stable sort *IN PLACE*.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.meta.**RemoveCBAsByOutOfImageFraction**(*fraction*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Remove coordinate-based augmentables exceeding an out of image fraction.

This augmenter inspects all coordinate-based augmentables (e.g. bounding boxes, line strings) within a given batch and removes any such augmentable which's out of image fraction is exactly a given value or greater than that. The out of image fraction denotes the fraction of the augmentable's area that is outside of the image, e.g. for a bounding box that has half of its area outside of the image it would be `0.5`.

Added in 0.4.0.

Supported dtypes:

- `uint8`: yes; fully tested
- `uint16`: yes; fully tested
- `uint32`: yes; fully tested
- `uint64`: yes; fully tested
- `int8`: yes; fully tested
- `int16`: yes; fully tested
- `int32`: yes; fully tested
- `int64`: yes; fully tested
- `float16`: yes; fully tested
- `float32`: yes; fully tested
- `float64`: yes; fully tested
- `float128`: yes; fully tested
- `bool`: yes; fully tested

Parameters

- **fraction** (*number*) – Remove any augmentable for which `fraction_{actual} >= fraction`, where `fraction_{actual}` denotes the estimated out of image fraction.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.Sequential([
>>>     iaa.Affine(translate_px={"x": (-100, 100)}),
>>>     iaa.RemoveCBAsByOutOfImageFraction(0.5)
>>> ])
```

Translate all inputs by `-100` to `100` pixels on the x-axis, then remove any coordinate-based augmentable (e.g. bounding boxes) which has at least `50%` of its area outside of the image plane.

```
>>> import imgaug as ia
>>> import imgaug.augmenters as iaa
>>> image = ia.quokka_square((100, 100))
>>> bb = ia.BoundingBox(x1=50-25, y1=0, x2=50+25, y2=100)
>>> bbsoi = ia.BoundingBoxesOnImage([bb], shape=image.shape)
>>> aug_without = iaa.Affine(translate_px={"x": 51})
>>> aug_with = iaa.Sequential([
>>>     iaa.Affine(translate_px={"x": 51}),
>>>     iaa.RemoveCBAsByOutOfImageFraction(0.5)
>>> ])
>>>
>>> image_without, bbsoi_without = aug_without(
>>>     image=image, bounding_boxes=bbsoi)
>>> image_with, bbsoi_with = aug_with(
>>>     image=image, bounding_boxes=bbsoi)
>>>
>>> assert len(bbsoi_without.bounding_boxes) == 1
>>> assert len(bbsoi_with.bounding_boxes) == 0
```

Create a bounding box on an example image, then translate the image so that `50%` of the bounding box's area is outside of the image and compare the effects and using `RemoveCBAsByOutOfImageFraction` with not using it.

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.

Continued on next page

Table 186 – continued from previous page

find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.meta.**Sequential**(*children=None*, *random_order=False*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*, list

List augmenter containing child augmenters to apply to inputs.

This augmenter is simply a list of other augmenters. To augment an image or any other data, it iterates over its children and applies each one of them independently to the data. (This also means that the second applied augmenter will already receive augmented input data and augment it further.)

This augmenter offers the option to apply its children in random order using the *random_order* parameter. This should often be activated as it greatly increases the space of possible augmentations.

Note: You are *not* forced to use *Sequential* in order to use other augmenters. Each augmenter can be used on its own, e.g the following defines an augmenter for horizontal flips and then augments a single image:

```
>>> import numpy as np
>>> import imgaug.augmenters as iaa
>>> image = np.zeros((32, 32, 3), dtype=np.uint8)
>>> aug = iaa.Fliplr(0.5)
>>> image_aug = aug.augment_image(image)
```

Supported dtypes:

- `uint8`: yes; fully tested
- `uint16`: yes; tested
- `uint32`: yes; tested
- `uint64`: yes; tested
- `int8`: yes; tested
- `int16`: yes; tested
- `int32`: yes; tested
- `int64`: yes; tested
- `float16`: yes; tested
- `float32`: yes; tested
- `float64`: yes; tested
- `float128`: yes; tested
- `bool`: yes; tested

Parameters

- **children** (*imgaug.augmenters.meta.Augmenter or list of imgaug.augmenters.meta.Augmenter or None, optional*) – The augmenters to apply to images.
- **random_order** (*bool, optional*) – Whether to apply the child augmenters in random order. If `True`, the order will be randomly sampled once per batch.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import numpy as np
>>> import imgaug.augmenters as iaa
>>> imgs = [np.random.rand(10, 10)]
>>> seq = iaa.Sequential([
>>>     iaa.Fliplr(0.5),
>>>     iaa.Flipud(0.5)
>>> ])
>>> imgs_aug = seq.augment_images(imgs)
```

Create a *Sequential* that always first applies a horizontal flip augmenter and then a vertical flip augmenter. Each of these two augmenters has a 50% probability of actually flipping the image.

```
>>> seq = iaa.Sequential([
>>>     iaa.Fliplr(0.5),
>>>     iaa.Flipud(0.5)
>>> ], random_order=True)
>>> imgs_aug = seq.augment_images(imgs)
```

Create a *Sequential* that sometimes first applies a horizontal flip augmenter (followed by a vertical flip augmenter) and sometimes first a vertical flip augmenter (followed by a horizontal flip augmenter). Again, each of them has a 50% probability of actually flipping the image.

Methods

__call__(self, *args, **kwargs)	Alias for *augment()*.
add(self, augmenter)	Add an augmenter to the list of child augmenters.
append(self, object, /)	Append object to the end of the list.
augment(self[, return_batch, hooks])	Augment a batch.
augment_batch(self, batch[, hooks])	**Deprecated**.
augment_batch_(self, batch[, parents, hooks])	Augment a single batch in-place.
augment_batches(self, batches[, hooks, ...])	Augment multiple batches.
augment_bounding_boxes(self, ...[, parents, ...])	Augment a batch of bounding boxes.
augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
clear(self, /)	Remove all items from list.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
count(self, value, /)	Return number of occurrences of value.
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
extend(self, iterable, /)	Extend list by appending elements from the iterable.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.

Continued on next page

Table 187 – continued from previous page

find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	See *get_children_lists()*.
get_parameters(self)	See *get_parameters()*.
index(self, value[, start, stop])	Return first index of value.
insert(self, index, object, /)	Insert object before index.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
pop(self[, index])	Remove and return item at index (default last).
remove(self, value, /)	Remove first occurrence of value.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
reverse(self, /)	Reverse *IN PLACE*.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
sort(self, /, *[, key, reverse])	Stable sort *IN PLACE*.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

add(*self*, *augmenter*)
Add an augmenter to the list of child augmenters.

Parameters **imgaug.augmenters.meta.Augmenter** – The augmenter to add.

get_children_lists(*self*)
See *get_children_lists()*.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.meta.**SomeOf**(*n=None*, *children=None*, *random_order=False*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)
Bases: *imgaug.augmenters.meta.Augmenter*, list

List augmenter that applies only some of its children to inputs.

This augmenter is similar to *Sequential*, but may apply only a fixed or random subset of its child augmenters to inputs. E.g. the augmenter could be initialized with a list of 20 child augmenters and then apply 5 randomly chosen child augmenters to images.

The subset of augmenters to apply (and their order) is sampled once *per image*. If *random_order* is True, the order will be sampled once *per batch* (similar to *Sequential*).

This augmenter currently does not support replacing (i.e. picking the same child multiple times) due to implementation difficulties in connection with deterministic augmenters.

Supported dtypes:

- `uint8`: yes; fully tested
- `uint16`: yes; tested
- `uint32`: yes; tested
- `uint64`: yes; tested
- `int8`: yes; tested
- `int16`: yes; tested
- `int32`: yes; tested
- `int64`: yes; tested
- `float16`: yes; tested
- `float32`: yes; tested
- `float64`: yes; tested
- `float128`: yes; tested
- `bool`: yes; tested

Parameters

- **n** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter or None, optional*) –

  Count of augmenters to apply.

  - If `int`, then exactly *n* of the child augmenters are applied to every image.
  - If tuple of two `int`s `(a, b)`, then a random value will be uniformly sampled per image from the discrete interval `[a..b]` and denote the number of child augmenters to pick and apply. `b` may be set to `None`, which is then equivalent to `(a..C)` with `C` denoting the number of children that the augmenter has.
  - If `StochasticParameter`, then `N` numbers will be sampled for `N` images. The parameter is expected to be discrete.
  - If `None`, then the total number of available children will be used (i.e. all children will be applied).

- **children** (*imgaug.augmenters.meta.Augmenter or list of imgaug.augmenters.meta.Augmenter or None, optional*) – The augmenters to apply to images. If this is a list of augmenters, it will be converted to a *Sequential*.
- **random_order** (*boolean, optional*) – Whether to apply the child augmenters in random order. If `True`, the order will be randomly sampled once per batch.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.

- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> imgs = [np.random.rand(10, 10)]
>>> seq = iaa.SomeOf(1, [
>>>     iaa.Fliplr(1.0),
>>>     iaa.Flipud(1.0)
>>> ])
>>> imgs_aug = seq.augment_images(imgs)
```

Apply either `Fliplr` or `Flipud` to images.

```
>>> seq = iaa.SomeOf((1, 3), [
>>>     iaa.Fliplr(1.0),
>>>     iaa.Flipud(1.0),
>>>     iaa.GaussianBlur(1.0)
>>> ])
>>> imgs_aug = seq.augment_images(imgs)
```

Apply one to three of the listed augmenters (`Fliplr`, `Flipud`, `GaussianBlur`) to images. They are always applied in the provided order, i.e. first `Fliplr`, second `Flipud`, third `GaussianBlur`.

```
>>> seq = iaa.SomeOf((1, None), [
>>>     iaa.Fliplr(1.0),
>>>     iaa.Flipud(1.0),
>>>     iaa.GaussianBlur(1.0)
>>> ], random_order=True)
>>> imgs_aug = seq.augment_images(imgs)
```

Apply one to all of the listed augmenters (`Fliplr`, `Flipud`, `GaussianBlur`) to images. They are applied in random order, i.e. sometimes `GaussianBlur` first, followed by `Fliplr`, sometimes `Fliplr` followed by `Flipud` followed by `Blur` etc. The order is sampled once per batch.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
add(self, augmenter)	Add an augmenter to the list of child augmenters.
`append`(self, object, /)	Append object to the end of the list.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.

Continued on next page

Table 188 – continued from previous page

augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
clear(self, /)	Remove all items from list.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
count(self, value, /)	Return number of occurrences of value.
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
extend(self, iterable, /)	Extend list by appending elements from the iterable.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	See *get_children_lists()*.
get_parameters(self)	See *get_parameters()*.
index(self, value[, start, stop])	Return first index of value.
insert(self, index, object, /)	Insert object before index.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
pop(self[, index])	Remove and return item at index (default last).
remove(self, value, /)	Remove first occurrence of value.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
reverse(self, /)	Reverse *IN PLACE*.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
sort(self, /, *[, key, reverse])	Stable sort *IN PLACE*.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

add(*self*, *augmenter*)
Add an augmenter to the list of child augmenters.

Parameters **augmenter** (*imgaug.augmenters.meta.Augmenter*) – The augmenter to add.

get_children_lists(*self*)
See *get_children_lists()*.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.meta.**Sometimes**(*p=0.5*, *then_list=None*, *else_list=None*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)
Bases: *imgaug.augmenters.meta.Augmenter*

Apply child augmenter(s) with a probability of *p*.

Let `C` be one or more child augmenters given to *Sometimes*. Let `p` be the fraction of images (or other data) to augment. Let `I` be the input images (or other data). Let `N` be the number of input images (or other entities). Then (on average) `p*N` images of `I` will be augmented using `C`.

Supported dtypes:

- `uint8`: yes; fully tested
- `uint16`: yes; tested
- `uint32`: yes; tested
- `uint64`: yes; tested
- `int8`: yes; tested
- `int16`: yes; tested
- `int32`: yes; tested
- `int64`: yes; tested
- `float16`: yes; tested
- `float32`: yes; tested
- `float64`: yes; tested
- `float128`: yes; tested
- `bool`: yes; tested

Parameters

- **p** (*float or imgaug.parameters.StochasticParameter, optional*) – Sets the probability with which the given augmenters will be applied to input images/data. E.g. a value of `0.5` will result in `50%` of all input images (or other augmentables) being augmented.
- **then_list** (*None or imgaug.augmenters.meta.Augmenter or list of imgaug.augmenters.meta.Augmenter, optional*) – Augmenter(s) to apply to *p%* percent of all images. If this is a list of augmenters, it will be converted to a *Sequential*.
- **else_list** (*None or imgaug.augmenters.meta.Augmenter or list of imgaug.augmenters.meta.Augmenter, optional*) – Augmenter(s) to apply to `(1-p)` percent of all images. These augmenters will be applied only when the ones in *then_list* are *not* applied (either-or-relationship). If this is a list of augmenters, it will be converted to a *Sequential*.

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.Sometimes(0.5, iaa.GaussianBlur(0.3))
```

Apply `GaussianBlur` to `50%` of all input images.

```
>>> aug = iaa.Sometimes(0.5, iaa.GaussianBlur(0.3), iaa.Fliplr(1.0))
```

Apply `GaussianBlur` to `50%` of all input images. Apply `Fliplr` to the other `50%` of all input images.

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.

Continued on next page

Table 189 – continued from previous page

copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	See *get_children_lists()*.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_children_lists(*self*)
See *get_children_lists()*.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.meta.**WithChannels**(*channels=None*, *children=None*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Apply child augmenters to specific channels.

Let `C` be one or more child augmenters given to this augmenter. Let `H` be a list of channels. Let `I` be the input images. Then this augmenter will pick the channels `H` from each image in `I` (resulting in new images) and apply `C` to them. The result of the augmentation will be merged back into the original images.

Supported dtypes:

- `uint8`: yes; fully tested
- `uint16`: yes; tested
- `uint32`: yes; tested
- `uint64`: yes; tested

- `int8`: yes; tested
- `int16`: yes; tested
- `int32`: yes; tested
- `int64`: yes; tested
- `float16`: yes; tested
- `float32`: yes; tested
- `float64`: yes; tested
- `float128`: yes; tested
- `bool`: yes; tested

Parameters

- **channels** (*None or int or list of int, optional*) – Sets the channels to be extracted from each image. If `None`, all channels will be used. Note that this is not stochastic - the extracted channels are always the same ones.
- **children** (*imgaug.augmenters.meta.Augmenter or list of imgaug.augmenters.meta.Augmenter or None, optional*) – One or more augmenters to apply to images, after the channels are extracted.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.WithChannels([0], iaa.Add(10))
```

Assuming input images are RGB, then this augmenter will add `10` only to the first channel, i.e. it will make images appear more red.

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.

Continued on next page

Table 190 – continued from previous page

augment_batch_(self, batch[, parents, hooks])	Augment a single batch in-place.
augment_batches(self, batches[, hooks, ...])	Augment multiple batches.
augment_bounding_boxes(self, ...[, parents, ...])	Augment a batch of bounding boxes.
augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	See *get_children_lists()*.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_children_lists(*self*)
See *get_children_lists()*.

get_parameters(*self*)
See *get_parameters()*.

imgaug.augmenters.meta.**clip_augmented_image**(*image*, *min_value*, *max_value*)
Deprecated. Use `imgaug.dtypes.clip_` instead.

Clip image.

imgaug.augmenters.meta.**clip_augmented_image_**(*image*, *min_value*, *max_value*)
Deprecated. Use `imgaug.dtypes.clip_` instead.

Clip image in-place.

imgaug.augmenters.meta.**clip_augmented_images**(*images*, *min_value*, *max_value*)
Deprecated. Use `imgaug.dtypes.clip_` instead.

Clip images.

imgaug.augmenters.meta.**clip_augmented_images_**(*images*, *min_value*, *max_value*)
Deprecated. Use `imgaug.dtypes.clip_` instead.

Clip images in-place.

imgaug.augmenters.meta.**copy_arrays**(*arrays*)
Copy the arrays of a single input array or list of input arrays.

imgaug.augmenters.meta.**estimate_max_number_of_channels**(*images*)
Compute the maximum number of image channels among a list of images.

imgaug.augmenters.meta.**handle_children_list**(*lst*, *augmenter_name*, *lst_name*, *default='sequential'*)
Normalize an augmenter list provided by a user.

imgaug.augmenters.meta.**invert_reduce_to_nonempty**(*objs*, *ids*, *objs_reduced*)
Inverse of *reduce_to_nonempty()*.

imgaug.augmenters.meta.**reduce_to_nonempty**(*objs*)
Remove from a list all objects that don't follow `obj.empty==True`.

imgaug.augmenters.meta.**shuffle_channels**(*image*, *random_state*, *channels=None*)
Randomize the order of (color) channels in an image.

Supported dtypes:

- `uint8`: yes; fully tested
- `uint16`: yes; indirectly tested (1)
- `uint32`: yes; indirectly tested (1)
- `uint64`: yes; indirectly tested (1)
- `int8`: yes; indirectly tested (1)
- `int16`: yes; indirectly tested (1)
- `int32`: yes; indirectly tested (1)
- `int64`: yes; indirectly tested (1)
- `float16`: yes; indirectly tested (1)
- `float32`: yes; indirectly tested (1)
- `float64`: yes; indirectly tested (1)
- `float128`: yes; indirectly tested (1)

- `bool`: yes; indirectly tested (1)
- (1) Indirectly tested via *ChannelShuffle*.

Parameters

- **image** (*((H,W,[C]) ndarray*) – Image of any dtype for which to shuffle the channels.
- **random_state** (*imgaug.random.RNG*) – The random state to use for this shuffling operation.
- **channels** (*None or imgaug.ALL or list of int, optional*) – Which channels are allowed to be shuffled with each other. If this is `None` or `imgaug.ALL`, then all channels may be shuffled. If it is a `list` of `int` s, then only the channels with indices in that list may be shuffled. (Values start at `0`. All channel indices in the list must exist in the image.)

Returns The input image with shuffled channels.

Return type ndarray

13.32 imgaug.augmenters.pillike

Augmenters that have identical outputs to well-known PIL functions.

The `like` in `pillike` indicates that the augmenters in this module have identical outputs and mostly identical inputs to corresponding PIL functions, but do not *have to* wrap these functions internally. They may use internally different (e.g. faster) techniques to produce these outputs.

Some of the augmenters in this module may also exist in other modules under similar name. These other augmenters may currently have the same outputs as the corresponding PIL functions, but that is not guaranteed for the future. Use the augmenters in this module if identical outputs to PIL are required.

List of augmenters:

- *Solarize*
- *Posterize*
- *Equalize*
- *Autocontrast*
- *EnhanceColor*
- *EnhanceContrast*
- *EnhanceBrightness*
- *EnhanceSharpness*
- *FilterBlur*
- *FilterSmooth*
- *FilterSmoothMore*
- *FilterEdgeEnhance*
- *FilterEdgeEnhanceMore*
- *FilterFindEdges*
- *FilterContour*
- *FilterEmboss*

- *FilterSharpen*
- *FilterDetail*
- *Affine*

Standard usage of these augmenters follows roughly the schema:

```
import numpy as np
import imgaug.augmenters as iaa

aug = iaa.pillike.Affine(translate_px={"x": (-5, 5)})
image = np.full((32, 32, 3), 255, dtype=np.uint8)

images_aug = aug(images=[image, image, image])
```

Added in 0.4.0.

class imgaug.augmenters.pillike.**Affine**(*scale=1.0*, *translate_percent=None*, *translate_px=None*, *rotate=0.0*, *shear=0.0*, *fillcolor=0*, *center=(0.5, 0.5)*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.geometric.Affine*

Apply PIL-like affine transformations to images.

This augmenter has identical outputs to `PIL.Image.transform` with parameter `method=PIL.Image.AFFINE`.

Warning: This augmenter can currently only transform image-data. Batches containing heatmaps, segmentation maps and coordinate-based augmentables will be rejected with an error. Use *Affine* if you have to transform such inputs.

Note: This augmenter uses the image center as the transformation center. This has to be explicitly enforced in PIL using corresponding translation matrices. Without such translation, PIL uses the image top left corner as the transformation center. To mirror that behaviour, use `center=(0.0, 0.0)`.

Added in 0.4.0.

Supported dtypes:

See *warp_affine()*.

Parameters

- **scale** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter or dict {"x": number/tuple/list/StochasticParameter, "y": number/tuple/list/StochasticParameter}, optional*) – See *Affine*.
- **translate_percent** (*None or number or tuple of number or list of number or imgaug.parameters.StochasticParameter or dict {"x": number/tuple/list/StochasticParameter, "y": number/tuple/list/StochasticParameter}, optional*) – See *Affine*.
- **translate_px** (*None or int or tuple of int or list of int or imgaug.parameters.StochasticParameter or dict {"x": int/tuple/list/StochasticParameter, "y": int/tuple/list/StochasticParameter}, optional*) – See *Affine*.

- **rotate** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – See *Affine*.
- **shear** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter or dict {"x": int/tuple/list/StochasticParameter, "y": int/tuple/list/StochasticParameter}, optional*) – See *Affine*.
- **fillcolor** (*number or tuple of number or list of number or imgaug.ALL or imgaug.parameters.StochasticParameter, optional*) – See parameter `cval` in *Affine*.
- **center** (*{'uniform', 'normal', 'center', 'left-top', 'left-center', 'left-bottom', 'center-top', 'center-center', 'center-bottom', 'right-top', 'right-center', 'right-bottom'} or tuple of float or StochasticParameter or tuple of StochasticParameter, optional*) – The center point of the affine transformation, given as relative xy-coordinates. Set this to `(0.0, 0.0)` or `left-top` to use the top left image corner as the transformation center. Set this to `(0.5, 0.5)` or `center-center` to use the image center as the transformation center. See also paramerer `position` in *PadToFixedSize* for details about valid datatypes of this parameter.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.pillike.Affine(scale={"x": (0.8, 1.2), "y": (0.5, 1.5)})
```

Create an augmenter that applies affine scaling (zoom in/out) to images. Along the x-axis they are scaled to 80-120% of their size, along the y-axis to 50-150% (both values randomly and uniformly chosen per image).

```
>>> aug = iaa.pillike.Affine(translate_px={"x": 0, "y": [-10, 10]},
>>>                          fillcolor=128)
```

Create an augmenter that translates images along the y-axis by either `-10px` or `10px`. Newly created pixels are always filled with the value `128` (along all channels).

```
>>> aug = iaa.pillike.Affine(rotate=(-20, 20), fillcolor=(0, 256))
```

Rotate an image by `-20` to `20` degress and fill up all newly created pixels with a random RGB color.

See the similar augmenter *Affine* for more examples.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`localize_random_state_`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`pool`(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
`remove_augmenters`(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
`remove_augmenters_`(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
`remove_augmenters_inplace`(self, func[, parents])	**Deprecated**.
`reseed`(self[, random_state, deterministic_too])	**Deprecated**.
`seed_`(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.

Continued on next page

Table 191 – continued from previous page

show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.pillike.**Autocontrast**(*cutoff=(0, 20), per_channel=False, seed=None, name=None, random_state='deprecated', deterministic='deprecated'*)

Bases: imgaug.augmenters.contrast._ContrastFuncWrapper

Adjust contrast by cutting off p% of lowest/highest histogram values.

This augmenter has identical outputs to PIL.ImageOps.autocontrast.

See *autocontrast()* for more details.

Added in 0.4.0.

Supported dtypes:

See *autocontrast()*.

Parameters

- **cutoff** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter, optional*) – Percentage of values to cut off from the low and high end of each image's histogram, before stretching it to [0, 255].
  - If int: The value will be used for all images.
  - If tuple (a, b): A value will be uniformly sampled from the discrete interval [a..b] per image.
  - If list: A random value will be sampled from the list per image.
  - If StochasticParameter: A value will be sampled from that parameter per image.
- **per_channel** (*bool or float, optional*) – Whether to use the same value for all channels (False) or to sample a new value for each channel (True). If this value is a float p, then for p percent of all images *per_channel* will be treated as True, otherwise as False.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See __init__().
- **name** (*None or str, optional*) – See __init__().
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method to_deterministic() for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.pillike.Autocontrast()
```

Modify the contrast of images by cutting off the `0` to `20%` lowest and highest values from the histogram, then stretching it to full length.

```
>>> aug = iaa.pillike.Autocontrast((10, 20), per_channel=True)
```

Modify the contrast of images by cutting off the `10` to `20%` lowest and highest values from the histogram, then stretching it to full length. The cutoff value is sampled per *channel* instead of per *image*.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.
`get_parameters`(self)	See *get_parameters()*.

Continued on next page

Table 192 – continued from previous page

localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.pillike.**EnhanceBrightness**(*factor=(0.5, 1.5), seed=None, name=None, random_state='deprecated', deterministic='deprecated'*)

Bases: imgaug.augmenters.pillike._EnhanceBase

Change the brightness of images.

This augmenter has identical outputs to PIL.ImageEnhance.Brightness.

Added in 0.4.0.

Supported dtypes:

See *enhance_brightness()*.

Parameters

- **factor** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Brightness of the image. Values below 1.0 decrease the brightness, leading to a black image around 0.0. Values above 1.0 increase the brightness. Sane values are roughly in [0.5, 1.5].
  - If number: The value will be used for all images.
  - If tuple (a, b): A value will be uniformly sampled per image from the interval [a, b).
  - If list: A random value will be picked from the list per image.
  - If StochasticParameter: Per batch of size N, the parameter will be queried once to return (N,) samples.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See __init__().
- **name** (*None or str, optional*) – See __init__().
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage

will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.

- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.pillike.EnhanceBrightness()
```

Create an augmenter that worsens the brightness of an image by a random factor.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.

Continued on next page

Table 193 – continued from previous page

get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.pillike.**EnhanceColor**(*factor=(0.0, 3.0)*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: imgaug.augmenters.pillike._EnhanceBase

Convert images to grayscale.

This augmenter has identical outputs to PIL.ImageEnhance.Color.

Added in 0.4.0.

Supported dtypes:

See *enhance_color()*.

Parameters

- **factor** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Colorfulness of the output image. Values close to 0.0 lead to grayscale images, values above 1.0 increase the strength of colors. Sane values are roughly in [0.0, 3.0].
  - If number: The value will be used for all images.
  - If tuple (a, b): A value will be uniformly sampled per image from the interval [a, b).
  - If list: A random value will be picked from the list per image.
  - If StochasticParameter: Per batch of size N, the parameter will be queried once to return (N,) samples.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See __init__().
- **name** (*None or str, optional*) – See __init__().
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage

will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.

- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.pillike.EnhanceColor()
```

Create an augmenter to remove a random fraction of color from input images.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.

Continued on next page

Table 194 – continued from previous page

get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.pillike.**EnhanceContrast**(*factor=(0.5, 1.5)*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: imgaug.augmenters.pillike._EnhanceBase

Change the contrast of images.

This augmenter has identical outputs to PIL.ImageEnhance.Contrast.

Added in 0.4.0.

Supported dtypes:

See *enhance_contrast()*.

Parameters

- **factor** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Strength of contrast in the image. Values below 1.0 decrease the contrast, leading to a gray image around 0.0. Values above 1.0 increase the contrast. Sane values are roughly in [0.5, 1.5].
  - If number: The value will be used for all images.
  - If tuple (a, b): A value will be uniformly sampled per image from the interval [a, b).
  - If list: A random value will be picked from the list per image.
  - If StochasticParameter: Per batch of size N, the parameter will be queried once to return (N,) samples.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See __init__().
- **name** (*None or str, optional*) – See __init__().
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or*

numpy.random.RandomState, optional) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.

- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.pillike.EnhanceContrast()
```

Create an augmenter that worsens the contrast of an image by a random factor.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.

Continued on next page

Table 195 – continued from previous page

get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.pillike.**EnhanceSharpness**(*factor=(0.0, 2.0)*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: imgaug.augmenters.pillike._EnhanceBase

Change the sharpness of images.

This augmenter has identical outputs to PIL.ImageEnhance.Sharpness.

Added in 0.4.0.

Supported dtypes:

See *enhance_sharpness()*.

Parameters

- **factor** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Sharpness of the image. Values below 1.0 decrease the sharpness, values above 1.0 increase it. Sane values are roughly in [0.0, 2.0].
  - If number: The value will be used for all images.
  - If tuple (a, b): A value will be uniformly sampled per image from the interval [a, b).
  - If list: A random value will be picked from the list per image.
  - If StochasticParameter: Per batch of size N, the parameter will be queried once to return (N,) samples.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See __init__().
- **name** (*None or str, optional*) – See __init__().

- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.pillike.EnhanceSharpness()
```

Create an augmenter that randomly decreases or increases the sharpness of an image.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.

Continued on next page

Table 196 – continued from previous page

get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.pillike.**Equalize**(*seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Equalize the image histogram.

This augmenter has identical outputs to `PIL.ImageOps.equalize`.

Added in 0.4.0.

Supported dtypes:

See *equalize_()*.

Parameters

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.pillike.Equalize()
```

Equalize the histograms of all input images.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`localize_random_state_`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`pool`(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.

Continued on next page

Table 197 – continued from previous page

remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.pillike.**FilterBlur**(*seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: imgaug.augmenters.pillike._FilterBase

Apply a blur filter kernel to images.

This augmenter has identical outputs to calling PIL.Image.filter with kernel PIL.ImageFilter.BLUR.

Added in 0.4.0.

Supported dtypes:

See *filter_blur()*.

Parameters

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See __init__().
- **name** (*None or str, optional*) – See __init__().
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method to_deterministic() for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.pillike.FilterBlur()
```

Create an augmenter that applies a blur filter kernel to images.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.
`get_parameters`(self)	See *get_parameters()*.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`localize_random_state_`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`pool`(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
`remove_augmenters`(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
`remove_augmenters_`(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
`remove_augmenters_inplace`(self, func[, parents])	**Deprecated**.
`reseed`(self[, random_state, deterministic_too])	**Deprecated**.
`seed_`(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.

Continued on next page

Table 198 – continued from previous page

show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.pillike.**FilterContour**(*seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: imgaug.augmenters.pillike._FilterBase

Apply a contour detection filter kernel to images.

This augmenter has identical outputs to calling PIL.Image.filter with kernel PIL.ImageFilter.CONTOUR.

Added in 0.4.0.

Supported dtypes:

See *filter_contour()*.

Parameters

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See __init__().
- **name** (*None or str, optional*) – See __init__().
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method to_deterministic() for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.pillike.FilterContour()
```

Create an augmenter that applies a contour detection filter kernel to images.

Methods

__call__(self, *args, **kwargs)	Alias for *augment()*.
augment(self[, return_batch, hooks])	Augment a batch.
augment_batch(self, batch[, hooks])	**Deprecated**.
augment_batch_(self, batch[, parents, hooks])	Augment a single batch in-place.
augment_batches(self, batches[, hooks, ...])	Augment multiple batches.

Continued on next page

Table 199 – continued from previous page

augment_bounding_boxes(self, ...[, parents, ...])	Augment a batch of bounding boxes.
augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.pillike.**FilterDetail**(*seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: imgaug.augmenters.pillike._FilterBase

Apply a detail enhancement filter kernel to images.

This augmenter has identical outputs to calling `PIL.Image.filter` with kernel `PIL.ImageFilter.DETAIL`.

Added in 0.4.0.

Supported dtypes:

See *filter_detail()*.

Parameters

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.pillike.FilterDetail()
```

Create an augmenter that applies a detail enhancement filter kernel to images.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.

Continued on next page

Table 200 – continued from previous page

augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.pillike.**FilterEdgeEnhance**(*seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: imgaug.augmenters.pillike._FilterBase

Apply an edge enhance filter kernel to images.

This augmenter has identical outputs to calling PIL.Image.filter with kernel PIL.ImageFilter.EDGE_ENHANCE.

Added in 0.4.0.

Supported dtypes:

See *filter_edge_enhance()*.

Parameters

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or*

numpy.random.RandomState, optional) – See `__init__()`.

- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.pillike.FilterEdgeEnhance()
```

Create an augmenter that applies a edge enhancement filter kernel to images.

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.

Continued on next page

Table 201 – continued from previous page

find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.pillike.**FilterEdgeEnhanceMore**(*seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: imgaug.augmenters.pillike._FilterBase

Apply a strong edge enhancement filter kernel to images.

This augmenter has identical outputs to calling PIL.Image.filter with kernel PIL.ImageFilter.EDGE_ENHANCE_MORE.

Added in 0.4.0.

Supported dtypes:

See *filter_edge_enhance_more()*.

Parameters

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See __init__().
- **name** (*None or str, optional*) – See __init__().
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method to_deterministic() for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.pillike.FilterEdgeEnhanceMore()
```

Create an augmenter that applies a strong edge enhancement filter kernel to images.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.
`get_parameters`(self)	See *get_parameters()*.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`localize_random_state_`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`pool`(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.

Continued on next page

Table 202 – continued from previous page

remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.pillike.**FilterEmboss**(*seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: imgaug.augmenters.pillike._FilterBase

Apply an emboss filter kernel to images.

This augmenter has identical outputs to calling PIL.Image.filter with kernel PIL.ImageFilter.EMBOSS.

Added in 0.4.0.

Supported dtypes:

See *filter_emboss()*.

Parameters

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See __init__().
- **name** (*None or str, optional*) – See __init__().
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method to_deterministic() for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.pillike.FilterEmboss()
```

Create an augmenter that applies an emboss filter kernel to images.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.
`get_parameters`(self)	See *get_parameters()*.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`localize_random_state_`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`pool`(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
`remove_augmenters`(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
`remove_augmenters_`(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
`remove_augmenters_inplace`(self, func[, parents])	**Deprecated**.
`reseed`(self[, random_state, deterministic_too])	**Deprecated**.
`seed_`(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.

Continued on next page

Table 203 – continued from previous page

show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.pillike.**FilterFindEdges**(*seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: imgaug.augmenters.pillike._FilterBase

Apply a edge detection kernel to images.

This augmenter has identical outputs to calling PIL.Image.filter with kernel PIL.ImageFilter.FIND_EDGES.

Added in 0.4.0.

Supported dtypes:

See *filter_find_edges()*.

Parameters

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See __init__().
- **name** (*None or str, optional*) – See __init__().
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method to_deterministic() for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.pillike.FilterFindEdges()
```

Create an augmenter that applies an edge detection filter kernel to images.

Methods

__call__(self, *args, **kwargs)	Alias for *augment()*.
augment(self[, return_batch, hooks])	Augment a batch.
augment_batch(self, batch[, hooks])	**Deprecated**.
augment_batch_(self, batch[, parents, hooks])	Augment a single batch in-place.
augment_batches(self, batches[, hooks, ...])	Augment multiple batches.

Continued on next page

Table 204 – continued from previous page

augment_bounding_boxes(self, ...[, parents, ...])	Augment a batch of bounding boxes.
augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.pillike.**FilterSharpen**(*seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: imgaug.augmenters.pillike._FilterBase

Apply a sharpening filter kernel to images.

This augmenter has identical outputs to calling `PIL.Image.filter` with kernel `PIL.ImageFilter.SHARPEN`.

Added in 0.4.0.

Supported dtypes:

See *filter_sharpen()*.

Parameters

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.pillike.FilterSharpen()
```

Create an augmenter that applies a sharpening filter kernel to images.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.

Continued on next page

Table 205 – continued from previous page

augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.pillike.**FilterSmooth**(*seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: imgaug.augmenters.pillike._FilterBase

Apply a smoothening filter kernel to images.

This augmenter has identical outputs to calling PIL.Image.filter with kernel PIL.ImageFilter.SMOOTH.

Added in 0.4.0.

Supported dtypes:

See *filter_smooth()*.

Parameters

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or*

numpy.random.RandomState, optional) – See `__init__()`.

- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.pillike.FilterSmooth()
```

Create an augmenter that applies a smoothening filter kernel to images.

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.

Continued on next page

Table 206 – continued from previous page

find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.pillike.**FilterSmoothMore**(*seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: imgaug.augmenters.pillike._FilterBase

Apply a strong smoothening filter kernel to images.

This augmenter has identical outputs to calling PIL.Image.filter with kernel PIL.ImageFilter.BLUR.

Added in 0.4.0.

Supported dtypes:

See *filter_smooth_more()*.

Parameters

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See __init__().
- **name** (*None or str, optional*) – See __init__().
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method to_deterministic() for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.pillike.FilterSmoothMore()
```

Create an augmenter that applies a strong smoothening filter kernel to images.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.
`get_parameters`(self)	See *get_parameters()*.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`localize_random_state_`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`pool`(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.

Continued on next page

Table 207 – continued from previous page

remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.pillike.**Posterize**(*nb_bits=(1, 8), from_colorspace='RGB', to_colorspace=None, max_size=None, interpolation='linear', seed=None, name=None, random_state='deprecated', deterministic='deprecated'*)

Bases: *imgaug.augmenters.color.Posterize*

Augmenter with identical outputs to PIL's `posterize()` function.

This augmenter quantizes each array component to `N` bits.

This class is currently an alias for *Posterize*, which again is an alias for *UniformColorQuantizationToNBits*, i.e. all three classes are right now guarantueed to have the same outputs as PIL's function.

Added in 0.4.0.

Supported dtypes:

See *Posterize*.

Methods

__call__(self, *args, **kwargs)	Alias for *augment()*.
augment(self[, return_batch, hooks])	Augment a batch.
augment_batch(self, batch[, hooks])	**Deprecated**.
augment_batch_(self, batch[, parents, hooks])	Augment a single batch in-place.
augment_batches(self, batches[, hooks, ...])	Augment multiple batches.
augment_bounding_boxes(self, ...[, parents, ...])	Augment a batch of bounding boxes.
augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.

Continued on next page

Table 208 – continued from previous page

augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.pillike.**Solarize**(*p=1.0*, *threshold=128*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.arithmetic.Invert*

Augmenter with identical outputs to PIL's `solarize()` function.

This augmenter inverts all pixel values above a threshold.

The outputs are identical to PIL's `solarize()`.

Added in 0.4.0.

Supported dtypes:

See ~imgaug.augmenters.arithmetic.invert_(min_value=None and max_value=None).

Parameters

- **p** (*float or imgaug.parameters.StochasticParameter, optional*) – See `Invert`.
- **threshold** (*None or number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – See `Invert`.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.Solarize(0.5, threshold=(32, 128))
```

Invert the colors in `50` percent of all images for pixels with a value between `32` and `128` or more. The threshold is sampled once per image. The thresholding operation happens per channel.

Methods

`__call__`(self, *args, **kwargs)	Alias for `augment()`.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.

Continued on next page

Table 209 – continued from previous page

copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

imgaug.augmenters.pillike.**autocontrast**(*image*, *cutoff=0*, *ignore=None*)

Maximize (normalize) image contrast.

This function calculates a histogram of the input image, removes **cutoff** percent of the lightest and darkest pixels from the histogram, and remaps the image so that the darkest pixel becomes black (0), and the lightest becomes white (255).

This function has identical outputs to PIL.ImageOps.autocontrast. The speed is almost identical.

Added in 0.4.0.

Supported dtypes:

- uint8: yes; fully tested
- uint16: no
- uint32: no
- uint64: no
- int8: no
- int16: no
- int32: no

- int64: no
- float16: no
- float32: no
- float64: no
- float128: no
- bool: no

Parameters

- **image** (*ndarray*) – The image for which to enhance the contrast.
- **cutoff** (*number*) – How many percent to cut off at the low and high end of the histogram. E.g. 20 will cut off the lowest and highest 20% of values. Expected value range is [0, 100].
- **ignore** (*None or int or iterable of int*) – Intensity values to ignore, i.e. to treat as background. If None, no pixels will be ignored. Otherwise exactly the given intensity value(s) will be ignored.

Returns Contrast-enhanced image.

Return type ndarray

imgaug.augmenters.pillike.**enhance_brightness**(*image*, *factor*)

Change the brightness of images.

This function has identical outputs to PIL.ImageEnhance.Brightness.

Added in 0.4.0.

Supported dtypes:

- uint8: yes; fully tested
- uint16: no
- uint32: no
- uint64: no
- int8: no
- int16: no
- int32: no
- int64: no
- float16: no
- float32: no
- float64: no
- float128: no
- bool: no

Parameters

- **image** (*ndarray*) – The image to modify.

- **factor** (*number*) – Brightness of the image. Values below `1.0` decrease the brightness, leading to a black image around `0.0`. Values above `1.0` increase the brightness. Sane values are roughly in `[0.5, 1.5]`.

Returns Brightness-modified image.

Return type ndarray

`imgaug.augmenters.pillike.`**`enhance_color`**(*image*, *factor*)

Change the strength of colors in an image.

This function has identical outputs to `PIL.ImageEnhance.Color`.

Added in 0.4.0.

Supported dtypes:

- `uint8`: yes; fully tested
- `uint16`: no
- `uint32`: no
- `uint64`: no
- `int8`: no
- `int16`: no
- `int32`: no
- `int64`: no
- `float16`: no
- `float32`: no
- `float64`: no
- `float128`: no
- `bool`: no

Parameters

- **image** (*ndarray*) – The image to modify.
- **factor** (*number*) – Colorfulness of the output image. Values close to `0.0` lead to grayscale images, values above `1.0` increase the strength of colors. Sane values are roughly in `[0.0, 3.0]`.

Returns Color-modified image.

Return type ndarray

`imgaug.augmenters.pillike.`**`enhance_contrast`**(*image*, *factor*)

Change the contrast of an image.

This function has identical outputs to `PIL.ImageEnhance.Contrast`.

Added in 0.4.0.

Supported dtypes:

- `uint8`: yes; fully tested
- `uint16`: no

- uint32: no
- uint64: no
- int8: no
- int16: no
- int32: no
- int64: no
- float16: no
- float32: no
- float64: no
- float128: no
- bool: no

Parameters

- **image** (*ndarray*) – The image to modify.
- **factor** (*number*) – Strength of contrast in the image. Values below 1.0 decrease the contrast, leading to a gray image around 0.0. Values above 1.0 increase the contrast. Sane values are roughly in [0.5, 1.5].

Returns Contrast-modified image.

Return type ndarray

imgaug.augmenters.pillike.**enhance_sharpness**(*image*, *factor*)

Change the sharpness of an image.

This function has identical outputs to PIL.ImageEnhance.Sharpness.

Added in 0.4.0.

Supported dtypes:

- uint8: yes; fully tested
- uint16: no
- uint32: no
- uint64: no
- int8: no
- int16: no
- int32: no
- int64: no
- float16: no
- float32: no
- float64: no
- float128: no
- bool: no

Parameters

- **image** (*ndarray*) – The image to modify.
- **factor** (*number*) – Sharpness of the image. Values below `1.0` decrease the sharpness, values above `1.0` increase it. Sane values are roughly in `[0.0, 2.0]`.

Returns Sharpness-modified image.

Return type ndarray

imgaug.augmenters.pillike.**equalize**(*image*, *mask=None*)

Equalize the image histogram.

See *equalize_()* for details.

This function is identical in inputs and outputs to `PIL.ImageOps.equalize`.

Added in 0.4.0.

Supported dtypes:

See *equalize_()*.

Parameters

- **image** (*ndarray*) – `uint8 (H,W,[C])` image to equalize.
- **mask** (*None or ndarray, optional*) – An optional mask. If given, only the pixels selected by the mask are included in the analysis.

Returns Equalized image.

Return type ndarray

imgaug.augmenters.pillike.**equalize_**(*image*, *mask=None*)

Equalize the image histogram in-place.

This function applies a non-linear mapping to the input image, in order to create a uniform distribution of grayscale values in the output image.

This function has identical outputs to `PIL.ImageOps.equalize`. It does however work in-place.

Added in 0.4.0.

Supported dtypes:

- `uint8`: yes; fully tested
- `uint16`: no
- `uint32`: no
- `uint64`: no
- `int8`: no
- `int16`: no
- `int32`: no
- `int64`: no
- `float16`: no
- `float32`: no
- `float64`: no

- float128: no
- bool: no

Parameters

- **image** (*ndarray*) – uint8 (H,W,[C]) image to equalize.
- **mask** (*None or ndarray, optional*) – An optional mask. If given, only the pixels selected by the mask are included in the analysis.

Returns Equalized image. *Might* have been modified in-place.

Return type ndarray

imgaug.augmenters.pillike.**filter_blur**(*image*)

Apply a blur filter kernel to the image.

This is the same as using PIL's PIL.ImageFilter.BLUR kernel.

Added in 0.4.0.

Supported dtypes:

- uint8: yes; fully tested
- uint16: no
- uint32: no
- uint64: no
- int8: no
- int16: no
- int32: no
- int64: no
- float16: no
- float32: no
- float64: no
- float128: no
- bool: no

Parameters **image** (*ndarray*) – The image to modify.

Returns Blurred image.

Return type ndarray

imgaug.augmenters.pillike.**filter_contour**(*image*)

Apply a contour filter kernel to the image.

This is the same as using PIL's PIL.ImageFilter.CONTOUR kernel.

Added in 0.4.0.

Supported dtypes:

- uint8: yes; fully tested
- uint16: no

- uint32: no
- uint64: no
- int8: no
- int16: no
- int32: no
- int64: no
- float16: no
- float32: no
- float64: no
- float128: no
- bool: no

Parameters **image** (*ndarray*) – The image to modify.

Returns Image with pronounced contours.

Return type ndarray

imgaug.augmenters.pillike.**filter_detail**(*image*)

Apply a detail enhancement filter kernel to the image.

This is the same as using PIL's PIL.ImageFilter.DETAIL kernel.

Added in 0.4.0.

Supported dtypes:

- uint8: yes; fully tested
- uint16: no
- uint32: no
- uint64: no
- int8: no
- int16: no
- int32: no
- int64: no
- float16: no
- float32: no
- float64: no
- float128: no
- bool: no

Parameters **image** (*ndarray*) – The image to modify.

Returns Image with enhanced details.

Return type ndarray

imgaug.augmenters.pillike.**filter_edge_enhance**(*image*)

Apply an edge enhancement filter kernel to the image.

This is the same as using PIL's `PIL.ImageFilter.EDGE_ENHANCE` kernel.

Added in 0.4.0.

Supported dtypes:

- `uint8`: yes; fully tested
- `uint16`: no
- `uint32`: no
- `uint64`: no
- `int8`: no
- `int16`: no
- `int32`: no
- `int64`: no
- `float16`: no
- `float32`: no
- `float64`: no
- `float128`: no
- `bool`: no

Parameters **image** (*ndarray*) – The image to modify.

Returns Image with enhanced edges.

Return type ndarray

imgaug.augmenters.pillike.**filter_edge_enhance_more**(*image*)

Apply a stronger edge enhancement filter kernel to the image.

This is the same as using PIL's `PIL.ImageFilter.EDGE_ENHANCE_MORE` kernel.

Added in 0.4.0.

Supported dtypes:

- `uint8`: yes; fully tested
- `uint16`: no
- `uint32`: no
- `uint64`: no
- `int8`: no
- `int16`: no
- `int32`: no
- `int64`: no
- `float16`: no
- `float32`: no

- `float64`: no
- `float128`: no
- `bool`: no

Parameters **image** (*ndarray*) – The image to modify.

Returns Smoothened image.

Return type ndarray

`imgaug.augmenters.pillike.`**`filter_emboss`**(*image*)

Apply an emboss filter kernel to the image.

This is the same as using PIL's `PIL.ImageFilter.EMBOSS` kernel.

Added in 0.4.0.

Supported dtypes:

- `uint8`: yes; fully tested
- `uint16`: no
- `uint32`: no
- `uint64`: no
- `int8`: no
- `int16`: no
- `int32`: no
- `int64`: no
- `float16`: no
- `float32`: no
- `float64`: no
- `float128`: no
- `bool`: no

Parameters **image** (*ndarray*) – The image to modify.

Returns Embossed image.

Return type ndarray

`imgaug.augmenters.pillike.`**`filter_find_edges`**(*image*)

Apply an edge detection filter kernel to the image.

This is the same as using PIL's `PIL.ImageFilter.FIND_EDGES` kernel.

Added in 0.4.0.

Supported dtypes:

- `uint8`: yes; fully tested
- `uint16`: no
- `uint32`: no

- uint64: no
- int8: no
- int16: no
- int32: no
- int64: no
- float16: no
- float32: no
- float64: no
- float128: no
- bool: no

Parameters **image** (*ndarray*) – The image to modify.

Returns Image with detected edges.

Return type ndarray

imgaug.augmenters.pillike.**filter_sharpen**(*image*)

Apply a sharpening filter kernel to the image.

This is the same as using PIL's PIL.ImageFilter.SHARPEN kernel.

Added in 0.4.0.

Supported dtypes:

- uint8: yes; fully tested
- uint16: no
- uint32: no
- uint64: no
- int8: no
- int16: no
- int32: no
- int64: no
- float16: no
- float32: no
- float64: no
- float128: no
- bool: no

Parameters **image** (*ndarray*) – The image to modify.

Returns Sharpened image.

Return type ndarray

imgaug.augmenters.pillike.**filter_smooth**(*image*)

Apply a smoothness filter kernel to the image.

This is the same as using PIL's `PIL.ImageFilter.SMOOTH` kernel.

Added in 0.4.0.

Supported dtypes:

- `uint8`: yes; fully tested
- `uint16`: no
- `uint32`: no
- `uint64`: no
- `int8`: no
- `int16`: no
- `int32`: no
- `int64`: no
- `float16`: no
- `float32`: no
- `float64`: no
- `float128`: no
- `bool`: no

Parameters **image** (*ndarray*) – The image to modify.

Returns Smoothened image.

Return type ndarray

imgaug.augmenters.pillike.**filter_smooth_more**(*image*)

Apply a strong smoothness filter kernel to the image.

This is the same as using PIL's `PIL.ImageFilter.SMOOTH_MORE` kernel.

Added in 0.4.0.

Supported dtypes:

- `uint8`: yes; fully tested
- `uint16`: no
- `uint32`: no
- `uint64`: no
- `int8`: no
- `int16`: no
- `int32`: no
- `int64`: no
- `float16`: no
- `float32`: no

- `float64`: no
- `float128`: no
- `bool`: no

Parameters **image** (*ndarray*) – The image to modify.

Returns Smoothened image.

Return type ndarray

`imgaug.augmenters.pillike.`**`posterize`**(*image*, *bits*)
Reduce the number of bits for each color channel.

This function has identical outputs to `PIL.ImageOps.posterize`.

Added in 0.4.0.

Supported dtypes:

See *quantize_uniform_to_n_bits()*.

Parameters

- **image** (*ndarray*) – Image array of shape `(H,W,[C])`.
- **bits** (*int*) – The number of bits to keep per component. Values in the interval `[1, 8]` are valid.

Returns Posterized image.

Return type ndarray

`imgaug.augmenters.pillike.`**`posterize_`**(*image*, *bits*)
Reduce the number of bits for each color channel in-place.

This function has identical outputs to `PIL.ImageOps.posterize`. It does however work in-place.

Added in 0.4.0.

Supported dtypes:

See *quantize_uniform_to_n_bits_()*.

Parameters

- **image** (*ndarray*) – Image array of shape `(H,W,[C])`.
- **bits** (*int*) – The number of bits to keep per component. Values in the interval `[1, 8]` are valid.

Returns Posterized image. This *can* be the same array as input in *image*, modified in-place.

Return type ndarray

`imgaug.augmenters.pillike.`**`solarize`**(*image*, *threshold=128*)
Invert all array components above a threshold.

This function has identical outputs to `PIL.ImageOps.solarize`.

Added in 0.4.0.

Supported dtypes:

See `~imgaug.augmenters.arithmetic.invert_(min_value=None and max_value=None)`.

Parameters

- **image** (*ndarray*) – Image array of shape `(H,W,[C])`.
- **threshold** (*int, optional*) – A threshold to use in order to invert only numbers above or below the threshold.

Returns Inverted image.

Return type ndarray

`imgaug.augmenters.pillike.`**`solarize_`**(*image*, *threshold=128*)

Invert all array components above a threshold in-place.

This function has identical outputs to `PIL.ImageOps.solarize`. It does however work in-place.

Added in 0.4.0.

Supported dtypes:

See `~imgaug.augmenters.arithmetic.invert_(min_value=None and max_value=None)`.

Parameters

- **image** (*ndarray*) – Image array of shape `(H,W,[C])`. The array *might* be modified in-place.
- **threshold** (*int, optional*) – A threshold to use in order to invert only numbers above or below the threshold.

Returns Inverted image. This *can* be the same array as input in *image*, modified in-place.

Return type ndarray

`imgaug.augmenters.pillike.`**`warp_affine`**(*image*, *scale_x=1.0*, *scale_y=1.0*, *translate_x_px=0*, *translate_y_px=0*, *rotate_deg=0*, *shear_x_deg=0*, *shear_y_deg=0*, *fillcolor=None*, *center=(0.5*, *0.5)*)

Apply an affine transformation to an image.

This function has identical outputs to `PIL.Image.transform` with `method=PIL.Image.AFFINE`.

Added in 0.4.0.

Supported dtypes:

- `uint8`: yes; fully tested
- `uint16`: no
- `uint32`: no
- `uint64`: no
- `int8`: no
- `int16`: no
- `int32`: no
- `int64`: no
- `float16`: no
- `float32`: no
- `float64`: no
- `float128`: no

- `bool`: no

Parameters

- **image** (*ndarray*) – The image to modify. Expected to be `uint8` with shape `(H,W)` or `(H,W,C)` with `C` being `3` or `4`.
- **scale_x** (*number, optional*) – Affine scale factor along the x-axis, where `1.0` denotes an identity transform and `2.0` is a strong zoom-in effect.
- **scale_y** (*number, optional*) – Affine scale factor along the y-axis, where `1.0` denotes an identity transform and `2.0` is a strong zoom-in effect.
- **translate_x_px** (*number, optional*) – Affine translation along the x-axis in pixels. Positive values translate the image towards the right.
- **translate_y_px** (*number, optional*) – Affine translation along the y-axis in pixels. Positive values translate the image towards the bottom.
- **rotate_deg** (*number, optional*) – Affine rotation in degrees *around the top left* of the image.
- **shear_x_deg** (*number, optional*) – Affine shearing in degrees along the x-axis with center point being the top-left of the image.
- **shear_y_deg** (*number, optional*) – Affine shearing in degrees along the y-axis with center point being the top-left of the image.
- **fillcolor** (*None or int or tuple of int, optional*) – Color tuple or intensity value to use when filling up newly created pixels. `None` fills with zeros. `int` will only fill the `0` th channel with that intensity value and all other channels with `0` (this is the default behaviour of PIL, use a tuple to fill all channels).
- **center** (*tuple of number, optional*) – Center xy-coordinate of the affine transformation, given as *relative* values, i.e. `(0.0, 0.0)` sets the transformation center to the top-left image corner, `(1.0, 0.0)` sets it to the the top-right image corner and `(0.5, 0.5)` sets it to the image center. The transformation center is relevant e.g. for rotations ("rotate around this center point"). PIL uses the image top-left corner as the transformation center if no centerization is included in the affine transformation matrix.

Returns Image after affine transformation.

Return type ndarray

13.33 imgaug.augmenters.pooling

Augmenters that apply pooling operations to images.

List of augmenters:

- *AveragePooling*
- *MaxPooling*
- *MinPooling*
- *MedianPooling*

class imgaug.augmenters.pooling.**AveragePooling**(*kernel_size=(1, 5), keep_size=True, seed=None, name=None, random_state='deprecated', deterministic='deprecated'*)

Bases: imgaug.augmenters.pooling._AbstractPoolingBase

Apply average pooling to images.

This augmenter pools images with kernel sizes `H x W` by averaging the pixel values within these windows. For e.g. `2 x 2` this halves the image size. Optionally, the augmenter will automatically re-upscale the image to the input size (by default this is activated).

Note that this augmenter is very similar to `AverageBlur`. `AverageBlur` applies averaging within windows of given kernel size *without* striding, while `AveragePooling` applies striding corresponding to the kernel size, with optional upscaling afterwards. The upscaling is configured to create "pixelated"/"blocky" images by default.

Note: During heatmap or segmentation map augmentation, the respective arrays are not changed, only the shapes of the underlying images are updated. This is because imgaug can handle maps/maks that are larger/smaller than their corresponding image.

Supported dtypes:

See *avg_pool()*.

Variables

- **kernel_size** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter or tuple of tuple of int or tuple of list of int or tuple of imgaug.parameters.StochasticParameter, optional*) – The kernel size of the pooling operation.
  - If an int, then that value will be used for all images for both kernel height and width.
  - If a tuple `(a, b)`, then a value from the discrete range `[a..b]` will be sampled per image.
  - If a list, then a random value will be sampled from that list per image and used for both kernel height and width.
  - If a StochasticParameter, then a value will be sampled per image from that parameter per image and used for both kernel height and width.
  - If a tuple of tuple of int given as `((a, b), (c, d))`, then two values will be sampled independently from the discrete ranges `[a..b]` and `[c..d]` per image and used as the kernel height and width.
  - If a tuple of lists of int, then two values will be sampled independently per image, one from the first list and one from the second, and used as the kernel height and width.
  - If a tuple of StochasticParameter, then two values will be sampled indepdently per image, one from the first parameter and one from the second, and used as the kernel height and width.
- **keep_size** (*bool, optional*) – After pooling, the result image will usually have a different height/width compared to the original input image. If this parameter is set to True, then the pooled image will be resized to the input image's size, i.e. the augmenter's output shape is always identical to the input shape.

- ***seed*** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.AveragePooling(2)
```

Create an augmenter that always pools with a kernel size of `2 x 2`.

```
>>> aug = iaa.AveragePooling(2, keep_size=False)
```

Create an augmenter that always pools with a kernel size of `2 x 2` and does *not* resize back to the input image size, i.e. the resulting images have half the resolution.

```
>>> aug = iaa.AveragePooling([2, 8])
```

Create an augmenter that always pools either with a kernel size of `2 x 2` or `8 x 8`.

```
>>> aug = iaa.AveragePooling((1, 7))
```

Create an augmenter that always pools with a kernel size of `1 x 1` (does nothing) to `7 x 7`. The kernel sizes are always symmetric.

```
>>> aug = iaa.AveragePooling(((1, 7), (1, 7)))
```

Create an augmenter that always pools with a kernel size of `H x W` where `H` and `W` are both sampled independently from the range `[1..7]`. E.g. resulting kernel sizes could be `3 x 7` or `5 x 1`.

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.

Continued on next page

Table 210 – continued from previous page

augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.pooling.**MaxPooling**(*kernel_size=(1, 5)*, *keep_size=True*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: imgaug.augmenters.pooling._AbstractPoolingBase

Apply max pooling to images.

This augmenter pools images with kernel sizes `H x W` by taking the maximum pixel value over windows. For e.g. `2 x 2` this halves the image size. Optionally, the augmenter will automatically re-upscale the image to the input size (by default this is activated).

The maximum within each pixel window is always taken channelwise..

Note: During heatmap or segmentation map augmentation, the respective arrays are not changed, only the shapes of the underlying images are updated. This is because imgaug can handle maps/maks that are larger/smaller than their corresponding image.

Supported dtypes:

See *max_pool()*.

Variables

- **kernel_size** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter or tuple of tuple of int or tuple of list of int or tuple of imgaug.parameters.StochasticParameter, optional*) – The kernel size of the pooling operation.
  - If an int, then that value will be used for all images for both kernel height and width.
  - If a tuple `(a, b)`, then a value from the discrete range `[a..b]` will be sampled per image.
  - If a list, then a random value will be sampled from that list per image and used for both kernel height and width.
  - If a StochasticParameter, then a value will be sampled per image from that parameter per image and used for both kernel height and width.
  - If a tuple of tuple of int given as `((a, b), (c, d))`, then two values will be sampled independently from the discrete ranges `[a..b]` and `[c..d]` per image and used as the kernel height and width.
  - If a tuple of lists of int, then two values will be sampled independently per image, one from the first list and one from the second, and used as the kernel height and width.
  - If a tuple of StochasticParameter, then two values will be sampled indepdently per image, one from the first parameter and one from the second, and used as the kernel height and width.
- **keep_size** (*bool, optional*) – After pooling, the result image will usually have a different height/width compared to the original input image. If this parameter is set to True, then the pooled image will be resized to the input image's size, i.e. the augmenter's output shape is always identical to the input shape.
- ***seed*** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.

- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.MaxPooling(2)
```

Create an augmenter that always pools with a kernel size of `2 x 2`.

```
>>> aug = iaa.MaxPooling(2, keep_size=False)
```

Create an augmenter that always pools with a kernel size of `2 x 2` and does *not* resize back to the input image size, i.e. the resulting images have half the resolution.

```
>>> aug = iaa.MaxPooling([2, 8])
```

Create an augmenter that always pools either with a kernel size of `2 x 2` or `8 x 8`.

```
>>> aug = iaa.MaxPooling((1, 7))
```

Create an augmenter that always pools with a kernel size of `1 x 1` (does nothing) to `7 x 7`. The kernel sizes are always symmetric.

```
>>> aug = iaa.MaxPooling(((1, 7), (1, 7)))
```

Create an augmenter that always pools with a kernel size of `H x W` where `H` and `W` are both sampled independently from the range `[1..7]`. E.g. resulting kernel sizes could be `3 x 7` or `5 x 1`.

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.

Continued on next page

Table 211 – continued from previous page

copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.pooling.**MedianPooling**(*kernel_size=(1, 5)*, *keep_size=True*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: imgaug.augmenters.pooling._AbstractPoolingBase

Apply median pooling to images.

This augmenter pools images with kernel sizes `H x W` by taking the median pixel value over windows. For e.g. `2 x 2` this halves the image size. Optionally, the augmenter will automatically re-upscale the image to the input size (by default this is activated).

The median within each pixel window is always taken channelwise.

Note: During heatmap or segmentation map augmentation, the respective arrays are not changed, only the shapes of the underlying images are updated. This is because imgaug can handle maps/maks that are larger/smaller than their corresponding image.

Supported dtypes:

See *median_pool()*.

Variables

- **kernel_size** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter or tuple of tuple of int or tuple of list of int or tuple of imgaug.parameters.StochasticParameter, optional*) – The kernel size of the pooling operation.
  - If an int, then that value will be used for all images for both kernel height and width.
  - If a tuple `(a, b)`, then a value from the discrete range `[a..b]` will be sampled per image.
  - If a list, then a random value will be sampled from that list per image and used for both kernel height and width.
  - If a StochasticParameter, then a value will be sampled per image from that parameter per image and used for both kernel height and width.
  - If a tuple of tuple of int given as `((a, b), (c, d))`, then two values will be sampled independently from the discrete ranges `[a..b]` and `[c..d]` per image and used as the kernel height and width.
  - If a tuple of lists of int, then two values will be sampled independently per image, one from the first list and one from the second, and used as the kernel height and width.
  - If a tuple of StochasticParameter, then two values will be sampled indepdently per image, one from the first parameter and one from the second, and used as the kernel height and width.
- **keep_size** (*bool, optional*) – After pooling, the result image will usually have a different height/width compared to the original input image. If this parameter is set to True, then the pooled image will be resized to the input image's size, i.e. the augmenter's output shape is always identical to the input shape.
- ***seed*** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.MedianPooling(2)
```

Create an augmenter that always pools with a kernel size of `2 x 2`.

```
>>> aug = iaa.MedianPooling(2, keep_size=False)
```

Create an augmenter that always pools with a kernel size of `2 x 2` and does *not* resize back to the input image size, i.e. the resulting images have half the resolution.

```
>>> aug = iaa.MedianPooling([2, 8])
```

Create an augmenter that always pools either with a kernel size of `2 x 2` or `8 x 8`.

```
>>> aug = iaa.MedianPooling((1, 7))
```

Create an augmenter that always pools with a kernel size of `1 x 1` (does nothing) to `7 x 7`. The kernel sizes are always symmetric.

```
>>> aug = iaa.MedianPooling(((1, 7), (1, 7)))
```

Create an augmenter that always pools with a kernel size of `H x W` where `H` and `W` are both sampled independently from the range `[1..7]`. E.g. resulting kernel sizes could be `3 x 7` or `5 x 1`.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.

Continued on next page

Table 212 – continued from previous page

find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.pooling.**MinPooling**(*kernel_size=(1, 5), keep_size=True, seed=None, name=None, random_state='deprecated', deterministic='deprecated'*)

Bases: imgaug.augmenters.pooling._AbstractPoolingBase

Apply minimum pooling to images.

This augmenter pools images with kernel sizes `H x W` by taking the minimum pixel value over windows. For e.g. `2 x 2` this halves the image size. Optionally, the augmenter will automatically re-upscale the image to the input size (by default this is activated).

The minimum within each pixel window is always taken channelwise.

Note: During heatmap or segmentation map augmentation, the respective arrays are not changed, only the shapes of the underlying images are updated. This is because imgaug can handle maps/maks that are larger/smaller than their corresponding image.

Supported dtypes:

See *min_pool()*.

Variables

- **kernel_size** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter or tuple of tuple of int or tuple of list of int or tuple of imgaug.parameters.StochasticParameter, optional*) – The kernel size of the pooling operation.
  - If an int, then that value will be used for all images for both kernel height and width.
  - If a tuple `(a, b)`, then a value from the discrete range `[a..b]` will be sampled per image.

  - If a list, then a random value will be sampled from that list per image and used for both kernel height and width.
  - If a StochasticParameter, then a value will be sampled per image from that parameter per image and used for both kernel height and width.
  - If a tuple of tuple of int given as `((a, b), (c, d))`, then two values will be sampled independently from the discrete ranges `[a..b]` and `[c..d]` per image and used as the kernel height and width.
  - If a tuple of lists of int, then two values will be sampled independently per image, one from the first list and one from the second, and used as the kernel height and width.
  - If a tuple of StochasticParameter, then two values will be sampled indepdently per image, one from the first parameter and one from the second, and used as the kernel height and width.

- **keep_size** (*bool, optional*) – After pooling, the result image will usually have a different height/width compared to the original input image. If this parameter is set to True, then the pooled image will be resized to the input image's size, i.e. the augmenter's output shape is always identical to the input shape.
- ***seed*** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.MinPooling(2)
```

Create an augmenter that always pools with a kernel size of `2 x 2`.

```
>>> aug = iaa.MinPooling(2, keep_size=False)
```

Create an augmenter that always pools with a kernel size of `2 x 2` and does *not* resize back to the input image size, i.e. the resulting images have half the resolution.

```
>>> aug = iaa.MinPooling([2, 8])
```

Create an augmenter that always pools either with a kernel size of `2 x 2` or `8 x 8`.

```
>>> aug = iaa.MinPooling((1, 7))
```

Create an augmenter that always pools with a kernel size of `1 x 1` (does nothing) to `7 x 7`. The kernel sizes are always symmetric.

```
>>> aug = iaa.MinPooling(((1, 7), (1, 7)))
```

Create an augmenter that always pools with a kernel size of `H x W` where `H` and `W` are both sampled independently from the range `[1..7]`. E.g. resulting kernel sizes could be `3 x 7` or `5 x 1`.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.
`get_parameters`(self)	See *get_parameters()*.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`localize_random_state_`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`pool`(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
`remove_augmenters`(self, func[, copy, ...])	Remove this augmenter or children that match a condition.

Continued on next page

Table 213 – continued from previous page

remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

13.34 imgaug.augmenters.segmentation

Augmenters that apply changes to images based on segmentation methods.

List of augmenters:

- *Superpixels*
- *Voronoi*
- *UniformVoronoi*
- *RegularGridVoronoi*
- *RelativeRegularGridVoronoi*

class imgaug.augmenters.segmentation.**DropoutPointsSampler**(*other_points_sampler*, *p_drop*)

Bases: *imgaug.augmenters.segmentation.IPointsSampler*

Remove a defined fraction of sampled points.

Parameters

- **other_points_sampler** (*IPointsSampler*) – Another point sampler that is queried to generate a list of points. The dropout operation will be applied to that list.
- **p_drop** (*number or tuple of number or imgaug.parameters.StochasticParameter*) – The probability that a coordinate will be removed from the list of all sampled coordinates. A value of `1.0` would mean that (on average) `100` percent of all coordinates will be dropped, while `0.0` denotes `0` percent. Note that this sampler will always ensure that at least one coordinate is left after the dropout operation, i.e. even `1.0` will only drop all *except one* coordinate.
  - If a `float`, then that value will be used for all images.
  - If a `tuple (a, b)`, then a value `p` will be sampled from the interval `[a, b]` per image.
  - If a `StochasticParameter`, then this parameter will be used to determine per coordinate whether it should be *kept* (sampled value of `>0.5`) or shouldn't be kept (sampled value of `<=0.5`). If you instead want to provide the probability as a stochastic parameter, you can usually do `imgaug.parameters.Binomial(1-p)` to convert parameter *p* to a 0/1 representation.

Examples

```
>>> import imgaug.augmenters as iaa
>>> sampler = iaa.DropoutPointsSampler(
>>>     iaa.RegularGridPointsSampler(10, 20),
>>>     0.2)
```

Create a point sampler that first generates points following a regular grid of `10` rows and `20` columns, then randomly drops `20` percent of these points.

Methods

sample_points(self, images, random_state)	Generate coordinates of points on images.

sample_points (*self*, *images*, *random_state*)
Generate coordinates of points on images.

Parameters

- **images** (*ndarray or list of ndarray*) – One or more images for which to generate points. If this is a `list` of arrays, each one of them is expected to have three dimensions. If this is an array, it must be four-dimensional and the first axis is expected to denote the image index. For `RGB` images the array would hence have to be of shape `(N, H, W, 3)`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState*) – A random state to use for any probabilistic function required during the point sampling. See *RNG()* for details.

Returns An `(N,2) float32` array containing `(x,y)` subpixel coordinates, all of which being within the intervals `[0.0, width]` and `[0.0, height]`.

Return type ndarray

class imgaug.augmenters.segmentation.**IPointsSampler**
Bases: `object`

Interface for all point samplers.

Point samplers return coordinate arrays of shape `Nx2`. These coordinates can be used in other augmenters, see e.g. *Voronoi*.

Methods

sample_points(self, images, random_state)	Generate coordinates of points on images.

sample_points (*self*, *images*, *random_state*)
Generate coordinates of points on images.

Parameters

- **images** (*ndarray or list of ndarray*) – One or more images for which to generate points. If this is a `list` of arrays, each one of them is expected to have three dimensions. If this is an array, it must be four-dimensional and the first axis is expected to denote the image index. For `RGB` images the array would hence have to be of shape `(N, H, W, 3)`.

- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState*) – A random state to use for any probabilistic function required during the point sampling. See *RNG()* for details.

Returns An `(N,2) float32` array containing `(x,y)` subpixel coordinates, all of which being within the intervals `[0.0, width]` and `[0.0, height]`.

Return type ndarray

class imgaug.augmenters.segmentation.**RegularGridPointsSampler**(*n_rows*, *n_cols*)

Bases: *imgaug.augmenters.segmentation.IPointsSampler*

Sampler that generates a regular grid of coordinates on an image.

'Regular grid' here means that on each axis all coordinates have the same distance from each other. Note that the distance may change between axis.

Parameters

- **n_rows** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter, optional*) – Number of rows of coordinates to place on each image, i.e. the number of coordinates on the y-axis. Note that for each image, the sampled value is clipped to the interval `[1..H]`, where `H` is the image height.
  - If a single `int`, then that value will always be used.
  - If a `tuple (a, b)`, then a value from the discrete interval `[a..b]` will be sampled per image.
  - If a `list`, then a random value will be sampled from that `list` per image.
  - If a `StochasticParameter`, then that parameter will be queried to draw one value per image.
- **n_cols** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter, optional*) – Number of columns of coordinates to place on each image, i.e. the number of coordinates on the x-axis. Note that for each image, the sampled value is clipped to the interval `[1..W]`, where `W` is the image width.
  - If a single `int`, then that value will always be used.
  - If a `tuple (a, b)`, then a value from the discrete interval `[a..b]` will be sampled per image.
  - If a `list`, then a random value will be sampled from that `list` per image.
  - If a `StochasticParameter`, then that parameter will be queried to draw one value per image.

Examples

```
>>> import imgaug.augmenters as iaa
>>> sampler = iaa.RegularGridPointsSampler(
>>>     n_rows=(5, 20),
>>>     n_cols=50)
```

Create a point sampler that generates regular grids of points. These grids contain `r` points on the y-axis, where `r` is sampled uniformly from the discrete interval `[5..20]` per image. On the x-axis, the grids always contain `50` points.

Methods

sample_points(self, images, random_state)	Generate coordinates of points on images.

sample_points (*self*, *images*, *random_state*)
Generate coordinates of points on images.

Parameters

- **images** (*ndarray or list of ndarray*) – One or more images for which to generate points. If this is a `list` of arrays, each one of them is expected to have three dimensions. If this is an array, it must be four-dimensional and the first axis is expected to denote the image index. For `RGB` images the array would hence have to be of shape `(N, H, W, 3)`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState*) – A random state to use for any probabilistic function required during the point sampling. See *RNG()* for details.

Returns An `(N,2) float32` array containing `(x,y)` subpixel coordinates, all of which being within the intervals `[0.0, width]` and `[0.0, height]`.

Return type ndarray

class imgaug.augmenters.segmentation.**RegularGridVoronoi**(*n_rows=(10, 30)*, *n_cols=(10, 30)*, *p_drop_points=(0.0, 0.5)*, *p_replace=(0.5, 1.0)*, *max_size=128*, *interpolation='linear'*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.segmentation.Voronoi*

Sample Voronoi cells from regular grids and color-average them.

This augmenter is a shortcut for the combination of *Voronoi*, *RegularGridPointsSampler* and *DropoutPointsSampler*. Hence, it generates a regular grid with `R` rows and `C` columns of coordinates on each image. Then, it drops `p` percent of the `R*C` coordinates to randomize the grid. Each image pixel then belongs to the voronoi cell with the closest coordinate.

Supported dtypes:

See *Voronoi*.

Parameters

- **n_rows** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter, optional*) – Number of rows of coordinates to place on each image, i.e. the number of coordinates on the y-axis. Note that for each image, the sampled value is clipped to the interval `[1..H]`, where `H` is the image height.
  - If a single `int`, then that value will always be used.
  - If a `tuple (a, b)`, then a value from the discrete interval `[a..b]` will be sampled per image.
  - If a `list`, then a random value will be sampled from that `list` per image.

  - If a `StochasticParameter`, then that parameter will be queried to draw one value per image.

- **n_cols** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter, optional*) – Number of columns of coordinates to place on each image, i.e. the number of coordinates on the x-axis. Note that for each image, the sampled value is clipped to the interval `[1..W]`, where `W` is the image width.

  - If a single `int`, then that value will always be used.
  - If a `tuple (a, b)`, then a value from the discrete interval `[a..b]` will be sampled per image.
  - If a `list`, then a random value will be sampled from that `list` per image.
  - If a `StochasticParameter`, then that parameter will be queried to draw one value per image.

- **p_drop_points** (*number or tuple of number or imgaug.parameters.StochasticParameter, optional*) – The probability that a coordinate will be removed from the list of all sampled coordinates. A value of `1.0` would mean that (on average) `100` percent of all coordinates will be dropped, while `0.0` denotes `0` percent. Note that this sampler will always ensure that at least one coordinate is left after the dropout operation, i.e. even `1.0` will only drop all *except one* coordinate.

  - If a `float`, then that value will be used for all images.
  - If a `tuple (a, b)`, then a value `p` will be sampled from the interval `[a, b]` per image.
  - If a `StochasticParameter`, then this parameter will be used to determine per coordinate whether it should be *kept* (sampled value of `>0.5`) or shouldn't be kept (sampled value of `<=0.5`). If you instead want to provide the probability as a stochastic parameter, you can usually do `imgaug.parameters.Binomial(1-p)` to convert parameter *p* to a 0/1 representation.

- **p_replace** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Defines for any segment the probability that the pixels within that segment are replaced by their average color (otherwise, the pixels are not changed). Examples:

  - A probability of `0.0` would mean, that the pixels in no segment are replaced by their average color (image is not changed at all).
  - A probability of `0.5` would mean, that around half of all segments are replaced by their average color.
  - A probability of `1.0` would mean, that all segments are replaced by their average color (resulting in a voronoi image).

  Behaviour based on chosen datatypes for this parameter:

  - If a `number`, then that number will always be used.
  - If `tuple (a, b)`, then a random probability will be sampled from the interval `[a, b]` per image.
  - If a `list`, then a random value will be sampled from that `list` per image.
  - If a `StochasticParameter`, it is expected to return values between `0.0` and `1.0` and will be queried *for each individual segment* to determine whether it is supposed to be averaged (`>0.5`) or not (`<=0.5`). Recommended to be some form of `Binomial(...)`.

- **max_size** (*int or None, optional*) – Maximum image size at which the augmentation is performed. If the width or height of an image exceeds this value, it will be downscaled before the augmentation so that the longest side matches *max_size*. This is done to speed up the process. The final output image has the same size as the input image. Note that in case *p_replace* is below `1.0`, the down-/upscaling will affect the not-replaced pixels too. Use `None` to apply no down-/upscaling.
- **interpolation** (*int or str, optional*) – Interpolation method to use during downscaling when *max_size* is exceeded. Valid methods are the same as in *imresize_single_image()*.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.RegularGridVoronoi(10, 20)
```

Place a regular grid of `10x20` (`height x width`) coordinates on each image. Randomly drop on average `20` percent of these points to create a less regular pattern. Then use the remaining coordinates to group the image pixels into voronoi cells and average the colors within them. The process is performed at an image size not exceeding `128` px on any side (default). If necessary, the downscaling is performed using `linear` interpolation (default).

```
>>> aug = iaa.RegularGridVoronoi(
>>>     (10, 30), 20, p_drop_points=0.0, p_replace=0.9, max_size=None)
```

Same as above, generates a grid with randomly `10` to `30` rows, drops none of the generates points, replaces only `90` percent of the voronoi cells with their average color (the pixels of the remaining `10` percent are not changed) and performs the transformation at the original image size (`max_size=None`).

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.

Continued on next page

Table 217 – continued from previous page

augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.segmentation.**RelativeRegularGridPointsSampler**(*n_rows_frac*, *n_cols_frac*)

Bases: *imgaug.augmenters.segmentation.IPointsSampler*

Regular grid coordinate sampler; places more points on larger images.

This is similar to RegularGridPointsSampler, but the number of rows and columns is given as fractions

of each image's height and width. Hence, more coordinates are generated for larger images.

Parameters

- **n_rows_frac** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Relative number of coordinates to place on the y-axis. For a value `y` and image height `H` the number of actually placed coordinates (i.e. computed rows) is given by `int(round(y*H))`. Note that for each image, the number of coordinates is clipped to the interval `[1,H]`, where `H` is the image height.
  - If a single `number`, then that value will always be used.
  - If a `tuple (a, b)`, then a value from the interval `[a, b]` will be sampled per image.
  - If a `list`, then a random value will be sampled from that `list` per image.
  - If a `StochasticParameter`, then that parameter will be queried to draw one value per image.
- **n_cols_frac** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Relative number of coordinates to place on the x-axis. For a value `x` and image height `W` the number of actually placed coordinates (i.e. computed columns) is given by `int(round(x*W))`. Note that for each image, the number of coordinates is clipped to the interval `[1,W]`, where `W` is the image width.
  - If a single `number`, then that value will always be used.
  - If a `tuple (a, b)`, then a value from the interval `[a, b]` will be sampled per image.
  - If a `list`, then a random value will be sampled from that `list` per image.
  - If a `StochasticParameter`, then that parameter will be queried to draw one value per image.

Examples

```
>>> import imgaug.augmenters as iaa
>>> sampler = iaa.RelativeRegularGridPointsSampler(
>>>     n_rows_frac=(0.01, 0.1),
>>>     n_cols_frac=0.2)
```

Create a point sampler that generates regular grids of points. These grids contain `round(y*H)` points on the y-axis, where `y` is sampled uniformly from the interval `[0.01, 0.1]` per image and `H` is the image height. On the x-axis, the grids always contain `0.2*W` points, where `W` is the image width.

Methods

sample_points(self, images, random_state)	Generate coordinates of points on images.

sample_points (*self*, *images*, *random_state*)

Generate coordinates of points on images.

Parameters

- **images** (*ndarray or list of ndarray*) – One or more images for which to generate points. If this is a `list` of arrays, each one of them is expected to have three dimensions. If this

is an array, it must be four-dimensional and the first axis is expected to denote the image index. For `RGB` images the array would hence have to be of shape `(N, H, W, 3)`.

- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState*) – A random state to use for any probabilistic function required during the point sampling. See *RNG()* for details.

Returns An `(N,2) float32` array containing `(x,y)` subpixel coordinates, all of which being within the intervals `[0.0, width]` and `[0.0, height]`.

Return type ndarray

class imgaug.augmenters.segmentation.**RelativeRegularGridVoronoi**(*n_rows_frac=(0.05, 0.15)*, *n_cols_frac=(0.05, 0.15)*, *p_drop_points=(0.0, 0.5)*, *p_replace=(0.5, 1.0)*, *max_size=None*, *interpolation='linear'*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.segmentation.Voronoi*

Sample Voronoi cells from image-dependent grids and color-average them.

This augmenter is a shortcut for the combination of *Voronoi*, *RegularGridPointsSampler* and *DropoutPointsSampler*. Hence, it generates a regular grid with `R` rows and `C` columns of coordinates on each image. Then, it drops `p` percent of the `R*C` coordinates to randomize the grid. Each image pixel then belongs to the voronoi cell with the closest coordinate.

Note: In contrast to the other voronoi augmenters, this one uses `None` as the default value for *max_size*, i.e. the color averaging is always performed at full resolution. This enables the augmenter to make use of the additional points on larger images. It does however slow down the augmentation process.

Supported dtypes:

See *Voronoi*.

Parameters

- **n_rows_frac** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Relative number of coordinates to place on the y-axis. For a value `y` and image height `H` the number of actually placed coordinates (i.e. computed rows) is given by `int(round(y*H))`. Note that for each image, the number of coordinates is clipped to the interval `[1,H]`, where `H` is the image height.
  - If a single `number`, then that value will always be used.
  - If a `tuple (a, b)`, then a value from the interval `[a, b]` will be sampled per image.

  - If a `list`, then a random value will be sampled from that `list` per image.
  - If a `StochasticParameter`, then that parameter will be queried to draw one value per image.

- **n_cols_frac** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Relative number of coordinates to place on the x-axis. For a value `x` and image height `W` the number of actually placed coordinates (i.e. computed columns) is given by `int(round(x*W))`. Note that for each image, the number of coordinates is clipped to the interval `[1,W]`, where `W` is the image width.
  - If a single `number`, then that value will always be used.
  - If a `tuple (a, b)`, then a value from the interval `[a, b]` will be sampled per image.
  - If a `list`, then a random value will be sampled from that `list` per image.
  - If a `StochasticParameter`, then that parameter will be queried to draw one value per image.

- **p_drop_points** (*number or tuple of number or imgaug.parameters.StochasticParameter, optional*) – The probability that a coordinate will be removed from the list of all sampled coordinates. A value of `1.0` would mean that (on average) `100` percent of all coordinates will be dropped, while `0.0` denotes `0` percent. Note that this sampler will always ensure that at least one coordinate is left after the dropout operation, i.e. even `1.0` will only drop all *except one* coordinate.
  - If a `float`, then that value will be used for all images.
  - If a `tuple (a, b)`, then a value `p` will be sampled from the interval `[a, b]` per image.
  - If a `StochasticParameter`, then this parameter will be used to determine per coordinate whether it should be *kept* (sampled value of `>0.5`) or shouldn't be kept (sampled value of `<=0.5`). If you instead want to provide the probability as a stochastic parameter, you can usually do `imgaug.parameters.Binomial(1-p)` to convert parameter *p* to a 0/1 representation.

- **p_replace** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Defines for any segment the probability that the pixels within that segment are replaced by their average color (otherwise, the pixels are not changed). Examples:
  - A probability of `0.0` would mean, that the pixels in no segment are replaced by their average color (image is not changed at all).
  - A probability of `0.5` would mean, that around half of all segments are replaced by their average color.
  - A probability of `1.0` would mean, that all segments are replaced by their average color (resulting in a voronoi image).

  Behaviour based on chosen datatypes for this parameter:
  - If a `number`, then that `number` will always be used.
  - If `tuple (a, b)`, then a random probability will be sampled from the interval `[a, b]` per image.
  - If a `list`, then a random value will be sampled from that `list` per image.

  - If a `StochasticParameter`, it is expected to return values between `0.0` and `1.0` and will be queried *for each individual segment* to determine whether it is supposed to be averaged (`>0.5`) or not (`<=0.5`). Recommended to be some form of `Binomial(...)`.

- **max_size** (*int or None, optional*) – Maximum image size at which the augmentation is performed. If the width or height of an image exceeds this value, it will be downscaled before the augmentation so that the longest side matches *max_size*. This is done to speed up the process. The final output image has the same size as the input image. Note that in case *p_replace* is below `1.0`, the down-/upscaling will affect the not-replaced pixels too. Use `None` to apply no down-/upscaling.
- **interpolation** (*int or str, optional*) – Interpolation method to use during downscaling when *max_size* is exceeded. Valid methods are the same as in *imresize_single_image()*.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.RelativeRegularGridVoronoi(0.1, 0.25)
```

Place a regular grid of `R x C` coordinates on each image, where `R` is the number of rows and computed as `R=0.1*H` with `H` being the height of the input image. `C` is the number of columns and analogously estimated from the image width `W` as `C=0.25*W`. Larger images will lead to larger `R` and `C` values. On average, `20` percent of these grid coordinates are randomly dropped to create a less regular pattern. Then, the remaining coordinates are used to group the image pixels into voronoi cells and the colors within them are averaged.

```
>>> aug = iaa.RelativeRegularGridVoronoi(
>>>     (0.03, 0.1), 0.1, p_drop_points=0.0, p_replace=0.9, max_size=512)
```

Same as above, generates a grid with randomly `R=r*H` rows, where `r` is sampled uniformly from the interval `[0.03, 0.1]` and `C=0.1*W` rows. No points are dropped. The augmenter replaces only `90` percent of the voronoi cells with their average color (the pixels of the remaining `10` percent are not changed). Images larger than `512` px are temporarily downscaled (*before* sampling the grid points) so that no side exceeds `512` px. This improves performance, but degrades the quality of the resulting image.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.

Continued on next page

Table 219 – continued from previous page

augment(self[, return_batch, hooks])	Augment a batch.
augment_batch(self, batch[, hooks])	**Deprecated**.
augment_batch_(self, batch[, parents, hooks])	Augment a single batch in-place.
augment_batches(self, batches[, hooks, ...])	Augment multiple batches.
augment_bounding_boxes(self, ...[, parents, ...])	Augment a batch of bounding boxes.
augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.segmentation.**SubsamplingPointsSampler**(*other_points_sampler*, *n_points_max*)

Bases: *imgaug.augmenters.segmentation.IPointsSampler*

Ensure that the number of sampled points is below a maximum.

This point sampler will sample points from another sampler and then – in case more points were generated than an allowed maximum – will randomly pick *n_points_max* of these.

Parameters

- **other_points_sampler** (*IPointsSampler*) – Another point sampler that is queried to generate a `list` of points. The dropout operation will be applied to that `list`.
- **n_points_max** (*int*) – Maximum number of allowed points. If *other_points_sampler* generates more points than this maximum, a random subset of size *n_points_max* will be selected.

Examples

```
>>> import imgaug.augmenters as iaa
>>> sampler = iaa.SubsamplingPointsSampler(
>>>     iaa.RelativeRegularGridPointsSampler(0.1, 0.2),
>>>     50
>>> )
```

Create a points sampler that places `y*H` points on the y-axis (with `y` being `0.1` and `H` being an image's height) and `x*W` on the x-axis (analogous). Then, if that number of placed points exceeds `50` (can easily happen for larger images), a random subset of `50` points will be picked and returned.

Methods

sample_points(self, images, random_state)	Generate coordinates of points on images.

sample_points(*self*, *images*, *random_state*)

Generate coordinates of points on images.

Parameters

- **images** (*ndarray or list of ndarray*) – One or more images for which to generate points. If this is a `list` of arrays, each one of them is expected to have three dimensions. If this is an array, it must be four-dimensional and the first axis is expected to denote the image index. For `RGB` images the array would hence have to be of shape `(N, H, W, 3)`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState*) – A random state to use for any probabilistic function required during the point sampling. See *RNG()* for details.

Returns An `(N,2) float32` array containing `(x,y)` subpixel coordinates, all of which being within the intervals `[0.0, width]` and `[0.0, height]`.

Return type ndarray

class imgaug.augmenters.segmentation.**Superpixels**(*p_replace=(0.5, 1.0), n_segments=(50, 120), max_size=128, interpolation='linear', seed=None, name=None, random_state='deprecated', deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Transform images parially/completely to their superpixel representation.

This implementation uses skimage's version of the SLIC algorithm.

Note: This augmenter is fairly slow. See *Performance*.

Supported dtypes:

if (image size <= max_size):

- `uint8`: yes; fully tested
- `uint16`: yes; tested
- `uint32`: yes; tested
- `uint64`: limited (1)
- `int8`: yes; tested
- `int16`: yes; tested
- `int32`: yes; tested
- `int64`: limited (1)
- `float16`: no (2)
- `float32`: no (2)
- `float64`: no (3)
- `float128`: no (2)
- `bool`: yes; tested

- (1) Superpixel mean intensity replacement requires computing these means as `float64` s. This can cause inaccuracies for large integer values.
- (2) Error in scikit-image.
- (3) Loss of resolution in scikit-image.

if (image size > max_size):

minimum of (`imgaug.augmenters.segmentation.Superpixels(image size <= max_size)`, `_ensure_image_max_size()`

)

Parameters

- **p_replace** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Defines for any segment the probability that the pixels within that segment are replaced by their average color (otherwise, the pixels are not changed). Examples:
  - A probability of `0.0` would mean, that the pixels in no segment are replaced by their average color (image is not changed at all).
  - A probability of `0.5` would mean, that around half of all segments are replaced by their average color.
  - A probability of `1.0` would mean, that all segments are replaced by their average color (resulting in a voronoi image).

  Behaviour based on chosen datatypes for this parameter:
  - If a `number`, then that `number` will always be used.
  - If `tuple (a, b)`, then a random probability will be sampled from the interval `[a, b]` per image.
  - If a `list`, then a random value will be sampled from that `list` per image.
  - If a `StochasticParameter`, it is expected to return values between `0.0` and `1.0` and will be queried *for each individual segment* to determine whether it is supposed to be averaged (`>0.5`) or not (`<=0.5`). Recommended to be some form of `Binomial(...)`.
- **n_segments** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter, optional*) – Rough target number of how many superpixels to generate (the algorithm may deviate from this number). Lower value will lead to coarser superpixels. Higher values are computationally more intensive and will hence lead to a slowdown.
  - If a single `int`, then that value will always be used as the number of segments.
  - If a `tuple (a, b)`, then a value from the discrete interval `[a..b]` will be sampled per image.
  - If a `list`, then a random value will be sampled from that `list` per image.
  - If a `StochasticParameter`, then that parameter will be queried to draw one value per image.
- **max_size** (*int or None, optional*) – Maximum image size at which the augmentation is performed. If the width or height of an image exceeds this value, it will be downscaled before the augmentation so that the longest side matches *max_size*. This is done to speed up the process. The final output image has the same size as the input image. Note that in case *p_replace* is below `1.0`, the down-/upscaling will affect the not-replaced pixels too. Use `None` to apply no down-/upscaling.
- **interpolation** (*int or str, optional*) – Interpolation method to use during downscaling when *max_size* is exceeded. Valid methods are the same as in *imresize_single_image()*.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage

will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.

- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.Superpixels(p_replace=1.0, n_segments=64)
```

Generate around `64` superpixels per image and replace all of them with their average color (standard superpixel image).

```
>>> aug = iaa.Superpixels(p_replace=0.5, n_segments=64)
```

Generate around `64` superpixels per image and replace half of them with their average color, while the other half are left unchanged (i.e. they still show the input image's content).

```
>>> aug = iaa.Superpixels(p_replace=(0.25, 1.0), n_segments=(16, 128))
```

Generate between `16` and `128` superpixels per image and replace `25` to `100` percent of them with their average color.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).

Continued on next page

Table 221 – continued from previous page

deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.segmentation.**UniformPointsSampler**(*n_points*)
Bases: *imgaug.augmenters.segmentation.IPointsSampler*

Sample points uniformly on images.

This point sampler generates *n_points* points per image. The x- and y-coordinates are both sampled from uniform distributions matching the respective image width and height.

Parameters **n_points** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter, optional*) –

Number of points to sample on each image.

- If a single `int`, then that value will always be used.
- If a `tuple (a, b)`, then a value from the discrete interval `[a..b]` will be sampled per image.
- If a `list`, then a random value will be sampled from that `list` per image.
- If a `StochasticParameter`, then that parameter will be queried to draw one value per image.

Examples

```
>>> import imgaug.augmenters as iaa
>>> sampler = iaa.UniformPointsSampler(500)
```

Create a point sampler that generates an array of `500` random points for each input image. The x- and y-coordinates of each point are sampled from uniform distributions.

Methods

sample_points(self, images, random_state)	Generate coordinates of points on images.

sample_points (*self*, *images*, *random_state*)
Generate coordinates of points on images.

Parameters

- **images** (*ndarray or list of ndarray*) – One or more images for which to generate points. If this is a `list` of arrays, each one of them is expected to have three dimensions. If this is an array, it must be four-dimensional and the first axis is expected to denote the image index. For `RGB` images the array would hence have to be of shape `(N, H, W, 3)`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState*) – A random state to use for any probabilistic function required during the point sampling. See *RNG()* for details.

Returns An `(N,2) float32` array containing `(x,y)` subpixel coordinates, all of which being within the intervals `[0.0, width]` and `[0.0, height]`.

Return type ndarray

class imgaug.augmenters.segmentation.**UniformVoronoi**(*n_points=(50, 500)*, *p_replace=(0.5, 1.0)*, *max_size=128*, *interpolation='linear'*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.segmentation.Voronoi*

Uniformly sample Voronoi cells on images and average colors within them.

This augmenter is a shortcut for the combination of *Voronoi* with *UniformPointsSampler*. Hence, it generates a fixed amount of `N` random coordinates of voronoi cells on each image. The cell coordinates are sampled uniformly using the image height and width as maxima.

Supported dtypes:

See *Voronoi*.

Parameters

- **n_points** (*int or tuple of int or list of int or imgaug.parameters.StochasticParameter, optional*) –

  Number of points to sample on each image.

  – If a single `int`, then that value will always be used.

  - If a `tuple (a, b)`, then a value from the discrete interval `[a..b]` will be sampled per image.
  - If a `list`, then a random value will be sampled from that `list` per image.
  - If a `StochasticParameter`, then that parameter will be queried to draw one value per image.

- **p_replace** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Defines for any segment the probability that the pixels within that segment are replaced by their average color (otherwise, the pixels are not changed). Examples:
  - A probability of `0.0` would mean, that the pixels in no segment are replaced by their average color (image is not changed at all).
  - A probability of `0.5` would mean, that around half of all segments are replaced by their average color.
  - A probability of `1.0` would mean, that all segments are replaced by their average color (resulting in a voronoi image).

  Behaviour based on chosen datatypes for this parameter:
  - If a `number`, then that `number` will always be used.
  - If `tuple (a, b)`, then a random probability will be sampled from the interval `[a, b]` per image.
  - If a `list`, then a random value will be sampled from that `list` per image.
  - If a `StochasticParameter`, it is expected to return values between `0.0` and `1.0` and will be queried *for each individual segment* to determine whether it is supposed to be averaged (`>0.5`) or not (`<=0.5`). Recommended to be some form of `Binomial(...)`.
- **max_size** (*int or None, optional*) – Maximum image size at which the augmentation is performed. If the width or height of an image exceeds this value, it will be downscaled before the augmentation so that the longest side matches *max_size*. This is done to speed up the process. The final output image has the same size as the input image. Note that in case *p_replace* is below `1.0`, the down-/upscaling will affect the not-replaced pixels too. Use `None` to apply no down-/upscaling.
- **interpolation** (*int or str, optional*) – Interpolation method to use during downscaling when *max_size* is exceeded. Valid methods are the same as in *imresize_single_image()*.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.UniformVoronoi((100, 500))
```

Sample for each image uniformly the number of voronoi cells `N` from the interval `[100, 500]`. Then generate `N` coordinates by sampling uniformly the x-coordinates from `[0, W]` and the y-coordinates from `[0, H]`, where `H` is the image height and `W` the image width. Then use these coordinates to group the image pixels into voronoi cells and average the colors within them. The process is performed at an image size not exceeding `128` px on any side (default). If necessary, the downscaling is performed using `linear` interpolation (default).

```
>>> aug = iaa.UniformVoronoi(250, p_replace=0.9, max_size=None)
```

Same as above, but always samples `N=250` cells, replaces only `90` percent of them with their average color (the pixels of the remaining `10` percent are not changed) and performs the transformation at the original image size (`max_size=None`).

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.

Continued on next page

Table 223 – continued from previous page

get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.segmentation.**Voronoi**(*points_sampler*, *p_replace=1.0*, *max_size=128*, *interpolation='linear'*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Average colors of an image within Voronoi cells.

This augmenter performs the following steps:

1. Query *points_sampler* to sample random coordinates of cell centers. On the image.
2. Estimate for each pixel to which voronoi cell (i.e. segment) it belongs. Each pixel belongs to the cell with the closest center coordinate (euclidean distance).
3. Compute for each cell the average color of the pixels within it.
4. Replace the pixels of *p_replace* percent of all cells by their average color. Do not change the pixels of `(1 - p_replace)` percent of all cells. (The percentages are average values over many images. Some images may get more/less cells replaced by their average color.)

This code is very loosely based on https://codegolf.stackexchange.com/questions/50299/draw-an-image-as-a-voronoi-map/50345#50345

Supported dtypes:

if (image size <= max_size):

- `uint8`: yes; fully tested
- `uint16`: no; not tested
- `uint32`: no; not tested
- `uint64`: no; not tested
- `int8`: no; not tested

- `int16`: no; not tested
- `int32`: no; not tested
- `int64`: no; not tested
- `float16`: no; not tested
- `float32`: no; not tested
- `float64`: no; not tested
- `float128`: no; not tested
- `bool`: no; not tested

if (image size > max_size):

minimum of (`imgaug.augmenters.segmentation.Voronoi(image size <= max_size)`, `_ensure_image_max_size()`

)

Parameters

- **points_sampler** (*IPointsSampler*) – A points sampler which will be queried per image to generate the coordinates of the centers of voronoi cells.
- **p_replace** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Defines for any segment the probability that the pixels within that segment are replaced by their average color (otherwise, the pixels are not changed). Examples:
  - A probability of `0.0` would mean, that the pixels in no segment are replaced by their average color (image is not changed at all).
  - A probability of `0.5` would mean, that around half of all segments are replaced by their average color.
  - A probability of `1.0` would mean, that all segments are replaced by their average color (resulting in a voronoi image).

  Behaviour based on chosen datatypes for this parameter:
  - If a `number`, then that `number` will always be used.
  - If `tuple (a, b)`, then a random probability will be sampled from the interval `[a, b]` per image.
  - If a `list`, then a random value will be sampled from that `list` per image.
  - If a `StochasticParameter`, it is expected to return values between `0.0` and `1.0` and will be queried *for each individual segment* to determine whether it is supposed to be averaged (`>0.5`) or not (`<=0.5`). Recommended to be some form of `Binomial(...)`.
- **max_size** (*int or None, optional*) – Maximum image size at which the augmentation is performed. If the width or height of an image exceeds this value, it will be downscaled before the augmentation so that the longest side matches *max_size*. This is done to speed up the process. The final output image has the same size as the input image. Note that in case *p_replace* is below `1.0`, the down-/upscaling will affect the not-replaced pixels too. Use `None` to apply no down-/upscaling.
- **interpolation** (*int or str, optional*) – Interpolation method to use during downscaling when *max_size* is exceeded. Valid methods are the same as in *imresize_single_image()*.

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> points_sampler = iaa.RegularGridPointsSampler(n_cols=20, n_rows=40)
>>> aug = iaa.Voronoi(points_sampler)
```

Create an augmenter that places a `20x40` (`HxW`) grid of cells on the image and replaces all pixels within each cell by the cell's average color. The process is performed at an image size not exceeding `128` px on any side (default). If necessary, the downscaling is performed using `linear` interpolation (default).

```
>>> points_sampler = iaa.DropoutPointsSampler(
>>>     iaa.RelativeRegularGridPointsSampler(
>>>         n_cols_frac=(0.05, 0.2),
>>>         n_rows_frac=0.1),
>>>     0.2)
>>> aug = iaa.Voronoi(points_sampler, p_replace=0.9, max_size=None)
```

Create a voronoi augmenter that generates a grid of cells dynamically adapted to the image size. Larger images get more cells. On the x-axis, the distance between two cells is `w * W` pixels, where `W` is the width of the image and `w` is always `0.1`. On the y-axis, the distance between two cells is `h * H` pixels, where `H` is the height of the image and `h` is sampled uniformly from the interval `[0.05, 0.2]`. To make the voronoi pattern less regular, about `20` percent of the cell coordinates are randomly dropped (i.e. the remaining cells grow in size). In contrast to the first example, the image is not resized (if it was, the sampling would happen *after* the resizing, which would affect `W` and `H`). Not all voronoi cells are replaced by their average color, only around `90` percent of them. The remaining `10` percent's pixels remain unchanged.

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.

Continued on next page

Table 224 – continued from previous page

augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See *get_parameters()*.

imgaug.augmenters.segmentation.**segment_voronoi**(*image*, *cell_coordinates*, *replace_mask=None*)

Average colors within voronoi cells of an image.

Parameters

- **image** (*ndarray*) – The image to convert to a voronoi image. May be HxW or HxWxC. Note

that for RGBA images the alpha channel will currently also by averaged.

- **cell_coordinates** (*ndarray*) – A Nx2 float array containing the center coordinates of voronoi cells on the image. Values are expected to be in the interval [0.0, height-1.0] for the y-axis (x-axis analogous). If this array contains no coordinate, the image will not be changed.
- **replace_mask** (*None or ndarray, optional*) – Boolean mask of the same length as *cell_coordinates*, denoting for each cell whether its pixels are supposed to be replaced by the cell's average color (True) or left untouched (False). If this is set to None, all cells will be replaced.

Returns Voronoi image.

Return type ndarray

13.35 imgaug.augmenters.size

Augmenters that somehow change the size of the images.

List of augmenters:

- *Resize*
- *CropAndPad*
- *Crop*
- *Pad*
- *PadToFixedSize*
- *CenterPadToFixedSize*
- *CropToFixedSize*
- *CenterCropToFixedSize*
- *CropToMultiplesOf*
- *CenterCropToMultiplesOf*
- *PadToMultiplesOf*
- *CenterPadToMultiplesOf*
- *CropToPowersOf*
- *CenterCropToPowersOf*
- *PadToPowersOf*
- *CenterPadToPowersOf*
- *CropToAspectRatio*
- *CenterCropToAspectRatio*
- *PadToAspectRatio*
- *CenterPadToAspectRatio*
- *CropToSquare*
- *CenterCropToSquare*

- *PadToSquare*
- *CenterPadToSquare*
- *KeepSizeByResize*

class imgaug.augmenters.size.**CenterCropToAspectRatio**(*aspect_ratio*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.size.CropToAspectRatio*

Crop images equally on all sides until they reach an aspect ratio.

This is the same as *CropToAspectRatio*, but uses `position="center"` by default, which spreads the crop amounts equally over all image sides, while *CropToAspectRatio* by default spreads them randomly.

Added in 0.4.0.

Supported dtypes:

See *CropToFixedSize*.

Parameters

- **aspect_ratio** (*number*) – See `CropToAspectRatio.__init__()`.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.CenterCropToAspectRatio(2.0)
```

Create an augmenter that crops each image until its aspect ratio is as close as possible to `2.0` (i.e. two times as many pixels along the x-axis than the y-axis). The rows to be cropped will be spread *equally* over the top and bottom sides (analogous for the left/right sides).

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.

Continued on next page

Table 225 – continued from previous page

augment_batches(self, batches[, hooks, ...])	Augment multiple batches.
augment_bounding_boxes(self, ...[, parents, ...])	Augment a batch of bounding boxes.
augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.size.**CenterCropToFixedSize**(*width*, *height*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.size.CropToFixedSize*

Take a crop from the center of each image.

This is an alias for *CropToFixedSize* with `position="center"`.

Note: If images already have a width and/or height below the provided width and/or height then this augmenter will do nothing for the respective axis. Hence, resulting images can be smaller than the provided axis sizes.

Added in 0.4.0.

Supported dtypes:

See *CropToFixedSize*.

Parameters

- **width** (*int or None*) – See `CropToFixedSize.__init__()`.
- **height** (*int or None*) – See `CropToFixedSize.__init__()`.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> crop = iaa.CenterCropToFixedSize(height=20, width=10)
```

Create an augmenter that takes `20x10` sized crops from the center of images.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.

Continued on next page

Table 226 – continued from previous page

augment_bounding_boxes(self, ...[, parents, ...])	Augment a batch of bounding boxes.
augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.size.**CenterCropToMultiplesOf**(*width_multiple*, *height_multiple*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.size.CropToMultiplesOf*

Crop images equally on all sides until H/W are multiples of given values.

This is the same as *CropToMultiplesOf*, but uses `position="center"` by default, which spreads the crop amounts equally over all image sides, while *CropToMultiplesOf* by default spreads them randomly.

Added in 0.4.0.

Supported dtypes:

See *CropToFixedSize*.

Parameters

- **width_multiple** (*int or None*) – See `CropToMultiplesOf.__init__()`.
- **height_multiple** (*int or None*) – See `CropToMultiplesOf.__init__()`.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.CenterCropToMultiplesOf(height_multiple=10, width_multiple=6)
```

Create an augmenter that crops images to multiples of `10` along the y-axis (i.e. 10, 20, 30, ...) and to multiples of `6` along the x-axis (i.e. 6, 12, 18, ...). The rows to be cropped will be spread *equally* over the top and bottom sides (analogous for the left/right sides).

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.

Continued on next page

Table 227 – continued from previous page

augment_bounding_boxes(self, ...[, parents, ...])	Augment a batch of bounding boxes.
augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.size.**CenterCropToPowersOf**(*width_base*, *height_base*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.size.CropToPowersOf*

Crop images equally on all sides until H/W is a power of a base.

This is the same as *CropToPowersOf*, but uses `position="center"` by default, which spreads the crop amounts equally over all image sides, while *CropToPowersOf* by default spreads them randomly.

Added in 0.4.0.

Supported dtypes:

See *CropToFixedSize*.

Parameters

- **width_base** (*int or None*) – See `CropToPowersOf.__init__()`.
- **height_base** (*int or None*) – See `CropToPowersOf.__init__()`.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.CropToPowersOf(height_base=3, width_base=2)
```

Create an augmenter that crops each image down to powers of `3` along the y-axis (i.e. 3, 9, 27, ...) and powers of `2` along the x-axis (i.e. 2, 4, 8, 16, ...). The rows to be cropped will be spread *equally* over the top and bottom sides (analogous for the left/right sides).

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.

Continued on next page

Table 228 – continued from previous page

augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.size.**CenterCropToSquare**(*seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.size.CropToSquare*

Crop images equally on all sides until their height/width are identical.

In contrast to *CropToSquare*, this augmenter always tries to spread the columns/rows to remove equally over both sides of the respective axis to be cropped. *CropToAspectRatio* by default spreads the croppings randomly.

This augmenter is identical to *CropToSquare* with `position="center"`, and thereby the same as *CropToAspectRatio* with `aspect_ratio=1.0, position="center"`.

Images with axis sizes of `0` will not be altered.

Added in 0.4.0.

Supported dtypes:

See *CropToFixedSize*.

Parameters

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.CenterCropToSquare()
```

Create an augmenter that crops each image until its square, i.e. height and width match. The rows to be cropped will be spread *equally* over the top and bottom sides (analogous for the left/right sides).

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.

Continued on next page

Table 229 – continued from previous page

augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.size.**CenterPadToAspectRatio**(*aspect_ratio*, *pad_mode='constant'*, *pad_cval=0*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.size.PadToAspectRatio*

Pad images equally on all sides until H/W matches an aspect ratio.

This is the same as *PadToAspectRatio*, but uses `position="center"` by default, which spreads the

pad amounts equally over all image sides, while *PadToAspectRatio* by default spreads them randomly.

Added in 0.4.0.

Supported dtypes:

See *PadToFixedSize*.

Parameters

- **aspect_ratio** (*number*) – See `PadToAspectRatio.__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **pad_mode** (*imgaug.ALL or str or list of str or imgaug.parameters.StochasticParameter, optional*) – See `__init__()`.
- **pad_cval** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – See `__init__()`.
- **deterministic** (*bool, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.PadToAspectRatio(2.0)
```

Create am augmenter that pads each image until its aspect ratio is as close as possible to `2.0` (i.e. two times as many pixels along the x-axis than the y-axis). The rows to be padded will be spread *equally* over the top and bottom sides (analogous for the left/right sides).

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.

Continued on next page

Table 230 – continued from previous page

augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.size.**CenterPadToFixedSize**(*width*, *height*, *pad_mode='constant'*, *pad_cval=0*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.size.PadToFixedSize*

Pad images equally on all sides up to given minimum heights/widths.

This is an alias for *PadToFixedSize* with `position="center"`. It spreads the pad amounts equally over all image sides, while *PadToFixedSize* by defaults spreads them randomly.

Added in 0.4.0.

Supported dtypes:

See *PadToFixedSize*.

Parameters

- **width** (*int or None*) – See `PadToFixedSize.__init__()`.
- **height** (*int or None*) – See `PadToFixedSize.__init__()`.
- **pad_mode** (*imgaug.ALL or str or list of str or imgaug.parameters.StochasticParameter, optional*) – See `PadToFixedSize.__init__()`.
- **pad_cval** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – See `PadToFixedSize.__init__()`.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.CenterPadToFixedSize(height=20, width=30)
```

Create an augmenter that pads images up to `20x30`, with the padded rows added *equally* on the top and bottom (analogous for the padded columns).

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.

Continued on next page

Table 231 – continued from previous page

augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.size.**CenterPadToMultiplesOf**(*width_multiple*, *height_multiple*, *pad_mode='constant'*, *pad_cval=0*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.size.PadToMultiplesOf*

Pad images equally on all sides until H/W are multiples of given values.

This is the same as *PadToMultiplesOf*, but uses `position="center"` by default, which spreads the pad amounts equally over all image sides, while *PadToMultiplesOf* by default spreads them randomly.

Added in 0.4.0.

Supported dtypes:

See *PadToFixedSize*.

Parameters

- **width_multiple** (*int or None*) – See `PadToMultiplesOf.__init__()`.
- **height_multiple** (*int or None*) – See `PadToMultiplesOf.__init__()`.
- **pad_mode** (*imgaug.ALL or str or list of str or imgaug.parameters.StochasticParameter, optional*) – See `__init__()`.
- **pad_cval** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – See `__init__()`.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.CenterPadToMultiplesOf(height_multiple=10, width_multiple=6)
```

Create an augmenter that pads images to multiples of `10` along the y-axis (i.e. 10, 20, 30, ...) and to multiples of `6` along the x-axis (i.e. 6, 12, 18, ...). The rows to be padded will be spread *equally* over the top and bottom sides (analogous for the left/right sides).

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.

Continued on next page

Table 232 – continued from previous page

augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.size.**CenterPadToPowersOf**(*width_base*, *height_base*, *pad_mode='constant'*, *pad_cval=0*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.size.PadToPowersOf*

Pad images equally on all sides until H/W is a power of a base.

This is the same as *PadToPowersOf*, but uses `position="center"` by default, which spreads the pad amounts equally over all image sides, while *PadToPowersOf* by default spreads them randomly.

Added in 0.4.0.

Supported dtypes:

See *PadToFixedSize*.

Parameters

- **width_base** (*int or None*) – See `PadToPowersOf.__init__()`.
- **height_base** (*int or None*) – See `PadToPowersOf.__init__()`.
- **pad_mode** (*imgaug.ALL or str or list of str or imgaug.parameters.StochasticParameter, optional*) – See `__init__()`.
- **pad_cval** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – See `__init__()`.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.CenterPadToPowersOf(height_base=5, width_base=2)
```

Create an augmenter that pads each image to powers of `3` along the y-axis (i.e. 3, 9, 27, ...) and powers of `2` along the x-axis (i.e. 2, 4, 8, 16, ...). The rows to be padded will be spread *equally* over the top and bottom sides (analogous for the left/right sides).

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.

Continued on next page

Table 233 – continued from previous page

augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.size.**CenterPadToSquare**(*pad_mode='constant'*, *pad_cval=0*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.size.PadToSquare*

Pad images equally on all sides until their height & width are identical.

This is the same as *PadToSquare*, but uses `position="center"` by default, which spreads the pad amounts equally over all image sides, while *PadToSquare* by default spreads them randomly. This augmenter is thus also identical to *PadToAspectRatio* with `aspect_ratio=1.0, position="center"`.

Added in 0.4.0.

Supported dtypes:

See *PadToFixedSize*.

Parameters

- **name** (*None or str, optional*) – See `__init__()`.
- **pad_mode** (*imgaug.ALL or str or list of str or imgaug.parameters.StochasticParameter, optional*) – See `__init__()`.
- **pad_cval** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – See `__init__()`.
- **deterministic** (*bool, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.CenterPadToSquare()
```

Create an augmenter that pads each image until its square, i.e. height and width match. The rows to be padded will be spread *equally* over the top and bottom sides (analogous for the left/right sides).

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.

Continued on next page

Table 234 – continued from previous page

find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.size.**Crop**(*px=None*, *percent=None*, *keep_size=True*, *sample_independently=True*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.size.CropAndPad*

Crop images, i.e. remove columns/rows of pixels at the sides of images.

This augmenter allows to extract smaller-sized subimages from given full-sized input images. The number of pixels to cut off may be defined in absolute values or as fractions of the image sizes.

This augmenter will never crop images below a height or width of `1`.

Supported dtypes:

See *CropAndPad*.

Parameters

- **px** (*None or int or imgaug.parameters.StochasticParameter or tuple, optional*) – The number of pixels to crop on each side of the image. Expected value range is `[0, inf)`. Either this or the parameter *percent* may be set, not both at the same time.
  - If `None`, then pixel-based cropping will not be used.
  - If `int`, then that exact number of pixels will always be cropped.
  - If `StochasticParameter`, then that parameter will be used for each image. Four samples will be drawn per image (top, right, bottom, left), unless *sample_independently* is set to `False`, as then only one value will be sampled per image and used for all sides.
  - If a `tuple` of two `int` s with values `a` and `b`, then each side will be cropped by a random amount sampled uniformly per image and side from the inteval `[a, b]`. If however

sample_independently is set to `False`, only one value will be sampled per image and used for all sides.

  - If a `tuple` of four entries, then the entries represent top, right, bottom, left. Each entry may be a single `int` (always crop by exactly that value), a `tuple` of two `int` s `a` and `b` (crop by an amount within `[a, b]`), a `list` of `int` s (crop by a random value that is contained in the `list`) or a `StochasticParameter` (sample the amount to crop from that parameter).

- **percent** (*None or int or float or imgaug.parameters.StochasticParameter or tuple, optional*) – The number of pixels to crop on each side of the image given as a *fraction* of the image height/width. E.g. if this is set to `0.1`, the augmenter will always crop `10%` of the image's height at both the top and the bottom (both `10%` each), as well as `10%` of the width at the right and left. Expected value range is `[0.0, 1.0)`. Either this or the parameter *px* may be set, not both at the same time.

  - If `None`, then fraction-based cropping will not be used.

  - If `number`, then that fraction will always be cropped.

  - If `StochasticParameter`, then that parameter will be used for each image. Four samples will be drawn per image (top, right, bottom, left). If however *sample_independently* is set to `False`, only one value will be sampled per image and used for all sides.

  - If a `tuple` of two `float` s with values `a` and `b`, then each side will be cropped by a random fraction sampled uniformly per image and side from the interval `[a, b]`. If however *sample_independently* is set to `False`, only one value will be sampled per image and used for all sides.

  - If a `tuple` of four entries, then the entries represent top, right, bottom, left. Each entry may be a single `float` (always crop by exactly that fraction), a `tuple` of two `float` s `a` and `b` (crop by a fraction from `[a, b]`), a `list` of `float` s (crop by a random value that is contained in the list) or a `StochasticParameter` (sample the percentage to crop from that parameter).

- **keep_size** (*bool, optional*) – After cropping, the result image will usually have a different height/width compared to the original input image. If this parameter is set to `True`, then the cropped image will be resized to the input image's size, i.e. the augmenter's output shape is always identical to the input shape.

- **sample_independently** (*bool, optional*) – If `False` *and* the values for *px*/*percent* result in exactly *one* probability distribution for all image sides, only one single value will be sampled from that probability distribution and used for all sides. I.e. the crop amount then is the same for all sides. If `True`, four values will be sampled independently, one per side.

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.

- **name** (*None or str, optional*) – See `__init__()`.

- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.

- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic

mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.Crop(px=(0, 10))
```

Crop each side by a random pixel value sampled uniformly per image and side from the discrete interval `[0..10]`.

```
>>> aug = iaa.Crop(px=(0, 10), sample_independently=False)
```

Crop each side by a random pixel value sampled uniformly once per image from the discrete interval `[0..10]`. Each sampled value is used for *all* sides of the corresponding image.

```
>>> aug = iaa.Crop(px=(0, 10), keep_size=False)
```

Crop each side by a random pixel value sampled uniformly per image and side from the discrete interval `[0..10]`. Afterwards, do **not** resize the cropped image back to the input image's size. This will decrease the image's height and width by a maximum of `20` pixels.

```
>>> aug = iaa.Crop(px=((0, 10), (0, 5), (0, 10), (0, 5)))
```

Crop the top and bottom by a random pixel value sampled uniformly from the discrete interval `[0..10]`. Crop the left and right analogously by a random value sampled from `[0..5]`. Each value is always sampled independently.

```
>>> aug = iaa.Crop(percent=(0, 0.1))
```

Crop each side by a random fraction sampled uniformly from the continuous interval `[0.0, 0.10]`. The fraction is sampled once per image and side. E.g. a sampled fraction of `0.1` for the top side would crop by `0.1*H`, where `H` is the height of the input image.

```
>>> aug = iaa.Crop(
>>>     percent=([0.05, 0.1], [0.05, 0.1], [0.05, 0.1], [0.05, 0.1]))
```

Crops each side by either `5%` or `10%`. The values are sampled once per side and image.

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.

Continued on next page

Table 235 – continued from previous page

augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.size.**CropAndPad**(*px=None*, *percent=None*, *pad_mode='constant'*, *pad_cval=0*, *keep_size=True*, *sample_independently=True*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Crop/pad images by pixel amounts or fractions of image sizes.

Cropping removes pixels at the sides (i.e. extracts a subimage from a given full image). Padding adds pixels to the sides (e.g. black pixels).

This augmenter will never crop images below a height or width of 1.

Note: This augmenter automatically resizes images back to their original size after it has augmented them. To deactivate this, add the parameter keep_size=False.

Supported dtypes:

if (keep_size=False):

- uint8: yes; fully tested
- uint16: yes; tested
- uint32: yes; tested
- uint64: yes; tested
- int8: yes; tested
- int16: yes; tested
- int32: yes; tested
- int64: yes; tested
- float16: yes; tested
- float32: yes; tested
- float64: yes; tested
- float128: yes; tested
- bool: yes; tested

if (keep_size=True):

minimum of (imgaug.augmenters.size.CropAndPad(keep_size=False), *imresize_many_images()*

)

Parameters

- **px** (*None or int or imgaug.parameters.StochasticParameter or tuple, optional*) – The number of pixels to crop (negative values) or pad (positive values) on each side of the image. Either this or the parameter *percent* may be set, not both at the same time.
  - If None, then pixel-based cropping/padding will not be used.
  - If int, then that exact number of pixels will always be cropped/padded.
  - If StochasticParameter, then that parameter will be used for each image. Four samples will be drawn per image (top, right, bottom, left), unless *sample_independently* is set to False, as then only one value will be sampled per image and used for all sides.
  - If a tuple of two int s with values a and b, then each side will be cropped/padded by a random amount sampled uniformly per image and side from the inteval [a, b]. If however *sample_independently* is set to False, only one value will be sampled per image and used for all sides.
  - If a tuple of four entries, then the entries represent top, right, bottom, left. Each entry may be a single int (always crop/pad by exactly that value), a tuple of two int s a and b (crop/pad by an amount within [a, b]), a list of int s (crop/pad by a random

value that is contained in the `list`) or a `StochasticParameter` (sample the amount to crop/pad from that parameter).

- **percent** (*None or number or imgaug.parameters.StochasticParameter or tuple, optional*) – The number of pixels to crop (negative values) or pad (positive values) on each side of the image given as a *fraction* of the image height/width. E.g. if this is set to `-0.1`, the augmenter will always crop away `10%` of the image's height at both the top and the bottom (both `10%` each), as well as `10%` of the width at the right and left. Expected value range is `(-1.0, inf)`. Either this or the parameter *px* may be set, not both at the same time.
  - If `None`, then fraction-based cropping/padding will not be used.
  - If `number`, then that fraction will always be cropped/padded.
  - If `StochasticParameter`, then that parameter will be used for each image. Four samples will be drawn per image (top, right, bottom, left). If however *sample_independently* is set to `False`, only one value will be sampled per image and used for all sides.
  - If a `tuple` of two `float` s with values `a` and `b`, then each side will be cropped/padded by a random fraction sampled uniformly per image and side from the interval `[a, b]`. If however *sample_independently* is set to `False`, only one value will be sampled per image and used for all sides.
  - If a `tuple` of four entries, then the entries represent top, right, bottom, left. Each entry may be a single `float` (always crop/pad by exactly that percent value), a `tuple` of two `float` s `a` and `b` (crop/pad by a fraction from `[a, b]`), a `list` of `float` s (crop/pad by a random value that is contained in the list) or a `StochasticParameter` (sample the percentage to crop/pad from that parameter).
- **pad_mode** (*imgaug.ALL or str or list of str or imgaug.parameters.StochasticParameter, optional*) – Padding mode to use. The available modes match the numpy padding modes, i.e. `constant`, `edge`, `linear_ramp`, `maximum`, `median`, `minimum`, `reflect`, `symmetric`, `wrap`. The modes `constant` and `linear_ramp` use extra values, which are provided by `pad_cval` when necessary. See *pad()* for more details.
  - If `imgaug.ALL`, then a random mode from all available modes will be sampled per image.
  - If a `str`, it will be used as the pad mode for all images.
  - If a `list` of `str`, a random one of these will be sampled per image and used as the mode.
  - If `StochasticParameter`, a random mode will be sampled from this parameter per image.
- **pad_cval** (*number or tuple of number list of number or imgaug.parameters.StochasticParameter, optional*) – The constant value to use if the pad mode is `constant` or the end value to use if the mode is `linear_ramp`. See *pad()* for more details.
  - If `number`, then that value will be used.
  - If a `tuple` of two `number` s and at least one of them is a `float`, then a random number will be uniformly sampled per image from the continuous interval `[a, b]` and used as the value. If both `number` s are `int` s, the interval is discrete.
  - If a `list` of `number`, then a random value will be chosen from the elements of the `list` and used as the value.

  - If `StochasticParameter`, a random value will be sampled from that parameter per image.

- **keep_size** (*bool, optional*) – After cropping and padding, the result image will usually have a different height/width compared to the original input image. If this parameter is set to `True`, then the cropped/padded image will be resized to the input image's size, i.e. the augmenter's output shape is always identical to the input shape.
- **sample_independently** (*bool, optional*) – If `False` *and* the values for *px*/*percent* result in exactly *one* probability distribution for all image sides, only one single value will be sampled from that probability distribution and used for all sides. I.e. the crop/pad amount then is the same for all sides. If `True`, four values will be sampled independently, one per side.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.CropAndPad(px=(-10, 0))
```

Crop each side by a random pixel value sampled uniformly per image and side from the discrete interval `[-10..0]`.

```
>>> aug = iaa.CropAndPad(px=(0, 10))
```

Pad each side by a random pixel value sampled uniformly per image and side from the discrete interval `[0..10]`. The padding happens by zero-padding, i.e. it adds black pixels (default setting).

```
>>> aug = iaa.CropAndPad(px=(0, 10), pad_mode="edge")
```

Pad each side by a random pixel value sampled uniformly per image and side from the discrete interval `[0..10]`. The padding uses the `edge` mode from numpy's pad function, i.e. the pixel colors around the image sides are repeated.

```
>>> aug = iaa.CropAndPad(px=(0, 10), pad_mode=["constant", "edge"])
```

Similar to the previous example, but uses zero-padding (`constant`) for half of the images and `edge` padding for the other half.

```
>>> aug = iaa.CropAndPad(px=(0, 10), pad_mode=ia.ALL, pad_cval=(0, 255))
```

Similar to the previous example, but uses any available padding mode. In case the padding mode ends up being `constant` or `linear_ramp`, and random intensity is uniformly sampled (once per image) from the discrete interval `[0..255]` and used as the intensity of the new pixels.

```
>>> aug = iaa.CropAndPad(px=(0, 10), sample_independently=False)
```

Pad each side by a random pixel value sampled uniformly once per image from the discrete interval `[0..10]`. Each sampled value is used for *all* sides of the corresponding image.

```
>>> aug = iaa.CropAndPad(px=(0, 10), keep_size=False)
```

Pad each side by a random pixel value sampled uniformly per image and side from the discrete interval `[0..10]`. Afterwards, do **not** resize the padded image back to the input image's size. This will increase the image's height and width by a maximum of `20` pixels.

```
>>> aug = iaa.CropAndPad(px=((0, 10), (0, 5), (0, 10), (0, 5)))
```

Pad the top and bottom by a random pixel value sampled uniformly from the discrete interval `[0..10]`. Pad the left and right analogously by a random value sampled from `[0..5]`. Each value is always sampled independently.

```
>>> aug = iaa.CropAndPad(percent=(0, 0.1))
```

Pad each side by a random fraction sampled uniformly from the continuous interval `[0.0, 0.10]`. The fraction is sampled once per image and side. E.g. a sampled fraction of `0.1` for the top side would pad by `0.1*H`, where `H` is the height of the input image.

```
>>> aug = iaa.CropAndPad(
>>>     percent=([0.05, 0.1], [0.05, 0.1], [0.05, 0.1], [0.05, 0.1]))
```

Pads each side by either `5%` or `10%`. The values are sampled once per side and image.

```
>>> aug = iaa.CropAndPad(px=(-10, 10))
```

Sample uniformly per image and side a value `v` from the discrete range `[-10..10]`. Then either crop (negative sample) or pad (positive sample) the side by `v` pixels.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.

Continued on next page

Table 236 – continued from previous page

augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.size.**CropToAspectRatio**(*aspect_ratio*, *position='uniform'*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.size.CropToFixedSize*

Crop images until their width/height matches an aspect ratio.

This augmenter removes either rows or columns until the image reaches the desired aspect ratio given in `width / height`. The cropping operation is stopped once the desired aspect ratio is reached or the image side to crop reaches a size of `1`. If any side of the image starts with a size of `0`, the image will not be changed.

Added in 0.4.0.

Supported dtypes:

See *CropToFixedSize*.

Parameters

- **aspect_ratio** (*number*) – The desired aspect ratio, given as `width/height`. E.g. a ratio of `2.0` denotes an image that is twice as wide as it is high.
- **position** (*{'uniform', 'normal', 'center', 'left-top', 'left-center', 'left-bottom', 'center-top', 'center-center', 'center-bottom', 'right-top', 'right-center', 'right-bottom'} or tuple of float or StochasticParameter or tuple of StochasticParameter, optional*) – See `CropToFixedSize.__init__()`.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.CropToAspectRatio(2.0)
```

Create an augmenter that crops each image until its aspect ratio is as close as possible to `2.0` (i.e. two times as many pixels along the x-axis than the y-axis). The rows to be cropped will be spread *randomly* over the top and bottom sides (analogous for the left/right sides).

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.

Continued on next page

Table 237 – continued from previous page

augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.size.**CropToFixedSize**(*width*, *height*, *position='uniform'*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Crop images down to a predefined maximum width and/or height.

If images are already at the maximum width/height or are smaller, they will not be cropped. Note that this also means that images will not be padded if they are below the required width/height.

The augmenter randomly decides per image how to distribute the required cropping amounts over the image axis. E.g. if 2px have to be cropped on the left or right to reach the required width, the augmenter will sometimes

remove 2px from the left and 0px from the right, sometimes remove 2px from the right and 0px from the left and sometimes remove 1px from both sides. Set *position* to `center` to prevent that.

Supported dtypes:

- `uint8`: yes; fully tested
- `uint16`: yes; tested
- `uint32`: yes; tested
- `uint64`: yes; tested
- `int8`: yes; tested
- `int16`: yes; tested
- `int32`: yes; tested
- `int64`: yes; tested
- `float16`: yes; tested
- `float32`: yes; tested
- `float64`: yes; tested
- `float128`: yes; tested
- `bool`: yes; tested

Parameters

- **width** (*int or None*) – Crop images down to this maximum width. If `None`, image widths will not be altered.
- **height** (*int or None*) – Crop images down to this maximum height. If `None`, image heights will not be altered.
- **position** (*{'uniform', 'normal', 'center', 'left-top', 'left-center', 'left-bottom', 'center-top', 'center-center', 'center-bottom', 'right-top', 'right-center', 'right-bottom'} or tuple of float or StochasticParameter or tuple of StochasticParameter, optional*) – Sets the center point of the cropping, which determines how the required cropping amounts are distributed to each side. For a `tuple (a, b)`, both `a` and `b` are expected to be in range `[0.0, 1.0]` and describe the fraction of cropping applied to the left/right (low/high values for `a`) and the fraction of cropping applied to the top/bottom (low/high values for `b`). A cropping position at `(0.5, 0.5)` would be the center of the image and distribute the cropping equally over all sides. A cropping position at `(1.0, 0.0)` would be the right-top and would apply 100% of the required cropping to the right and top sides of the image.
  - If string `uniform` then the share of cropping is randomly and uniformly distributed over each side. Equivalent to `(Uniform(0.0, 1.0), Uniform(0.0, 1.0))`.
  - If string `normal` then the share of cropping is distributed based on a normal distribution, leading to a focus on the center of the images. Equivalent to `(Clip(Normal(0.5, 0.45/2), 0, 1), Clip(Normal(0.5, 0.45/2), 0, 1))`.
  - If string `center` then center point of the cropping is identical to the image center. Equivalent to `(0.5, 0.5)`.
  - If a string matching regex `^(left|center|right)-(top|center|bottom)$`, e.g. `left-top` or `center-bottom` then sets the center point of the cropping to the X-Y position matching that description.

  - If a tuple of float, then expected to have exactly two entries between `0.0` and `1.0`, which will always be used as the combination the position matching (x, y) form.
  - If a `StochasticParameter`, then that parameter will be queried once per call to `augment_*()` to get `Nx2` center positions in `(x, y)` form (with `N` the number of images).
  - If a `tuple` of `StochasticParameter`, then expected to have exactly two entries that will both be queried per call to `augment_*()`, each for `(N,)` values, to get the center positions. First parameter is used for `x` coordinates, second for `y` coordinates.

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.CropToFixedSize(width=100, height=100)
```

For image sides larger than `100` pixels, crop to `100` pixels. Do nothing for the other sides. The cropping amounts are randomly (and uniformly) distributed over the sides of the image.

```
>>> aug = iaa.CropToFixedSize(width=100, height=100, position="center")
```

For sides larger than `100` pixels, crop to `100` pixels. Do nothing for the other sides. The cropping amounts are always equally distributed over the left/right sides of the image (and analogously for top/bottom).

```
>>> aug = iaa.Sequential([
>>>     iaa.PadToFixedSize(width=100, height=100),
>>>     iaa.CropToFixedSize(width=100, height=100)
>>> ])
```

Pad images smaller than `100x100` until they reach `100x100`. Analogously, crop images larger than `100x100` until they reach `100x100`. The output images therefore have a fixed size of `100x100`.

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.

Continued on next page

Table 238 – continued from previous page

augment_batches(self, batches[, hooks, ...])	Augment multiple batches.
augment_bounding_boxes(self, ...[, parents, ...])	Augment a batch of bounding boxes.
augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.size.**CropToMultiplesOf**(*width_multiple*, *height_multiple*, *position='uniform'*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.size.CropToFixedSize*

Crop images down until their height/width is a multiple of a value.

Note: For a given axis size `A` and multiple `M`, if `A` is in the interval `[0 .. M]`, the axis will not be changed. As a result, this augmenter can still produce axis sizes that are not multiples of the given values.

Added in 0.4.0.

Supported dtypes:

See *CropToFixedSize*.

Parameters

- **width_multiple** (*int or None*) – Multiple for the width. Images will be cropped down until their width is a multiple of this value. If `None`, image widths will not be altered.
- **height_multiple** (*int or None*) – Multiple for the height. Images will be cropped down until their height is a multiple of this value. If `None`, image heights will not be altered.
- **position** (*{'uniform', 'normal', 'center', 'left-top', 'left-center', 'left-bottom', 'center-top', 'center-center', 'center-bottom', 'right-top', 'right-center', 'right-bottom'} or tuple of float or StochasticParameter or tuple of StochasticParameter, optional*) – See `CropToFixedSize.__init__()`.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.CropToMultiplesOf(height_multiple=10, width_multiple=6)
```

Create an augmenter that crops images to multiples of `10` along the y-axis (i.e. 10, 20, 30, ...) and to multiples of `6` along the x-axis (i.e. 6, 12, 18, ...). The rows to be cropped will be spread *randomly* over the top and bottom sides (analogous for the left/right sides).

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`localize_random_state_`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`pool`(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
`remove_augmenters`(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
`remove_augmenters_`(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
`remove_augmenters_inplace`(self, func[, parents])	**Deprecated**.
`reseed`(self[, random_state, deterministic_too])	**Deprecated**.
`seed_`(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.

Continued on next page

Table 239 – continued from previous page

show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.size.**CropToPowersOf**(*width_base*, *height_base*, *position='uniform'*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.size.CropToFixedSize*

Crop images until their height/width is a power of a base.

This augmenter removes pixels from an axis with size `S` leading to the new size `S'` until `S' = B^E` is fulfilled, where `B` is a provided base (e.g. `2`) and `E` is an exponent from the discrete interval `[1 .. inf)`.

Note: This augmenter does nothing for axes with size less than `B^1 = B`. If you have images with `S < B^1`, it is recommended to combine this augmenter with a padding augmenter that pads each axis up to `B`.

Added in 0.4.0.

Supported dtypes:

See *CropToFixedSize*.

Parameters

- **width_base** (*int or None*) – Base for the width. Images will be cropped down until their width fulfills `width' = width_base ^ E` with `E` being any natural number. If `None`, image widths will not be altered.
- **height_base** (*int or None*) – Base for the height. Images will be cropped down until their height fulfills `height' = height_base ^ E` with `E` being any natural number. If `None`, image heights will not be altered.
- **position** (*{'uniform', 'normal', 'center', 'left-top', 'left-center', 'left-bottom', 'center-top', 'center-center', 'center-bottom', 'right-top', 'right-center', 'right-bottom'} or tuple of float or StochasticParameter or tuple of StochasticParameter, optional*) – See `CropToFixedSize.__init__()`.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.CropToPowersOf(height_base=3, width_base=2)
```

Create an augmenter that crops each image down to powers of 3 along the y-axis (i.e. 3, 9, 27, ...) and powers of 2 along the x-axis (i.e. 2, 4, 8, 16, ...). The rows to be cropped will be spread *randomly* over the top and bottom sides (analogous for the left/right sides).

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`localize_random_state_`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.

Continued on next page

Table 240 – continued from previous page

pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.size.**CropToSquare**(*position='uniform'*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.size.CropToAspectRatio*

Crop images until their width and height are identical.

This is identical to *CropToAspectRatio* with `aspect_ratio=1.0`.

Images with axis sizes of `0` will not be altered.

Added in 0.4.0.

Supported dtypes:

See *CropToFixedSize*.

Parameters

- **position** (*{'uniform', 'normal', 'center', 'left-top', 'left-center', 'left-bottom', 'center-top', 'center-center', 'center-bottom', 'right-top', 'right-center', 'right-bottom'} or tuple of float or StochasticParameter or tuple of StochasticParameter, optional*) – See `CropToFixedSize.__init__()`.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.CropToSquare()
```

Create an augmenter that crops each image until its square, i.e. height and width match. The rows to be cropped will be spread *randomly* over the top and bottom sides (analogous for the left/right sides).

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.
`get_parameters`(self)	See *get_parameters()*.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`localize_random_state_`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`pool`(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.

Continued on next page

Table 241 – continued from previous page

remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.size.**KeepSizeByResize**(*children*, *interpolation='cubic'*, *interpolation_heatmaps='SAME_AS_IMAGES'*, *interpolation_segmaps='nearest'*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Resize images back to their input sizes after applying child augmenters.

Combining this with e.g. a cropping augmenter as the child will lead to images being resized back to the input size after the crop operation was applied. Some augmenters have a `keep_size` argument that achieves the same goal (if set to `True`), though this augmenter offers control over the interpolation mode and which augmentables to resize (images, heatmaps, segmentation maps).

Supported dtypes:

See *imresize_many_images()*.

Parameters

- **children** (*Augmenter or list of imgaug.augmenters.meta.Augmenter or None, optional*) – One or more augmenters to apply to images. These augmenters may change the image size.
- **interpolation** (*KeepSizeByResize.NO_RESIZE or {'nearest', 'linear', 'area', 'cubic'} or {cv2.INTER_NEAREST, cv2.INTER_LINEAR, cv2.INTER_AREA, cv2.INTER_CUBIC} or list of str or list of int or StochasticParameter, optional*) – The interpolation mode to use when resizing images. Can take any value that *imresize_single_image()* accepts, e.g. `cubic`.
  - If this is `KeepSizeByResize.NO_RESIZE` then images will not be resized.
  - If this is a single `str`, it is expected to have one of the following values: `nearest`, `linear`, `area`, `cubic`.
  - If this is a single integer, it is expected to have a value identical to one of: `cv2.INTER_NEAREST`, `cv2.INTER_LINEAR`, `cv2.INTER_AREA`, `cv2.INTER_CUBIC`.
  - If this is a `list` of `str` or `int`, it is expected that each `str`/`int` is one of the above mentioned valid ones. A random one of these values will be sampled per image.
  - If this is a `StochasticParameter`, it will be queried once per call to `_augment_images()` and must return `N` `str` s or `int` s (matching the above mentioned ones) for `N` images.

- **interpolation_heatmaps** (*KeepSizeByResize.SAME_AS_IMAGES or KeepSizeByResize.NO_RESIZE or {'nearest', 'linear', 'area', 'cubic'} or {cv2.INTER_NEAREST, cv2.INTER_LINEAR, cv2.INTER_AREA, cv2.INTER_CUBIC} or list of str or list of int or StochasticParameter, optional*) – The interpolation mode to use when resizing heatmaps. Meaning and valid values are similar to *interpolation*. This parameter may also take the value `KeepSizeByResize.SAME_AS_IMAGES`, which will lead to copying the interpolation modes used for the corresponding images. The value may also be returned on a per-image basis if *interpolation_heatmaps* is provided as a `StochasticParameter` or may be one possible value if it is provided as a `list` of `str`.
- **interpolation_segmaps** (*KeepSizeByResize.SAME_AS_IMAGES or KeepSizeByResize.NO_RESIZE or {'nearest', 'linear', 'area', 'cubic'} or {cv2.INTER_NEAREST, cv2.INTER_LINEAR, cv2.INTER_AREA, cv2.INTER_CUBIC} or list of str or list of int or StochasticParameter, optional*) – The interpolation mode to use when resizing segmentation maps. Similar to *interpolation_heatmaps*. **Note**: For segmentation maps, only `NO_RESIZE` or nearest neighbour interpolation (i.e. `nearest`) make sense in the vast majority of all cases.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.KeepSizeByResize(
>>>     iaa.Crop((20, 40), keep_size=False)
>>> )
```

Apply random cropping to input images, then resize them back to their original input sizes. The resizing is done using this augmenter instead of the corresponding internal resizing operation in `Crop`.

```
>>> aug = iaa.KeepSizeByResize(
>>>     iaa.Crop((20, 40), keep_size=False),
>>>     interpolation="nearest"
>>> )
```

Same as in the previous example, but images are now always resized using nearest neighbour interpolation.

```
>>> aug = iaa.KeepSizeByResize(
>>>     iaa.Crop((20, 40), keep_size=False),
>>>     interpolation=["nearest", "cubic"],
>>>     interpolation_heatmaps=iaa.KeepSizeByResize.SAME_AS_IMAGES,
>>>     interpolation_segmaps=iaa.KeepSizeByResize.NO_RESIZE
>>> )
```

Similar to the previous example, but images are now sometimes resized using linear interpolation and sometimes using nearest neighbour interpolation. Heatmaps are resized using the same interpolation as was used for the corresponding image. Segmentation maps are not resized and will therefore remain at their size after cropping.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	See *get_children_lists()*.
get_parameters(self)	See *get_parameters()*.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`localize_random_state_`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`pool`(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
`remove_augmenters`(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
`remove_augmenters_`(self, func[, parents])	Remove in-place children of this augmenter that match a condition.

Continued on next page

Table 242 – continued from previous page

remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

NO_RESIZE = 'NO_RESIZE'

SAME_AS_IMAGES = 'SAME_AS_IMAGES'

get_children_lists(*self*)
See *get_children_lists()*.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.size.**Pad**(*px=None*, *percent=None*, *pad_mode='constant'*, *pad_cval=0*, *keep_size=True*, *sample_independently=True*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.size.CropAndPad*

Pad images, i.e. adds columns/rows of pixels to them.

Supported dtypes:

See *CropAndPad*.

Parameters

- **px** (*None or int or imgaug.parameters.StochasticParameter or tuple, optional*) – The number of pixels to pad on each side of the image. Expected value range is `[0, inf)`. Either this or the parameter *percent* may be set, not both at the same time.
  - If `None`, then pixel-based padding will not be used.
  - If `int`, then that exact number of pixels will always be padded.
  - If `StochasticParameter`, then that parameter will be used for each image. Four samples will be drawn per image (top, right, bottom, left), unless *sample_independently* is set to `False`, as then only one value will be sampled per image and used for all sides.
  - If a `tuple` of two `int` s with values `a` and `b`, then each side will be padded by a random amount sampled uniformly per image and side from the inteval `[a, b]`. If however *sample_independently* is set to `False`, only one value will be sampled per image and used for all sides.
  - If a `tuple` of four entries, then the entries represent top, right, bottom, left. Each entry may be a single `int` (always pad by exactly that value), a `tuple` of two `int` s `a` and `b` (pad by an amount within `[a, b]`), a `list` of `int` s (pad by a random value that is contained in the `list`) or a `StochasticParameter` (sample the amount to pad from that parameter).
- **percent** (*None or int or float or imgaug.parameters.StochasticParameter or tuple, optional*) – The number of pixels to pad on each side of the image given as a *fraction* of the image height/width. E.g. if this is set to `0.1`, the augmenter will always pad `10%` of the image's height at both the top and the bottom (both `10%` each), as well as `10%` of the width at the

right and left. Expected value range is `[0.0, inf)`. Either this or the parameter *px* may be set, not both at the same time.

- If `None`, then fraction-based padding will not be used.
- If `number`, then that fraction will always be padded.
- If `StochasticParameter`, then that parameter will be used for each image. Four samples will be drawn per image (top, right, bottom, left). If however *sample_independently* is set to `False`, only one value will be sampled per image and used for all sides.
- If a `tuple` of two `float` s with values `a` and `b`, then each side will be padded by a random fraction sampled uniformly per image and side from the interval `[a, b]`. If however *sample_independently* is set to `False`, only one value will be sampled per image and used for all sides.
- If a `tuple` of four entries, then the entries represent top, right, bottom, left. Each entry may be a single `float` (always pad by exactly that fraction), a `tuple` of two `float` s `a` and `b` (pad by a fraction from `[a, b]`), a `list` of `float` s (pad by a random value that is contained in the list) or a `StochasticParameter` (sample the percentage to pad from that parameter).

- **pad_mode** (*imgaug.ALL or str or list of str or imgaug.parameters.StochasticParameter, optional*) – Padding mode to use. The available modes match the numpy padding modes, i.e. `constant`, `edge`, `linear_ramp`, `maximum`, `median`, `minimum`, `reflect`, `symmetric`, `wrap`. The modes `constant` and `linear_ramp` use extra values, which are provided by `pad_cval` when necessary. See *pad()* for more details.
  - If `imgaug.ALL`, then a random mode from all available modes will be sampled per image.
  - If a `str`, it will be used as the pad mode for all images.
  - If a `list` of `str`, a random one of these will be sampled per image and used as the mode.
  - If `StochasticParameter`, a random mode will be sampled from this parameter per image.
- **pad_cval** (*number or tuple of number list of number or imgaug.parameters.StochasticParameter, optional*) – The constant value to use if the pad mode is `constant` or the end value to use if the mode is `linear_ramp`. See *pad()* for more details.
  - If `number`, then that value will be used.
  - If a `tuple` of two `number` s and at least one of them is a `float`, then a random number will be uniformly sampled per image from the continuous interval `[a, b]` and used as the value. If both `number` s are `int` s, the interval is discrete.
  - If a `list` of `number`, then a random value will be chosen from the elements of the `list` and used as the value.
  - If `StochasticParameter`, a random value will be sampled from that parameter per image.
- **keep_size** (*bool, optional*) – After padding, the result image will usually have a different height/width compared to the original input image. If this parameter is set to `True`, then the padded image will be resized to the input image's size, i.e. the augmenter's output shape is always identical to the input shape.

- **sample_independently** (*bool, optional*) – If `False` *and* the values for *px*/*percent* result in exactly *one* probability distribution for all image sides, only one single value will be sampled from that probability distribution and used for all sides. I.e. the pad amount then is the same for all sides. If `True`, four values will be sampled independently, one per side.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.Pad(px=(0, 10))
```

Pad each side by a random pixel value sampled uniformly per image and side from the discrete interval `[0..10]`. The padding happens by zero-padding, i.e. it adds black pixels (default setting).

```
>>> aug = iaa.Pad(px=(0, 10), pad_mode="edge")
```

Pad each side by a random pixel value sampled uniformly per image and side from the discrete interval `[0..10]`. The padding uses the `edge` mode from numpy's pad function, i.e. the pixel colors around the image sides are repeated.

```
>>> aug = iaa.Pad(px=(0, 10), pad_mode=["constant", "edge"])
```

Similar to the previous example, but uses zero-padding (`constant`) for half of the images and `edge` padding for the other half.

```
>>> aug = iaa.Pad(px=(0, 10), pad_mode=ia.ALL, pad_cval=(0, 255))
```

Similar to the previous example, but uses any available padding mode. In case the padding mode ends up being `constant` or `linear_ramp`, and random intensity is uniformly sampled (once per image) from the discrete interval `[0..255]` and used as the intensity of the new pixels.

```
>>> aug = iaa.Pad(px=(0, 10), sample_independently=False)
```

Pad each side by a random pixel value sampled uniformly once per image from the discrete interval `[0..10]`. Each sampled value is used for *all* sides of the corresponding image.

```
>>> aug = iaa.Pad(px=(0, 10), keep_size=False)
```

Pad each side by a random pixel value sampled uniformly per image and side from the discrete interval `[0..10]`. Afterwards, do **not** resize the padded image back to the input image's size. This will increase the image's height and width by a maximum of `20` pixels.

```
>>> aug = iaa.Pad(px=((0, 10), (0, 5), (0, 10), (0, 5)))
```

Pad the top and bottom by a random pixel value sampled uniformly from the discrete interval `[0..10]`. Pad the left and right analogously by a random value sampled from `[0..5]`. Each value is always sampled independently.

```
>>> aug = iaa.Pad(percent=(0, 0.1))
```

Pad each side by a random fraction sampled uniformly from the continuous interval `[0.0, 0.10]`. The fraction is sampled once per image and side. E.g. a sampled fraction of `0.1` for the top side would pad by `0.1*H`, where `H` is the height of the input image.

```
>>> aug = iaa.Pad(
>>>     percent=([0.05, 0.1], [0.05, 0.1], [0.05, 0.1], [0.05, 0.1]))
```

Pads each side by either `5%` or `10%`. The values are sampled once per side and image.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.

Continued on next page

Table 243 – continued from previous page

get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.size.**PadToAspectRatio**(*aspect_ratio*, *pad_mode='constant'*, *pad_cval=0*, *position='uniform'*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.size.PadToFixedSize*

Pad images until their width/height matches an aspect ratio.

This augmenter adds either rows or columns until the image reaches the desired aspect ratio given in `width / height`.

Added in 0.4.0.

Supported dtypes:

See *PadToFixedSize*.

Parameters

- **aspect_ratio** (*number*) – The desired aspect ratio, given as `width/height`. E.g. a ratio of `2.0` denotes an image that is twice as wide as it is high.
- **position** (*{'uniform', 'normal', 'center', 'left-top', 'left-center', 'left-bottom', 'center-top', 'center-center', 'center-bottom', 'right-top', 'right-center', 'right-bottom'} or tuple of float or StochasticParameter or tuple of StochasticParameter, optional*) – See `PadToFixedSize.__init__()`.
- **pad_mode** (*imgaug.ALL or str or list of str or imgaug.parameters.StochasticParameter, optional*) – See `__init__()`.
- **pad_cval** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – See `__init__()`.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.

- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.PadToAspectRatio(2.0)
```

Create an augmenter that pads each image until its aspect ratio is as close as possible to `2.0` (i.e. two times as many pixels along the x-axis than the y-axis). The rows to be padded will be spread *randomly* over the top and bottom sides (analogous for the left/right sides).

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.

Continued on next page

Table 244 – continued from previous page

find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.size.**PadToFixedSize**(*width*, *height*, *pad_mode='constant'*, *pad_cval=0*, *position='uniform'*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Pad images to a predefined minimum width and/or height.

If images are already at the minimum width/height or are larger, they will not be padded. Note that this also means that images will not be cropped if they exceed the required width/height.

The augmenter randomly decides per image how to distribute the required padding amounts over the image axis. E.g. if 2px have to be padded on the left or right to reach the required width, the augmenter will sometimes add 2px to the left and 0px to the right, sometimes add 2px to the right and 0px to the left and sometimes add 1px to both sides. Set *position* to `center` to prevent that.

Supported dtypes:

See *pad()*.

Parameters

- **width** (*int or None*) – Pad images up to this minimum width. If `None`, image widths will not be altered.
- **height** (*int or None*) – Pad images up to this minimum height. If `None`, image heights will not be altered.

- **pad_mode** (*imgaug.ALL or str or list of str or imgaug.parameters.StochasticParameter, optional*) – See `__init__()`.
- **pad_cval** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – See `__init__()`.
- **position** (*{'uniform', 'normal', 'center', 'left-top', 'left-center', 'left-bottom', 'center-top', 'center-center', 'center-bottom', 'right-top', 'right-center', 'right-bottom'} or tuple of float or StochasticParameter or tuple of StochasticParameter, optional*) – Sets the center point of the padding, which determines how the required padding amounts are distributed to each side. For a `tuple (a, b)`, both `a` and `b` are expected to be in range `[0.0, 1.0]` and describe the fraction of padding applied to the left/right (low/high values for `a`) and the fraction of padding applied to the top/bottom (low/high values for `b`). A padding position at `(0.5, 0.5)` would be the center of the image and distribute the padding equally to all sides. A padding position at `(0.0, 1.0)` would be the left-bottom and would apply 100% of the required padding to the bottom and left sides of the image so that the bottom left corner becomes more and more the new image center (depending on how much is padded).
  - If string `uniform` then the share of padding is randomly and uniformly distributed over each side. Equivalent to `(Uniform(0.0, 1.0), Uniform(0.0, 1.0))`.
  - If string `normal` then the share of padding is distributed based on a normal distribution, leading to a focus on the center of the images. Equivalent to `(Clip(Normal(0.5, 0.45/2), 0, 1), Clip(Normal(0.5, 0.45/2), 0, 1))`.
  - If string `center` then center point of the padding is identical to the image center. Equivalent to `(0.5, 0.5)`.
  - If a string matching regex `^(left|center|right)-(top|center|bottom)$`, e.g. `left-top` or `center-bottom` then sets the center point of the padding to the X-Y position matching that description.
  - If a tuple of float, then expected to have exactly two entries between `0.0` and `1.0`, which will always be used as the combination the position matching (x, y) form.
  - If a `StochasticParameter`, then that parameter will be queried once per call to `augment_*()` to get `Nx2` center positions in `(x, y)` form (with `N` the number of images).
  - If a `tuple` of `StochasticParameter`, then expected to have exactly two entries that will both be queried per call to `augment_*()`, each for `(N,)` values, to get the center positions. First parameter is used for `x` coordinates, second for `y` coordinates.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.PadToFixedSize(width=100, height=100)
```

For image sides smaller than `100` pixels, pad to `100` pixels. Do nothing for the other edges. The padding is randomly (uniformly) distributed over the sides, so that e.g. sometimes most of the required padding is applied to the left, sometimes to the right (analogous top/bottom).

```
>>> aug = iaa.PadToFixedSize(width=100, height=100, position="center")
```

For image sides smaller than `100` pixels, pad to `100` pixels. Do nothing for the other image sides. The padding is always equally distributed over the left/right and top/bottom sides.

```
>>> aug = iaa.PadToFixedSize(width=100, height=100, pad_mode=ia.ALL)
```

For image sides smaller than `100` pixels, pad to `100` pixels and use any possible padding mode for that. Do nothing for the other image sides. The padding is always equally distributed over the left/right and top/bottom sides.

```
>>> aug = iaa.Sequential([
>>>     iaa.PadToFixedSize(width=100, height=100),
>>>     iaa.CropToFixedSize(width=100, height=100)
>>> ])
```

Pad images smaller than `100x100` until they reach `100x100`. Analogously, crop images larger than `100x100` until they reach `100x100`. The output images therefore have a fixed size of `100x100`.

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.

Continued on next page

Table 245 – continued from previous page

copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.size.**PadToMultiplesOf**(*width_multiple*, *height_multiple*, *pad_mode='constant'*, *pad_cval=0*, *position='uniform'*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.size.PadToFixedSize*

Pad images until their height/width is a multiple of a value.

Added in 0.4.0.

Supported dtypes:

See *PadToFixedSize*.

Parameters

- **width_multiple** (*int or None*) – Multiple for the width. Images will be padded until their width is a multiple of this value. If None, image widths will not be altered.
- **height_multiple** (*int or None*) – Multiple for the height. Images will be padded until their height is a multiple of this value. If None, image heights will not be altered.

- **pad_mode** (*imgaug.ALL or str or list of str or imgaug.parameters.StochasticParameter, optional*) – See `__init__()`.
- **pad_cval** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – See `__init__()`.
- **position** (*{'uniform', 'normal', 'center', 'left-top', 'left-center', 'left-bottom', 'center-top', 'center-center', 'center-bottom', 'right-top', 'right-center', 'right-bottom'} or tuple of float or StochasticParameter or tuple of StochasticParameter, optional*) – See `PadToFixedSize.__init__()`.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.PadToMultiplesOf(height_multiple=10, width_multiple=6)
```

Create an augmenter that pads images to multiples of `10` along the y-axis (i.e. 10, 20, 30, ...) and to multiples of `6` along the x-axis (i.e. 6, 12, 18, ...). The rows to be padded will be spread *randomly* over the top and bottom sides (analogous for the left/right sides).

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.

Continued on next page

Table 246 – continued from previous page

augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.size.**PadToPowersOf**(*width_base*, *height_base*, *pad_mode='constant'*, *pad_cval=0*, *position='uniform'*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.size.PadToFixedSize*

Pad images until their height/width is a power of a base.

This augmenter adds pixels to an axis with size `S` leading to the new size `S'` until `S' = B^E` is fulfilled, where `B` is a provided base (e.g. `2`) and `E` is an exponent from the discrete interval `[1 .. inf)`.

Added in 0.4.0.

Supported dtypes:

See *PadToFixedSize*.

Parameters

- **width_base** (*int or None*) – Base for the width. Images will be padded down until their width fulfills `width' = width_base ^ E` with `E` being any natural number. If `None`, image widths will not be altered.
- **height_base** (*int or None*) – Base for the height. Images will be padded until their height fulfills `height' = height_base ^ E` with `E` being any natural number. If `None`, image heights will not be altered.
- **pad_mode** (*imgaug.ALL or str or list of str or imgaug.parameters.StochasticParameter, optional*) – See `__init__()`.
- **pad_cval** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – See `__init__()`.
- **position** (*{'uniform', 'normal', 'center', 'left-top', 'left-center', 'left-bottom', 'center-top', 'center-center', 'center-bottom', 'right-top', 'right-center', 'right-bottom'} or tuple of float or StochasticParameter or tuple of StochasticParameter, optional*) – See `PadToFixedSize.__init__()`.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.PadToPowersOf(height_base=3, width_base=2)
```

Create an augmenter that pads each image to powers of `3` along the y-axis (i.e. 3, 9, 27, ...) and powers of `2` along the x-axis (i.e. 2, 4, 8, 16, ...). The rows to be padded will be spread *randomly* over the top and bottom sides (analogous for the left/right sides).

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.

Continued on next page

Table 247 – continued from previous page

augment_batches(self, batches[, hooks, ...])	Augment multiple batches.
augment_bounding_boxes(self, ...[, parents, ...])	Augment a batch of bounding boxes.
augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.size.**PadToSquare**(*pad_mode='constant'*, *pad_cval=0*, *position='uniform'*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.size.PadToAspectRatio*

Pad images until their height and width are identical.

This augmenter is identical to *PadToAspectRatio* with `aspect_ratio=1.0`.

Added in 0.4.0.

Supported dtypes:

See *PadToFixedSize*.

Parameters

- **position** (*{'uniform', 'normal', 'center', 'left-top', 'left-center', 'left-bottom', 'center-top', 'center-center', 'center-bottom', 'right-top', 'right-center', 'right-bottom'} or tuple of float or StochasticParameter or tuple of StochasticParameter, optional*) – See `PadToFixedSize.__init__()`.
- **pad_mode** (*imgaug.ALL or str or list of str or imgaug.parameters.StochasticParameter, optional*) – See `__init__()`.
- **pad_cval** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – See `__init__()`.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.PadToSquare()
```

Create an augmenter that pads each image until its square, i.e. height and width match. The rows to be padded will be spread *randomly* over the top and bottom sides (analogous for the left/right sides).

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.

Continued on next page

Table 248 – continued from previous page

augment_batch(self, batch[, hooks])	**Deprecated**.
augment_batch_(self, batch[, parents, hooks])	Augment a single batch in-place.
augment_batches(self, batches[, hooks, ...])	Augment multiple batches.
augment_bounding_boxes(self, ...[, parents, ...])	Augment a batch of bounding boxes.
augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.size.**Resize**(*size*, *interpolation='cubic'*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Augmenter that resizes images to specified heights and widths.

Supported dtypes:

See *imresize_many_images()*.

Parameters

- **size** (*'keep' or int or float or tuple of int or tuple of float or list of int or list of float or imgaug.parameters.StochasticParameter or dict*) –

  The new size of the images.

  - If this has the string value `keep`, the original height and width values will be kept (image is not resized).
  - If this is an `int`, this value will always be used as the new height and width of the images.
  - If this is a `float` `v`, then per image the image's height `H` and width `W` will be changed to `H*v` and `W*v`.
  - If this is a `tuple`, it is expected to have two entries `(a, b)`. If at least one of these are `float` s, a value will be sampled from range `[a, b]` and used as the `float` value to resize the image (see above). If both are `int` s, a value will be sampled from the discrete range `[a..b]` and used as the integer value to resize the image (see above).
  - If this is a `list`, a random value from the `list` will be picked to resize the image. All values in the `list` must be `int` s or `float` s (no mixture is possible).
  - If this is a `StochasticParameter`, then this parameter will first be queried once per image. The resulting value will be used for both height and width.
  - If this is a `dict`, it may contain the keys `height` and `width` or the keys `shorter-side` and `longer-side`. Each key may have the same datatypes as above and describes the scaling on x and y-axis or the shorter and longer axis, respectively. Both axis are sampled independently. Additionally, one of the keys may have the value `keep-aspect-ratio`, which means that the respective side of the image will be resized so that the original aspect ratio is kept. This is useful when only resizing one image size by a pixel value (e.g. resize images to a height of `64` pixels and resize the width so that the overall aspect ratio is maintained).

- **interpolation** (*imgaug.ALL or int or str or list of int or list of str or imgaug.parameters.StochasticParameter, optional*) –

  Interpolation to use.

  - If `imgaug.ALL`, then a random interpolation from `nearest`, `linear`, `area` or `cubic` will be picked (per image).
  - If `int`, then this interpolation will always be used. Expected to be any of the following: `cv2.INTER_NEAREST`, `cv2.INTER_LINEAR`, `cv2.INTER_AREA`, `cv2.INTER_CUBIC`
  - If string, then this interpolation will always be used. Expected to be any of the following: `nearest`, `linear`, `area`, `cubic`
  - If `list` of `int` / `str`, then a random one of the values will be picked per image as the interpolation.

  - If a `StochasticParameter`, then this parameter will be queried per image and is expected to return an `int` or `str`.

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.Resize(32)
```

Resize all images to `32x32` pixels.

```
>>> aug = iaa.Resize(0.5)
```

Resize all images to `50` percent of their original size.

```
>>> aug = iaa.Resize((16, 22))
```

Resize all images to a random height and width within the discrete interval `[16..22]` (uniformly sampled per image).

```
>>> aug = iaa.Resize((0.5, 0.75))
```

Resize all any input image so that its height (`H`) and width (`W`) become `H*v` and `W*v`, where `v` is uniformly sampled from the interval `[0.5, 0.75]`.

```
>>> aug = iaa.Resize([16, 32, 64])
```

Resize all images either to `16x16`, `32x32` or `64x64` pixels.

```
>>> aug = iaa.Resize({"height": 32})
```

Resize all images to a height of `32` pixels and keeps the original width.

```
>>> aug = iaa.Resize({"height": 32, "width": 48})
```

Resize all images to a height of `32` pixels and a width of `48`.

```
>>> aug = iaa.Resize({"height": 32, "width": "keep-aspect-ratio"})
```

Resize all images to a height of `32` pixels and resizes the x-axis (width) so that the aspect ratio is maintained.

```
>>> aug = iaa.Resize(
>>>     {"shorter-side": 224, "longer-side": "keep-aspect-ratio"})
```

Resize all images to a height/width of `224` pixels, depending on which axis is shorter and resize the other axis so that the aspect ratio is maintained.

```
>>> aug = iaa.Resize({"height": (0.5, 0.75), "width": [16, 32, 64]})
```

Resize all images to a height of `H*v`, where `H` is the original height and `v` is a random value sampled from the interval `[0.5, 0.75]`. The width/x-axis of each image is resized to either `16` or `32` or `64` pixels.

```
>>> aug = iaa.Resize(32, interpolation=["linear", "cubic"])
```

Resize all images to `32x32` pixels. Randomly use either `linear` or `cubic` interpolation.

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.

Continued on next page

Table 249 – continued from previous page

get_parameters(self)	See get_parameters().
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See get_parameters().

imgaug.augmenters.size.**Scale**(**args*, ***kwargs*)
Deprecated. Use Resize instead. Resize has the exactly same interface as Scale.

Augmenter that resizes images to specified heights and widths.

imgaug.augmenters.size.**compute_croppings_to_reach_aspect_ratio**(*arr*, *aspect_ratio*)
Compute crop amounts required to fulfill an aspect ratio.

"Crop amounts" here denotes the number of pixels that have to be removed from each side to fulfill the desired constraint.

The aspect ratio is given as ratio = width / height. Depending on which dimension is smaller (height or width), only the corresponding sides (top/bottom or left/right) will be cropped.

The axis-wise padding amounts are always distributed equally over the sides of the respective axis (i.e. left and right, top and bottom). For odd pixel amounts, one pixel will be left over after the equal distribution and could be added to either side of the axis. This function will always add such a left over pixel to the bottom (y-axis) or right (x-axis) side.

If an aspect ratio cannot be reached exactly, this function will return rather one pixel too few than one pixel too many.

Added in 0.4.0.

Parameters

- **arr** (*(H,W) ndarray or (H,W,C) ndarray or tuple of int*) – Image-like array or shape tuple for which to compute crop amounts.
- **aspect_ratio** (*float*) – Target aspect ratio, given as width/height. E.g. 2.0 denotes the image having twice as much width as height.

Returns Required cropping amounts to reach the target aspect ratio, given as a tuple of the form (top, right, bottom, left).

Return type tuple of int

imgaug.augmenters.size.**compute_croppings_to_reach_multiples_of**(*arr*, *height_multiple*, *width_multiple*)

Compute croppings to reach multiples of given heights/widths.

See *compute_paddings_for_aspect_ratio()* for an explanation of how the required cropping amounts are distributed per image axis.

Added in 0.4.0.

Parameters

- **arr** (*(H,W) ndarray or (H,W,C) ndarray or tuple of int*) – Image-like array or shape tuple for which to compute crop amounts.
- **height_multiple** (*None or int*) – The desired multiple of the height. The computed croppings will reflect a crop operation that decreases the y axis size until it is a multiple of this value.
- **width_multiple** (*None or int*) – The desired multiple of the width. The computed croppings amount will reflect a crop operation that decreases the x axis size until it is a multiple of this value.

Returns Required cropping amounts to reach multiples of the provided values, given as a `tuple` of the form `(top, right, bottom, left)`.

Return type tuple of int

imgaug.augmenters.size.**compute_croppings_to_reach_powers_of**(*arr*, *height_base*, *width_base*, *allow_zero_exponent=False*)

Compute croppings to reach powers of given base values.

For given axis size `S`, cropped size `S'` (`S' <= S`) and base `B` this function computes croppings that fulfill `S' = B^E`, where `E` is any exponent from the discrete interval `[0 .. inf)`.

See *compute_paddings_for_aspect_ratio()* for an explanation of how the required cropping amounts are distributed per image axis.

Note: For axes where `S == 0`, this function alwayws returns zeros as croppings.

For axes where `1 <= S < B` see parameter *allow_zero_exponent*.

Added in 0.4.0.

Parameters

- **arr** (*(H,W) ndarray or (H,W,C) ndarray or tuple of int*) – Image-like array or shape tuple for which to compute crop amounts.
- **height_base** (*None or int*) – The desired base of the height.
- **width_base** (*None or int*) – The desired base of the width.
- **allow_zero_exponent** (*bool*) – Whether `E=0` in `S'=B^E` is a valid value. If `True`, axes with size `1 <= S < B` will be cropped to size `B^0=1`. If `False`, axes with sizes `S < B` will not be changed.

Returns Required cropping amounts to fulfill `S' = B^E` given as a `tuple` of the form `(top, right, bottom, left)`.

Return type tuple of int

imgaug.augmenters.size.**compute_paddings_to_reach_aspect_ratio**(*arr*, *aspect_ratio*)

Compute pad amounts required to fulfill an aspect ratio.

"Pad amounts" here denotes the number of pixels that have to be added to each side to fulfill the desired constraint.

The aspect ratio is given as `ratio = width / height`. Depending on which dimension is smaller (height or width), only the corresponding sides (top/bottom or left/right) will be padded.

The axis-wise padding amounts are always distributed equally over the sides of the respective axis (i.e. left and right, top and bottom). For odd pixel amounts, one pixel will be left over after the equal distribution and could be added to either side of the axis. This function will always add such a left over pixel to the bottom (y-axis) or right (x-axis) side.

Added in 0.4.0. (Previously named `imgaug.imgaug.compute_paddings_to_reach_aspect_ratio()`.)

Parameters

- **arr** (*(H,W) ndarray or (H,W,C) ndarray or tuple of int*) – Image-like array or shape tuple for which to compute pad amounts.
- **aspect_ratio** (*float*) – Target aspect ratio, given as width/height. E.g. `2.0` denotes the image having twice as much width as height.

Returns Required padding amounts to reach the target aspect ratio, given as a `tuple` of the form `(top, right, bottom, left)`.

Return type tuple of int

imgaug.augmenters.size.**compute_paddings_to_reach_multiples_of**(*arr*, *height_multiple*, *width_multiple*)

Compute pad amounts until img height/width are multiples of given values.

See *compute_paddings_for_aspect_ratio()* for an explanation of how the required padding amounts are distributed per image axis.

Added in 0.4.0. (Previously named `imgaug.imgaug.compute_paddings_to_reach_multiples_of()`.)

Parameters

- **arr** (*(H,W) ndarray or (H,W,C) ndarray or tuple of int*) – Image-like array or shape tuple for which to compute pad amounts.
- **height_multiple** (*None or int*) – The desired multiple of the height. The computed padding amount will reflect a padding that increases the y axis size until it is a multiple of this value.
- **width_multiple** (*None or int*) – The desired multiple of the width. The computed padding amount will reflect a padding that increases the x axis size until it is a multiple of this value.

Returns Required padding amounts to reach multiples of the provided values, given as a `tuple` of the form `(top, right, bottom, left)`.

Return type tuple of int

imgaug.augmenters.size.**compute_paddings_to_reach_powers_of**(*arr*, *height_base*, *width_base*, *allow_zero_exponent=False*)

Compute paddings to reach powers of given base values.

For given axis size `S`, padded size `S'` (`S' >= S`) and base `B` this function computes paddings that fulfill `S' = B^E`, where `E` is any exponent from the discrete interval `[0 .. inf)`.

See *compute_paddings_for_aspect_ratio()* for an explanation of how the required padding amounts are distributed per image axis.

Added in 0.4.0. (Previously named `imgaug.imgaug.compute_paddings_to_reach_exponents_of()`.)

Parameters

- **arr** (*(H,W) ndarray or (H,W,C) ndarray or tuple of int*) – Image-like array or shape tuple for which to compute pad amounts.
- **height_base** (*None or int*) – The desired base of the height.
- **width_base** (*None or int*) – The desired base of the width.
- **allow_zero_exponent** (*bool, optional*) – Whether `E=0` in `S'=B^E` is a valid value. If `True`, axes with size `0` or `1` will be padded up to size `B^0=1` and axes with size `1 < S <= B` will be padded up to `B^1=B`. If `False`, the minimum output axis size is always at least `B`.

Returns Required padding amounts to fulfill `S' = B^E` given as a `tuple` of the form `(top, right, bottom, left)`.

Return type tuple of int

`imgaug.augmenters.size.`**`pad`**(*arr*, *top=0*, *right=0*, *bottom=0*, *left=0*, *mode='constant'*, *cval=0*)

Pad an image-like array on its top/right/bottom/left side.

This function is a wrapper around `numpy.pad()`.

Added in 0.4.0. (Previously named `imgaug.imgaug.pad()`.)

Supported dtypes:

- `uint8`: yes; fully tested (1)
- `uint16`: yes; fully tested (1)
- `uint32`: yes; fully tested (2) (3)
- `uint64`: yes; fully tested (2) (3)
- `int8`: yes; fully tested (1)
- `int16`: yes; fully tested (1)
- `int32`: yes; fully tested (1)
- `int64`: yes; fully tested (2) (3)
- `float16`: yes; fully tested (2) (3)
- `float32`: yes; fully tested (1)
- `float64`: yes; fully tested (1)
- `float128`: yes; fully tested (2) (3)
- `bool`: yes; tested (2) (3)

- (1) Uses `cv2` if *mode* is one of: `"constant"`, `"edge"`, `"reflect"`, `"symmetric"`. Otherwise uses `numpy`.
- (2) Uses `numpy`.
- (3) Rejected by `cv2`.

Parameters

- **arr** (*(H,W) ndarray or (H,W,C) ndarray*) – Image-like array to pad.

- **top** (*int, optional*) – Amount of pixels to add to the top side of the image. Must be `0` or greater.
- **right** (*int, optional*) – Amount of pixels to add to the right side of the image. Must be `0` or greater.
- **bottom** (*int, optional*) – Amount of pixels to add to the bottom side of the image. Must be `0` or greater.
- **left** (*int, optional*) – Amount of pixels to add to the left side of the image. Must be `0` or greater.
- **mode** (*str, optional*) – Padding mode to use. See `numpy.pad()` for details. In case of mode `constant`, the parameter *cval* will be used as the `constant_values` parameter to `numpy.pad()`. In case of mode `linear_ramp`, the parameter *cval* will be used as the `end_values` parameter to `numpy.pad()`.
- **cval** (*number or iterable of number, optional*) – Value to use for padding if *mode* is `constant`. See `numpy.pad()` for details. The cval is expected to match the input array's dtype and value range. If an iterable is used, it is expected to contain one value per channel. The number of values and number of channels are expected to match.

Returns Padded array with height `H'=H+top+bottom` and width `W'=W+left+right`.

Return type (H',W') ndarray or (H',W',C) ndarray

`imgaug.augmenters.size.`**`pad_to_aspect_ratio`**(*arr*, *aspect_ratio*, *mode='constant'*, *cval=0*, *return_pad_amounts=False*)

Pad an image array on its sides so that it matches a target aspect ratio.

See *compute_paddings_for_aspect_ratio()* for an explanation of how the required padding amounts are distributed per image axis.

Added in 0.4.0. (Previously named `imgaug.imgaug.pad_to_aspect_ratio()`.)

Supported dtypes:

See *pad()*.

Parameters

- **arr** (*(H,W) ndarray or (H,W,C) ndarray*) – Image-like array to pad.
- **aspect_ratio** (*float*) – Target aspect ratio, given as width/height. E.g. `2.0` denotes the image having twice as much width as height.
- **mode** (*str, optional*) – Padding mode to use. See *pad()* for details.
- **cval** (*number, optional*) – Value to use for padding if *mode* is `constant`. See `numpy.pad()` for details.
- **return_pad_amounts** (*bool, optional*) – If `False`, then only the padded image will be returned. If `True`, a `tuple` with two entries will be returned, where the first entry is the padded image and the second entry are the amounts by which each image side was padded. These amounts are again a `tuple` of the form `(top, right, bottom, left)`, with each value being an `int`.

Returns

- *(H',W') ndarray or (H',W',C) ndarray* – Padded image as `(H',W')` or `(H',W',C)` ndarray, fulfilling the given *aspect_ratio*.

- *tuple of int* – Amounts by which the image was padded on each side, given as a `tuple (top, right, bottom, left)`. This `tuple` is only returned if *return_pad_amounts* was set to `True`.

`imgaug.augmenters.size.`**`pad_to_multiples_of`**(*arr*, *height_multiple*, *width_multiple*, *mode='constant'*, *cval=0*, *return_pad_amounts=False*)

Pad an image array until its side lengths are multiples of given values.

See *compute_paddings_for_aspect_ratio()* for an explanation of how the required padding amounts are distributed per image axis.

Added in 0.4.0. (Previously named `imgaug.imgaug.pad_to_multiples_of()`.)

Supported dtypes:

See *pad()*.

Parameters

- **arr** (*(H,W) ndarray or (H,W,C) ndarray*) – Image-like array to pad.
- **height_multiple** (*None or int*) – The desired multiple of the height. The computed padding amount will reflect a padding that increases the y axis size until it is a multiple of this value.
- **width_multiple** (*None or int*) – The desired multiple of the width. The computed padding amount will reflect a padding that increases the x axis size until it is a multiple of this value.
- **mode** (*str, optional*) – Padding mode to use. See *pad()* for details.
- **cval** (*number, optional*) – Value to use for padding if *mode* is `constant`. See `numpy.pad()` for details.
- **return_pad_amounts** (*bool, optional*) – If `False`, then only the padded image will be returned. If `True`, a `tuple` with two entries will be returned, where the first entry is the padded image and the second entry are the amounts by which each image side was padded. These amounts are again a `tuple` of the form `(top, right, bottom, left)`, with each value being an integer.

Returns

- *(H',W') ndarray or (H',W',C) ndarray* – Padded image as `(H',W')` or `(H',W',C)` ndarray.
- *tuple of int* – Amounts by which the image was padded on each side, given as a `tuple (top, right, bottom, left)`. This `tuple` is only returned if *return_pad_amounts* was set to `True`.

13.36 imgaug.augmenters.weather

Augmenters that create weather effects.

List of augmenters:

- *FastSnowyLandscape*
- *CloudLayer*
- *Clouds*
- *Fog*
- *SnowflakesLayer*

- *Snowflakes*
- *RainLayer*
- *Rain*

class imgaug.augmenters.weather.**CloudLayer**(*intensity_mean*, *intensity_freq_exponent*, *intensity_coarse_scale*, *alpha_min*, *alpha_multiplier*, *alpha_size_px_max*, *alpha_freq_exponent*, *sparsity*, *density_multiplier*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Add a single layer of clouds to an image.

Supported dtypes:

- `uint8`: yes; indirectly tested (1)
- `uint16`: no
- `uint32`: no
- `uint64`: no
- `int8`: no
- `int16`: no
- `int32`: no
- `int64`: no
- `float16`: yes; not tested
- `float32`: yes; not tested
- `float64`: yes; not tested
- `float128`: yes; not tested (2)
- `bool`: no

- (1) Indirectly tested via tests for `Clouds` and *Fog*
- (2) Note that random values are usually sampled as `int64` or `float64`, which `float128` images would exceed. Note also that random values might have to upscaled, which is done via *imresize_many_images()* and has its own limited dtype support (includes however floats up to `64bit`).

Parameters

- **intensity_mean** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – Mean intensity of the clouds (i.e. mean color). Recommended to be in the interval `[190, 255]`.
  - If a `number`, then that value will always be used.
  - If a `tuple (a, b)`, then a value will be uniformly sampled per image from the interval `[a, b]`.
  - If a `list`, then a random value will be sampled from that `list` per image.

  - If a `StochasticParameter`, then a value will be sampled per image from that parameter.

- **intensity_freq_exponent** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – Exponent of the frequency noise used to add fine intensity to the mean intensity. Recommended to be in the interval `[-2.5, -1.5]`. See `__init__()` for details.

- **intensity_coarse_scale** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – Standard deviation of the gaussian distribution used to add more localized intensity to the mean intensity. Sampled in low resolution space, i.e. affects final intensity on a coarse level. Recommended to be in the interval `(0, 10]`.

  - If a `number`, then that value will always be used.

  - If a `tuple (a, b)`, then a value will be uniformly sampled per image from the interval `[a, b]`.

  - If a `list`, then a random value will be sampled from that `list` per image.

  - If a `StochasticParameter`, then a value will be sampled per image from that parameter.

- **alpha_min** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – Minimum alpha when blending cloud noise with the image. High values will lead to clouds being "everywhere". Recommended to usually be at around `0.0` for clouds and `>0` for fog.

  - If a `number`, then that value will always be used.

  - If a `tuple (a, b)`, then a value will be uniformly sampled per image from the interval `[a, b]`.

  - If a `list`, then a random value will be sampled from that `list` per image.

  - If a `StochasticParameter`, then a value will be sampled per image from that parameter.

- **alpha_multiplier** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – Multiplier for the sampled alpha values. High values will lead to denser clouds wherever they are visible. Recommended to be in the interval `[0.3, 1.0]`. Note that this parameter currently overlaps with *density_multiplier*, which is applied a bit later to the alpha mask.

  - If a `number`, then that value will always be used.

  - If a `tuple (a, b)`, then a value will be uniformly sampled per image from the interval `[a, b]`.

  - If a `list`, then a random value will be sampled from that `list` per image.

  - If a `StochasticParameter`, then a value will be sampled per image from that parameter.

- **alpha_size_px_max** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – Controls the image size at which the alpha mask is sampled. Lower values will lead to coarser alpha masks and hence larger clouds (and empty areas). See `__init__()` for details.

- **alpha_freq_exponent** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – Exponent of the frequency noise used to sample the alpha mask. Similarly to *alpha_size_max_px*, lower values will lead to coarser alpha

patterns. Recommended to be in the interval `[-4.0, -1.5]`. See `__init__()` for details.

- **sparsity** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – Exponent applied late to the alpha mask. Lower values will lead to coarser cloud patterns, higher values to finer patterns. Recommended to be somewhere around `1.0`. Do not deviate far from that value, otherwise the alpha mask might get weird patterns with sudden fall-offs to zero that look very unnatural.
  - If a `number`, then that value will always be used.
  - If a `tuple (a, b)`, then a value will be uniformly sampled per image from the interval `[a, b]`.
  - If a `list`, then a random value will be sampled from that `list` per image.
  - If a `StochasticParameter`, then a value will be sampled per image from that parameter.
- **density_multiplier** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – Late multiplier for the alpha mask, similar to *alpha_multiplier*. Set this higher to get "denser" clouds wherever they are visible. Recommended to be around `[0.5, 1.5]`.
  - If a `number`, then that value will always be used.
  - If a `tuple (a, b)`, then a value will be uniformly sampled per image from the interval `[a, b]`.
  - If a `list`, then a random value will be sampled from that `list` per image.
  - If a `StochasticParameter`, then a value will be sampled per image from that parameter.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.

Continued on next page

Table 250 – continued from previous page

augment_bounding_boxes(self, ...[, parents, ...])	Augment a batch of bounding boxes.
augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

draw_on_image	
generate_maps	

draw_on_image(*self*, *image*, *random_state*)

generate_maps(*self*, *image*, *random_state*)

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.weather.**Clouds**(*seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.SomeOf*

Add clouds to images.

This is a wrapper around *CloudLayer*. It executes 1 to 2 layers per image, leading to varying densities and frequency patterns of clouds.

This augmenter seems to be fairly robust w.r.t. the image size. Tested with `96x128`, `192x256` and `960x1280`.

Supported dtypes:

- `uint8`: yes; tested
- `uint16`: no (1)
- `uint32`: no (1)
- `uint64`: no (1)
- `int8`: no (1)
- `int16`: no (1)
- `int32`: no (1)
- `int64`: no (1)
- `float16`: no (1)
- `float32`: no (1)
- `float64`: no (1)
- `float128`: no (1)
- `bool`: no (1)

- (1) Parameters of this augmenter are optimized for the value range of `uint8`. While other dtypes may be accepted, they will lead to images augmented in ways inappropriate for the respective dtype.

Parameters

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.

- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.Clouds()
```

Create an augmenter that adds clouds to images.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`add`(self, augmenter)	Add an augmenter to the list of child augmenters.
`append`(self, object, /)	Append object to the end of the list.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`clear`(self, /)	Remove all items from list.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`count`(self, value, /)	Return number of occurrences of value.
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`extend`(self, iterable, /)	Extend list by appending elements from the iterable.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.

Continued on next page

Table 251 – continued from previous page

get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	See *get_children_lists()*.
get_parameters(self)	See *get_parameters()*.
index(self, value[, start, stop])	Return first index of value.
insert(self, index, object, /)	Insert object before index.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
pop(self[, index])	Remove and return item at index (default last).
remove(self, value, /)	Remove first occurrence of value.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
reverse(self, /)	Reverse *IN PLACE*.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
sort(self, /, *[, key, reverse])	Stable sort *IN PLACE*.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.weather.**FastSnowyLandscape**(*lightness_threshold=(100, 255)*, *lightness_multiplier=(1.0, 4.0)*, *from_colorspace='RGB'*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Convert non-snowy landscapes to snowy ones.

This augmenter expects to get an image that roughly shows a landscape.

This augmenter is based on the method proposed in https://medium.freecodecamp.org/image-augmentation-make-it-rain-make-it-snow-how-to-modify-a-photo-with-machine-learning-163c0cb3843f?gi=bca4a13e634c

Supported dtypes:

- uint8: yes; fully tested
- uint16: no (1)
- uint32: no (1)
- uint64: no (1)
- int8: no (1)
- int16: no (1)
- int32: no (1)

- `int64`: no (1)
- `float16`: no (1)
- `float32`: no (1)
- `float64`: no (1)
- `float128`: no (1)
- `bool`: no (1)

- (1) This augmenter is based on a colorspace conversion to HLS. Hence, only RGB `uint8` inputs are sensible.

Parameters

- **lightness_threshold** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – All pixels with lightness in HLS colorspace that is below this value will have their lightness increased by *lightness_multiplier*.
  - If a `number`, then that value will always be used.
  - If a `tuple (a, b)`, then a value will be uniformly sampled per image from the discrete interval `[a..b]`.
  - If a `list`, then a random value will be sampled from that `list` per image.
  - If a `StochasticParameter`, then a value will be sampled per image from that parameter.
- **lightness_multiplier** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter, optional*) – Multiplier for pixel's lightness value in HLS colorspace. Affects all pixels selected via *lightness_threshold*.
  - If a `number`, then that value will always be used.
  - If a `tuple (a, b)`, then a value will be uniformly sampled per image from the discrete interval `[a..b]`.
  - If a `list`, then a random value will be sampled from that `list` per image.
  - If a `StochasticParameter`, then a value will be sampled per image from that parameter.
- **from_colorspace** (*str, optional*) – The source colorspace of the input images. See `__init__()`.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.FastSnowyLandscape(
>>>     lightness_threshold=140,
>>>     lightness_multiplier=2.5
>>> )
```

Search for all pixels in the image with a lightness value in HLS colorspace of less than `140` and increase their lightness by a factor of `2.5`.

```
>>> aug = iaa.FastSnowyLandscape(
>>>     lightness_threshold=[128, 200],
>>>     lightness_multiplier=(1.5, 3.5)
>>> )
```

Search for all pixels in the image with a lightness value in HLS colorspace of less than `128` or less than `200` (one of these values is picked per image) and multiply their lightness by a factor of `x` with `x` being sampled from `uniform(1.5, 3.5)` (once per image).

```
>>> aug = iaa.FastSnowyLandscape(
>>>     lightness_threshold=(100, 255),
>>>     lightness_multiplier=(1.0, 4.0)
>>> )
```

Similar to the previous example, but the lightness threshold is sampled from `uniform(100, 255)` (per image) and the multiplier from `uniform(1.0, 4.0)` (per image). This seems to produce good and varied results.

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.

Continued on next page

Table 252 – continued from previous page

copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

get_parameters(*self*)
See *get_parameters()*.

class imgaug.augmenters.weather.**Fog**(*seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.weather.CloudLayer*

Add fog to images.

This is a wrapper around *CloudLayer*. It executes a single layer per image with a configuration leading to fairly dense clouds with low-frequency patterns.

This augmenter seems to be fairly robust w.r.t. the image size. Tested with 96x128, 192x256 and 960x1280.

Supported dtypes:

- uint8: yes; tested
- uint16: no (1)
- uint32: no (1)
- uint64: no (1)

- `int8`: no (1)
- `int16`: no (1)
- `int32`: no (1)
- `int64`: no (1)
- `float16`: no (1)
- `float32`: no (1)
- `float64`: no (1)
- `float128`: no (1)
- `bool`: no (1)

- (1) Parameters of this augmenter are optimized for the value range of `uint8`. While other dtypes may be accepted, they will lead to images augmented in ways inappropriate for the respective dtype.

Parameters

- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.Fog()
```

Create an augmenter that adds fog to images.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.

Continued on next page

Table 253 – continued from previous page

augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

draw_on_image	
generate_maps	

class imgaug.augmenters.weather.**Rain**(*nb_iterations=(1, 3)*, *drop_size=(0.01, 0.02)*, *speed=(0.04, 0.2)*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.SomeOf*

Add falling snowflakes to images.

This is a wrapper around *RainLayer*. It executes 1 to 3 layers per image.

Note: This augmenter currently seems to work best for medium-sized images around `192x256`. For smaller images, you may want to increase the *speed* value to e.g. `(0.1, 0.3)`, otherwise the drops tend to look like snowflakes. For larger images, you may want to increase the *drop_size* to e.g. `(0.10, 0.20)`.

Added in 0.4.0.

Supported dtypes:

- `uint8`: yes; tested
- `uint16`: no (1)
- `uint32`: no (1)
- `uint64`: no (1)
- `int8`: no (1)
- `int16`: no (1)
- `int32`: no (1)
- `int64`: no (1)
- `float16`: no (1)
- `float32`: no (1)
- `float64`: no (1)
- `float128`: no (1)
- `bool`: no (1)

- (1) Parameters of this augmenter are optimized for the value range of `uint8`. While other dtypes may be accepted, they will lead to images augmented in ways inappropriate for the respective dtype.

Parameters

- **drop_size** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – See *RainLayer*.
- **speed** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – See *RainLayer*.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage

will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.

- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.Rain(speed=(0.1, 0.3))
```

Add rain to small images (around `96x128`).

```
>>> aug = iaa.Rain()
```

Add rain to medium sized images (around `192x256`).

```
>>> aug = iaa.Rain(drop_size=(0.10, 0.20))
```

Add rain to large images (around `960x1280`).

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`add`(self, augmenter)	Add an augmenter to the list of child augmenters.
`append`(self, object, /)	Append object to the end of the list.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`clear`(self, /)	Remove all items from list.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).

Continued on next page

Table 254 – continued from previous page

count(self, value, /)	Return number of occurrences of value.
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
extend(self, iterable, /)	Extend list by appending elements from the iterable.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	See *get_children_lists()*.
get_parameters(self)	See *get_parameters()*.
index(self, value[, start, stop])	Return first index of value.
insert(self, index, object, /)	Insert object before index.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
pop(self[, index])	Remove and return item at index (default last).
remove(self, value, /)	Remove first occurrence of value.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
reverse(self, /)	Reverse *IN PLACE*.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
sort(self, /, *[, key, reverse])	Stable sort *IN PLACE*.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.weather.**RainLayer**(*density*, *density_uniformity*, *drop_size*, *drop_size_uniformity*, *angle*, *speed*, *blur_sigma_fraction*, *blur_sigma_limits=(0.5, 3.75)*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.weather.SnowflakesLayer*

Add a single layer of falling raindrops to images.

Added in 0.4.0.

Supported dtypes:

- uint8: yes; indirectly tested (1)
- uint16: no

- `uint32`: no
- `uint64`: no
- `int8`: no
- `int16`: no
- `int32`: no
- `int64`: no
- `float16`: no
- `float32`: no
- `float64`: no
- `float128`: no
- `bool`: no

- (1) indirectly tested via tests for *Rain*

Parameters

- **density** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – Same as in *SnowflakesLayer*.
- **density_uniformity** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – Same as in *SnowflakesLayer*.
- **drop_size** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – Same as *flake_size* in *SnowflakesLayer*.
- **drop_size_uniformity** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – Same as *flake_size_uniformity* in *SnowflakesLayer*.
- **angle** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – Same as in *SnowflakesLayer*.
- **speed** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – Same as in *SnowflakesLayer*.
- **blur_sigma_fraction** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – Same as in *SnowflakesLayer*.
- **blur_sigma_limits** (*tuple of float, optional*) – Same as in *SnowflakesLayer*.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Methods

`__call__`(self, *args, **kwargs)	Alias for *augment()*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
`copy_random_state_`(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
`deepcopy`(self)	Create a deep copy of this Augmenter instance.
`draw_grid`(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
`find_augmenters`(self, func[, parents, flat])	Find augmenters that match a condition.
`find_augmenters_by_name`(self, name[, regex, ...])	Find augmenter(s) by name.
`find_augmenters_by_names`(self, names[, ...])	Find augmenter(s) by names.
`get_all_children`(self[, flat])	Get all children of this augmenter as a list.
`get_children_lists`(self)	Get a list of lists of children of this augmenter.
`get_parameters`(self)	See *get_parameters()*.
`localize_random_state`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`localize_random_state_`(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
`pool`(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
`remove_augmenters`(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
`remove_augmenters_`(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
`remove_augmenters_inplace`(self, func[, parents])	**Deprecated**.
`reseed`(self[, random_state, deterministic_too])	**Deprecated**.
`seed_`(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.

Continued on next page

Table 255 – continued from previous page

show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

draw_on_image

class imgaug.augmenters.weather.**Snowflakes**(*density=(0.005, 0.075)*, *density_uniformity=(0.3, 0.9)*, *flake_size=(0.2, 0.7)*, *flake_size_uniformity=(0.4, 0.8)*, *angle=(-30, 30)*, *speed=(0.007, 0.03)*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.SomeOf*

Add falling snowflakes to images.

This is a wrapper around *SnowflakesLayer*. It executes 1 to 3 layers per image.

Supported dtypes:

- uint8: yes; tested
- uint16: no (1)
- uint32: no (1)
- uint64: no (1)
- int8: no (1)
- int16: no (1)
- int32: no (1)
- int64: no (1)
- float16: no (1)
- float32: no (1)
- float64: no (1)
- float128: no (1)
- bool: no (1)

- (1) Parameters of this augmenter are optimized for the value range of uint8. While other dtypes may be accepted, they will lead to images augmented in ways inappropriate for the respective dtype.

Parameters

- **density** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – Density of the snowflake layer, as a probability of each pixel in low resolution space to be a snowflake. Valid values are in the interval [0.0, 1.0]. Recommended to be in the interval [0.01, 0.075].
  - If a number, then that value will always be used.

  - If a `tuple (a, b)`, then a value will be uniformly sampled per image from the interval `[a, b]`.
  - If a `list`, then a random value will be sampled from that `list` per image.
  - If a `StochasticParameter`, then a value will be sampled per image from that parameter.

- **density_uniformity** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – Size uniformity of the snowflakes. Higher values denote more similarly sized snowflakes. Valid values are in the interval `[0.0, 1.0]`. Recommended to be around `0.5`.
  - If a `number`, then that value will always be used.
  - If a `tuple (a, b)`, then a value will be uniformly sampled per image from the interval `[a, b]`.
  - If a `list`, then a random value will be sampled from that `list` per image.
  - If a `StochasticParameter`, then a value will be sampled per image from that parameter.

- **flake_size** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – Size of the snowflakes. This parameter controls the resolution at which snowflakes are sampled. Higher values mean that the resolution is closer to the input image's resolution and hence each sampled snowflake will be smaller (because of the smaller pixel size).

  Valid values are in the interval `(0.0, 1.0]`. Recommended values:
  - On `96x128` a value of `(0.1, 0.4)` worked well.
  - On `192x256` a value of `(0.2, 0.7)` worked well.
  - On `960x1280` a value of `(0.7, 0.95)` worked well.

  Datatype behaviour:
  - If a `number`, then that value will always be used.
  - If a `tuple (a, b)`, then a value will be uniformly sampled per image from the interval `[a, b]`.
  - If a `list`, then a random value will be sampled from that `list` per image.
  - If a `StochasticParameter`, then a value will be sampled per image from that parameter.

- **flake_size_uniformity** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – Controls the size uniformity of the snowflakes. Higher values mean that the snowflakes are more similarly sized. Valid values are in the interval `[0.0, 1.0]`. Recommended to be around `0.5`.
  - If a `number`, then that value will always be used.
  - If a `tuple (a, b)`, then a value will be uniformly sampled per image from the interval `[a, b]`.
  - If a `list`, then a random value will be sampled from that `list` per image.
  - If a `StochasticParameter`, then a value will be sampled per image from that parameter.

- **angle** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – Angle in degrees of motion blur applied to the snowflakes, where `0.0` is motion blur that points straight upwards. Recommended to be in the interval `[-30, 30]`. See also `__init__()`.
  - If a `number`, then that value will always be used.
  - If a `tuple (a, b)`, then a value will be uniformly sampled per image from the interval `[a, b]`.
  - If a `list`, then a random value will be sampled from that `list` per image.
  - If a `StochasticParameter`, then a value will be sampled per image from that parameter.
- **speed** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – Perceived falling speed of the snowflakes. This parameter controls the motion blur's kernel size. It follows roughly the form `kernel_size = image_size * speed`. Hence, values around `1.0` denote that the motion blur should "stretch" each snowflake over the whole image.

  Valid values are in the interval `[0.0, 1.0]`. Recommended values:
  - On `96x128` a value of `(0.01, 0.05)` worked well.
  - On `192x256` a value of `(0.007, 0.03)` worked well.
  - On `960x1280` a value of `(0.001, 0.03)` worked well.

  Datatype behaviour:
  - If a `number`, then that value will always be used.
  - If a `tuple (a, b)`, then a value will be uniformly sampled per image from the interval `[a, b]`.
  - If a `list`, then a random value will be sampled from that `list` per image.
  - If a `StochasticParameter`, then a value will be sampled per image from that parameter.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Examples

```
>>> import imgaug.augmenters as iaa
>>> aug = iaa.Snowflakes(flake_size=(0.1, 0.4), speed=(0.01, 0.05))
```

Add snowflakes to small images (around 96x128).

```
>>> aug = iaa.Snowflakes(flake_size=(0.2, 0.7), speed=(0.007, 0.03))
```

Add snowflakes to medium-sized images (around 192x256).

```
>>> aug = iaa.Snowflakes(flake_size=(0.7, 0.95), speed=(0.001, 0.03))
```

Add snowflakes to large images (around 960x1280).

Methods

__call__(self, *args, **kwargs)	Alias for *augment()*.
add(self, augmenter)	Add an augmenter to the list of child augmenters.
append(self, object, /)	Append object to the end of the list.
augment(self[, return_batch, hooks])	Augment a batch.
augment_batch(self, batch[, hooks])	**Deprecated**.
augment_batch_(self, batch[, parents, hooks])	Augment a single batch in-place.
augment_batches(self, batches[, hooks, ...])	Augment multiple batches.
augment_bounding_boxes(self, ...[, parents, ...])	Augment a batch of bounding boxes.
augment_heatmaps(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
augment_image(self, image[, hooks])	Augment a single image.
augment_images(self, images[, parents, hooks])	Augment a batch of images.
augment_keypoints(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
augment_line_strings(self, ...[, parents, hooks])	Augment a batch of line strings.
augment_polygons(self, polygons_on_images[, ...])	Augment a batch of polygons.
augment_segmentation_maps(self, segmaps[, ...])	Augment a batch of segmentation maps.
clear(self, /)	Remove all items from list.
copy(self)	Create a shallow copy of this Augmenter instance.
copy_random_state(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.
copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
count(self, value, /)	Return number of occurrences of value.
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
extend(self, iterable, /)	Extend list by appending elements from the iterable.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	See *get_children_lists()*.

Continued on next page

Table 256 – continued from previous page

get_parameters(self)	See *get_parameters()*.
index(self, value[, start, stop])	Return first index of value.
insert(self, index, object, /)	Insert object before index.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
pop(self[, index])	Remove and return item at index (default last).
remove(self, value, /)	Remove first occurrence of value.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
reverse(self, /)	Reverse *IN PLACE*.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
sort(self, /, *[, key, reverse])	Stable sort *IN PLACE*.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

class imgaug.augmenters.weather.**SnowflakesLayer**(*density*, *density_uniformity*, *flake_size*, *flake_size_uniformity*, *angle*, *speed*, *blur_sigma_fraction*, *blur_sigma_limits=(0.5, 3.75)*, *seed=None*, *name=None*, *random_state='deprecated'*, *deterministic='deprecated'*)

Bases: *imgaug.augmenters.meta.Augmenter*

Add a single layer of falling snowflakes to images.

Supported dtypes:

- uint8: yes; indirectly tested (1)
- uint16: no
- uint32: no
- uint64: no
- int8: no
- int16: no
- int32: no
- int64: no
- float16: no
- float32: no
- float64: no

- `float128`: no
- `bool`: no
- (1) indirectly tested via tests for *Snowflakes*

Parameters

- **density** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – Density of the snowflake layer, as a probability of each pixel in low resolution space to be a snowflake. Valid values are in the interval `[0.0, 1.0]`. Recommended to be in the interval `[0.01, 0.075]`.
  - If a `number`, then that value will always be used.
  - If a `tuple (a, b)`, then a value will be uniformly sampled per image from the interval `[a, b]`.
  - If a `list`, then a random value will be sampled from that `list` per image.
  - If a `StochasticParameter`, then a value will be sampled per image from that parameter.
- **density_uniformity** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – Size uniformity of the snowflakes. Higher values denote more similarly sized snowflakes. Valid values are in the interval `[0.0, 1.0]`. Recommended to be around `0.5`.
  - If a `number`, then that value will always be used.
  - If a `tuple (a, b)`, then a value will be uniformly sampled per image from the interval `[a, b]`.
  - If a `list`, then a random value will be sampled from that `list` per image.
  - If a `StochasticParameter`, then a value will be sampled per image from that parameter.
- **flake_size** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – Size of the snowflakes. This parameter controls the resolution at which snowflakes are sampled. Higher values mean that the resolution is closer to the input image's resolution and hence each sampled snowflake will be smaller (because of the smaller pixel size).

  Valid values are in the interval `(0.0, 1.0]`. Recommended values:
  - On 96x128 a value of `(0.1, 0.4)` worked well.
  - On 192x256 a value of `(0.2, 0.7)` worked well.
  - On 960x1280 a value of `(0.7, 0.95)` worked well.

  Datatype behaviour:
  - If a `number`, then that value will always be used.
  - If a `tuple (a, b)`, then a value will be uniformly sampled per image from the interval `[a, b]`.
  - If a `list`, then a random value will be sampled from that `list` per image.
  - If a `StochasticParameter`, then a value will be sampled per image from that parameter.

- **flake_size_uniformity** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – Controls the size uniformity of the snowflakes. Higher values mean that the snowflakes are more similarly sized. Valid values are in the interval `[0.0, 1.0]`. Recommended to be around `0.5`.
  - If a `number`, then that value will always be used.
  - If a `tuple (a, b)`, then a value will be uniformly sampled per image from the interval `[a, b]`.
  - If a `list`, then a random value will be sampled from that `list` per image.
  - If a `StochasticParameter`, then a value will be sampled per image from that parameter.
- **angle** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – Angle in degrees of motion blur applied to the snowflakes, where `0.0` is motion blur that points straight upwards. Recommended to be in the interval `[-30, 30]`. See also `__init__()`.
  - If a `number`, then that value will always be used.
  - If a `tuple (a, b)`, then a value will be uniformly sampled per image from the interval `[a, b]`.
  - If a `list`, then a random value will be sampled from that `list` per image.
  - If a `StochasticParameter`, then a value will be sampled per image from that parameter.
- **speed** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – Perceived falling speed of the snowflakes. This parameter controls the motion blur's kernel size. It follows roughly the form `kernel_size = image_size * speed`. Hence, values around `1.0` denote that the motion blur should "stretch" each snowflake over the whole image.

  Valid values are in the interval `[0.0, 1.0]`. Recommended values:
  - On 96x128 a value of `(0.01, 0.05)` worked well.
  - On 192x256 a value of `(0.007, 0.03)` worked well.
  - On 960x1280 a value of `(0.001, 0.03)` worked well.

  Datatype behaviour:
  - If a `number`, then that value will always be used.
  - If a `tuple (a, b)`, then a value will be uniformly sampled per image from the interval `[a, b]`.
  - If a `list`, then a random value will be sampled from that `list` per image.
  - If a `StochasticParameter`, then a value will be sampled per image from that parameter.
- **blur_sigma_fraction** (*number or tuple of number or list of number or imgaug.parameters.StochasticParameter*) – Standard deviation (as a fraction of the image size) of gaussian blur applied to the snowflakes. Valid values are in the interval `[0.0, 1.0]`. Recommended to be in the interval `[0.0001, 0.001]`. May still require tinkering based on image size.
  - If a `number`, then that value will always be used.

  - If a `tuple (a, b)`, then a value will be uniformly sampled per image from the interval `[a, b]`.
  - If a `list`, then a random value will be sampled from that `list` per image.
  - If a `StochasticParameter`, then a value will be sampled per image from that parameter.

- **blur_sigma_limits** (*tuple of float, optional*) – Controls allowed min and max values of *blur_sigma_fraction* after(!) multiplication with the image size. First value is the minimum, second value is the maximum. Values outside of that range will be clipped to be within that range. This prevents extreme values for very small or large images.
- **seed** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – See `__init__()`.
- **name** (*None or str, optional*) – See `__init__()`.
- **random_state** (*None or int or imgaug.random.RNG or numpy.random.Generator or numpy.random.BitGenerator or numpy.random.SeedSequence or numpy.random.RandomState, optional*) – Old name for parameter *seed*. Its usage will not yet cause a deprecation warning, but it is still recommended to use *seed* now. Outdated since 0.4.0.
- **deterministic** (*bool, optional*) – Deprecated since 0.4.0. See method `to_deterministic()` for an alternative and for details about what the "deterministic mode" actually does.

Methods

`__call__`(self, *args, **kwargs)	Alias for *`augment()`*.
`augment`(self[, return_batch, hooks])	Augment a batch.
`augment_batch`(self, batch[, hooks])	**Deprecated**.
`augment_batch_`(self, batch[, parents, hooks])	Augment a single batch in-place.
`augment_batches`(self, batches[, hooks, ...])	Augment multiple batches.
`augment_bounding_boxes`(self, ...[, parents, ...])	Augment a batch of bounding boxes.
`augment_heatmaps`(self, heatmaps[, parents, ...])	Augment a batch of heatmaps.
`augment_image`(self, image[, hooks])	Augment a single image.
`augment_images`(self, images[, parents, hooks])	Augment a batch of images.
`augment_keypoints`(self, keypoints_on_images)	Augment a batch of keypoints/landmarks.
`augment_line_strings`(self, ...[, parents, hooks])	Augment a batch of line strings.
`augment_polygons`(self, polygons_on_images[, ...])	Augment a batch of polygons.
`augment_segmentation_maps`(self, segmaps[, ...])	Augment a batch of segmentation maps.
`copy`(self)	Create a shallow copy of this Augmenter instance.
`copy_random_state`(self, source[, recursive, ...])	Copy the RNGs from a source augmenter sequence.

Continued on next page

Table 257 – continued from previous page

copy_random_state_(self, source[, ...])	Copy the RNGs from a source augmenter sequence (in-place).
deepcopy(self)	Create a deep copy of this Augmenter instance.
draw_grid(self, images, rows, cols)	Augment images and draw the results as a single grid-like image.
find_augmenters(self, func[, parents, flat])	Find augmenters that match a condition.
find_augmenters_by_name(self, name[, regex, ...])	Find augmenter(s) by name.
find_augmenters_by_names(self, names[, ...])	Find augmenter(s) by names.
get_all_children(self[, flat])	Get all children of this augmenter as a list.
get_children_lists(self)	Get a list of lists of children of this augmenter.
get_parameters(self)	See *get_parameters()*.
localize_random_state(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
localize_random_state_(self[, recursive])	Assign augmenter-specific RNGs to this augmenter and its children.
pool(self[, processes, maxtasksperchild, seed])	Create a pool used for multicore augmentation.
remove_augmenters(self, func[, copy, ...])	Remove this augmenter or children that match a condition.
remove_augmenters_(self, func[, parents])	Remove in-place children of this augmenter that match a condition.
remove_augmenters_inplace(self, func[, parents])	**Deprecated**.
reseed(self[, random_state, deterministic_too])	**Deprecated**.
seed_(self[, entropy, deterministic_too])	Seed this augmenter and all of its children.
show_grid(self, images, rows, cols)	Augment images and plot the results as a single grid-like image.
to_deterministic(self[, n])	Convert this augmenter from a stochastic to a deterministic one.

draw_on_image	

draw_on_image(*self*, *image*, *random_state*)

get_parameters(*self*)
See *get_parameters()*.

See modindex for API.

CHAPTER 14

Indices and tables

- genindex
- modindex
- search

Python Module Index

i

imgaug.augmentables.base, 409
imgaug.augmentables.batches, 409
imgaug.augmentables.bbs, 416
imgaug.augmentables.heatmaps, 434
imgaug.augmentables.kps, 439
imgaug.augmentables.lines, 452
imgaug.augmentables.normalization, 471
imgaug.augmentables.polys, 472
imgaug.augmentables.segmaps, 491
imgaug.augmentables.utils, 494
imgaug.augmenters.arithmetic, 497
imgaug.augmenters.artistic, 568
imgaug.augmenters.base, 572
imgaug.augmenters.blend, 572
imgaug.augmenters.blur, 626
imgaug.augmenters.collections, 645
imgaug.augmenters.color, 648
imgaug.augmenters.contrast, 703
imgaug.augmenters.convolutional, 729
imgaug.augmenters.debug, 741
imgaug.augmenters.edges, 745
imgaug.augmenters.flip, 750
imgaug.augmenters.geometric, 756
imgaug.augmenters.imgcorruptlike, 810
imgaug.augmenters.meta, 855
imgaug.augmenters.pillike, 912
imgaug.augmenters.pooling, 962
imgaug.augmenters.segmentation, 974
imgaug.augmenters.size, 998
imgaug.augmenters.weather, 1065
imgaug.dtypes, 393
imgaug.imgaug, 343
imgaug.multicore, 389
imgaug.parameters, 362
imgaug.random, 394
imgaug.validation, 408

Index

A

Absolute (*class in imgaug.parameters*), 362
Add (*class in imgaug.augmenters.arithmetic*), 498
Add (*class in imgaug.parameters*), 363
add() (*imgaug.augmenters.meta.Sequential method*), 902
add() (*imgaug.augmenters.meta.SomeOf method*), 906
add_elementwise() (*in module imgaug.augmenters.arithmetic*), 559
add_scalar() (*in module imgaug.augmenters.arithmetic*), 560
AddElementwise (*class in imgaug.augmenters.arithmetic*), 500
AdditiveGaussianNoise (*class in imgaug.augmenters.arithmetic*), 502
AdditiveLaplaceNoise (*class in imgaug.augmenters.arithmetic*), 505
AdditivePoissonNoise (*class in imgaug.augmenters.arithmetic*), 508
AddToBrightness (*class in imgaug.augmenters.color*), 649
AddToHue (*class in imgaug.augmenters.color*), 651
AddToHueAndSaturation (*class in imgaug.augmenters.color*), 653
AddToSaturation (*class in imgaug.augmenters.color*), 656
adjust_contrast_gamma() (*in module imgaug.augmenters.contrast*), 725
adjust_contrast_linear() (*in module imgaug.augmenters.contrast*), 726
adjust_contrast_log() (*in module imgaug.augmenters.contrast*), 727
adjust_contrast_sigmoid() (*in module imgaug.augmenters.contrast*), 728
advance_() (*imgaug.random.RNG method*), 397
advance_generator_() (*in module imgaug.random*), 403
Affine (*class in imgaug.augmenters.geometric*), 756
Affine (*class in imgaug.augmenters.pillike*), 913
AffineCv2 (*class in imgaug.augmenters.geometric*), 766
all_finished() (*imgaug.multicore.BackgroundAugmenter method*), 390
all_finished() (*imgaug.multicore.BatchLoader method*), 391
AllChannelsCLAHE (*class in imgaug.augmenters.contrast*), 704
AllChannelsHistogramEqualization (*class in imgaug.augmenters.contrast*), 707
ALLOW_DTYPES_CUSTOM_MINMAX (*imgaug.augmenters.arithmetic.Invert attribute*), 538
almost_equals() (*imgaug.augmentables.bbs.BoundingBox method*), 418
almost_equals() (*imgaug.augmentables.kps.Keypoint method*), 440
almost_equals() (*imgaug.augmentables.lines.LineString method*), 453
almost_equals() (*imgaug.augmentables.polys.Polygon method*), 474
Alpha() (*in module imgaug.augmenters.blend*), 572
AlphaElementwise() (*in module imgaug.augmenters.blend*), 573
angle_between_vectors() (*in module imgaug.imgaug*), 346
apply_brightness() (*in module imgaug.augmenters.imgcorruptlike*), 848
apply_contrast() (*in module imgaug.augmenters.imgcorruptlike*), 849
apply_defocus_blur() (*in module imgaug.augmenters.imgcorruptlike*), 849
apply_elastic_transform() (*in module imgaug.augmenters.imgcorruptlike*), 849
apply_fog() (*in module im-*

gaug.augmenters.imgcorruptlike), 850
apply_frost() (*in module imgaug.augmenters.imgcorruptlike*), 850
apply_gaussian_blur() (*in module imgaug.augmenters.imgcorruptlike*), 850
apply_gaussian_noise() (*in module imgaug.augmenters.imgcorruptlike*), 851
apply_glass_blur() (*in module imgaug.augmenters.imgcorruptlike*), 851
apply_impulse_noise() (*in module imgaug.augmenters.imgcorruptlike*), 851
apply_jigsaw() (*in module imgaug.augmenters.geometric*), 809
apply_jigsaw_to_coords() (*in module imgaug.augmenters.geometric*), 809
apply_jpeg_compression() (*in module imgaug.augmenters.imgcorruptlike*), 852
apply_lut() (*in module imgaug.imgaug*), 346
apply_lut_() (*in module imgaug.imgaug*), 347
apply_motion_blur() (*in module imgaug.augmenters.imgcorruptlike*), 852
apply_pixelate() (*in module imgaug.augmenters.imgcorruptlike*), 852
apply_saturate() (*in module imgaug.augmenters.imgcorruptlike*), 853
apply_shot_noise() (*in module imgaug.augmenters.imgcorruptlike*), 853
apply_snow() (*in module imgaug.augmenters.imgcorruptlike*), 853
apply_spatter() (*in module imgaug.augmenters.imgcorruptlike*), 854
apply_speckle_noise() (*in module imgaug.augmenters.imgcorruptlike*), 854
apply_zoom_blur() (*in module imgaug.augmenters.imgcorruptlike*), 854
area (*imgaug.augmentables.bbs.BoundingBox attribute*), 418
area (*imgaug.augmentables.polys.Polygon attribute*), 474
assert_is_iterable_of() (*in module imgaug.validation*), 408
AssertLambda (*class in imgaug.augmenters.meta*), 856
AssertShape (*class in imgaug.augmenters.meta*), 859
augment() (*imgaug.augmenters.meta.Augmenter method*), 864
augment_batch() (*imgaug.augmenters.meta.Augmenter method*), 867
augment_batch_() (*imgaug.augmenters.meta.Augmenter method*), 868
augment_batches() (*imgaug.augmenters.meta.Augmenter method*), 868
augment_bounding_boxes() (*imgaug.augmenters.meta.Augmenter method*), 869
augment_heatmaps() (*imgaug.augmenters.meta.Augmenter method*), 869
augment_image() (*imgaug.augmenters.meta.Augmenter method*), 870
augment_images() (*imgaug.augmenters.meta.Augmenter method*), 870
augment_keypoints() (*imgaug.augmenters.meta.Augmenter method*), 871
augment_line_strings() (*imgaug.augmenters.meta.Augmenter method*), 871
augment_polygons() (*imgaug.augmenters.meta.Augmenter method*), 872
augment_segmentation_maps() (*imgaug.augmenters.meta.Augmenter method*), 873
Augmenter (*class in imgaug.augmenters.meta*), 862
Autocontrast (*class in imgaug.augmenters.pillike*), 916
autocontrast() (*in module imgaug.augmenters.pillike*), 949
AverageBlur (*class in imgaug.augmenters.blur*), 626
AveragePooling (*class in imgaug.augmenters.pooling*), 962
avg_pool() (*imgaug.augmentables.heatmaps.HeatmapsOnImage method*), 435
avg_pool() (*in module imgaug.imgaug*), 347

B

BackgroundAugmenter (*class in imgaug.multicore*), 389
BackgroundAugmenter() (*in module imgaug.imgaug*), 343
Batch (*class in imgaug.augmentables.batches*), 409
Batch() (*in module imgaug.imgaug*), 343
BatchLoader (*class in imgaug.multicore*), 390
BatchLoader() (*in module imgaug.imgaug*), 343
Beta (*class in imgaug.parameters*), 364
beta() (*imgaug.random.RNG method*), 398
BGR (*imgaug.augmenters.color.ChangeColorspace attribute*), 662
BGR (*imgaug.augmenters.contrast.CLAHE attribute*), 713
BGR (*imgaug.augmenters.contrast.HistogramEqualization attribute*), 718

BilateralBlur (*class in imgaug.augmenters.blur*), 629
Binomial (*class in imgaug.parameters*), 365
binomial() (*imgaug.random.RNG method*), 398
blend_alpha() (*in module imgaug.augmenters.blend*), 625
BlendAlpha (*class in imgaug.augmenters.blend*), 573
BlendAlphaBoundingBoxes (*class in imgaug.augmenters.blend*), 576
BlendAlphaCheckerboard (*class in imgaug.augmenters.blend*), 579
BlendAlphaElementwise (*class in imgaug.augmenters.blend*), 582
BlendAlphaFrequencyNoise (*class in imgaug.augmenters.blend*), 585
BlendAlphaHorizontalLinearGradient (*class in imgaug.augmenters.blend*), 590
BlendAlphaMask (*class in imgaug.augmenters.blend*), 593
BlendAlphaRegularGrid (*class in imgaug.augmenters.blend*), 595
BlendAlphaSegMapClassIds (*class in imgaug.augmenters.blend*), 598
BlendAlphaSimplexNoise (*class in imgaug.augmenters.blend*), 601
BlendAlphaSomeColors (*class in imgaug.augmenters.blend*), 606
BlendAlphaVerticalLinearGradient (*class in imgaug.augmenters.blend*), 609
blur_gaussian_() (*in module imgaug.augmenters.blur*), 641
blur_mean_shift_() (*in module imgaug.augmenters.blur*), 644
both_np_float_if_one_is_float() (*in module imgaug.parameters*), 389
bounding_boxes (*imgaug.augmentables.batches.Batch attribute*), 410
BoundingBox (*class in imgaug.augmentables.bbs*), 416
BoundingBox() (*in module imgaug.imgaug*), 343
BoundingBoxesMaskGen (*class in imgaug.augmenters.blend*), 612
BoundingBoxesOnImage (*class in imgaug.augmentables.bbs*), 427
BoundingBoxesOnImage() (*in module imgaug.imgaug*), 343
Brightness (*class in imgaug.augmenters.imgcorruptlike*), 811
bytes() (*imgaug.random.RNG method*), 398

C

caller_name() (*in module imgaug.imgaug*), 348
Canny (*class in imgaug.augmenters.edges*), 745
Cartoon (*class in imgaug.augmenters.artistic*), 568
center_x (*imgaug.augmentables.bbs.BoundingBox attribute*), 418
center_y (*imgaug.augmentables.bbs.BoundingBox attribute*), 418
CenterCropToAspectRatio (*class in imgaug.augmenters.size*), 999
CenterCropToFixedSize (*class in imgaug.augmenters.size*), 1001
CenterCropToMultiplesOf (*class in imgaug.augmenters.size*), 1003
CenterCropToPowersOf (*class in imgaug.augmenters.size*), 1005
CenterCropToSquare (*class in imgaug.augmenters.size*), 1006
CenterPadToAspectRatio (*class in imgaug.augmenters.size*), 1008
CenterPadToFixedSize (*class in imgaug.augmenters.size*), 1010
CenterPadToMultiplesOf (*class in imgaug.augmenters.size*), 1012
CenterPadToPowersOf (*class in imgaug.augmenters.size*), 1014
CenterPadToSquare (*class in imgaug.augmenters.size*), 1016
change_color_temperature() (*in module imgaug.augmenters.color*), 697
change_color_temperatures_() (*in module imgaug.augmenters.color*), 697
change_colorspace_() (*in module imgaug.augmenters.color*), 698
change_colorspaces_() (*in module imgaug.augmenters.color*), 699
change_dtype_() (*in module imgaug.dtypes*), 393
change_dtypes_() (*in module imgaug.dtypes*), 393
change_first_point_by_coords() (*imgaug.augmentables.polys.Polygon method*), 474
change_first_point_by_index() (*imgaug.augmentables.polys.Polygon method*), 475
change_normalization() (*imgaug.augmentables.heatmaps.HeatmapsOnImage class method*), 435
ChangeColorspace (*class in imgaug.augmenters.color*), 660
ChangeColorTemperature (*class in imgaug.augmenters.color*), 658
ChannelShuffle (*class in imgaug.augmenters.meta*), 881
CheckerboardMaskGen (*class in imgaug.augmenters.blend*), 613
ChiSquare (*class in imgaug.parameters*), 365
chisquare() (*imgaug.random.RNG method*), 398

Choice (*class in imgaug.parameters*), 366
choice() (*imgaug.random.RNG method*), 398
CIE (*imgaug.augmenters.color.ChangeColorspace attribute*), 662
CLAHE (*class in imgaug.augmenters.contrast*), 709
Clip (*class in imgaug.parameters*), 367
clip_() (*in module imgaug.dtypes*), 393
clip_augmented_image() (*in module imgaug.augmenters.meta*), 911
clip_augmented_image_() (*in module imgaug.augmenters.meta*), 911
clip_augmented_images() (*in module imgaug.augmenters.meta*), 911
clip_augmented_images_() (*in module imgaug.augmenters.meta*), 911
clip_out_of_image() (*imgaug.augmentables.bbs.BoundingBox method*), 418
clip_out_of_image() (*imgaug.augmentables.bbs.BoundingBoxesOnImage method*), 429
clip_out_of_image() (*imgaug.augmentables.kps.KeypointsOnImage method*), 445
clip_out_of_image() (*imgaug.augmentables.lines.LineString method*), 454
clip_out_of_image() (*imgaug.augmentables.lines.LineStringsOnImage method*), 466
clip_out_of_image() (*imgaug.augmentables.polys.Polygon method*), 475
clip_out_of_image() (*imgaug.augmentables.polys.PolygonsOnImage method*), 484
clip_out_of_image_() (*imgaug.augmentables.bbs.BoundingBox method*), 418
clip_out_of_image_() (*imgaug.augmentables.bbs.BoundingBoxesOnImage method*), 429
clip_out_of_image_() (*imgaug.augmentables.kps.KeypointsOnImage method*), 446
clip_out_of_image_() (*imgaug.augmentables.lines.LineStringsOnImage method*), 466
clip_out_of_image_() (*imgaug.augmentables.polys.PolygonsOnImage method*), 485
clip_to_dtype_value_range_() (*in module imgaug.dtypes*), 393
ClipCBAsToImagePlanes (*class in imgaug.augmenters.meta*), 884
close() (*imgaug.multicore.Pool method*), 392
CloudLayer (*class in imgaug.augmenters.weather*), 1066
Clouds (*class in imgaug.augmenters.weather*), 1070
CoarseDropout (*class in imgaug.augmenters.arithmetic*), 511
CoarsePepper (*class in imgaug.augmenters.arithmetic*), 514
CoarseSalt (*class in imgaug.augmenters.arithmetic*), 517
CoarseSaltAndPepper (*class in imgaug.augmenters.arithmetic*), 520
colorize() (*imgaug.augmenters.edges.IBinaryImageColorizer method*), 749
colorize() (*imgaug.augmenters.edges.RandomColorsBinaryImageColo method*), 750
COLORSPACES (*imgaug.augmenters.color.ChangeColorspace attribute*), 662
compress_jpeg() (*in module imgaug.augmenters.arithmetic*), 560
compute_croppings_to_reach_aspect_ratio() (*in module imgaug.augmenters.size*), 1060
compute_croppings_to_reach_multiples_of() (*in module imgaug.augmenters.size*), 1060
compute_croppings_to_reach_powers_of() (*in module imgaug.augmenters.size*), 1061
compute_distance() (*imgaug.augmentables.lines.LineString method*), 454
compute_geometric_median() (*in module imgaug.augmentables.kps*), 451
compute_geometric_median() (*in module imgaug.imgaug*), 348
compute_line_intersection_point() (*in module imgaug.imgaug*), 348
compute_neighbour_distances() (*imgaug.augmentables.lines.LineString method*), 454
compute_out_of_image_area() (*imgaug.augmentables.bbs.BoundingBox method*), 418
compute_out_of_image_area() (*imgaug.augmentables.polys.Polygon method*), 475
compute_out_of_image_fraction() (*imgaug.augmentables.bbs.BoundingBox method*), 419
compute_out_of_image_fraction() (*imgaug.augmentables.kps.Keypoint method*), 440
compute_out_of_image_fraction() (*imgaug.augmentables.lines.LineString method*), 454

compute_out_of_image_fraction() (*imgaug.augmentables.polys.Polygon method*), 476
compute_paddings_for_aspect_ratio() (*in module imgaug.imgaug*), 349
compute_paddings_to_reach_aspect_ratio() (*in module imgaug.augmenters.size*), 1061
compute_paddings_to_reach_exponents_of() (*in module imgaug.imgaug*), 349
compute_paddings_to_reach_multiples_of() (*in module imgaug.augmenters.size*), 1062
compute_paddings_to_reach_multiples_of() (*in module imgaug.imgaug*), 349
compute_paddings_to_reach_powers_of() (*in module imgaug.augmenters.size*), 1062
compute_pointwise_distances() (*imgaug.augmentables.lines.LineString method*), 454
concatenate() (*imgaug.augmentables.lines.LineString method*), 455
contains() (*imgaug.augmentables.bbs.BoundingBox method*), 419
contains() (*imgaug.augmentables.lines.LineString method*), 455
Contrast (*class in imgaug.augmenters.imgcorruptlike*), 814
ContrastNormalization() (*in module imgaug.augmenters.arithmetic*), 523
convert_cbaois_to_kpsois() (*in module imgaug.augmentables.utils*), 494
convert_iterable_to_string_of_types() (*in module imgaug.validation*), 408
convert_seed_sequence_to_generator() (*in module imgaug.random*), 403
convert_seed_to_generator() (*in module imgaug.random*), 404
Convolve (*class in imgaug.augmenters.convolutional*), 729
coords (*imgaug.augmentables.bbs.BoundingBox attribute*), 419
coords (*imgaug.augmentables.kps.Keypoint attribute*), 441
coords (*imgaug.augmentables.polys.Polygon attribute*), 476
coords_almost_equals() (*imgaug.augmentables.bbs.BoundingBox method*), 419
coords_almost_equals() (*imgaug.augmentables.kps.Keypoint method*), 441
coords_almost_equals() (*imgaug.augmentables.lines.LineString method*), 455
coords_almost_equals() (*imgaug.augmentables.polys.Polygon method*), 476
copy() (*imgaug.augmentables.bbs.BoundingBox method*), 420
copy() (*imgaug.augmentables.bbs.BoundingBoxesOnImage method*), 429
copy() (*imgaug.augmentables.heatmaps.HeatmapsOnImage method*), 436
copy() (*imgaug.augmentables.kps.Keypoint method*), 441
copy() (*imgaug.augmentables.kps.KeypointsOnImage method*), 446
copy() (*imgaug.augmentables.lines.LineString method*), 455
copy() (*imgaug.augmentables.lines.LineStringsOnImage method*), 466
copy() (*imgaug.augmentables.polys.Polygon method*), 476
copy() (*imgaug.augmentables.polys.PolygonsOnImage method*), 485
copy() (*imgaug.augmentables.segmaps.SegmentationMapsOnImage method*), 491
copy() (*imgaug.augmenters.meta.Augmenter method*), 873
copy() (*imgaug.parameters.StochasticParameter method*), 384
copy() (*imgaug.random.RNG method*), 398
copy_arrays() (*in module imgaug.augmenters.meta*), 911
copy_augmentables() (*in module imgaug.augmentables.utils*), 495
copy_dtypes_for_restore() (*in module imgaug.dtypes*), 393
copy_generator() (*in module imgaug.random*), 404
copy_generator_unless_global_generator() (*in module imgaug.random*), 404
copy_random_state() (*imgaug.augmenters.meta.Augmenter method*), 874
copy_random_state() (*in module imgaug.imgaug*), 349
copy_random_state_() (*imgaug.augmenters.meta.Augmenter method*), 874
copy_unless_global_rng() (*imgaug.random.RNG method*), 398
count_workers_alive() (*imgaug.multicore.BatchLoader method*), 391
create_for_noise() (*imgaug.parameters.Sigmoid static method*), 382
create_fully_random() (*imgaug.random.RNG class method*), 398
create_fully_random_generator() (*in mod-*

ule imgaug.random), 404
create_pseudo_random_() (imgaug.random.RNG class method), 398
create_pseudo_random_generator_() (in module imgaug.random), 404
Crop (class in imgaug.augmenters.size), 1018
CropAndPad (class in imgaug.augmenters.size), 1021
CropToAspectRatio (class in imgaug.augmenters.size), 1026
CropToFixedSize (class in imgaug.augmenters.size), 1028
CropToMultiplesOf (class in imgaug.augmenters.size), 1031
CropToPowersOf (class in imgaug.augmenters.size), 1034
CropToSquare (class in imgaug.augmenters.size), 1036
current_random_state() (in module imgaug.imgaug), 349
cut_out_of_image() (imgaug.augmentables.bbs.BoundingBox method), 420
cut_out_of_image() (imgaug.augmentables.bbs.BoundingBoxesOnImage method), 429
cut_out_of_image() (imgaug.augmentables.polys.Polygon method), 477
Cutout (class in imgaug.augmenters.arithmetic), 524
cutout() (in module imgaug.augmenters.arithmetic), 561
cutout_() (in module imgaug.augmenters.arithmetic), 562
CV_VARS (imgaug.augmenters.color.ChangeColorspace attribute), 662

D

deepcopy() (imgaug.augmentables.batches.Batch method), 410
deepcopy() (imgaug.augmentables.bbs.BoundingBox method), 420
deepcopy() (imgaug.augmentables.bbs.BoundingBoxesOnImage method), 429
deepcopy() (imgaug.augmentables.heatmaps.HeatmapsOnImage method), 436
deepcopy() (imgaug.augmentables.kps.Keypoint method), 441
deepcopy() (imgaug.augmentables.kps.KeypointsOnImage method), 446
deepcopy() (imgaug.augmentables.lines.LineString method), 456
deepcopy() (imgaug.augmentables.lines.LineStringsOnImage method), 466
deepcopy() (imgaug.augmentables.polys.Polygon method), 477
deepcopy() (imgaug.augmentables.polys.PolygonsOnImage method), 485
deepcopy() (imgaug.augmentables.segmaps.SegmentationMapsOnImage method), 492
deepcopy() (imgaug.augmenters.meta.Augmenter method), 874
deepcopy() (imgaug.parameters.StochasticParameter method), 384
deepcopy_fast() (in module imgaug.augmentables.utils), 495
DEFAULT_SEGMENT_COLORS (imgaug.augmentables.segmaps.SegmentationMapsOnImage attribute), 491
DefocusBlur (class in imgaug.augmenters.imgcorruptlike), 816
deprecated (class in imgaug.imgaug), 349
DeprecationWarning, 343
derive_generator_() (in module imgaug.random), 404
derive_generators_() (in module imgaug.random), 405
derive_random_state() (in module imgaug.imgaug), 350
derive_random_states() (in module imgaug.imgaug), 350
derive_rng_() (imgaug.random.RNG method), 398
derive_rngs_() (imgaug.random.RNG method), 399
Deterministic (class in imgaug.parameters), 367
DeterministicList (class in imgaug.parameters), 368
DirectedEdgeDetect (class in imgaug.augmenters.convolutional), 732
dirichlet() (imgaug.random.RNG method), 399
DiscreteUniform (class in imgaug.parameters), 368
Discretize (class in imgaug.parameters), 369
Divide (class in imgaug.parameters), 370
do_assert() (in module imgaug.imgaug), 350
draw() (imgaug.augmentables.heatmaps.HeatmapsOnImage method), 436
draw() (imgaug.augmentables.segmaps.SegmentationMapsOnImage method), 492
draw_box_on_image() (imgaug.augmentables.bbs.BoundingBox method), 420
draw_debug_image() (in module imgaug.augmenters.debug), 743
draw_distribution_graph() (imgaug.parameters.StochasticParameter method), 384
draw_distributions_grid() (in module imgaug.parameters), 389
draw_grid() (imgaug.augmenters.meta.Augmenter

method), 875
draw_grid() (*in module imgaug.imgaug*), 350
draw_heatmap_array() (*imgaug.augmentables.lines.LineString method*), 456
draw_label_on_image() (*imgaug.augmentables.bbs.BoundingBox method*), 421
draw_lines_heatmap_array() (*imgaug.augmentables.lines.LineString method*), 456
draw_lines_on_image() (*imgaug.augmentables.lines.LineString method*), 457
draw_mask() (*imgaug.augmentables.lines.LineString method*), 457
draw_masks() (*imgaug.augmenters.blend.BoundingBoxesMaskGen method*), 613
draw_masks() (*imgaug.augmenters.blend.CheckerboardMaskGen method*), 614
draw_masks() (*imgaug.augmenters.blend.IBatchwiseMaskGenerator method*), 617
draw_masks() (*imgaug.augmenters.blend.InvertMaskGen method*), 617
draw_masks() (*imgaug.augmenters.blend.RegularGridMaskGen method*), 618
draw_masks() (*imgaug.augmenters.blend.SegMapClassIdsMaskGen method*), 620
draw_masks() (*imgaug.augmenters.blend.SomeColorsMaskGen method*), 622
draw_masks() (*imgaug.augmenters.blend.StochasticParameterMaskGen method*), 623
draw_on_image() (*imgaug.augmentables.bbs.BoundingBox method*), 422
draw_on_image() (*imgaug.augmentables.bbs.BoundingBoxesOnImage method*), 429
draw_on_image() (*imgaug.augmentables.heatmaps.HeatmapsOnImage method*), 436
draw_on_image() (*imgaug.augmentables.kps.Keypoint method*), 442
draw_on_image() (*imgaug.augmentables.kps.KeypointsOnImage method*), 446
draw_on_image() (*imgaug.augmentables.lines.LineString method*), 457
draw_on_image() (*imgaug.augmentables.lines.LineStringsOnImage method*), 467
draw_on_image() (*imgaug.augmentables.polys.Polygon method*), 477
draw_on_image() (*imgaug.augmentables.polys.PolygonsOnImage method*), 486
draw_on_image() (*imgaug.augmentables.segmaps.SegmentationMapsOnImage method*), 492
draw_on_image() (*imgaug.augmenters.weather.CloudLayer method*), 1069
draw_on_image() (*imgaug.augmenters.weather.SnowflakesLayer method*), 1091
draw_points_heatmap_array() (*imgaug.augmentables.lines.LineString method*), 458
draw_points_on_image() (*imgaug.augmentables.lines.LineString method*), 458
draw_sample() (*imgaug.parameters.StochasticParameter method*), 385
draw_samples() (*imgaug.parameters.StochasticParameter method*), 385
duplicate() (*in module imgaug.imgaug*), 351
Dropout (*class in imgaug.augmenters.arithmetic*), 528
Dropout2d (*class in imgaug.augmenters.arithmetic*), 530
DropoutPointsSampler (*class in imgaug.augmenters.segmentation*), 974
dummy_random_state() (*in module imgaug.imgaug*), 352
duplicate() (*imgaug.random.RNG method*), 399

E

EdgeDetect (*class in imgaug.augmenters.convolutional*), 734
ElasticTransform (*class in imgaug.augmenters.imgcorruptlike*), 818
ElasticTransformation (*class in imgaug.augmenters.geometric*), 772
Emboss (*class in imgaug.augmenters.convolutional*), 736
empty (*imgaug.augmentables.bbs.BoundingBoxesOnImage attribute*), 430
empty (*imgaug.augmentables.kps.KeypointsOnImage attribute*), 447
empty (*imgaug.augmentables.lines.LineStringsOnImage attribute*), 467
empty (*imgaug.augmentables.polys.PolygonsOnImage attribute*), 487

enhance_brightness() (*in module imgaug.augmenters.pillike*), 950
enhance_color() (*in module imgaug.augmenters.pillike*), 951
enhance_contrast() (*in module imgaug.augmenters.pillike*), 951
enhance_sharpness() (*in module imgaug.augmenters.pillike*), 952
EnhanceBrightness (*class in imgaug.augmenters.pillike*), 918
EnhanceColor (*class in imgaug.augmenters.pillike*), 920
EnhanceContrast (*class in imgaug.augmenters.pillike*), 922
EnhanceSharpness (*class in imgaug.augmenters.pillike*), 924
Equalize (*class in imgaug.augmenters.pillike*), 926
equalize() (*in module imgaug.augmenters.pillike*), 953
equalize_() (*in module imgaug.augmenters.pillike*), 953
equals() (*imgaug.random.RNG method*), 399
equals_global_rng() (*imgaug.random.RNG method*), 399
estimate_bounding_boxes_norm_type() (*in module imgaug.augmentables.normalization*), 471
estimate_heatmaps_norm_type() (*in module imgaug.augmentables.normalization*), 471
estimate_keypoints_norm_type() (*in module imgaug.augmentables.normalization*), 471
estimate_line_strings_norm_type() (*in module imgaug.augmentables.normalization*), 471
estimate_max_number_of_channels() (*in module imgaug.augmenters.meta*), 911
estimate_normalization_type() (*in module imgaug.augmentables.normalization*), 471
estimate_polygons_norm_type() (*in module imgaug.augmentables.normalization*), 471
estimate_segmaps_norm_type() (*in module imgaug.augmentables.normalization*), 471
exponential() (*imgaug.random.RNG method*), 399
extend() (*imgaug.augmentables.bbs.BoundingBox method*), 422
extend_() (*imgaug.augmentables.bbs.BoundingBox method*), 422
exterior_almost_equals() (*imgaug.augmentables.polys.Polygon method*), 478
extract_from_image() (*imgaug.augmentables.bbs.BoundingBox method*), 423
extract_from_image() (*imgaug.augmentables.lines.LineString method*), 459
extract_from_image() (*imgaug.augmentables.polys.Polygon method*), 478

F

f() (*imgaug.random.RNG method*), 399
factor (*imgaug.augmenters.blend.BlendAlphaElementwise attribute*), 585
FastSnowyLandscape (*class in imgaug.augmenters.weather*), 1072
fill_from_augmented_normalized_batch() (*imgaug.augmentables.batches.UnnormalizedBatch method*), 415
fill_from_augmented_normalized_batch_() (*imgaug.augmentables.batches.UnnormalizedBatch method*), 415
fill_from_batch_in_augmentation_() (*imgaug.augmentables.batches.Batch method*), 412
fill_from_xy_array_() (*imgaug.augmentables.bbs.BoundingBoxesOnImage method*), 430
fill_from_xy_array_() (*imgaug.augmentables.kps.KeypointsOnImage method*), 447
fill_from_xy_array_() (*imgaug.augmentables.lines.LineStringsOnImage method*), 468
fill_from_xy_array_() (*imgaug.augmentables.polys.PolygonsOnImage method*), 487
fill_from_xyxy_array_() (*imgaug.augmentables.bbs.BoundingBoxesOnImage method*), 430
filter_blur() (*in module imgaug.augmenters.pillike*), 954
filter_contour() (*in module imgaug.augmenters.pillike*), 954
filter_detail() (*in module imgaug.augmenters.pillike*), 955
filter_edge_enhance() (*in module imgaug.augmenters.pillike*), 955
filter_edge_enhance_more() (*in module imgaug.augmenters.pillike*), 956
filter_emboss() (*in module imgaug.augmenters.pillike*), 957
filter_find_edges() (*in module imgaug.augmenters.pillike*), 957
filter_sharpen() (*in module imgaug.augmenters.pillike*), 958
filter_smooth() (*in module imgaug.augmenters.pillike*), 958

filter_smooth_more() (*in module imgaug.augmenters.pillike*), 959
FilterBlur (*class in imgaug.augmenters.pillike*), 928
FilterContour (*class in imgaug.augmenters.pillike*), 930
FilterDetail (*class in imgaug.augmenters.pillike*), 931
FilterEdgeEnhance (*class in imgaug.augmenters.pillike*), 933
FilterEdgeEnhanceMore (*class in imgaug.augmenters.pillike*), 935
FilterEmboss (*class in imgaug.augmenters.pillike*), 937
FilterFindEdges (*class in imgaug.augmenters.pillike*), 939
FilterSharpen (*class in imgaug.augmenters.pillike*), 940
FilterSmooth (*class in imgaug.augmenters.pillike*), 942
FilterSmoothMore (*class in imgaug.augmenters.pillike*), 944
find_augmenters() (*imgaug.augmenters.meta.Augmenter method*), 875
find_augmenters_by_name() (*imgaug.augmenters.meta.Augmenter method*), 876
find_augmenters_by_names() (*imgaug.augmenters.meta.Augmenter method*), 876
find_closest_point_index() (*imgaug.augmentables.polys.Polygon method*), 479
find_first_nonempty() (*in module imgaug.augmentables.normalization*), 471
find_intersections_with() (*imgaug.augmentables.lines.LineString method*), 460
flatten() (*in module imgaug.imgaug*), 352
Fliplr (*class in imgaug.augmenters.flip*), 750
fliplr() (*in module imgaug.augmenters.flip*), 754
Flipud (*class in imgaug.augmenters.flip*), 752
flipud() (*in module imgaug.augmenters.flip*), 755
Fog (*class in imgaug.augmenters.imgcorruptlike*), 820
Fog (*class in imgaug.augmenters.weather*), 1075
force_np_float_dtype() (*in module imgaug.parameters*), 389
ForceSign (*class in imgaug.parameters*), 371
forward_random_state() (*in module imgaug.imgaug*), 352
FrequencyNoise (*class in imgaug.parameters*), 371
FrequencyNoiseAlpha() (*in module imgaug.augmenters.blend*), 615
from_0to1() (*imgaug.augmentables.heatmaps.HeatmapsOnImage method*), 1070
static method), 437
from_coords_array() (*imgaug.augmentables.kps.KeypointsOnImage static method*), 447
from_distance_maps() (*imgaug.augmentables.kps.KeypointsOnImage static method*), 447
from_keypoint_image() (*imgaug.augmentables.kps.KeypointsOnImage static method*), 448
from_point_soup() (*imgaug.augmentables.bbs.BoundingBox class method*), 423
from_point_soups() (*imgaug.augmentables.bbs.BoundingBoxesOnImage class method*), 430
from_shapely() (*imgaug.augmentables.polys.MultiPolygon static method*), 472
from_shapely() (*imgaug.augmentables.polys.Polygon static method*), 479
from_uint8() (*imgaug.augmentables.heatmaps.HeatmapsOnImage static method*), 437
from_xy_array() (*imgaug.augmentables.kps.KeypointsOnImage class method*), 448
from_xy_arrays() (*imgaug.augmentables.lines.LineStringsOnImage class method*), 468
from_xyxy_array() (*imgaug.augmentables.bbs.BoundingBoxesOnImage class method*), 431
FromLowerResolution (*class in imgaug.parameters*), 373
Frost (*class in imgaug.augmenters.imgcorruptlike*), 821

G

gamma() (*imgaug.random.RNG method*), 399
GammaContrast (*class in imgaug.augmenters.contrast*), 713
gate_dtypes() (*in module imgaug.dtypes*), 393
GaussianBlur (*class in imgaug.augmenters.blur*), 632
GaussianBlur (*class in imgaug.augmenters.imgcorruptlike*), 823
GaussianNoise (*class in imgaug.augmenters.imgcorruptlike*), 825
generate_jigsaw_destinations() (*in module imgaug.augmenters.geometric*), 810
generate_maps() (*imgaug.augmenters.weather.CloudLayer*

generate_mask() (imgaug.augmenters.blend.BoundingBoxesMaskGen class method), 613
generate_mask() (imgaug.augmenters.blend.CheckerboardMaskGen class method), 614
generate_mask() (imgaug.augmenters.blend.HorizontalLinearGradientMaskGen class method), 616
generate_mask() (imgaug.augmenters.blend.RegularGridMaskGen class method), 618
generate_mask() (imgaug.augmenters.blend.SegMapClassIdsMaskGen class method), 620
generate_mask() (imgaug.augmenters.blend.SomeColorsMaskGen class method), 622
generate_mask() (imgaug.augmenters.blend.VerticalLinearGradientMaskGen class method), 624
generate_seed_() (imgaug.random.RNG method), 399
generate_seed_() (in module imgaug.random), 405
generate_seeds_() (imgaug.random.RNG method), 400
generate_seeds_() (in module imgaug.random), 405
generate_similar_points_manhattan() (imgaug.augmentables.kps.Keypoint method), 442
geometric() (imgaug.random.RNG method), 400
get_all_children() (imgaug.augmenters.meta.Augmenter method), 876
get_arr() (imgaug.augmentables.heatmaps.HeatmapsOnImage method), 437
get_arr() (imgaug.augmentables.segmaps.SegmentationMapsOnImage method), 493
get_arr_int() (imgaug.augmentables.segmaps.SegmentationMapsOnImage method), 493
get_batch() (imgaug.multicore.BackgroundAugmenter method), 390
get_children_lists() (imgaug.augmenters.blend.BlendAlpha method), 576
get_children_lists() (imgaug.augmenters.blend.BlendAlphaMask method), 595
get_children_lists() (imgaug.augmenters.color.WithBrightnessChannels method), 692
get_children_lists() (imgaug.augmenters.color.WithColorspace method), 694
get_children_lists() (imgaug.augmenters.color.WithHueAndSaturation method), 697
get_children_lists() (imgaug.augmenters.geometric.WithPolarWarping method), 809
get_children_lists() (imgaug.augmenters.meta.Augmenter method), 876
get_children_lists() (imgaug.augmenters.meta.Sequential method), 902
get_children_lists() (imgaug.augmenters.meta.SomeOf method), 906
get_children_lists() (imgaug.augmenters.meta.Sometimes method), 908
get_children_lists() (imgaug.augmenters.meta.WithChannels method), 910
get_children_lists() (imgaug.augmenters.size.KeepSizeByResize method), 1041
get_column_names() (imgaug.augmentables.batches.Batch method), 412
get_column_names() (imgaug.augmentables.batches.UnnormalizedBatch method), 415
get_coords_array() (imgaug.augmentables.kps.KeypointsOnImage method), 449
get_corruption_names() (in module imgaug.augmenters.imgcorruptlike), 855
get_generator_state() (in module imgaug.random), 405
get_global_rng() (in module imgaug.random), 405
get_minimal_dtype() (in module imgaug.dtypes), 393
get_parameters() (imgaug.augmenters.arithmetic.Add method), 500
get_parameters() (imgaug.augmenters.arithmetic.AddElementwise method), 502
get_parameters() (imgaug.augmenters.arithmetic.Cutout method), 528
get_parameters() (imgaug.augmenters.arithmetic.Dropout2d method), 533
get_parameters() (im-

gaug.augmenters.arithmetic.Invert method), 538
get_parameters() (*imgaug.augmenters.arithmetic.JpegCompression method*), 540
get_parameters() (*imgaug.augmenters.arithmetic.Multiply method*), 543
get_parameters() (*imgaug.augmenters.arithmetic.MultiplyElementwise method*), 545
get_parameters() (*imgaug.augmenters.arithmetic.ReplaceElementwise method*), 550
get_parameters() (*imgaug.augmenters.arithmetic.TotalDropout method*), 559
get_parameters() (*imgaug.augmenters.artistic.Cartoon method*), 571
get_parameters() (*imgaug.augmenters.blend.BlendAlpha method*), 576
get_parameters() (*imgaug.augmenters.blend.BlendAlphaMask method*), 595
get_parameters() (*imgaug.augmenters.blur.AverageBlur method*), 629
get_parameters() (*imgaug.augmenters.blur.BilateralBlur method*), 632
get_parameters() (*imgaug.augmenters.blur.GaussianBlur method*), 634
get_parameters() (*imgaug.augmenters.blur.MeanShiftBlur method*), 636
get_parameters() (*imgaug.augmenters.blur.MedianBlur method*), 639
get_parameters() (*imgaug.augmenters.collections.RandAugment method*), 648
get_parameters() (*imgaug.augmenters.color.AddToHueAndSaturation method*), 656
get_parameters() (*imgaug.augmenters.color.ChangeColorspace method*), 662
get_parameters() (*imgaug.augmenters.color.ChangeColorTemperature method*), 660
get_parameters() (*imgaug.augmenters.color.WithBrightnessChannels method*), 692
get_parameters() (*imgaug.augmenters.color.WithColorspace method*), 694
get_parameters() (*imgaug.augmenters.color.WithHueAndSaturation method*), 697
get_parameters() (*imgaug.augmenters.contrast.AllChannelsCLAHE method*), 707
get_parameters() (*imgaug.augmenters.contrast.AllChannelsHistogramEqualization method*), 709
get_parameters() (*imgaug.augmenters.contrast.CLAHE method*), 713
get_parameters() (*imgaug.augmenters.contrast.HistogramEqualization method*), 718
get_parameters() (*imgaug.augmenters.convolutional.Convolve method*), 732
get_parameters() (*imgaug.augmenters.debug.SaveDebugImageEveryNBatches method*), 743
get_parameters() (*imgaug.augmenters.edges.Canny method*), 749
get_parameters() (*imgaug.augmenters.flip.Fliplr method*), 752
get_parameters() (*imgaug.augmenters.flip.Flipud method*), 754
get_parameters() (*imgaug.augmenters.geometric.Affine method*), 766
get_parameters() (*imgaug.augmenters.geometric.AffineCv2 method*), 772
get_parameters() (*imgaug.augmenters.geometric.ElasticTransformation method*), 777
get_parameters() (*imgaug.augmenters.geometric.Jigsaw method*), 780
get_parameters() (*imgaug.augmenters.geometric.PerspectiveTransform method*), 784
get_parameters() (*imgaug.augmenters.geometric.PiecewiseAffine method*), 787
get_parameters() (*imgaug.augmenters.geometric.Rot90 method*), 790

get_parameters() (imgaug.augmenters.geometric.WithPolarWarping method), 809
get_parameters() (imgaug.augmenters.meta.Augmenter method), 877
get_parameters() (imgaug.augmenters.meta.ChannelShuffle method), 884
get_parameters() (imgaug.augmenters.meta.ClipCBAsToImagePlanes method), 886
get_parameters() (imgaug.augmenters.meta.Identity method), 888
get_parameters() (imgaug.augmenters.meta.Lambda method), 892
get_parameters() (imgaug.augmenters.meta.RemoveCBAsByOutOfImageFraction method), 899
get_parameters() (imgaug.augmenters.meta.Sequential method), 902
get_parameters() (imgaug.augmenters.meta.SomeOf method), 906
get_parameters() (imgaug.augmenters.meta.Sometimes method), 908
get_parameters() (imgaug.augmenters.meta.WithChannels method), 910
get_parameters() (imgaug.augmenters.pillike.Affine method), 916
get_parameters() (imgaug.augmenters.pillike.Equalize method), 928
get_parameters() (imgaug.augmenters.segmentation.Superpixels method), 990
get_parameters() (imgaug.augmenters.segmentation.Voronoi method), 997
get_parameters() (imgaug.augmenters.size.CropAndPad method), 1026
get_parameters() (imgaug.augmenters.size.CropToAspectRatio method), 1028
get_parameters() (imgaug.augmenters.size.CropToFixedSize method), 1031
get_parameters() (imgaug.augmenters.size.CropToMultiplesOf method), 1034
get_parameters() (imgaug.augmenters.size.CropToPowersOf method), 1036
get_parameters() (imgaug.augmenters.size.KeepSizeByResize method), 1041
get_parameters() (imgaug.augmenters.size.PadToAspectRatio method), 1047
get_parameters() (imgaug.augmenters.size.PadToFixedSize method), 1050
get_parameters() (imgaug.augmenters.size.PadToMultiplesOf method), 1052
get_parameters() (imgaug.augmenters.size.PadToPowersOf method), 1054
get_parameters() (imgaug.augmenters.size.Resize method), 1060
get_parameters() (imgaug.augmenters.weather.CloudLayer method), 1070
get_parameters() (imgaug.augmenters.weather.FastSnowyLandscape method), 1075
get_parameters() (imgaug.augmenters.weather.SnowflakesLayer method), 1091
get_pointwise_inside_image_mask() (imgaug.augmentables.lines.LineString method), 460
get_value_range_of_dtype() (in module imgaug.dtypes), 393
GlassBlur (class in imgaug.augmenters.imgcorruptlike), 827
GRAY (imgaug.augmenters.color.ChangeColorspace attribute), 662
Grayscale (class in imgaug.augmenters.color), 662
gumbel() (imgaug.random.RNG method), 400

H

handle_categorical_string_param() (in module imgaug.parameters), 389
handle_children_list() (in module imgaug.augmenters.meta), 911
handle_continuous_param() (in module imgaug.parameters), 389
handle_discrete_kernel_size_param() (in module imgaug.parameters), 389

handle_discrete_param() (*in module imgaug.parameters*), 389
handle_probability_param() (*in module imgaug.parameters*), 389
heatmaps (*imgaug.augmentables.batches.Batch attribute*), 412
HeatmapsOnImage (*class in imgaug.augmentables.heatmaps*), 434
HeatmapsOnImage() (*in module imgaug.imgaug*), 343
height (*imgaug.augmentables.bbs.BoundingBox attribute*), 423
height (*imgaug.augmentables.bbs.BoundingBoxesOnImage attribute*), 431
height (*imgaug.augmentables.kps.KeypointsOnImage attribute*), 449
height (*imgaug.augmentables.lines.LineString attribute*), 460
height (*imgaug.augmentables.polys.Polygon attribute*), 479
HistogramEqualization (*class in imgaug.augmenters.contrast*), 715
HLS (*imgaug.augmenters.color.ChangeColorspace attribute*), 662
HLS (*imgaug.augmenters.contrast.CLAHE attribute*), 713
HLS (*imgaug.augmenters.contrast.HistogramEqualization attribute*), 718
HooksHeatmaps (*class in imgaug.imgaug*), 343
HooksImages (*class in imgaug.imgaug*), 344
HooksKeypoints (*class in imgaug.imgaug*), 346
HorizontalFlip() (*in module imgaug.augmenters.flip*), 754
HorizontalLinearGradientMaskGen (*class in imgaug.augmenters.blend*), 615
HSV (*imgaug.augmenters.color.ChangeColorspace attribute*), 662
HSV (*imgaug.augmenters.contrast.CLAHE attribute*), 713
HSV (*imgaug.augmenters.contrast.HistogramEqualization attribute*), 718
hypergeometric() (*imgaug.random.RNG method*), 400

I

IAugmentable (*class in imgaug.augmentables.base*), 409
IBatchwiseMaskGenerator (*class in imgaug.augmenters.blend*), 616
IBinaryImageColorizer (*class in imgaug.augmenters.edges*), 749
Identity (*class in imgaug.augmenters.meta*), 886
images (*imgaug.augmentables.batches.Batch attribute*), 412
imap_batches() (*imgaug.multicore.Pool method*), 392
imap_batches_unordered() (*imgaug.multicore.Pool method*), 392
imgaug.augmentables.base (*module*), 409
imgaug.augmentables.batches (*module*), 409
imgaug.augmentables.bbs (*module*), 416
imgaug.augmentables.heatmaps (*module*), 434
imgaug.augmentables.kps (*module*), 439
imgaug.augmentables.lines (*module*), 452
imgaug.augmentables.normalization (*module*), 471
imgaug.augmentables.polys (*module*), 472
imgaug.augmentables.segmaps (*module*), 491
imgaug.augmentables.utils (*module*), 494
imgaug.augmenters.arithmetic (*module*), 497
imgaug.augmenters.artistic (*module*), 568
imgaug.augmenters.base (*module*), 572
imgaug.augmenters.blend (*module*), 572
imgaug.augmenters.blur (*module*), 626
imgaug.augmenters.collections (*module*), 645
imgaug.augmenters.color (*module*), 648
imgaug.augmenters.contrast (*module*), 703
imgaug.augmenters.convolutional (*module*), 729
imgaug.augmenters.debug (*module*), 741
imgaug.augmenters.edges (*module*), 745
imgaug.augmenters.flip (*module*), 750
imgaug.augmenters.geometric (*module*), 756
imgaug.augmenters.imgcorruptlike (*module*), 810
imgaug.augmenters.meta (*module*), 855
imgaug.augmenters.pillike (*module*), 912
imgaug.augmenters.pooling (*module*), 962
imgaug.augmenters.segmentation (*module*), 974
imgaug.augmenters.size (*module*), 998
imgaug.augmenters.weather (*module*), 1065
imgaug.dtypes (*module*), 393
imgaug.imgaug (*module*), 343
imgaug.multicore (*module*), 389
imgaug.parameters (*module*), 362
imgaug.random (*module*), 394
imgaug.validation (*module*), 408
ImpulseNoise (*class in imgaug.augmenters.arithmetic*), 533
ImpulseNoise (*class in imgaug.augmenters.imgcorruptlike*), 829
imresize_many_images() (*in module imgaug.imgaug*), 352
imresize_single_image() (*in module imgaug.imgaug*), 354
imshow() (*in module imgaug.imgaug*), 354

InColorspace() (*in module imgaug.augmenters.color*), 665
increase_array_resolutions_() (*in module imgaug.dtypes*), 393
increase_itemsize_of_dtype() (*in module imgaug.dtypes*), 394
integers() (*imgaug.random.RNG method*), 400
interpolate_point_pair() (*in module imgaug.augmentables.utils*), 495
interpolate_points() (*in module imgaug.augmentables.utils*), 495
interpolate_points_by_max_distance() (*in module imgaug.augmentables.utils*), 495
intersection() (*imgaug.augmentables.bbs.BoundingBox method*), 424
Invert (*class in imgaug.augmenters.arithmetic*), 535
invert() (*imgaug.augmentables.heatmaps.HeatmapsOnImage method*), 438
invert() (*in module imgaug.augmenters.arithmetic*), 562
invert_() (*in module imgaug.augmenters.arithmetic*), 563
invert_convert_cbaois_to_kpsois_() (*in module imgaug.augmentables.utils*), 496
invert_normalize_bounding_boxes() (*in module imgaug.augmentables.normalization*), 471
invert_normalize_heatmaps() (*in module imgaug.augmentables.normalization*), 471
invert_normalize_images() (*in module imgaug.augmentables.normalization*), 471
invert_normalize_keypoints() (*in module imgaug.augmentables.normalization*), 471
invert_normalize_line_strings() (*in module imgaug.augmentables.normalization*), 471
invert_normalize_polygons() (*in module imgaug.augmentables.normalization*), 471
invert_normalize_segmentation_maps() (*in module imgaug.augmentables.normalization*), 472
invert_reduce_to_nonempty() (*in module imgaug.augmenters.meta*), 911
invert_to_keypoints_on_image_() (*imgaug.augmentables.bbs.BoundingBoxesOnImage method*), 431
invert_to_keypoints_on_image_() (*imgaug.augmentables.kps.KeypointsOnImage method*), 449
invert_to_keypoints_on_image_() (*imgaug.augmentables.lines.LineStringsOnImage method*), 468
invert_to_keypoints_on_image_() (*imgaug.augmentables.polys.PolygonsOnImage method*), 487
InvertMaskGen (*class in imgaug.augmenters.blend*), 617
iou() (*imgaug.augmentables.bbs.BoundingBox method*), 424
IPointsSampler (*class in imgaug.augmenters.segmentation*), 975
is_activated() (*imgaug.imgaug.HooksImages method*), 345
is_callable() (*in module imgaug.imgaug*), 355
is_float_array() (*in module imgaug.imgaug*), 355
is_fully_within_image() (*imgaug.augmentables.bbs.BoundingBox method*), 424
is_fully_within_image() (*imgaug.augmentables.lines.LineString method*), 460
is_fully_within_image() (*imgaug.augmentables.polys.Polygon method*), 479
is_generator() (*in module imgaug.imgaug*), 355
is_generator_equal_to() (*in module imgaug.random*), 406
is_global_rng() (*imgaug.random.RNG method*), 400
is_integer_array() (*in module imgaug.imgaug*), 355
is_iterable() (*in module imgaug.imgaug*), 355
is_iterable_of() (*in module imgaug.validation*), 409
is_np_array() (*in module imgaug.imgaug*), 355
is_np_scalar() (*in module imgaug.imgaug*), 355
is_out_of_image() (*imgaug.augmentables.bbs.BoundingBox method*), 424
is_out_of_image() (*imgaug.augmentables.kps.Keypoint method*), 442
is_out_of_image() (*imgaug.augmentables.lines.LineString method*), 460
is_out_of_image() (*imgaug.augmentables.polys.Polygon method*), 479
is_partly_within_image() (*imgaug.augmentables.bbs.BoundingBox method*), 425
is_partly_within_image() (*imgaug.augmentables.lines.LineString method*), 461
is_partly_within_image() (*imgaug.augmentables.polys.Polygon method*), 480
is_propagating() (*imgaug.imgaug.HooksImages*

method), 345
is_single_bool() (*in module imgaug.imgaug*), 355
is_single_float() (*in module imgaug.imgaug*), 356
is_single_integer() (*in module imgaug.imgaug*), 356
is_single_number() (*in module imgaug.imgaug*), 356
is_string() (*in module imgaug.imgaug*), 356
is_valid (*imgaug.augmentables.polys.Polygon attribute*), 480
items (*imgaug.augmentables.bbs.BoundingBoxesOnImage attribute*), 431
items (*imgaug.augmentables.kps.KeypointsOnImage attribute*), 449
items (*imgaug.augmentables.lines.LineStringsOnImage attribute*), 468
items (*imgaug.augmentables.polys.PolygonsOnImage attribute*), 487
IterativeNoiseAggregator (*class in imgaug.parameters*), 374

J

Jigsaw (*class in imgaug.augmenters.geometric*), 777
join() (*imgaug.multicore.Pool method*), 392
JpegCompression (*class in imgaug.augmenters.arithmetic*), 538
JpegCompression (*class in imgaug.augmenters.imgcorruptlike*), 831

K

KeepSizeByResize (*class in imgaug.augmenters.size*), 1038
Keypoint (*class in imgaug.augmentables.kps*), 439
Keypoint() (*in module imgaug.imgaug*), 346
KEYPOINT_AUG_ALPHA_THRESH (*imgaug.augmenters.geometric.ElasticTransformation attribute*), 777
KEYPOINT_AUG_SIGMA_THRESH (*imgaug.augmenters.geometric.ElasticTransformation attribute*), 777
keypoints (*imgaug.augmentables.batches.Batch attribute*), 412
KeypointsOnImage (*class in imgaug.augmentables.kps*), 444
KeypointsOnImage() (*in module imgaug.imgaug*), 346
KMeansColorQuantization (*class in imgaug.augmenters.color*), 665

L

Lab (*imgaug.augmenters.color.ChangeColorspace attribute*), 662
Lab (*imgaug.augmenters.contrast.CLAHE attribute*), 713
Lab (*imgaug.augmenters.contrast.HistogramEqualization attribute*), 718
Lambda (*class in imgaug.augmenters.meta*), 888
Laplace (*class in imgaug.parameters*), 375
laplace() (*imgaug.random.RNG method*), 400
length (*imgaug.augmentables.lines.LineString attribute*), 461
LinearContrast (*class in imgaug.augmenters.contrast*), 718
LineString (*class in imgaug.augmentables.lines*), 452
LineStringsOnImage (*class in imgaug.augmentables.lines*), 464
localize_random_state() (*imgaug.augmenters.meta.Augmenter method*), 877
localize_random_state_() (*imgaug.augmenters.meta.Augmenter method*), 877
LogContrast (*class in imgaug.augmenters.contrast*), 721
logistic() (*imgaug.random.RNG method*), 400
lognormal() (*imgaug.random.RNG method*), 400
logseries() (*imgaug.random.RNG method*), 400
Luv (*imgaug.augmenters.color.ChangeColorspace attribute*), 662

M

map_batches() (*imgaug.multicore.Pool method*), 392
map_batches_async() (*imgaug.multicore.Pool method*), 393
max_pool() (*imgaug.augmentables.heatmaps.HeatmapsOnImage method*), 438
max_pool() (*in module imgaug.imgaug*), 356
MaxPooling (*class in imgaug.augmenters.pooling*), 965
MeanShiftBlur (*class in imgaug.augmenters.blur*), 634
median_pool() (*in module imgaug.imgaug*), 357
MedianBlur (*class in imgaug.augmenters.blur*), 636
MedianPooling (*class in imgaug.augmenters.pooling*), 968
min_pool() (*in module imgaug.imgaug*), 357
MinPooling (*class in imgaug.augmenters.pooling*), 971
MotionBlur (*class in imgaug.augmenters.blur*), 639
MotionBlur (*class in imgaug.augmenters.imgcorruptlike*), 833
multinomial() (*imgaug.random.RNG method*), 400
Multiply (*class in imgaug.augmenters.arithmetic*), 540
Multiply (*class in imgaug.parameters*), 376

multiply_elementwise() (*in module imgaug.augmenters.arithmetic*), 564
multiply_scalar() (*in module imgaug.augmenters.arithmetic*), 565
MultiplyAndAddToBrightness (*class in imgaug.augmenters.color*), 668
MultiplyBrightness (*class in imgaug.augmenters.color*), 670
MultiplyElementwise (*class in imgaug.augmenters.arithmetic*), 543
MultiplyHue (*class in imgaug.augmenters.color*), 672
MultiplyHueAndSaturation (*class in imgaug.augmenters.color*), 675
MultiplySaturation (*class in imgaug.augmenters.color*), 678
MultiPolygon (*class in imgaug.augmentables.polys*), 472
MultiPolygon() (*in module imgaug.imgaug*), 346
multivariate_normal() (*imgaug.random.RNG method*), 400

N

n_colors (*imgaug.augmenters.color.KMeansColorQuantization attribute*), 668
n_colors (*imgaug.augmenters.color.UniformColorQuantization attribute*), 686
nb_cols (*imgaug.augmenters.blend.CheckerboardMaskGen attribute*), 614
NB_NEIGHBOURING_KEYPOINTS (*imgaug.augmenters.geometric.ElasticTransformation attribute*), 777
nb_rows (*imgaug.augmenters.blend.CheckerboardMaskGen attribute*), 614
Negative() (*in module imgaug.parameters*), 377
negative_binomial() (*imgaug.random.RNG method*), 400
NEIGHBOURING_KEYPOINTS_DISTANCE (*imgaug.augmenters.geometric.ElasticTransformation attribute*), 777
new_random_state() (*in module imgaug.imgaug*), 358
NO_RESIZE (*imgaug.augmenters.size.KeepSizeByResize attribute*), 1041
noncentral_chisquare() (*imgaug.random.RNG method*), 400
noncentral_f() (*imgaug.random.RNG method*), 400
Noop (*class in imgaug.augmenters.meta*), 892
Normal (*class in imgaug.parameters*), 378
normal() (*imgaug.random.RNG method*), 401
normalize_bounding_boxes() (*in module imgaug.augmentables.normalization*), 472
normalize_dtype() (*in module imgaug.dtypes*), 394
normalize_dtypes() (*in module imgaug.dtypes*), 394
normalize_generator() (*in module imgaug.random*), 406
normalize_generator_() (*in module imgaug.random*), 406
normalize_heatmaps() (*in module imgaug.augmentables.normalization*), 472
normalize_images() (*in module imgaug.augmentables.normalization*), 472
normalize_keypoints() (*in module imgaug.augmentables.normalization*), 472
normalize_line_strings() (*in module imgaug.augmentables.normalization*), 472
normalize_polygons() (*in module imgaug.augmentables.normalization*), 472
normalize_random_state() (*in module imgaug.imgaug*), 358
normalize_segmentation_maps() (*in module imgaug.augmentables.normalization*), 472
normalize_shape() (*in module imgaug.augmentables.utils*), 496

O

on() (*imgaug.augmentables.bbs.BoundingBoxesOnImage method*), 431
on() (*imgaug.augmentables.kps.KeypointsOnImage method*), 449
on() (*imgaug.augmentables.lines.LineStringsOnImage method*), 468
on() (*imgaug.augmentables.polys.PolygonsOnImage method*), 487
on_() (*imgaug.augmentables.bbs.BoundingBoxesOnImage method*), 432
on_() (*imgaug.augmentables.kps.KeypointsOnImage method*), 449
on_() (*imgaug.augmentables.lines.LineStringsOnImage method*), 469
on_() (*imgaug.augmentables.polys.PolygonsOnImage method*), 488
OneOf (*class in imgaug.augmenters.meta*), 894

P

Pad (*class in imgaug.augmenters.size*), 1041
pad() (*imgaug.augmentables.heatmaps.HeatmapsOnImage method*), 438
pad() (*imgaug.augmentables.segmaps.SegmentationMapsOnImage method*), 493
pad() (*in module imgaug.augmenters.size*), 1063
pad() (*in module imgaug.imgaug*), 358
pad_to_aspect_ratio() (*imgaug.augmentables.heatmaps.HeatmapsOnImage method*), 438
pad_to_aspect_ratio() (*imgaug.augmentables.segmaps.SegmentationMapsOnImage method*), 494

pad_to_aspect_ratio() (*in module imgaug.augmenters.size*), 1064
pad_to_aspect_ratio() (*in module imgaug.imgaug*), 358
pad_to_multiples_of() (*in module imgaug.augmenters.size*), 1065
pad_to_multiples_of() (*in module imgaug.imgaug*), 358
PadToAspectRatio (*class in imgaug.augmenters.size*), 1045
PadToFixedSize (*class in imgaug.augmenters.size*), 1047
PadToMultiplesOf (*class in imgaug.augmenters.size*), 1050
PadToPowersOf (*class in imgaug.augmenters.size*), 1052
PadToSquare (*class in imgaug.augmenters.size*), 1054
pareto() (*imgaug.random.RNG method*), 401
Pepper (*class in imgaug.augmenters.arithmetic*), 545
permutation() (*imgaug.random.RNG method*), 401
PerspectiveTransform (*class in imgaug.augmenters.geometric*), 780
PiecewiseAffine (*class in imgaug.augmenters.geometric*), 784
Pixelate (*class in imgaug.augmenters.imgcorruptlike*), 835
Poisson (*class in imgaug.parameters*), 379
poisson() (*imgaug.random.RNG method*), 401
polyfill_integers() (*in module imgaug.random*), 407
polyfill_random() (*in module imgaug.random*), 407
Polygon (*class in imgaug.augmentables.polys*), 472
Polygon() (*in module imgaug.imgaug*), 346
PolygonsOnImage (*class in imgaug.augmentables.polys*), 483
PolygonsOnImage() (*in module imgaug.imgaug*), 346
Pool (*class in imgaug.multicore*), 391
pool (*imgaug.multicore.Pool attribute*), 393
pool() (*imgaug.augmenters.meta.Augmenter method*), 877
pool() (*in module imgaug.imgaug*), 358
Positive() (*in module imgaug.parameters*), 379
Posterize (*class in imgaug.augmenters.color*), 680
Posterize (*class in imgaug.augmenters.pillike*), 946
posterize() (*in module imgaug.augmenters.color*), 699
posterize() (*in module imgaug.augmenters.pillike*), 960
posterize_() (*in module imgaug.augmenters.pillike*), 960
postprocess() (*imgaug.imgaug.HooksImages method*), 345
Power (*class in imgaug.parameters*), 380
power() (*imgaug.random.RNG method*), 401
preprocess() (*imgaug.imgaug.HooksImages method*), 345
project() (*imgaug.augmentables.bbs.BoundingBox method*), 425
project() (*imgaug.augmentables.kps.Keypoint method*), 443
project() (*imgaug.augmentables.lines.LineString method*), 461
project() (*imgaug.augmentables.polys.Polygon method*), 480
project_() (*imgaug.augmentables.bbs.BoundingBox method*), 425
project_() (*imgaug.augmentables.kps.Keypoint method*), 443
project_() (*imgaug.augmentables.lines.LineString method*), 461
project_() (*imgaug.augmentables.polys.Polygon method*), 480
project_coords() (*in module imgaug.augmentables.utils*), 496
project_coords_() (*in module imgaug.augmentables.utils*), 497
promote_array_dtypes_() (*in module imgaug.dtypes*), 394

Q

quantize_colors_kmeans() (*in module imgaug.augmenters.color*), 700
quantize_colors_uniform() (*in module imgaug.augmenters.color*), 700
quantize_kmeans() (*in module imgaug.augmenters.color*), 700
quantize_uniform() (*in module imgaug.augmenters.color*), 701
quantize_uniform_() (*in module imgaug.augmenters.color*), 701
quantize_uniform_to_n_bits() (*in module imgaug.augmenters.color*), 702
quantize_uniform_to_n_bits_() (*in module imgaug.augmenters.color*), 703
quokka() (*in module imgaug.imgaug*), 359
quokka_bounding_boxes() (*in module imgaug.imgaug*), 360
quokka_heatmap() (*in module imgaug.imgaug*), 360
quokka_keypoints() (*in module imgaug.imgaug*), 360
quokka_polygons() (*in module imgaug.imgaug*), 360
quokka_segmentation_map() (*in module imgaug.imgaug*), 361
quokka_square() (*in module imgaug.imgaug*), 361

R

Rain (*class in imgaug.augmenters.weather*), 1077
RainLayer (*class in imgaug.augmenters.weather*), 1080
rand() (*imgaug.random.RNG method*), 401
RandAugment (*class in imgaug.augmenters.collections*), 645
randint() (*imgaug.random.RNG method*), 401
randn() (*imgaug.random.RNG method*), 401
random() (*imgaug.random.RNG method*), 401
random_integers() (*imgaug.random.RNG method*), 401
random_sample() (*imgaug.random.RNG method*), 402
RandomColorsBinaryImageColorizer (*class in imgaug.augmenters.edges*), 749
RandomSign (*class in imgaug.parameters*), 381
rayleigh() (*imgaug.random.RNG method*), 402
recover_psois_() (*in module imgaug.augmentables.polys*), 490
reduce_to_nonempty() (*in module imgaug.augmenters.meta*), 911
RegularGridMaskGen (*class in imgaug.augmenters.blend*), 617
RegularGridPointsSampler (*class in imgaug.augmenters.segmentation*), 976
RegularGridVoronoi (*class in imgaug.augmenters.segmentation*), 977
RelativeRegularGridPointsSampler (*class in imgaug.augmenters.segmentation*), 980
RelativeRegularGridVoronoi (*class in imgaug.augmenters.segmentation*), 982
remove_augmenters() (*imgaug.augmenters.meta.Augmenter method*), 878
remove_augmenters_() (*imgaug.augmenters.meta.Augmenter method*), 879
remove_augmenters_inplace() (*imgaug.augmenters.meta.Augmenter method*), 880
remove_out_of_image() (*imgaug.augmentables.bbs.BoundingBoxesOnImage method*), 432
remove_out_of_image() (*imgaug.augmentables.lines.LineStringsOnImage method*), 469
remove_out_of_image() (*imgaug.augmentables.polys.PolygonsOnImage method*), 488
remove_out_of_image_() (*imgaug.augmentables.bbs.BoundingBoxesOnImage method*), 432
remove_out_of_image_() (*imgaug.augmentables.lines.LineStringsOnImage method*), 469
remove_out_of_image_() (*imgaug.augmentables.polys.PolygonsOnImage method*), 488
remove_out_of_image_fraction() (*imgaug.augmentables.bbs.BoundingBoxesOnImage method*), 432
remove_out_of_image_fraction() (*imgaug.augmentables.kps.KeypointsOnImage method*), 449
remove_out_of_image_fraction() (*imgaug.augmentables.lines.LineStringsOnImage method*), 469
remove_out_of_image_fraction() (*imgaug.augmentables.polys.PolygonsOnImage method*), 488
remove_out_of_image_fraction_() (*imgaug.augmentables.bbs.BoundingBoxesOnImage method*), 433
remove_out_of_image_fraction_() (*imgaug.augmentables.kps.KeypointsOnImage method*), 450
remove_out_of_image_fraction_() (*imgaug.augmentables.lines.LineStringsOnImage method*), 469
remove_out_of_image_fraction_() (*imgaug.augmentables.polys.PolygonsOnImage method*), 489
RemoveCBAsByOutOfImageFraction (*class in imgaug.augmenters.meta*), 896
RemoveSaturation (*class in imgaug.augmenters.color*), 681
replace_elementwise_() (*in module imgaug.augmenters.arithmetic*), 566
ReplaceElementwise (*class in imgaug.augmenters.arithmetic*), 547
reseed() (*imgaug.augmenters.meta.Augmenter method*), 880
reset_cache_() (*imgaug.random.RNG method*), 402
reset_generator_cache_() (*in module imgaug.random*), 407
Resize (*class in imgaug.augmenters.size*), 1057
resize() (*imgaug.augmentables.heatmaps.HeatmapsOnImage method*), 439
resize() (*imgaug.augmentables.segmaps.SegmentationMapsOnImage method*), 494
restore_dtype_and_merge() (*in module imgaug.augmentables.normalization*), 472
restore_dtypes_() (*in module imgaug.dtypes*), 394
RGB (*imgaug.augmenters.color.ChangeColorspace attribute*), 662
RGB (*imgaug.augmenters.contrast.CLAHE attribute*), 713

RGB (*imgaug.augmenters.contrast.HistogramEqualization attribute*), 718
RNG (*class in imgaug.random*), 394
Rot90 (*class in imgaug.augmenters.geometric*), 787
Rotate (*class in imgaug.augmenters.geometric*), 790

S

Salt (*class in imgaug.augmenters.arithmetic*), 550
SaltAndPepper (*class in imgaug.augmenters.arithmetic*), 552
SAME_AS_IMAGES (*imgaug.augmenters.size.KeepSizeByResize attribute*), 1041
sample_points() (*imgaug.augmenters.segmentation.DropoutPointsSampler method*), 975
sample_points() (*imgaug.augmenters.segmentation.IPointsSampler method*), 975
sample_points() (*imgaug.augmenters.segmentation.RegularGridPointsSampler method*), 977
sample_points() (*imgaug.augmenters.segmentation.RelativeRegularGridPointsSampler method*), 981
sample_points() (*imgaug.augmenters.segmentation.SubsamplingPointsSampler method*), 986
sample_points() (*imgaug.augmenters.segmentation.UniformPointsSampler method*), 991
Saturate (*class in imgaug.augmenters.imgcorruptlike*), 837
SaveDebugImageEveryNBatches (*class in imgaug.augmenters.debug*), 741
scale() (*imgaug.augmentables.heatmaps.HeatmapsOnImage method*), 439
scale() (*imgaug.augmentables.segmaps.SegmentationMapsOnImage method*), 494
Scale() (*in module imgaug.augmenters.size*), 1060
ScaleX (*class in imgaug.augmenters.geometric*), 792
ScaleY (*class in imgaug.augmenters.geometric*), 795
seed() (*in module imgaug.imgaug*), 361
seed() (*in module imgaug.random*), 407
seed_() (*imgaug.augmenters.meta.Augmenter method*), 880
SegMapClassIdsMaskGen (*class in imgaug.augmenters.blend*), 619
segment_voronoi() (*in module imgaug.augmenters.segmentation*), 997
segmentation_maps (*imgaug.augmentables.batches.Batch attribute*), 412
SegmentationMapOnImage() (*in module imgaug.augmentables.segmaps*), 491
SegmentationMapsOnImage (*class in imgaug.augmentables.segmaps*), 491
SegmentationMapsOnImage() (*in module imgaug.imgaug*), 346
Sequential (*class in imgaug.augmenters.meta*), 899
set_generator_state_() (*in module imgaug.random*), 408
set_state_() (*imgaug.random.RNG method*), 402
Sharpen (*class in imgaug.augmenters.convolutional*), 739
ShearX (*class in imgaug.augmenters.geometric*), 797
ShearY (*class in imgaug.augmenters.geometric*), 799
shift() (*imgaug.augmentables.bbs.BoundingBox method*), 425
shift() (*imgaug.augmentables.bbs.BoundingBoxesOnImage method*), 433
shift() (*imgaug.augmentables.kps.Keypoint method*), 443
shift() (*imgaug.augmentables.kps.KeypointsOnImage method*), 450
shift() (*imgaug.augmentables.lines.LineString method*), 462
shift() (*imgaug.augmentables.lines.LineStringsOnImage method*), 470
shift() (*imgaug.augmentables.polys.Polygon method*), 481
shift() (*imgaug.augmentables.polys.PolygonsOnImage method*), 489
shift_() (*imgaug.augmentables.bbs.BoundingBox method*), 426
shift_() (*imgaug.augmentables.bbs.BoundingBoxesOnImage method*), 433
shift_() (*imgaug.augmentables.kps.Keypoint method*), 443
shift_() (*imgaug.augmentables.kps.KeypointsOnImage method*), 450
shift_() (*imgaug.augmentables.lines.LineString method*), 462
shift_() (*imgaug.augmentables.lines.LineStringsOnImage method*), 470
shift_() (*imgaug.augmentables.polys.Polygon method*), 481
shift_() (*imgaug.augmentables.polys.PolygonsOnImage method*), 489
ShotNoise (*class in imgaug.augmenters.imgcorruptlike*), 839
show_distributions_grid() (*in module imgaug.parameters*), 389
show_grid() (*imgaug.augmenters.meta.Augmenter method*), 881
show_grid() (*in module imgaug.imgaug*), 361
shuffle() (*imgaug.random.RNG method*), 402

shuffle_channels() (*in module imgaug.augmenters.meta*), 911
Sigmoid (*class in imgaug.parameters*), 381
SigmoidContrast (*class in imgaug.augmenters.contrast*), 723
SimplexNoise (*class in imgaug.parameters*), 383
SimplexNoiseAlpha() (*in module imgaug.augmenters.blend*), 620
Snow (*class in imgaug.augmenters.imgcorruptlike*), 841
Snowflakes (*class in imgaug.augmenters.weather*), 1083
SnowflakesLayer (*class in imgaug.augmenters.weather*), 1087
Solarize (*class in imgaug.augmenters.arithmetic*), 554
Solarize (*class in imgaug.augmenters.pillike*), 947
solarize() (*in module imgaug.augmenters.arithmetic*), 567
solarize() (*in module imgaug.augmenters.pillike*), 960
solarize_() (*in module imgaug.augmenters.arithmetic*), 567
solarize_() (*in module imgaug.augmenters.pillike*), 961
SomeColorsMaskGen (*class in imgaug.augmenters.blend*), 620
SomeOf (*class in imgaug.augmenters.meta*), 902
Sometimes (*class in imgaug.augmenters.meta*), 906
Spatter (*class in imgaug.augmenters.imgcorruptlike*), 843
SpeckleNoise (*class in imgaug.augmenters.imgcorruptlike*), 845
standard_cauchy() (*imgaug.random.RNG method*), 402
standard_exponential() (*imgaug.random.RNG method*), 402
standard_gamma() (*imgaug.random.RNG method*), 402
standard_normal() (*imgaug.random.RNG method*), 402
standard_t() (*imgaug.random.RNG method*), 402
state (*imgaug.random.RNG attribute*), 402
StochasticParameter (*class in imgaug.parameters*), 384
StochasticParameterMaskGen (*class in imgaug.augmenters.blend*), 623
stylize_cartoon() (*in module imgaug.augmenters.artistic*), 571
subdivide() (*imgaug.augmentables.lines.LineString method*), 462
subdivide() (*imgaug.augmentables.polys.Polygon method*), 481
subdivide() (*imgaug.augmentables.polys.PolygonsOnImage method*), 489
subdivide_() (*imgaug.augmentables.polys.Polygon method*), 482
subdivide_() (*imgaug.augmentables.polys.PolygonsOnImage method*), 490
SubsamplingPointsSampler (*class in imgaug.augmenters.segmentation*), 986
Subtract (*class in imgaug.parameters*), 385
Superpixels (*class in imgaug.augmenters.segmentation*), 986
supports_new_numpy_rng_style() (*in module imgaug.random*), 408
SuspiciousMultiImageShapeWarning, 572
SuspiciousSingleImageShapeWarning, 572

T

temporary_numpy_seed (*class in imgaug.random*), 408
terminate() (*imgaug.multicore.BackgroundAugmenter method*), 390
terminate() (*imgaug.multicore.BatchLoader method*), 391
terminate() (*imgaug.multicore.Pool method*), 393
to_batch_in_augmentation() (*imgaug.augmentables.batches.Batch method*), 412
to_bounding_box() (*imgaug.augmentables.lines.LineString method*), 463
to_bounding_box() (*imgaug.augmentables.polys.Polygon method*), 482
to_deterministic() (*imgaug.augmenters.meta.Augmenter method*), 881
to_distance_maps() (*imgaug.augmentables.kps.KeypointsOnImage method*), 450
to_heatmap() (*imgaug.augmentables.lines.LineString method*), 463
to_keypoint_image() (*imgaug.augmentables.kps.KeypointsOnImage method*), 451
to_keypoints() (*imgaug.augmentables.bbs.BoundingBox method*), 426
to_keypoints() (*imgaug.augmentables.lines.LineString method*), 463
to_keypoints() (*imgaug.augmentables.polys.Polygon method*), 482
to_keypoints_on_image() (*imgaug.augmentables.bbs.BoundingBoxesOnImage method*), 434

to_keypoints_on_image() (imgaug.augmentables.kps.KeypointsOnImage method), 451
to_keypoints_on_image() (imgaug.augmentables.lines.LineStringsOnImage method), 470
to_keypoints_on_image() (imgaug.augmentables.polys.PolygonsOnImage method), 490
to_line_string() (imgaug.augmentables.polys.Polygon method), 482
to_normalized_batch() (imgaug.augmentables.batches.Batch method), 412
to_normalized_batch() (imgaug.augmentables.batches.UnnormalizedBatch method), 416
to_polygon() (imgaug.augmentables.bbs.BoundingBox method), 426
to_polygon() (imgaug.augmentables.lines.LineString method), 463
to_polygons_on_image() (imgaug.augmentables.bbs.BoundingBoxesOnImage method), 434
to_segmentation_map() (imgaug.augmentables.lines.LineString method), 463
to_shapely_line_string() (imgaug.augmentables.polys.Polygon method), 482
to_shapely_polygon() (imgaug.augmentables.polys.Polygon method), 482
to_uint8() (imgaug.augmentables.heatmaps.HeatmapsOnImage method), 439
to_xy_array() (imgaug.augmentables.bbs.BoundingBoxesOnImage method), 434
to_xy_array() (imgaug.augmentables.kps.KeypointsOnImage method), 451
to_xy_array() (imgaug.augmentables.lines.LineStringsOnImage method), 471
to_xy_array() (imgaug.augmentables.polys.PolygonsOnImage method), 490
to_xy_arrays() (imgaug.augmentables.lines.LineStringsOnImage method), 471
to_xyxy_array() (imgaug.augmentables.bbs.BoundingBoxesOnImage method), 434
tomaxint() (imgaug.random.RNG method), 403
TotalDropout (class in imgaug.augmenters.arithmetic), 557
TranslateX (class in imgaug.augmenters.geometric), 801
TranslateY (class in imgaug.augmenters.geometric), 803
triangular() (imgaug.random.RNG method), 403
TruncatedNormal (class in imgaug.parameters), 386

U

Uniform (class in imgaug.parameters), 387
uniform() (imgaug.random.RNG method), 403
UniformColorQuantization (class in imgaug.augmenters.color), 683
UniformColorQuantizationToNBits (class in imgaug.augmenters.color), 686
UniformPointsSampler (class in imgaug.augmenters.segmentation), 990
UniformVoronoi (class in imgaug.augmenters.segmentation), 991
union() (imgaug.augmentables.bbs.BoundingBox method), 426
UnnormalizedBatch (class in imgaug.augmentables.batches), 413
use_state_of_() (imgaug.random.RNG method), 403

V

VerticalFlip() (in module imgaug.augmenters.flip), 754
VerticalLinearGradientMaskGen (class in imgaug.augmenters.blend), 623
vonmises() (imgaug.random.RNG method), 403
Voronoi (class in imgaug.augmenters.segmentation), 994

W

wald() (imgaug.random.RNG method), 403
warn() (in module imgaug.imgaug), 362
warn_deprecated() (in module imgaug.imgaug), 362
warp_affine() (in module imgaug.augmenters.pillike), 961
warpPolarCoords() (imgaug.augmenters.geometric.WithPolarWarping class method), 809
Weibull (class in imgaug.parameters), 388
weibull() (imgaug.random.RNG method), 403
width (imgaug.augmentables.bbs.BoundingBox attribute), 427
width (imgaug.augmentables.bbs.BoundingBoxesOnImage attribute), 434

width (*imgaug.augmentables.kps.KeypointsOnImage attribute*), 451
width (*imgaug.augmentables.lines.LineString attribute*), 464
width (*imgaug.augmentables.polys.Polygon attribute*), 482
WithBrightnessChannels (*class in imgaug.augmenters.color*), 690
WithChannels (*class in imgaug.augmenters.meta*), 908
WithColorspace (*class in imgaug.augmenters.color*), 692
WithHueAndSaturation (*class in imgaug.augmenters.color*), 694
WithPolarWarping (*class in imgaug.augmenters.geometric*), 805

X

x1_int (*imgaug.augmentables.bbs.BoundingBox attribute*), 427
x2_int (*imgaug.augmentables.bbs.BoundingBox attribute*), 427
x_int (*imgaug.augmentables.kps.Keypoint attribute*), 444
xx (*imgaug.augmentables.lines.LineString attribute*), 464
xx (*imgaug.augmentables.polys.Polygon attribute*), 483
xx_int (*imgaug.augmentables.lines.LineString attribute*), 464
xx_int (*imgaug.augmentables.polys.Polygon attribute*), 483
xy (*imgaug.augmentables.kps.Keypoint attribute*), 444
xy_int (*imgaug.augmentables.kps.Keypoint attribute*), 444

Y

y1_int (*imgaug.augmentables.bbs.BoundingBox attribute*), 427
y2_int (*imgaug.augmentables.bbs.BoundingBox attribute*), 427
y_int (*imgaug.augmentables.kps.Keypoint attribute*), 444
YCrCb (*imgaug.augmenters.color.ChangeColorspace attribute*), 662
yy (*imgaug.augmentables.lines.LineString attribute*), 464
yy (*imgaug.augmentables.polys.Polygon attribute*), 483
yy_int (*imgaug.augmentables.lines.LineString attribute*), 464
yy_int (*imgaug.augmentables.polys.Polygon attribute*), 483

Z

zipf() (*imgaug.random.RNG method*), 403
ZoomBlur (*class in imgaug.augmenters.imgcorruptlike*), 847