ChainerCV Documentation

Release 0.12.0

Preferred Networks, inc.

Feb 12, 2019

Contents

1 Installation Guide 3

2 ChainerCV Tutorial 5

3 ChainerCV Reference Manual 19

4 Naming Conventions 135

5 License 139

6 Indices and tables 141

Bibliography 143

Python Module Index 145

ChainerCV is a **deep learning based computer vision library** built on top of Chainer.

CHAPTER 1

Installation Guide

1.1 Pip

You can install ChainerCV using *pip*.

```
pip install -U numpy
pip install chainercv
```

1.2 Anaconda

Build instruction using Anaconda is as follows.

```
# For python 3
# wget https://repo.continuum.io/miniconda/Miniconda3-latest-Linux-x86_64.sh -O␣
↪miniconda.sh
wget https://repo.continuum.io/miniconda/Miniconda2-latest-Linux-x86_64.sh -O␣
↪miniconda.sh

bash miniconda.sh -b -p $HOME/miniconda
export PATH="$HOME/miniconda/bin:$PATH"
conda config --set always_yes yes --set changeps1 no
conda update -q conda

# Download ChainerCV and go to the root directory of ChainerCV
git clone https://github.com/chainer/chainercv
cd chainercv
conda env create -f environment.yml
source activate chainercv

# Install ChainerCV
pip install -e .
```

(continues on next page)

(continued from previous page)

```
# Try our demos at examples/* !
```

CHAPTER 2

ChainerCV Tutorial

2.1 Object Detection Tutorial

This tutorial will walk you through the features related to object detection that ChainerCV supports. We assume that readers have a basic understanding of Chainer framework (e.g. understand `chainer.Link`). For users new to Chainer, please first read Introduction to Chainer.

In ChainerCV, we define the object detection task as a problem of, given an image, bounding box based localization and categorization of objects. ChainerCV supports the task by providing the following features:

- Visualization
- BboxDataset
- Detection Link
- DetectionEvaluator
- Training script for various detection models

Here is a short example that conducts inference and visualizes output. Please download an image from a link below, and name it as `sample.jpg`. https://cloud.githubusercontent.com/assets/2062128/26187667/9cb236da-3bd5-11e7-8bcf-7dbd4302e2dc.jpg

```
# In the rest of the tutorial, we assume that the `plt`
# is imported before every code snippet.
import matplotlib.pyplot as plt

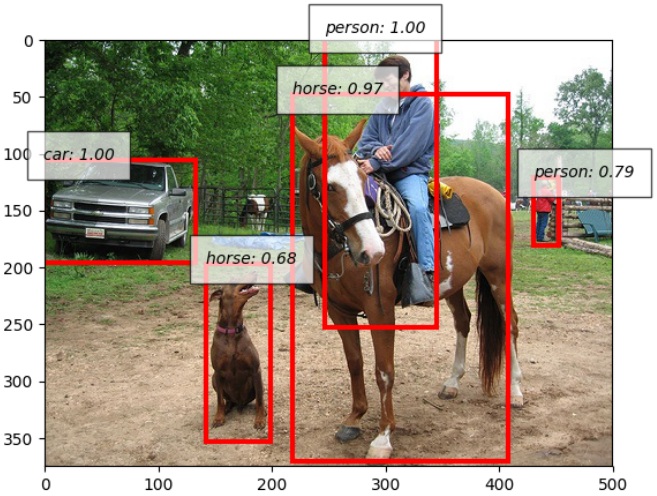
from chainercv.datasets import voc_bbox_label_names
from chainercv.links import SSD300
from chainercv.utils import read_image
from chainercv.visualizations import vis_bbox

# Read an RGB image and return it in CHW format.
img = read_image('sample.jpg')
model = SSD300(pretrained_model='voc0712')
```

(continues on next page)

(continued from previous page)

```
bboxes, labels, scores = model.predict([img])
vis_bbox(img, bboxes[0], labels[0], scores[0],
         label_names=voc_bbox_label_names)
plt.show()
```



2.1.1 Bounding boxes in ChainerCV

Bounding boxes in an image are represented as a two-dimensional array of shape $(R, 4)$, where R is the number of bounding boxes and the second axis corresponds to the coordinates of bounding boxes. The coordinates are ordered in the array by `(y_min, x_min, y_max, x_max)`, where `(y_min, x_min)` and `(y_max, x_max)` are the `(y, x)` coordinates of the top left and the bottom right vertices. Notice that ChainerCV orders coordinates in `yx` order, which is the opposite of the convention used by other libraries such as OpenCV. This convention is adopted because it is more consistent with the memory order of an image that follows row-column order. Also, the `dtype` of bounding box array is `numpy.float32`.

Here is an example with a simple toy data.

```
from chainercv.visualizations import vis_bbox
import numpy as np

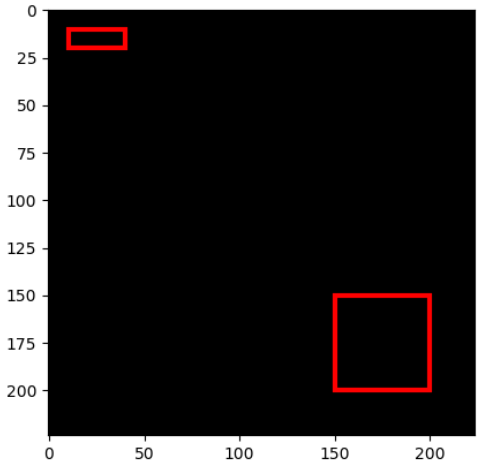
img = np.zeros((3, 224, 224), dtype=np.float32)
# We call a variable/array of bounding boxes as `bbox` throughout the library
bbox = np.array([[10, 10, 20, 40], [150, 150, 200, 200]], dtype=np.float32)

vis_bbox(img, bbox)
plt.show()
```

In this example, two bounding boxes are displayed on top of a black image. *`vis_bbox()`* is a utility function that visualizes bounding boxes and an image together.

2.1.2 Bounding Box Dataset

ChainerCV supports dataset loaders, which can be used to easily index examples with list-like interfaces. Dataset classes whose names end with `BboxDataset` contain annotations of where objects locate in an image and which



categories they are assigned to. These datasets can be indexed to return a tuple of an image, bounding boxes and labels. The labels are stored in an `np.int32` array of shape $(R,)$. Each element corresponds to a label of an object in the corresponding bounding box.

A mapping between an integer label and a category differs between datasets. This mapping can be obtained from objects whose names end with `label_names`, such as `voc_bbox_label_names`. These mappings become helpful when bounding boxes need to be visualized with label names. In the next example, the interface of `BboxDataset` and the functionality of *`vis_bbox()`* to visualize label names are illustrated.

```
from chainercv.datasets import VOCBboxDataset
from chainercv.datasets import voc_bbox_label_names
from chainercv.visualizations import vis_bbox

dataset = VOCBboxDataset(year='2012')
img, bbox, label = dataset[0]
print(bbox.shape)  # (2, 4)
print(label.shape)  # (2,)
vis_bbox(img, bbox, label, label_names=voc_bbox_label_names)
plt.show()
```

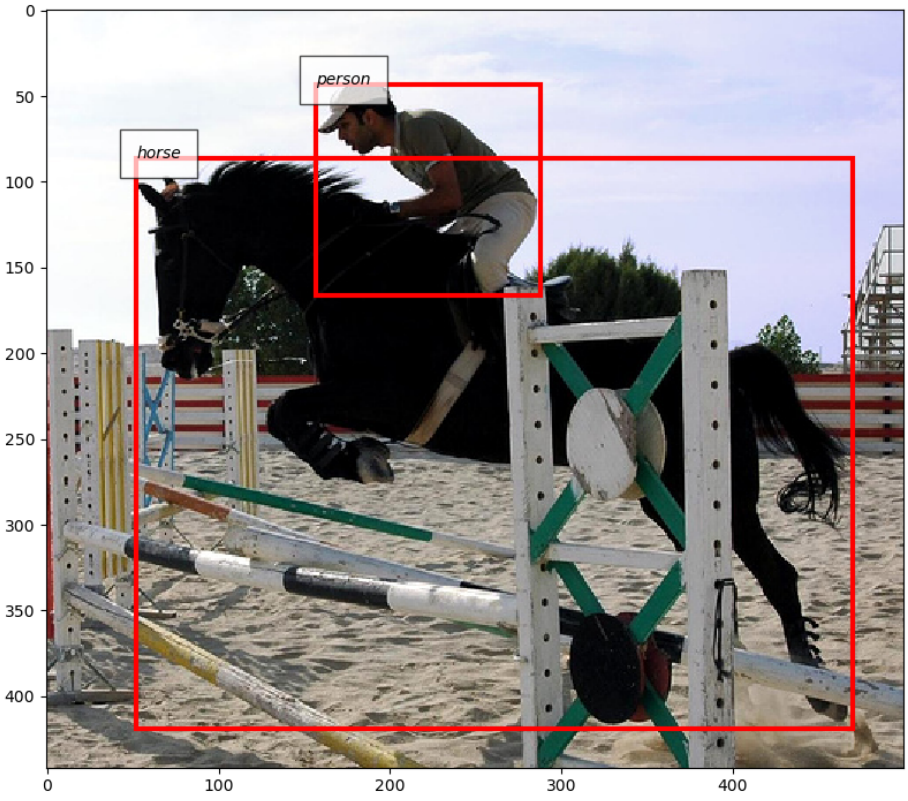
Note that the example downloads VOC 2012 dataset at runtime when it is used for the first time on the machine.

2.1.3 Detection Link

ChainerCV provides several network implementations that carry out object detection. For example, Single Shot Multi-Box Detector (SSD) *[Liu16]* and Faster R-CNN *[Ren15]* are supported. Despite the difference between the models in how prediction is carried out internally, they support the common method for prediction called *`predict()`*. This method takes a list of images and returns prediction result, which is a tuple of lists `bboxes, labels, scores`. The more description can be found here (*`predict()`*). Inference on these models runs smoothly by downloading necessary pre-trained weights from the internet automatically.

```
from chainercv.datasets import VOCBboxDataset
from chainercv.datasets import voc_bbox_label_names
from chainercv.links import SSD300
from chainercv.visualizations import vis_bbox
```

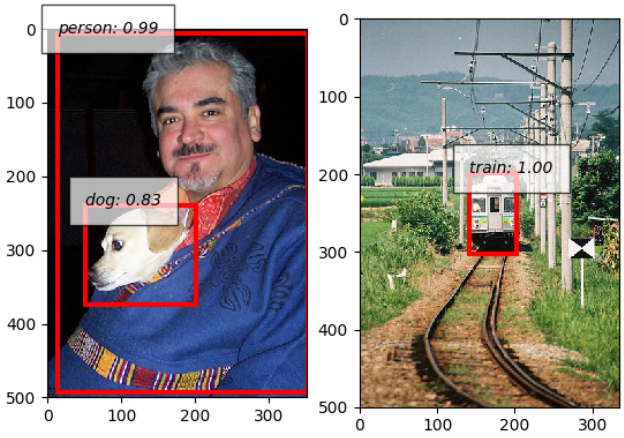
(continues on next page)


person
horse
0
50
100
150
200
250
300
350
400
0
100
200
300
400

(continued from previous page)

```
dataset = VOCBboxDataset(year='2007', split='test')
img_0, _, _ = dataset[0]
img_1, _, _ = dataset[1]
model = SSD300(pretrained_model='voc0712')
# Note that `predict` takes a list of images.
bboxes, labels, scores = model.predict([img_0, img_1])

# Visualize output of the first image on the left and
# the second image on the right.
fig = plt.figure()
ax1 = fig.add_subplot(1, 2, 1)
ax2 = fig.add_subplot(1, 2, 2)
vis_bbox(img_0, bboxes[0], labels[0], scores[0],
         label_names=voc_bbox_label_names, ax=ax1)
vis_bbox(img_1, bboxes[1], labels[1], scores[1],
         label_names=voc_bbox_label_names, ax=ax2)
plt.show()
```



The above example puts together functionality of detection link. It instantiates SSD300 model with weights trained on VOC 2007 and VOC 2012 datasets. The model runs prediction using *predict()*, and the outputs are visualized using *vis_bbox()*. Note that in this case, confidence scores are visualized together with other data.

Many detection algorithms post-process bounding box proposals calculated from the output of neural networks by removing unnecessary ones. Faster R-CNN and SSD use non-maximum suppression to remove overlapping bounding boxes. Also, they remove bounding boxes with low confidence scores. These two models have attributes `nms_thresh` and `score_thresh`, which configure the post-processing. In the following example, the algorithm runs with a very low `score_thresh` so that bounding boxes with low scores are kept. It is known that lower `score_thresh` produces higher mAP.

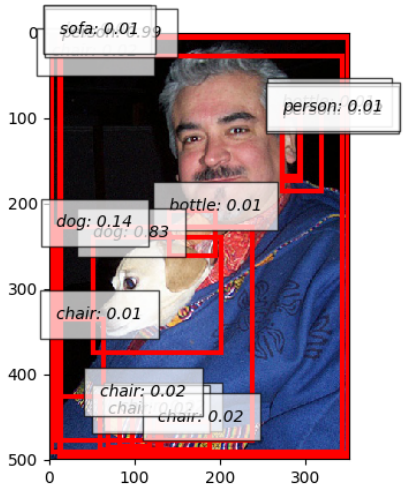
```
from chainercv.datasets import VOCBboxDataset
from chainercv.datasets import voc_bbox_label_names
from chainercv.links import SSD300
from chainercv.visualizations import vis_bbox

dataset = VOCBboxDataset(year='2007', split='test')
```

(continues on next page)

(continued from previous page)

```
img, _, _ = dataset[0]
model = SSD300(pretrained_model='voc0712')
# Alternatively, you can use predefined parameters by
# model.use_preset('evaluate')
model.score_thresh = 0.01
bboxes, labels, scores = model.predict([img])
vis_bbox(img, bboxes[0], labels[0], scores[0],
         label_names=voc_bbox_label_names)
plt.show()
```



2.1.4 Detection Evaluator

ChainerCV provides functionalities that make evaluating detection links easy. They are provided at two levels: evaluator extensions and evaluation functions.

Evaluator extensions such as *DetectionVOCEvaluator* inherit from `Evaluator`, and have similar interface. They are initialized by taking an iterator and a network that carries out prediction with method *predict()*. When this class is called (i.e. `__call__()` of *DetectionVOCEvaluator*), several actions are taken. First, it iterates over a dataset based on an iterator. Second, the network makes prediction using the images collected from the dataset. Last, an evaluation function is called with the ground truth annotations and the prediction results.

In contrast to evaluators that hide details, evaluation functions such as *eval_detection_voc()* are provided for those who need a finer level of control. These functions take the ground truth annotations and prediction results as arguments and return measured performance.

Here is a simple example that uses a detection evaluator.

```
from chainer.iterators import SerialIterator
from chainer.datasets import SubDataset
from chainercv.datasets import VOCBboxDataset
from chainercv.datasets import voc_bbox_label_names
from chainercv.extensions import DetectionVOCEvaluator
from chainercv.links import SSD300

# Only use subset of dataset so that evaluation finishes quickly.
```

(continues on next page)

(continued from previous page)

```
dataset = VOCBboxDataset(year='2007', split='test')
dataset = dataset[:6]
it = SerialIterator(dataset, 2, repeat=False, shuffle=False)
model = SSD300(pretrained_model='voc0712')
evaluator = DetectionVOCEvaluator(it, model,
                                  label_names=voc_bbox_label_names)
# result is a dictionary of evaluation scores. Print it and check it.
result = evaluator()
```

2.1.5 Training Detection Links

By putting together all the functions and utilities, training scripts can be easily written. Please check training scripts contained in the examples. Also, ChainerCV posts the performance achieved through running the training script in README.

- Faster R-CNN examples
- SSD examples

2.1.6 References

2.2 Tips using Links

2.2.1 Fine-tuning

Models in ChainerCV support the argument `pretrained_model` to load pretrained weights. This functionality is limited in the case when fine-tuning pretrained weights. In that circumstance, the layers specific to the classes of the original dataset may need to be randomly initialized. In this section, we give a procedure to cope with this problem.

Copying a subset of weights in a chain can be done in few lines of code. Here is a block of code that does this.

```
# src is a model with pretrained weights
# dst is a model randomly initialized
# ignore_names is the name of parameters to skip
# For the case of VGG16, this should be ['/fc7/W', '/fc7/b']
ignore_names = []
src_params = {p[0]: p[1] for p in src.namedparams()}
for dst_named_param in dst.namedparams():
    name = dst_named_param[0]
    if name not in ignore_names:
        dst_named_param[1].array[:] = src_params[name].array[:]
```

Fine-tuning to a dataset with a different number of classes

When the number of classes of the target dataset is different from the source dataset during fine-tuning, the names of the weights to skip can be found automatically with the following method.

```
def get_shape_mismatch_names(src, dst):
    # all parameters are assumed to be initialized
    mismatch_names = []
    src_params = {p[0]: p[1] for p in src.namedparams()}
```

(continues on next page)

(continued from previous page)

```
    for dst_named_param in dst.namedparams():
        name = dst_named_param[0]
        dst_param = dst_named_param[1]
        src_param = src_params[name]
        if src_param.shape != dst_param.shape:
            mismatch_names.append(name)
    return mismatch_names
```

Finally, this is a complete example using SSD300.

```
from chainercv.links import SSD300
import numpy as np

src = SSD300(pretrained_model='voc0712')
# the number of classes in VOC is different from 50
dst = SSD300(n_fg_class=50)
# initialized weights
dst(np.zeros((1, 3, dst.insize, dst.insize), dtype=np.float32))

# the method described above
ignore_names = get_shape_mismatch_names(src, dst)
src_params = {p[0]: p[1] for p in src.namedparams()}
for dst_named_param in dst.namedparams():
    name = dst_named_param[0]
    if name not in ignore_names:
        dst_named_param[1].array[:] = src_params[name].array[:]

# check that weights are transfered
np.testing.assert_equal(dst.extractor.conv1_1.W.data,
                        src.extractor.conv1_1.W.data)
# the names of the weights that are skipped
print(ignore_names)
```

2.3 Sliceable Dataset

This tutorial will walk you through the features related to sliceable dataset. We assume that readers have a basic understanding of Chainer dataset (e.g. understand `chainer.dataset.DatasetMixin`).

In ChainerCV, we introduce *sliceable* feature to datasets. Sliceable datasets support `slice()` that returns a view of the dataset.

This example that shows the basic usage.

```
# VOCBboxDataset supports sliceable feature
from chainercv.datasets import VOCBboxDataset
dataset = VOCBboxDataset()

# keys returns the names of data
print(dataset.keys)  # ('img', 'bbox', 'label')
# we can get an example by []
img, bbox, label = dataset[0]

# get a view of the first 100 examples
view = dataset.slice[:100]
print(len(view))  # 100
```

(continues on next page)

(continued from previous page)

```
# get a view of image and label
view = dataset.slice[:, ('img', 'label')]
# the view also supports sliceable, so that we can call keys
print(view.keys)  # ('img', 'label')
# we can get an example by []
img, label = view[0]
```

2.3.1 Motivation

`slice()` returns a view of the dataset without conducting data loading, where `DatasetMixin.__getitem__()` conducts `get_example()` for all required examples. Users can write efficient code by this view.

This example counts the number of images that contain dogs. With the sliceable feature, we can access the label information without loading images from disk.. Therefore, the first case becomes faster.

```
import time

from chainercv.datasets import VOCBboxDataset
from chainercv.datasets import voc_bbox_label_names

dataset = VOCBboxDataset()
dog_lb = voc_bbox_label_names.index('dog')

# with slice
t = time.time()
count = 0
# get a view of label
view = dataset.slice[:, 'label']
for i in range(len(view)):
    # we can focus on label
    label = view[i]
    if dog_lb in label:
        count += 1
print('w/ slice: {} secs'.format(time.time() - t))
print('{} images contain dogs'.format(count))
print()

# without slice
t = time.time()
count = 0
for i in range(len(dataset)):
    # img and bbox are loaded but not needed
    img, bbox, label = dataset[i]
    if dog_lb in label:
        count += 1
print('w/o slice: {} secs'.format(time.time() - t))
print('{} images contain dogs'.format(count))
print()
```

2.3.2 Usage: slice along with the axis of examples

`slice()` takes indices of examples as its first argument.

```
from chainercv.datasets import VOCBboxDataset
dataset = VOCBboxDataset()

# the view of the first 100 examples
view = dataset.slice[:100]

# the view of the last 100 examples
view = dataset.slice[-100:]

# the view of the 3rd, 5th, and 7th examples
view = dataset.slice[3:8:2]

# the view of the 3rd, 1st, and 4th examples
view = dataset.slice[[3, 1, 4]]
```

Also, it can take a list of booleans as its first argument. Note that the length of the list should be the same as `len(dataset)`.

```
from chainercv.datasets import VOCBboxDataset
dataset = VOCBboxDataset()

# make booleans
bboxes = dataset.slice[:, 'bbox']
booleans = [len(bbox) >= 3 for bbox in bboxes]

# a collection of samples that contain at least three bounding boxes
view = dataset.slice[booleans]
```

2.3.3 Usage: slice along with the axis of data

`slice()` takes names or indices of data as its second argument. `keys` returns all available names.

```
from chainercv.datasets import VOCBboxDataset
dataset = VOCBboxDataset()

# the view of image
# note that : of the first argument means all examples
view = dataset.slice[:, 'img']
print(view.keys)  # 'img'
img = view[0]

# the view of image and label
view = dataset.slice[:, ('img', 'label')]
print(view.keys)  # ('img', 'label')
img, label = view[0]

# the view of image (returns a tuple)
view = dataset.slice[:, ('img',)]
print(view.keys)  # ('img',)
img, = view[0]

# use an index instead of a name
view = dataset.slice[:, 1]
print(view.keys)  # 'bbox'
bbox = view[0]
```

(continues on next page)

(continued from previous page)

```
# mixture of names and indices
view = dataset.slice[:, (1, 'label')]
print(view.keys)  # ('bbox', 'label')
bbox, label = view[0]

# use booleans
# note that the number of booleans should be the same as len(dataset.keys)
view = dataset.slice[:, (True, True, False)]
print(view.keys)  # ('img', 'bbox')
img, bbox = view[0]
```

2.3.4 Usage: slice along with both axes

```
from chainercv.datasets import VOCBboxDataset
dataset = VOCBboxDataset()

# the view of the labels of the first 100 examples
view = dataset.slice[:100, 'label']
```

2.3.5 Concatenate and transform

ChainerCV provides *ConcatenatedDataset* and *TransformDataset*. The difference from `chainer.datasets.ConcatenatedDataset` and `chainer.datasets.TransformDataset` is that they take sliceable dataset(s) and return a sliceable dataset.

```
from chainercv.chainer_experimental.datasets.sliceable import ConcatenatedDataset
from chainercv.chainer_experimental.datasets.sliceable import TransformDataset
from chainercv.datasets import VOCBboxDataset
from chainercv.datasets import voc_bbox_label_names

dataset_07 = VOCBboxDataset(year='2007')
print('07:', dataset_07.keys, len(dataset_07))  # 07: ('img', 'bbox', 'label') 2501

dataset_12 = VOCBboxDataset(year='2012')
print('12:', dataset_12.keys, len(dataset_12))  # 12: ('img', 'bbox', 'label') 5717

# concatenate
dataset_0712 = ConcatenatedDataset(dataset_07, dataset_12)
print('0712:', dataset_0712.keys, len(dataset_0712))  # 0712: ('img', 'bbox', 'label
↪') 8218

# transform
def transform(in_data):
    img, bbox, label = in_data

    dog_lb = voc_bbox_label_names.index('dog')
    bbox_dog = bbox[label == dog_lb]

    return img, bbox_dog

# we need to specify the names of data that the transform function returns
dataset_0712_dog = TransformDataset(dataset_0712, ('img', 'bbox_dog'), transform)
print('0712_dog:', dataset_0712_dog.keys, len(dataset_0712_dog))  # 0712_dog: ('img',
↪'bbox_dog') 8218
```

(continues on next page)

(continued from previous page)

2.3.6 Make your own dataset

ChainerCV provides *GetterDataset* to construct a new sliceable dataset.

This example implements a sliceable bounding box dataset.

```
import numpy as np

from chainercv.chainer_experimental.datasets.sliceable import GetterDataset
from chainercv.utils import generate_random_bbox

class SampleBboxDataset(GetterDataset):
    def __init__(self):
        super(SampleBboxDataset, self).__init__()

        # register getter method for image
        self.add_getter('img', self.get_image)
        # register getter method for bbox and label
        self.add_getter(('bbox', 'label'), self.get_annotation)

    def __len__(self):
        return 20

    def get_image(self, i):
        print('get_image({})'.format(i))
        # generate dummy image
        img = np.random.uniform(0, 255, size=(3, 224, 224)).astype(np.float32)
        return img

    def get_annotation(self, i):
        print('get_annotation({})'.format(i))
        # generate dummy annotations
        bbox = generate_random_bbox(10, (224, 224), 10, 224)
        label = np.random.randint(0, 9, size=10).astype(np.int32)
        return bbox, label

dataset = SampleBboxDataset()
img, bbox, label = dataset[0]  # get_image(0) and get_annotation(0)

view = dataset.slice[:, 'label']
label = view[1]  # get_annotation(1)
```

If you have arrays of data, you can use *TupleDataset*.

```
import numpy as np

from chainercv.chainer_experimental.datasets.sliceable import TupleDataset
from chainercv.utils import generate_random_bbox

n = 20
imgs = np.random.uniform(0, 255, size=(n, 3, 224, 224)).astype(np.float32)
bboxes = [generate_random_bbox(10, (224, 224), 10, 224) for _ in range(n)]
labels = np.random.randint(0, 9, size=(n, 10)).astype(np.int32)
```

(continues on next page)

(continued from previous page)

```
dataset = TupleDataset(('img', imgs), ('bbox', bboxes), ('label', labels))

print(dataset.keys)  # ('img', 'bbox', 'label')
view = dataset.slice[:, 'label']
label = view[1]
```

CHAPTER 3

ChainerCV Reference Manual

3.1 Chainer Experimental

This module contains WIP modules of Chainer. After they are merged into chainer, these modules will be removed from ChainerCV.

3.1.1 Datasets

Sliceable

Sliceable

This module support sliceable feature. Please note that this module will be removed after Chainer implements sliceable feature.

See also:

https://github.com/chainer/chainercv/pull/454

ConcatenatedDataset

class chainercv.chainer_experimental.datasets.sliceable.**ConcatenatedDataset**(**datasets*)

A sliceable version of `chainer.datasets.ConcatenatedDataset`.

Here is an example.

```
>>> dataset_a = TupleDataset([0, 1, 2], [0, 1, 4])
>>> dataset_b = TupleDataset([3, 4, 5], [9, 16, 25])
>>>
>>> dataset = ConcatenatedDataset(dataset_a, dataset_b)
>>> dataset.slice[:, 0][:]  # [0, 1, 2, 3, 4, 5]
```

Parameters **datasets** – The underlying datasets. Each dataset should inherit `Sliceabledataset` and should have the same keys.

get_example_by_keys(*index*, *key_indices*)
Return data of an example by keys

Parameters

- **index** (*int*) – An index of an example.
- **key_indices** (*tuple of ints*) – A tuple of indices of requested keys.

Returns tuple of data

GetterDataset

class chainercv.chainer_experimental.datasets.sliceable.**GetterDataset**
A sliceable dataset class that is defined with getters.

This is a dataset class with getters. Please refer to the tutorial for more detailed explanation.

Here is an example.

```
>>> class SliceableLabeledImageDataset(GetterDataset):
>>>     def __init__(self, pairs, root='.'):
>>>         super(SliceableLabeledImageDataset, self).__init__()
>>>         with open(pairs) as f:
>>>             self._pairs = [l.split() for l in f]
>>>         self._root = root
>>>
>>>         self.add_getter('img', self.get_image)
>>>         self.add_getter('label', self.get_label)
>>>
>>>     def __len__(self):
>>>         return len(self._pairs)
>>>
>>>     def get_image(self, i):
>>>         path, _ = self._pairs[i]
>>>         return read_image(os.path.join(self._root, path))
>>>
>>>     def get_label(self, i):
>>>         _, label = self._pairs[i]
>>>         return np.int32(label)
>>>
>>> dataset = SliceableLabeledImageDataset('list.txt')
>>>
>>> # get a subset with label = 0, 1, 2
>>> # no images are loaded
>>> indices = [i for i, label in
...            enumerate(dataset.slice[:, 'label']) if label in {0, 1, 2}]
>>> dataset_012 = dataset.slice[indices]
```

add_getter(*keys*, *getter*)
Register a getter function

Parameters

- **keys** (*int or string or tuple of strings*) – The number or name(s) of data that the getter function returns.

- **getter** (*callable*) – A getter function that takes an index and returns data of the corresponding example.

get_example_by_keys(*index*, *key_indices*)
Return data of an example by keys

Parameters

- **index** (*int*) – An index of an example.
- **key_indices** (*tuple of ints*) – A tuple of indices of requested keys.

Returns tuple of data

TupleDataset

class chainercv.chainer_experimental.datasets.sliceable.**TupleDataset**(**datasets*)
A sliceable version of `chainer.datasets.TupleDataset`.

Here is an example.

```
>>> # omit keys
>>> dataset = TupleDataset([0, 1, 2], [0, 1, 4])
>>> dataset.keys  # (None, None)
>>> dataset.slice[:, 0][:]  # [0, 1, 2]
>>>
>>> dataset_more = TupleDataset(dataset, [0, 1, 8])
>>> dataset_more.keys  # (None, None, None)
>>> dataset_more.slice[:, [1, 2]][:]  # [(0, 0), (1, 1), (4, 8)]
>>>
>>> # specify the name of a key
>>> named_dataset = TupleDataset(('feat0', [0, 1, 2]), [0, 1, 4])
>>> named_dataset.keys  # ('feat0', None)
>>> # slice takes both key and index (or their mixture)
>>> named_dataset.slice[:, ['feat0', 1]][:]  # [(0, 0), (1, 1), (2, 4)]
```

Parameters **datasets** – The underlying datasets. The following datasets are acceptable.

- An inheritance of :class:~chainer.datasets.sliceable.SliceableDataset‘.
- A tuple of a name and a data array. The data array should be list or `numpy.ndarray`.
- A data array. In this case, the name of key is `None`.

get_example_by_keys(*index*, *key_indices*)
Return data of an example by keys

Parameters

- **index** (*int*) – An index of an example.
- **key_indices** (*tuple of ints*) – A tuple of indices of requested keys.

Returns tuple of data

TransformDataset

class chainercv.chainer_experimental.datasets.sliceable.**TransformDataset**(*dataset*, *keys*, *transform*)

A sliceable version of `chainer.datasets.TransformDataset`.

Note that it reuqires `keys` to determine the names of returned values.

Here is an example.

```
>>> def transfrom(in_data):
>>>     img, bbox, label = in_data
>>>     ...
>>>     return new_img, new_label
>>>
>>> dataset = TramsformDataset(dataset, ('img', 'label'), transform)
>>> dataset.keys  # ('img', 'label')
```

Parameters

- **dataset** – The underlying dataset. This dataset should have `__len__()` and `__getitem__()`.
- **keys** (*int or string or tuple of strings*) – The number or name(s) of data that the transform function returns.
- **transform** (*callable*) – A function that is called to transform values returned by the underlying dataset's `__getitem__()`.

3.1.2 Training

Extensions

Extensions

make_shift

chainercv.chainer_experimental.training.extensions.**make_shift**(*attr*, *optimizer=None*)

Decorator to make shift extensions.

This decorator wraps a function and makes a shift extension. Base function should take `trainer` and return a new value of `attr`.

Here is an example.

```
>>> # define an extension that updates 'lr' attribute
>>> @make_shift('lr')
>>> def warmup(trainer):
>>>     base_lr = 0.01
>>>     rate = 0.1
>>>
>>>     iteration = trainer.updater.iteration
>>>     if iteration < 1000:
```

(continues on next page)

(continued from previous page)

```
>>>         return base_lr * (rate + (1 - rate) * iteraion / 1000)
>>>     else:
>>>         return base_lr
>>>
>>> # use the extension
>>> trainer.extend(warmup)
```

Parameters

- **attr** (*str*) – Name of the attribute to shift.
- **optimizer** (*Optimizer*) – Target optimizer to adjust the attribute. If it is `None`, the main optimizer of the updater is used.

3.2 Datasets

3.2.1 General datasets

DirectoryParsingLabelDataset

class chainercv.datasets.**DirectoryParsingLabelDataset**(*root*, *check_img_file=None*, *color=True*, *numerical_sort=False*)

A label dataset whose label names are the names of the subdirectories.

The label names are the names of the directories that locate a layer below the root directory. All images locating under the subdirectoies will be categorized to classes with subdirectory names. An image is parsed only when the function `check_img_file` returns `True` by taking the path to the image as an argument. If `check_img_file` is `None`, the path with any image extensions will be parsed.

Example

A directory structure should be one like below.

```
root
|-- class_0
|   |-- img_0.png
|   |-- img_1.png
|
--- class_1
    |-- img_0.png
```

```
>>> from chainercv.datasets import DirectoryParsingLabelDataset
>>> dataset = DirectoryParsingLabelDataset('root')
>>> dataset.paths
['root/class_0/img_0.png', 'root/class_0/img_1.png',
'root_class_1/img_0.png']
>>> dataset.labels
array([0, 0, 1])
```

Parameters

- **root** (*string*) – The root directory.
- **check_img_file** (*callable*) – A function to determine if a file should be included in the dataset.
- **color** (*bool*) – If `True`, this dataset read images as color images. The default value is `True`.
- **numerical_sort** (*bool*) – Label names are sorted numerically. This means that label `2` is before label `10`, which is not the case when string sort is used. Regardless of this option, string sort is used for the order of files with the same label. The default value is `False`.

This dataset returns the following data.

name	shape	dtype	format
`img`	$(3, H, W)$[1]	`float32`	RGB, $[0, 255]$
`label`	scalar	`int32`	$[0, \#class - 1]$

directory_parsing_label_names

`chainercv.datasets.`**`directory_parsing_label_names`**(*root*, *numerical_sort=False*)

Get label names from the directories that are named by them.

The label names are the names of the directories that locate a layer below the root directory.

The label names can be used together with *DirectoryParsingLabelDataset*. The index of a label name corresponds to the label id that is used by the dataset to refer the label.

Parameters

- **root** (*string*) – The root directory.
- **numerical_sort** (*bool*) – Label names are sorted numerically. This means that label `2` is before label `10`, which is not the case when string sort is used. The default value is `False`.

Returns Sorted names of classes.

Return type list of strings

MixUpSoftLabelDataset

class `chainercv.datasets.`**`MixUpSoftLabelDataset`**(*dataset*, *n_class*)

Dataset which returns mixed images and labels for mixup learning[2].

MixUpSoftLabelDataset mixes two pairs of labeled images fetched from the base dataset.

Unlike *LabeledImageDatasets*, label is a one-dimensional float array with at most two nonnegative weights (i.e. soft label). The sum of the two weights is one.

Example

We construct a mixup dataset from MNIST.

[1] $(1, H, W)$ if `color = False`.

[2] Hongyi Zhang, Moustapha Cisse, Yann N. Dauphin, David Lopez-Paz. mixup: Beyond Empirical Risk Minimization. arXiv 2017.

```
>>> from chainer.datasets import get_mnist
>>> from chainercv.datasets import SiameseDataset
>>> from chainercv.datasets import MixUpSoftLabelDataset
>>> mnist, _ = get_mnist()
>>> base_dataset = SiameseDataset(mnist, mnist)
>>> dataset = MixUpSoftLabelDataset(base_dataset, 10)
>>> mixed_image, mixed_label = dataset[0]
>>> mixed_label.shape
(10,)
>>> mixed_label.dtype
dtype('float32')
```

Parameters

- **dataset** – The underlying dataset. The dataset returns `img_0, label_0, img_1, label_1`, which is a tuple containing two pairs of an image and a label. Typically, dataset is *SiameseDataset*.

  The shapes of images and labels should be constant.
- **n_class** (*int*) – The number of classes in the base dataset.

This dataset returns the following data.

name	shape	dtype	format
`img`	[3]	[3]	[3]
`label`	$(\#class,)$	`float32`	$[0, 1]$

SiameseDataset

class chainercv.datasets.**SiameseDataset**(*dataset_0*, *dataset_1*, *pos_ratio=None*, *length=None*, *labels_0=None*, *labels_1=None*)

A dataset that returns samples fetched from two datasets.

The dataset returns samples from the two base datasets. If `pos_ratio` is not `None`, *SiameseDataset* can be configured to return positive pairs at the ratio of `pos_ratio` and negative pairs at the ratio of `1 - pos_ratio`. In this mode, the base datasets are assumed to be label datasets that return an image and a label as a sample.

Example

We construct a siamese dataset from MNIST.

```
>>> from chainer.datasets import get_mnist
>>> from chainercv.datasets import SiameseDataset
>>> mnist, _ = get_mnist()
>>> dataset = SiameseDataset(mnist, mnist, pos_ratio=0.3)
# The probability of the two samples having the same label
# is 0.3 as specified by pos_ratio.
>>> img_0, label_0, img_1, label_1 = dataset[0]
# The returned examples may change in the next
# call even if the index is the same as before
```

(continues on next page)

[3] Same as `dataset`.

(continued from previous page)

```
# because SiameseDataset picks examples randomly
# (e.g., img_0_new may differ from img_0).
>>> img_0_new, label_0_new, img_1_new, label_1_new = dataset[0]
```

Parameters

- **dataset_0** – The first base dataset.
- **dataset_1** – The second base dataset.
- **pos_ratio** (*float*) – If this is not None, this dataset tries to construct positive pairs at the given rate. If None, this dataset randomly samples examples from the base datasets. The default value is None.
- **length** (*int*) – The length of this dataset. If None, the length of the first base dataset is the length of this dataset.
- **labels_0** (*numpy.ndarray*) – The labels associated to the first base dataset. The length should be the same as the length of the first dataset. If this is None, the labels are automatically fetched using the following line of code: `[ex[1] for ex in dataset_0]`. By setting `labels_0` and skipping the fetching iteration, the computation cost can be reduced. Also, if `pos_ratio` is None, this value is ignored. The default value is None. If `labels_1` is spcified and `dataset_0` and `dataset_1` are the same, `labels_0` can be skipped.
- **labels_1** (*numpy.ndarray*) – The labels associated to the second base dataset. If `labels_0` is spcified and `dataset_0` and `dataset_1` are the same, `labels_1` can be skipped. Please consult the explanation for `labels_0`.

This dataset returns the following data.

name	shape	dtype	format
img_0	[4]	[4]	[4]
label_0	scalar	int32	$[0, \#class - 1]$
img_1	[5]	[5]	[5]
label_1	scalar	int32	$[0, \#class - 1]$

3.2.2 ADE20K

ADE20KSemanticSegmentationDataset

class chainercv.datasets.**ADE20KSemanticSegmentationDataset**(*data_dir='auto'*, *split='train'*)

Semantic segmentation dataset for *ADE20K*.

This is ADE20K dataset distributed in MIT Scene Parsing Benchmark website. It has 20,210 training images and 2,000 validation images.

Parameters

- **data_dir** (*string*) – Path to the dataset directory. The directory should contain the `ADEChallengeData2016` directory. And that directory should contain at least `images`

[4] Same as `dataset_0`.
[5] Same as `dataset_1`.

and `annotations` directries. If `auto` is given, the dataset is automatically downloaded into `$CHAINER_DATASET_ROOT/pfnet/chainercv/ade20k`.

- **`split`** (*`{'train', 'val'}`*) – Select from dataset splits used in MIT Scene Parsing Benchmark dataset (ADE20K).

This dataset returns the following data.

name	shape	dtype	format
`img`	$(3, H, W)$	`float32`	RGB, $[0, 255]$
`label`	(H, W)	`int32`	$[-1, \#class - 1]$

ADE20KTestImageDataset

`class` `chainercv.datasets.`**`ADE20KTestImageDataset`**(*data_dir='auto'*)

Image dataset for test split of *ADE20K*.

This is an image dataset of test split in ADE20K dataset distributed at MIT Scene Parsing Benchmark website. It has 3,352 test images.

Parameters **`data_dir`** (*`string`*) – Path to the dataset directory. The directory should contain the `release_test` dir. If `auto` is given, the dataset is automatically downloaded into `$CHAINER_DATASET_ROOT/pfnet/chainercv/ade20k`.

This dataset returns the following data.

name	shape	dtype	format
`img`	$(3, H, W)$	`float32`	RGB, $[0, 255]$

3.2.3 CamVid

CamVidDataset

`class` `chainercv.datasets.`**`CamVidDataset`**(*data_dir='auto'*, *split='train'*)

Semantic segmentation dataset for CamVid.

Parameters

- **`data_dir`** (*`string`*) – Path to the root of the training data. If this is `auto`, this class will automatically download data for you under `$CHAINER_DATASET_ROOT/pfnet/chainercv/camvid`.
- **`split`** (*`{'train', 'val', 'test'}`*) – Select from dataset splits used in CamVid Dataset.

This dataset returns the following data.

name	shape	dtype	format
`img`	$(3, H, W)$	`float32`	RGB, $[0, 255]$
`label`	(H, W)	`int32`	$[-1, \#class - 1]$

3.2.4 Cityscapes

CityscapesSemanticSegmentationDataset

class chainercv.datasets.**CityscapesSemanticSegmentationDataset**(*data_dir='auto'*, *label_resolution=None*, *split='train'*, *ignore_labels=True*)

Semantic segmentation dataset for Cityscapes dataset.

Note: Please manually download the data because it is not allowed to re-distribute Cityscapes dataset.

Parameters

- **data_dir** (*string*) – Path to the dataset directory. The directory should contain at least two directories, `leftImg8bit` and either `gtFine` or `gtCoarse`. If `auto` is given, it uses `$CHAINER_DATSET_ROOT/pfnet/chainercv/cityscapes` by default.
- **label_resolution** (*{'fine', 'coarse'}*) – The resolution of the labels. It should be either `fine` or `coarse`.
- **split** (*{'train', 'val'}*) – Select from dataset splits used in Cityscapes dataset.
- **ignore_labels** (*bool*) – If `True`, the labels marked `ignoreInEval` defined in the original cityscapesScripts will be replaced with `-1` in the `get_example()` method. The default value is `True`.

This dataset returns the following data.

name	shape	dtype	format
`img`	$(3, H, W)$	`float32`	RGB, $[0, 255]$
`label`	(H, W)	`int32`	$[-1, \#class - 1]$

CityscapesTestImageDataset

class chainercv.datasets.**CityscapesTestImageDataset**(*data_dir='auto'*)

Image dataset for test split of Cityscapes dataset.

Note: Please manually download the data because it is not allowed to re-distribute Cityscapes dataset.

Parameters **data_dir** (*string*) – Path to the dataset directory. The directory should contain the `leftImg8bit` directory. If `auto` is given, it uses `$CHAINER_DATSET_ROOT/pfnet/chainercv/cityscapes` by default.

This dataset returns the following data.

name	shape	dtype	format
`img`	$(3, H, W)$	`float32`	RGB, $[0, 255]$

3.2.5 CUB

CUBLabelDataset

class chainercv.datasets.**CUBLabelDataset**(*data_dir='auto'*, *return_bb=False*, *prob_map_dir='auto'*, *return_prob_map=False*)

Caltech-UCSD Birds-200-2011 dataset with annotated class labels.

Parameters

- **data_dir** (*string*) – Path to the root of the training data. If this is `auto`, this class will automatically download data for you under `$CHAINER_DATASET_ROOT/pfnet/chainercv/cub`.
- **return_bb** (*bool*) – If `True`, this returns a bounding box around a bird. The default value is `False`.
- **prob_map_dir** (*string*) – Path to the root of the probability maps. If this is `auto`, this class will automatically download data for you under `$CHAINER_DATASET_ROOT/pfnet/chainercv/cub`.
- **return_prob_map** (*bool*) – Decide whether to include a probability map of the bird in a tuple served for a query. The default value is `False`.

This dataset returns the following data.

name	shape	dtype	format
`img`	$(3, H, W)$	`float32`	RGB, $[0, 255]$
`label`	scalar	`int32`	$[0, \#class - 1]$
`bb`[6]	$(4,)$	`float32`	$(y_{min}, x_{min}, y_{max}, x_{max})$
`prob_map`[7]	(H, W)	`float32`	$[0, 1]$

CUBPointDataset

class chainercv.datasets.**CUBPointDataset**(*data_dir='auto'*, *return_bb=False*, *prob_map_dir='auto'*, *return_prob_map=False*)

Caltech-UCSD Birds-200-2011 dataset with annotated points.

Parameters

- **data_dir** (*string*) – Path to the root of the training data. If this is `auto`, this class will automatically download data for you under `$CHAINER_DATASET_ROOT/pfnet/chainercv/cub`.
- **return_bb** (*bool*) – If `True`, this returns a bounding box around a bird. The default value is `False`.
- **prob_map_dir** (*string*) – Path to the root of the probability maps. If this is `auto`, this class will automatically download data for you under `$CHAINER_DATASET_ROOT/pfnet/chainercv/cub`.
- **return_prob_map** (*bool*) – Decide whether to include a probability map of the bird in a tuple served for a query. The default value is `False`.

[6] `bb` indicates the location of a bird. It is available if `return_bb = True`.
[7] `prob_map` indicates how likey a bird is located at each the pixel. It is available if `return_prob_map = True`.

This dataset returns the following data.

name	shape	dtype	format
`img`	$(3, H, W)$	`float32`	RGB, $[0, 255]$
`point`	$(P, 2)$	`float32`	(y, x)
`mask`	$(P,)$	`bool`	–
`bb`[8]	$(4,)$	`float32`	$(y_{min}, x_{min}, y_{max}, x_{max})$
`prob_map`[9]	(H, W)	`float32`	$[0, 1]$

3.2.6 MS COCO

COCOBboxDataset

class chainercv.datasets.**COCOBboxDataset**(*data_dir='auto'*, *split='train'*, *year='2017'*, *use_crowded=False*, *return_area=False*, *return_crowded=False*)

Bounding box dataset for MS COCO.

Parameters

- **data_dir** (*string*) – Path to the root of the training data. If this is `auto`, this class will automatically download data for you under `$CHAINER_DATASET_ROOT/pfnet/chainercv/coco`.
- **split** (*{'train', 'val', 'minival', 'valminusminival'}*) – Select a split of the dataset.
- **year** (*{'2014', '2017'}*) – Use a dataset released in `year`. Splits `minival` and `valminusminival` are only supported in year `2014`.
- **use_crowded** (*bool*) – If true, use bounding boxes that are labeled as crowded in the original annotation. The default value is `False`.
- **return_area** (*bool*) – If true, this dataset returns areas of masks around objects. The default value is `False`.
- **return_crowded** (*bool*) – If true, this dataset returns a boolean array that indicates whether bounding boxes are labeled as crowded or not. The default value is `False`.

This dataset returns the following data.

name	shape	dtype	format
`img`	$(3, H, W)$	`float32`	RGB, $[0, 255]$
`bbox`[10]	$(R, 4)$	`float32`	$(y_{min}, x_{min}, y_{max}, x_{max})$
`label`[10]	$(R,)$	`int32`	$[0, \#fg_class - 1]$
`area`[10][11]	$(R,)$	`float32`	–
`crowded`[12]	$(R,)$	`bool`	–

When there are more than ten objects from the same category, bounding boxes correspond to crowd of instances instead of individual instances. Please see more detail in the Fig. 12 (e) of the summary paper[13].

[8] `bb` indicates the location of a bird. It is available if `return_bb = True`.
[9] `prob_map` indicates how likey a bird is located at each the pixel. It is available if `return_prob_map = True`.
[10] If `use_crowded = True`, `bbox`, `label` and `area` contain crowded instances.
[11] `area` is available if `return_area = True`.
[12] `crowded` is available if `return_crowded = True`.
[13] Tsung-Yi Lin, Michael Maire, Serge Belongie, Lubomir Bourdev, Ross Girshick, James Hays, Pietro Perona, Deva Ramanan, C. Lawrence

COCOInstanceSegmentationDataset

class chainercv.datasets.**COCOInstanceSegmentationDataset**(*data_dir='auto'*, *split='train'*, *year='2017'*, *use_crowded=False*, *return_crowded=False*, *return_area=False*)

Instance segmentation dataset for MS COCO.

Parameters

- **data_dir** (*string*) – Path to the root of the training data. If this is `auto`, this class will automatically download data for you under `$CHAINER_DATASET_ROOT/pfnet/chainercv/coco`.
- **split** (*{'train', 'val', 'minival', 'valminusminival'}*) – Select a split of the dataset.
- **year** (*{'2014', '2017'}*) – Use a dataset released in `year`. Splits `minival` and `valminusminival` are only supported in year `2014`.
- **use_crowded** (*bool*) – If true, use masks that are labeled as crowded in the original annotation.
- **return_crowded** (*bool*) – If true, this dataset returns a boolean array that indicates whether masks are labeled as crowded or not. The default value is `False`.
- **return_area** (*bool*) – If true, this dataset returns areas of masks around objects.

This dataset returns the following data.

name	shape	dtype	format
`img`	$(3, H, W)$	`float32`	RGB, $[0, 255]$
`mask`[14]	(R, H, W)	`bool`	–
`label`[14]	$(R,)$	`int32`	$[0, \#fg_class - 1]$
`area`[14][15]	$(R,)$	`float32`	–
`crowded`[16]	$(R,)$	`bool`	–

When there are more than ten objects from the same category, masks correspond to crowd of instances instead of individual instances. Please see more detail in the Fig. 12 (e) of the summary paper[17].

COCOSemanticSegmentationDataset

class chainercv.datasets.**COCOSemanticSegmentationDataset**(*data_dir='auto'*, *split='train'*)

Semantic segmentation dataset for MS COCO.

Semantic segmentations are generated from panoptic segmentations as done in the official toolkit.

Parameters

Zitnick, Piotr Dollar. Microsoft COCO: Common Objects in Context. arXiv 2014.

[14] If `use_crowded = True`, `mask`, `label` and `area` contain crowded instances.

[15] `area` is available if `return_area = True`.

[16] `crowded` is available if `return_crowded = True`.

[17] Tsung-Yi Lin, Michael Maire, Serge Belongie, Lubomir Bourdev, Ross Girshick, James Hays, Pietro Perona, Deva Ramanan, C. Lawrence Zitnick, Piotr Dollar. Microsoft COCO: Common Objects in Context. arXiv 2014.

- **data_dir** (*string*) – Path to the root of the training data. If this is `auto`, this class will automatically download data for you under `$CHAINER_DATASET_ROOT/pfnet/chainercv/coco`.
- **split** (*{'train', 'val'}*) – Select a split of the dataset.

This dataset returns the following data.

name	shape	dtype	format
`img`	$(3, H, W)$	`float32`	RGB, $[0, 255]$
`label`	(H, W)	`int32`	$[-1, \#class - 1]$

3.2.7 OnlineProducts

OnlineProductsDataset

class chainercv.datasets.**OnlineProductsDataset**(*data_dir='auto'*, *split='train'*)

Dataset class for Stanford Online Products Dataset.

The `split` selects train and test split of the dataset as done in[18]. The train split contains the first 11318 classes and the test split contains the remaining 11316 classes.

Parameters

- **data_dir** (*string*) – Path to the root of the training data. If this is `auto`, this class will automatically download data for you under `$CHAINER_DATASET_ROOT/pfnet/chainercv/online_products`.
- **split** (*{'train', 'test'}*) – Select a split of the dataset.

This dataset returns the following data.

name	shape	dtype	format
`img`	$(3, H, W)$	`float32`	RGB, $[0, 255]$
`label`	scalar	`int32`	$[0, \#class - 1]$
`super_label`	scalar	`int32`	$[0, \#super_class - 1]$

3.2.8 PASCAL VOC

VOCBboxDataset

class chainercv.datasets.**VOCBboxDataset**(*data_dir='auto'*, *split='train'*, *year='2012'*, *use_difficult=False*, *return_difficult=False*)

Bounding box dataset for PASCAL VOC.

Parameters

- **data_dir** (*string*) – Path to the root of the training data. If this is `auto`, this class will automatically download data for you under `$CHAINER_DATASET_ROOT/pfnet/chainercv/voc`.
- **split** (*{'train', 'val', 'trainval', 'test'}*) – Select a split of the dataset. `test` split is only available for 2007 dataset.
- **year** (*{'2007', '2012'}*) – Use a dataset prepared for a challenge held in `year`.

[18] Hyun Oh Song, Yu Xiang, Stefanie Jegelka, Silvio Savarese. Deep Metric Learning via Lifted Structured Feature Embedding. arXiv 2015.

- **use_difficult** (*bool*) – If `True`, use images that are labeled as difficult in the original annotation.
- **return_difficult** (*bool*) – If `True`, this dataset returns a boolean array that indicates whether bounding boxes are labeled as difficult or not. The default value is `False`.

This dataset returns the following data.

name	shape	dtype	format
`img`	$(3, H, W)$	`float32`	RGB, $[0, 255]$
`bbox`[19]	$(R, 4)$	`float32`	$(y_{min}, x_{min}, y_{max}, x_{max})$
`label`[19]	$(R,)$	`int32`	$[0, \#fg_class - 1]$
`difficult` (optional[20])	$(R,)$	`bool`	–

VOCInstanceSegmentationDataset

class chainercv.datasets.**VOCInstanceSegmentationDataset**(*data_dir='auto'*, *split='train'*)

Instance segmentation dataset for PASCAL VOC2012.

Parameters

- **data_dir** (*string*) – Path to the root of the training data. If this is `auto`, this class will automatically download data for you under `$CHAINER_DATASET_ROOT/pfnet/chainercv/voc`.
- **split** (*{'train', 'val', 'trainval'}*) – Select a split of the dataset.

This dataset returns the following data.

name	shape	dtype	format
`img`	$(3, H, W)$	`float32`	RGB, $[0, 255]$
`mask`	(R, H, W)	`bool`	–
`label`	$(R,)$	`int32`	$[0, \#fg_class - 1]$

VOCSemanticSegmentationDataset

class chainercv.datasets.**VOCSemanticSegmentationDataset**(*data_dir='auto'*, *split='train'*)

Semantic segmentation dataset for PASCAL VOC2012.

Parameters

- **data_dir** (*string*) – Path to the root of the training data. If this is `auto`, this class will automatically download data for you under `$CHAINER_DATASET_ROOT/pfnet/chainercv/voc`.
- **split** (*{'train', 'val', 'trainval'}*) – Select a split of the dataset.

This dataset returns the following data.

name	shape	dtype	format
`img`	$(3, H, W)$	`float32`	RGB, $[0, 255]$
`label`	(H, W)	`int32`	$[-1, \#class - 1]$

[19] If `use_difficult = True`, `bbox` and `label` contain difficult instances.
[20] `difficult` is available if `return_difficult = True`.

3.2.9 Semantic Boundaries Dataset

SBDInstanceSegmentationDataset

class chainercv.datasets.**SBDInstanceSegmentationDataset**(*data_dir='auto'*, *split='train'*)

Instance segmentation dataset for Semantic Boundaries Dataset SBD.

Parameters

- **data_dir** (*string*) – Path to the root of the training data. If this is `auto`, this class will automatically download data for you under `$CHAINER_DATASET_ROOT/pfnet/chainercv/sbd`.
- **split** (*{'train', 'val', 'trainval'}*) – Select a split of the dataset.

This dataset returns the following data.

name	shape	dtype	format
`img`	$(3, H, W)$	`float32`	RGB, $[0, 255]$
`mask`	(R, H, W)	`bool`	–
`label`	$(R,)$	`int32`	$[0, \#fg_class - 1]$

3.3 Evaluations

3.3.1 Detection COCO

eval_detection_coco

chainercv.evaluations.**eval_detection_coco**(*pred_bboxes*, *pred_labels*, *pred_scores*, *gt_bboxes*, *gt_labels*, *gt_areas=None*, *gt_crowdeds=None*)

Evaluate detections based on evaluation code of MS COCO.

This function evaluates predicted bounding boxes obtained from a dataset by using average precision for each class. The code is based on the evaluation code used in MS COCO.

Parameters

- **pred_bboxes** (*iterable of numpy.ndarray*) – See the table below.
- **pred_labels** (*iterable of numpy.ndarray*) – See the table below.
- **pred_scores** (*iterable of numpy.ndarray*) – See the table below.
- **gt_bboxes** (*iterable of numpy.ndarray*) – See the table below.
- **gt_labels** (*iterable of numpy.ndarray*) – See the table below.
- **gt_areas** (*iterable of numpy.ndarray*) – See the table below. If `None`, some scores are not returned.
- **gt_crowdeds** (*iterable of numpy.ndarray*) – See the table below.

name	shape	dtype	format
pred_bboxes	$[(R, 4)]$	float32	$(y_{min}, x_{min}, y_{max}, x_{max})$
pred_labels	$[(R,)]$	int32	$[0, \#fg_class - 1]$
pred_scores	$[(R,)]$	float32	–
gt_bboxes	$[(R, 4)]$	float32	$(y_{min}, x_{min}, y_{max}, x_{max})$
gt_labels	$[(R,)]$	int32	$[0, \#fg_class - 1]$
gt_areas	$[(R,)]$	float32	–
gt_crowdeds	$[(R,)]$	bool	–

All inputs should have the same length. For more detailed explanation of the inputs, please refer to *chainercv.datasets.COCOBboxDataset*.

See also:

chainercv.datasets.COCOBboxDataset.

Returns

The keys, value-types and the description of the values are listed below. The APs and ARs calculated with different iou thresholds, sizes of objects, and numbers of detections per image. For more details on the 12 patterns of evaluation metrics, please refer to COCO's official evaluation page.

key	type	description
ap/iou=0.50:0.95/area=all/max_dets=100	*numpy.ndarray*	[1]
ap/iou=0.50/area=all/max_dets=100	*numpy.ndarray*	[1]
ap/iou=0.75/area=all/max_dets=100	*numpy.ndarray*	[1]
ap/iou=0.50:0.95/area=small/max_dets=100	*numpy.ndarray*	[1][5]
ap/iou=0.50:0.95/area=medium/max_dets=100	*numpy.ndarray*	[1][5]
ap/iou=0.50:0.95/area=large/max_dets=100	*numpy.ndarray*	[1][5]
ar/iou=0.50:0.95/area=all/max_dets=1	*numpy.ndarray*	[2]
ar/iou=0.50/area=all/max_dets=10	*numpy.ndarray*	[2]
ar/iou=0.75/area=all/max_dets=100	*numpy.ndarray*	[2]
ar/iou=0.50:0.95/area=small/max_dets=100	*numpy.ndarray*	[2][5]
ar/iou=0.50:0.95/area=medium/max_dets=100	*numpy.ndarray*	[2][5]
ar/iou=0.50:0.95/area=large/max_dets=100	*numpy.ndarray*	[2][5]
map/iou=0.50:0.95/area=all/max_dets=100	*float*	[3]
map/iou=0.50/area=all/max_dets=100	*float*	[3]
map/iou=0.75/area=all/max_dets=100	*float*	[3]
map/iou=0.50:0.95/area=small/max_dets=100	*float*	[3][5]
map/iou=0.50:0.95/area=medium/max_dets=100	*float*	[3][5]
map/iou=0.50:0.95/area=large/max_dets=100	*float*	[3][5]
mar/iou=0.50:0.95/area=all/max_dets=1	*float*	[4]
mar/iou=0.50/area=all/max_dets=10	*float*	[4]
mar/iou=0.75/area=all/max_dets=100	*float*	[4]
mar/iou=0.50:0.95/area=small/max_dets=100	*float*	[4][5]
mar/iou=0.50:0.95/area=medium/max_dets=100	*float*	[4][5]
mar/iou=0.50:0.95/area=large/max_dets=100	*float*	[4][5]
coco_eval	*pycocotools.cocoeval.COCOeval*	result from `pycocotools`
existent_labels	*numpy.ndarray*	used labels

[1] An array of average precisions. The l-th value corresponds to the average precision for class l. If class l does not exist in either `pred_labels`

Return type dict

3.3.2 Detection VOC

eval_detection_voc

chainercv.evaluations.**eval_detection_voc**(*pred_bboxes*, *pred_labels*, *pred_scores*, *gt_bboxes*, *gt_labels*, *gt_difficults=None*, *iou_thresh=0.5*, *use_07_metric=False*)

Calculate average precisions based on evaluation code of PASCAL VOC.

This function evaluates predicted bounding boxes obtained from a dataset which has N images by using average precision for each class. The code is based on the evaluation code used in PASCAL VOC Challenge.

Parameters

- **pred_bboxes** (*iterable of numpy.ndarray*) – An iterable of N sets of bounding boxes. Its index corresponds to an index for the base dataset. Each element of `pred_bboxes` is a set of coordinates of bounding boxes. This is an array whose shape is $(R, 4)$, where R corresponds to the number of bounding boxes, which may vary among boxes. The second axis corresponds to $y_{min}, x_{min}, y_{max}, x_{max}$ of a bounding box.
- **pred_labels** (*iterable of numpy.ndarray*) – An iterable of labels. Similar to `pred_bboxes`, its index corresponds to an index for the base dataset. Its length is N.
- **pred_scores** (*iterable of numpy.ndarray*) – An iterable of confidence scores for predicted bounding boxes. Similar to `pred_bboxes`, its index corresponds to an index for the base dataset. Its length is N.
- **gt_bboxes** (*iterable of numpy.ndarray*) – An iterable of ground truth bounding boxes whose length is N. An element of `gt_bboxes` is a bounding box whose shape is $(R, 4)$. Note that the number of bounding boxes in each image does not need to be same as the number of corresponding predicted boxes.
- **gt_labels** (*iterable of numpy.ndarray*) – An iterable of ground truth labels which are organized similarly to `gt_bboxes`.
- **gt_difficults** (*iterable of numpy.ndarray*) – An iterable of boolean arrays which is organized similarly to `gt_bboxes`. This tells whether the corresponding ground truth bounding box is difficult or not. By default, this is `None`. In that case, this function considers all bounding boxes to be not difficult.
- **iou_thresh** (*float*) – A prediction is correct if its Intersection over Union with the ground truth is above this value.
- **use_07_metric** (*bool*) – Whether to use PASCAL VOC 2007 evaluation metric for calculating average precision. The default value is `False`.

Returns

The keys, value-types and the description of the values are listed below.

or `gt_labels`, the corresponding value is set to `numpy.nan`.

[5] Skip if `gt_areas` is `None`.

[2] An array of average recalls. The l-th value corresponds to the average precision for class l. If class l does not exist in either `pred_labels` or `gt_labels`, the corresponding value is set to `numpy.nan`.

[3] The average of average precisions over classes.

[4] The average of average recalls over classes.

- **ap** (*numpy.ndarray*): An array of average precisions. The l-th value corresponds to the average precision for class l. If class l does not exist in either `pred_labels` or `gt_labels`, the corresponding value is set to `numpy.nan`.
- **map** (*float*): The average of Average Precisions over classes.

Return type dict

calc_detection_voc_ap

chainercv.evaluations.**calc_detection_voc_ap**(*prec*, *rec*, *use_07_metric=False*)

Calculate average precisions based on evaluation code of PASCAL VOC.

This function calculates average precisions from given precisions and recalls. The code is based on the evaluation code used in PASCAL VOC Challenge.

Parameters

- **prec** (*list of numpy.array*) – A list of arrays. `prec[l]` indicates precision for class l. If `prec[l]` is `None`, this function returns `numpy.nan` for class l.
- **rec** (*list of numpy.array*) – A list of arrays. `rec[l]` indicates recall for class l. If `rec[l]` is `None`, this function returns `numpy.nan` for class l.
- **use_07_metric** (*bool*) – Whether to use PASCAL VOC 2007 evaluation metric for calculating average precision. The default value is `False`.

Returns This function returns an array of average precisions. The l-th value corresponds to the average precision for class l. If `prec[l]` or `rec[l]` is `None`, the corresponding value is set to `numpy.nan`.

Return type ndarray

calc_detection_voc_prec_rec

chainercv.evaluations.**calc_detection_voc_prec_rec**(*pred_bboxes*, *pred_labels*, *pred_scores*, *gt_bboxes*, *gt_labels*, *gt_difficults=None*, *iou_thresh=0.5*)

Calculate precision and recall based on evaluation code of PASCAL VOC.

This function calculates precision and recall of predicted bounding boxes obtained from a dataset which has N images. The code is based on the evaluation code used in PASCAL VOC Challenge.

Parameters

- **pred_bboxes** (*iterable of numpy.ndarray*) – An iterable of N sets of bounding boxes. Its index corresponds to an index for the base dataset. Each element of `pred_bboxes` is a set of coordinates of bounding boxes. This is an array whose shape is $(R, 4)$, where R corresponds to the number of bounding boxes, which may vary among boxes. The second axis corresponds to $y_{min}, x_{min}, y_{max}, x_{max}$ of a bounding box.
- **pred_labels** (*iterable of numpy.ndarray*) – An iterable of labels. Similar to `pred_bboxes`, its index corresponds to an index for the base dataset. Its length is N.
- **pred_scores** (*iterable of numpy.ndarray*) – An iterable of confidence scores for predicted bounding boxes. Similar to `pred_bboxes`, its index corresponds to an index for the base dataset. Its length is N.

- **gt_bboxes** (*iterable of numpy.ndarray*) – An iterable of ground truth bounding boxes whose length is N. An element of `gt_bboxes` is a bounding box whose shape is $(R, 4)$. Note that the number of bounding boxes in each image does not need to be same as the number of corresponding predicted boxes.
- **gt_labels** (*iterable of numpy.ndarray*) – An iterable of ground truth labels which are organized similarly to `gt_bboxes`.
- **gt_difficults** (*iterable of numpy.ndarray*) – An iterable of boolean arrays which is organized similarly to `gt_bboxes`. This tells whether the corresponding ground truth bounding box is difficult or not. By default, this is `None`. In that case, this function considers all bounding boxes to be not difficult.
- **iou_thresh** (*float*) – A prediction is correct if its Intersection over Union with the ground truth is above this value..

Returns

This function returns two lists: `prec` and `rec`.

- `prec`: A list of arrays. `prec[l]` is precision for class l. If class l does not exist in either `pred_labels` or `gt_labels`, `prec[l]` is set to `None`.
- `rec`: A list of arrays. `rec[l]` is recall for class l. If class l that is not marked as difficult does not exist in `gt_labels`, `rec[l]` is set to `None`.

Return type tuple of two lists

3.3.3 Instance Segmentation COCO

eval_instance_segmentation_coco

chainercv.evaluations.**eval_instance_segmentation_coco**(*pred_masks*, *pred_labels*, *pred_scores*, *gt_masks*, *gt_labels*, *gt_areas=None*, *gt_crowdeds=None*)

Evaluate instance segmentations based on evaluation code of MS COCO.

This function evaluates predicted instance segmentations obtained from a dataset by using average precision for each class. The code is based on the evaluation code used in MS COCO.

Parameters

- **pred_masks** (*iterable of numpy.ndarray*) – See the table below.
- **pred_labels** (*iterable of numpy.ndarray*) – See the table below.
- **pred_scores** (*iterable of numpy.ndarray*) – See the table below.
- **gt_masks** (*iterable of numpy.ndarray*) – See the table below.
- **gt_labels** (*iterable of numpy.ndarray*) – See the table below.
- **gt_areas** (*iterable of numpy.ndarray*) – See the table below. If `None`, some scores are not returned.
- **gt_crowdeds** (*iterable of numpy.ndarray*) – See the table below.

name	shape	dtype	format
pred_masks	$[(R, H, W)]$	bool	–
pred_labels	$[(R,)]$	int32	$[0, \#fg_class - 1]$
pred_scores	$[(R,)]$	float32	–
gt_masks	$[(R, H, W)]$	bool	–
gt_labels	$[(R,)]$	int32	$[0, \#fg_class - 1]$
gt_areas	$[(R,)]$	float32	–
gt_crowdeds	$[(R,)]$	bool	–

All inputs should have the same length. For more detailed explanation of the inputs, please refer to *chainercv.datasets.COCOInstanceSegmentationDataset*.

See also:

chainercv.datasets.COCOInstanceSegmentationDataset.

Returns

The keys, value-types and the description of the values are listed below. The APs and ARs calculated with different iou thresholds, sizes of objects, and numbers of detections per image. For more details on the 12 patterns of evaluation metrics, please refer to COCO's official evaluation page.

key	type	description
ap/iou=0.50:0.95/area=all/max_dets=100	*numpy.ndarray*	[6]
ap/iou=0.50/area=all/max_dets=100	*numpy.ndarray*	[6]
ap/iou=0.75/area=all/max_dets=100	*numpy.ndarray*	[6]
ap/iou=0.50:0.95/area=small/max_dets=100	*numpy.ndarray*	[6][10]
ap/iou=0.50:0.95/area=medium/max_dets=100	*numpy.ndarray*	[6][10]
ap/iou=0.50:0.95/area=large/max_dets=100	*numpy.ndarray*	[6][10]
ar/iou=0.50:0.95/area=all/max_dets=1	*numpy.ndarray*	[7]
ar/iou=0.50/area=all/max_dets=10	*numpy.ndarray*	[7]
ar/iou=0.75/area=all/max_dets=100	*numpy.ndarray*	[7]
ar/iou=0.50:0.95/area=small/max_dets=100	*numpy.ndarray*	[7][10]
ar/iou=0.50:0.95/area=medium/max_dets=100	*numpy.ndarray*	[7][10]
ar/iou=0.50:0.95/area=large/max_dets=100	*numpy.ndarray*	[7][10]
map/iou=0.50:0.95/area=all/max_dets=100	*float*	[8]
map/iou=0.50/area=all/max_dets=100	*float*	[8]
map/iou=0.75/area=all/max_dets=100	*float*	[8]
map/iou=0.50:0.95/area=small/max_dets=100	*float*	[8][10]
map/iou=0.50:0.95/area=medium/max_dets=100	*float*	[8][10]
map/iou=0.50:0.95/area=large/max_dets=100	*float*	[8][10]
mar/iou=0.50:0.95/area=all/max_dets=1	*float*	[9]
mar/iou=0.50/area=all/max_dets=10	*float*	[9]
mar/iou=0.75/area=all/max_dets=100	*float*	[9]
mar/iou=0.50:0.95/area=small/max_dets=100	*float*	[9][10]
mar/iou=0.50:0.95/area=medium/max_dets=100	*float*	[9][10]
mar/iou=0.50:0.95/area=large/max_dets=100	*float*	[9][10]
coco_eval	*pycocotools.cocoeval.COCOeval*	result from pycocotools
existent_labels	*numpy.ndarray*	used labels

[6] An array of average precisions. The l-th value corresponds to the average precision for class l. If class l does not exist in either pred_labels

Return type dict

3.3.4 Instance Segmentation VOC

eval_instance_segmentation_voc

chainercv.evaluations.**eval_instance_segmentation_voc**(*pred_masks*, *pred_labels*, *pred_scores*, *gt_masks*, *gt_labels*, *iou_thresh=0.5*, *use_07_metric=False*)

Calculate average precisions based on evaluation code of PASCAL VOC.

This function evaluates predicted masks obtained from a dataset which has N images by using average precision for each class. The code is based on the evaluation code used in FCIS.

Parameters

- **pred_masks** (*iterable of numpy.ndarray*) – An iterable of N sets of masks. Its index corresponds to an index for the base dataset. Each element of `pred_masks` is an object mask and is an array whose shape is (R, H, W), where R corresponds to the number of masks, which may vary among images.
- **pred_labels** (*iterable of numpy.ndarray*) – An iterable of labels. Similar to `pred_masks`, its index corresponds to an index for the base dataset. Its length is N.
- **pred_scores** (*iterable of numpy.ndarray*) – An iterable of confidence scores for predicted masks. Similar to `pred_masks`, its index corresponds to an index for the base dataset. Its length is N.
- **gt_masks** (*iterable of numpy.ndarray*) – An iterable of ground truth masks whose length is N. An element of `gt_masks` is an object mask whose shape is (R, H, W). Note that the number of masks R in each image does not need to be same as the number of corresponding predicted masks.
- **gt_labels** (*iterable of numpy.ndarray*) – An iterable of ground truth labels which are organized similarly to `gt_masks`. Its length is N.
- **iou_thresh** (*float*) – A prediction is correct if its Intersection over Union with the ground truth is above this value.
- **use_07_metric** (*bool*) – Whether to use PASCAL VOC 2007 evaluation metric for calculating average precision. The default value is `False`.

Returns

The keys, value-types and the description of the values are listed below.

- **ap** (*numpy.ndarray*): An array of average precisions. The l-th value corresponds to the average precision for class l. If class l does not exist in either `pred_labels` or `gt_labels`, the corresponding value is set to `numpy.nan`.
- **map** (*float*): The average of Average Precisions over classes.

Return type dict

or `gt_labels`, the corresponding value is set to `numpy.nan`.

[10] Skip if `gt_areas` is `None`.

[7] An array of average recalls. The l-th value corresponds to the average precision for class l. If class l does not exist in either `pred_labels` or `gt_labels`, the corresponding value is set to `numpy.nan`.

[8] The average of average precisions over classes.

[9] The average of average recalls over classes.

calc_instance_segmentation_voc_prec_rec

chainercv.evaluations.**calc_instance_segmentation_voc_prec_rec**(*pred_masks*, *pred_labels*, *pred_scores*, *gt_masks*, *gt_labels*, *iou_thresh*)

Calculate precision and recall based on evaluation code of PASCAL VOC.

This function calculates precision and recall of predicted masks obtained from a dataset which has N images. The code is based on the evaluation code used in FCIS.

Parameters

- **pred_masks** (*iterable of numpy.ndarray*) – An iterable of N sets of masks. Its index corresponds to an index for the base dataset. Each element of `pred_masks` is an object mask and is an array whose shape is (R, H, W), where R corresponds to the number of masks, which may vary among images.
- **pred_labels** (*iterable of numpy.ndarray*) – An iterable of labels. Similar to `pred_masks`, its index corresponds to an index for the base dataset. Its length is N.
- **pred_scores** (*iterable of numpy.ndarray*) – An iterable of confidence scores for predicted masks. Similar to `pred_masks`, its index corresponds to an index for the base dataset. Its length is N.
- **gt_masks** (*iterable of numpy.ndarray*) – An iterable of ground truth masks whose length is N. An element of `gt_masks` is an object mask whose shape is (R, H, W). Note that the number of masks R in each image does not need to be same as the number of corresponding predicted masks.
- **gt_labels** (*iterable of numpy.ndarray*) – An iterable of ground truth labels which are organized similarly to `gt_masks`. Its length is N.
- **iou_thresh** (*float*) – A prediction is correct if its Intersection over Union with the ground truth is above this value.

Returns

This function returns two lists: `prec` and `rec`.

- `prec`: A list of arrays. `prec[l]` is precision for class l. If class l does not exist in either `pred_labels` or `gt_labels`, `prec[l]` is set to `None`.
- `rec`: A list of arrays. `rec[l]` is recall for class l. If class l that is not marked as difficult does not exist in `gt_labels`, `rec[l]` is set to `None`.

Return type tuple of two lists

3.3.5 Semantic Segmentation IoU

eval_semantic_segmentation

chainercv.evaluations.**eval_semantic_segmentation**(*pred_labels*, *gt_labels*)

Evaluate metrics used in Semantic Segmentation.

This function calculates Intersection over Union (IoU), Pixel Accuracy and Class Accuracy for the task of semantic segmentation.

The definition of metrics calculated by this function is as follows, where N_{ij} is the number of pixels that are labeled as class i by the ground truth and class j by the prediction.

- IoU of the i-th class $= \frac{N_{ii}}{\sum_{j=1}^{k} N_{ij} + \sum_{j=1}^{k} N_{ji} - N_{ii}}$
- mIoU $= \frac{1}{k} \sum_{i=1}^{k} \frac{N_{ii}}{\sum_{j=1}^{k} N_{ij} + \sum_{j=1}^{k} N_{ji} - N_{ii}}$
- Pixel Accuracy $= \frac{\sum_{i=1}^{k} N_{ii}}{\sum_{i=1}^{k} \sum_{j=1}^{k} N_{ij}}$
- Class Accuracy $= \frac{N_{ii}}{\sum_{j=1}^{k} N_{ij}}$
- Mean Class Accuracy $= \frac{1}{k} \sum_{i=1}^{k} \frac{N_{ii}}{\sum_{j=1}^{k} N_{ij}}$

The more detailed description of the above metrics can be found in a review on semantic segmentation[11].

The number of classes n_class is $max(pred_labels, gt_labels) + 1$, which is the maximum class id of the inputs added by one.

Parameters

- **pred_labels** (*iterable of numpy.ndarray*) – A collection of predicted labels. The shape of a label array is (H, W). H and W are height and width of the label. For example, this is a list of labels `[label_0, label_1, ...]`, where `label_i.shape = (H_i, W_i)`.
- **gt_labels** (*iterable of numpy.ndarray*) – A collection of ground truth labels. The shape of a ground truth label array is (H, W), and its corresponding prediction label should have the same shape. A pixel with value `-1` will be ignored during evaluation.

Returns

The keys, value-types and the description of the values are listed below.

- **iou** (*numpy.ndarray*): An array of IoUs for the n_class classes. Its shape is $(n_class,)$.
- **miou** (*float*): The average of IoUs over classes.
- **pixel_accuracy** (*float*): The computed pixel accuracy.
- **class_accuracy** (*numpy.ndarray*): An array of class accuracies for the n_class classes. Its shape is $(n_class,)$.
- **mean_class_accuracy** (*float*): The average of class accuracies.

Return type dict

calc_semantic_segmentation_confusion

chainercv.evaluations.**calc_semantic_segmentation_confusion**(*pred_labels*, *gt_labels*)

Collect a confusion matrix.

The number of classes n_class is $max(pred_labels, gt_labels) + 1$, which is the maximum class id of the inputs added by one.

Parameters

- **pred_labels** (*iterable of numpy.ndarray*) – A collection of predicted labels. The shape of a label array is (H, W). H and W are height and width of the label.

[11] Alberto Garcia-Garcia, Sergio Orts-Escolano, Sergiu Oprea, Victor Villena-Martinez, Jose Garcia-Rodriguez. A Review on Deep Learning Techniques Applied to Semantic Segmentation. arXiv 2017.

- **gt_labels** (*iterable of numpy.ndarray*) – A collection of ground truth labels. The shape of a ground truth label array is (H, W), and its corresponding prediction label should have the same shape. A pixel with value `-1` will be ignored during evaluation.

Returns A confusion matrix. Its shape is (n_class, n_class). The (i, j) th element corresponds to the number of pixels that are labeled as class i by the ground truth and class j by the prediction.

Return type numpy.ndarray

calc_semantic_segmentation_iou

chainercv.evaluations.**calc_semantic_segmentation_iou**(*confusion*)

Calculate Intersection over Union with a given confusion matrix.

The definition of Intersection over Union (IoU) is as follows, where N_{ij} is the number of pixels that are labeled as class i by the ground truth and class j by the prediction.

- IoU of the i-th class $= \frac{N_{ii}}{\sum_{j=1}^{k} N_{ij} + \sum_{j=1}^{k} N_{ji} - N_{ii}}$

Parameters **confusion** (*numpy.ndarray*) – A confusion matrix. Its shape is (n_class, n_class). The (i, j) th element corresponds to the number of pixels that are labeled as class i by the ground truth and class j by the prediction.

Returns An array of IoUs for the n_class classes. Its shape is $(n_class,)$.

Return type numpy.ndarray

3.4 Experimental

3.4.1 Links

Detection

Detection links share a common method `predict()` to detect objects in images.

YOLO

Object Detection Link

YOLOv2Tiny

class chainercv.experimental.links.model.yolo.**YOLOv2Tiny**(*n_fg_class=None*, *pretrained_model=None*)

YOLOv2 tiny.

This is a model of YOLOv2 tiny a.k.a. Tiny YOLO. This model uses `DarknetExtractor` as its feature extractor.

Parameters

- **n_fg_class** (*int*) – The number of classes excluding the background.
- **pretrained_model** (*string*) – The weight file to be loaded. This can take `'voc0712'`, *filepath* or `None`. The default value is `None`.

- `'voc0712'`: Load weights trained on trainval split of PASCAL VOC 2007 and 2012. The weight file is downloaded and cached automatically. `n_fg_class` must be `20` or `None`. These weights were converted from the darknet model provided by the original implementation. The conversion code is *chainercv/examples/yolo/darknet2npz.py*.
- *filepath*: A path of npz file. In this case, `n_fg_class` must be specified properly.
- `None`: Do not load weights.

Utility

DarknetExtractor

class chainercv.experimental.links.model.yolo.**DarknetExtractor**

A Darknet based feature extractor for YOLOv2Tiny.

This is a feature extractor for `YOLOv2Tiny`

__call__(*x*)

Compute a feature map from a batch of images.

Parameters **x** (*ndarray*) – An array holding a batch of images. The images should be resized to 416×416.

Returns

Return type Variable

Semantic Segmentation

Semantic segmentation links share a common method `predict()` to conduct semantic segmentation of images.

PSPNet

Semantic Segmentation Link

PSPNetResNet101

class chainercv.experimental.links.model.pspnet.**PSPNetResNet101**(*n_class=None*, *pretrained_model=None*, *input_size=None*, *initialW=None*)

PSPNet with Dilated ResNet101 as the feature extractor.

See also:

chainercv.experimental.links.model.pspnet.PSPNet

PSPNetResNet50

class chainercv.experimental.links.model.pspnet.**PSPNetResNet50**(*n_class=None*, *pretrained_model=None*, *input_size=None*, *initialW=None*)

PSPNet with Dilated ResNet50 as the feature extractor.

See also:

chainercv.experimental.links.model.pspnet.PSPNet

Utility

convolution_crop

chainercv.experimental.links.model.pspnet.**convolution_crop**(*img*, *size*, *stride*, *return_param=False*)

Strided cropping.

This extracts cropped images from the input. The cropped images are extracted from the entire image, while taking a constant steps between neighboring patches.

Parameters

- **img** (*ndarray*) – An image array to be cropped. This is in CHW format.
- **size** (*tuple*) – The size of output image after cropping. This value is $(height, width)$.
- **stride** (*tuple*) – The stride between crops. This contains two values: stride in the vertical and horizontal directions.
- **return_param** (*bool*) – If `True`, this function returns information of slices.

Returns

If `return_param = False`, returns an array `crop_imgs` that is a stack of cropped images.

If `return_param = True`, returns a tuple whose elements are `crop_imgs, param`. `param` is a dictionary of intermediate parameters whose contents are listed below with key, value-type and the description of the value.

- **y_slices** (*list slices*): Slices used to crop the input image. The relation below holds together with `x_slices`.
- **x_slices** (*list of slices*): Similar to `y_slices`.
- **crop_y_slices** (*list of slices*): This indicates the region of the cropped image that is actually extracted from the input. This is relevant only when borders of the input are cropped.
- **crop_x_slices** (*list of slices*): Similar to `crop_y_slices`.

```
crop_img = crop_imgs[i][:, crop_y_slices[i], crop_x_slices[i]]
crop_img == img[:, y_slices[i], x_slices[i]]
```

Return type ndarray or (ndarray, dict)

Examples

```
>>> import numpy as np
>>> from chainercv.datasets import VOCBboxDataset
>>> from chainercv.transforms import resize
>>> from chainercv.experimental.links.model.pspnet import          ...     ␣
↪convolution_crop
>>>
>>> img, _, _ = VOCBboxDataset(year='2007')[0]
>>> img = resize(img, (300, 300))
>>> imgs, param = convolution_crop(
>>>     img, (128, 128), (96, 96), return_param=True)
>>> # Restore the original image from the cropped images.
>>> output = np.zeros((3, 300, 300))
>>> count = np.zeros((300, 300))
>>> for i in range(len(imgs)):
>>>     crop_y_slice = param['crop_y_slices'][i]
>>>     crop_x_slice = param['crop_x_slices'][i]
>>>     y_slice = param['y_slices'][i]
>>>     x_slice = param['x_slices'][i]
>>>     output[:, y_slice, x_slice] +=        ...         imgs[i][:, crop_y_slice,
↪ crop_x_slice]
>>>     count[y_slice, x_slice] += 1
>>> output = output / count[None]
>>> np.testing.assert_equal(output, img)
>>>
>>> # Visualization of the cropped images
>>> import matplotlib.pyplot as plt
>>> from chainercv.utils import tile_images
>>> from chainercv.visualizations import vis_image
>>> v_imgs = tile_images(imgs, 5, fill=122.5)
>>> vis_image(v_imgs)
>>> plt.show()
```

PSPNet

class chainercv.experimental.links.model.pspnet.**PSPNet**(*n_class=None*, *pretrained_model=None*, *input_size=None*, *initialW=None*)

Pyramid Scene Parsing Network.

This is a PSPNet[1] model for semantic segmentation. This is based on the implementation found here.

Parameters

- **n_class** (*int*) – The number of channels in the last convolution layer.
- **pretrained_model** (*string*) – The weight file to be loaded. This can take `'cityscapes'`, *filepath* or `None`. The default value is `None`. * `'cityscapes'`: Load weights trained on the train split of Cityscapes dataset. `n_class` must be `19` or `None`. * `'ade20k'`: Load weights trained on the train split of ADE20K dataset. `n_class` must be `150` or `None`. * `'imagenet'`: Load ImageNet pretrained weights for the extractor. * *filepath*: A path of npz file. In this case, `n_class` must be specified properly. * `None`: Do not load weights.

[1] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang Jiaya Jia "Pyramid Scene Parsing Network" CVPR, 2017

- **input_size** (*tuple*) – The size of the input. This value is $(height, width)$.
- **initialW** (*callable*) – Initializer for the weights of convolution kernels.

predict(*imgs*)

Conduct semantic segmentation from images.

Parameters **imgs** (*iterable of numpy.ndarray*) – Arrays holding images. All images are in CHW and RGB format and the range of their values are $[0, 255]$.

Returns List of integer labels predicted from each image in the input list.

Return type list of numpy.ndarray

Instance Segmentation

Instance segmentation share a common method `predict()` to detect objects in images. For more details, please read `FCIS.predict()`.

FCIS

Instance Segmentation Link

FCISResNet101

class chainercv.experimental.links.model.fcis.**FCISResNet101**(*n_fg_class=None, pretrained_model=None, min_size=600, max_size=1000, roi_size=21, group_size=7, ratios=[0.5, 1, 2], anchor_scales=[8, 16, 32], loc_normalize_mean=(0.0, 0.0, 0.0, 0.0), loc_normalize_std=(0.2, 0.2, 0.5, 0.5), iter2=True, resnet_initialW=None, rpn_initialW=None, head_initialW=None, proposal_creator_params=None*)

FCIS based on ResNet101.

When you specify the path of a pre-trained chainer model serialized as a `npz` file in the constructor, this chain model automatically initializes all the parameters with it. When a string in prespecified set is provided, a pretrained model is loaded from weights distributed on the Internet. The list of pretrained models supported are as follows:

- `sbd`: Loads weights trained with the trainval split of Semantic Boundaries Dataset.

For descriptions on the interface of this model, please refer to *FCIS*.

FCISResNet101 supports finer control on random initializations of weights by arguments `resnet_initialW`, `rpn_initialW` and `head_initialW`. It accepts a callable that takes an array and edits its values. If `None` is passed as an initializer, the default initializer is used.

Parameters

- **n_fg_class** (*int*) – The number of classes excluding the background.
- **pretrained_model** (*str*) – The destination of the pre-trained chainer model serialized as a `npz` file. If this is one of the strings described above, it automatically loads weights stored under a directory `$CHAINER_DATASET_ROOT/pfnet/chainercv/models/`, where `$CHAINER_DATASET_ROOT` is set as `$HOME/.chainer/dataset` unless you specify another value by modifying the environment variable.
- **min_size** (*int*) – A preprocessing paramter for `prepare()`.
- **max_size** (*int*) – A preprocessing paramter for `prepare()`.
- **roi_size** (*int*) – Height and width of the feature maps after Position Sensitive RoI pooling.
- **group_size** (*int*) – Group height and width for Position Sensitive ROI pooling.
- **ratios** (*list of floats*) – This is ratios of width to height of the anchors.
- **anchor_scales** (*list of numbers*) – This is areas of anchors. Those areas will be the product of the square of an element in `anchor_scales` and the original area of the reference window.
- **loc_normalize_mean** (*tuple of four floats*) – Mean values of localization estimates.
- **loc_normalize_std** (*tupler of four floats*) – Standard deviation of localization estimates.
- **iter2** (*bool*) – if the value is set `True`, Position Sensitive ROI pooling is executed twice. In the second time, Position Sensitive ROI pooling uses improved ROIs by the localization parameters calculated in the first time.
- **resnet_initialW** (*callable*) – Initializer for the layers corresponding to the ResNet101 layers.
- **rpn_initialW** (*callable*) – Initializer for Region Proposal Network layers.
- **head_initialW** (*callable*) – Initializer for the head layers.
- **proposal_creator_params** (*dict*) – Key valued paramters for *ProposalCreator*.

Utility

FCIS

class chainercv.experimental.links.model.fcis.**FCIS**(*extractor*, *rpn*, *head*, *mean*, *min_size*, *max_size*, *loc_normalize_mean*, *loc_normalize_std*)

Base class for FCIS.

This is a base class for FCIS links supporting instance segmentation API[1]. The following three stages constitute FCIS.

1. **Feature extraction**: Images are taken and their feature maps are calculated.
2. **Region Proposal Networks**: Given the feature maps calculated in the previous stage, produce set of RoIs around objects.
3. **Localization, Segmentation and Classification Heads**: Using feature maps that belong to the proposed RoIs, segment regions of the objects, classify the categories of the objects in the RoIs and improve localizations.

Each stage is carried out by one of the callable `chainer.Chain` objects `feature`, `rpn` and `head`. There are two functions *predict()* and *__call__()* to conduct instance segmentation. *predict()* takes images and returns masks, object labels and their scores. *__call__()* is provided for a scnerario when intermediate outputs are needed, for instance, for training and debugging.

Links that support instance segmentation API have method *predict()* with the same interface. Please refer to *predict()* for further details.

Parameters

- **extractor** (*callable Chain*) – A callable that takes a BCHW image array and returns feature maps.
- **rpn** (*callable Chain*) – A callable that has the same interface as *RegionProposalNetwork*. Please refer to the documentation found there.
- **head** (*callable Chain*) – A callable that takes a BCHW array, RoIs and batch indices for RoIs. This returns class-agnostic segmentation scores, class-agnostic localization parameters, class scores, improved RoIs and batch indices for RoIs.
- **mean** (*numpy.ndarray*) – A value to be subtracted from an image in *prepare()*.
- **min_size** (*int*) – A preprocessing parameter for *prepare()*. Please refer to a docstring found for *prepare()*.
- **max_size** (*int*) – A preprocessing parameter for *prepare()*.
- **loc_normalize_mean** (*tuple of four floats*) – Mean values of localization estimates.
- **loc_normalize_std** (*tupler of four floats*) – Standard deviation of localization estimates.

__call__(*x*, *scale=1.0*)

Forward FCIS.

Scaling paramter `scale` is used by RPN to determine the threshold to select small objects, which are going to be rejected irrespective of their confidence scores.

Here are notations used.

- N is the number of batch size
- R' is the total number of RoIs produced across batches. Given R_i proposed RoIs from the i th image, $R' = \sum_{i=1}^{N} R_i$.
- L is the number of classes excluding the background.
- RH is the height of pooled image by Position Sensitive ROI pooling.
- RW is the height of pooled image by Position Sensitive ROI pooling.

[1] Yi Li, Haozhi Qi, Jifeng Dai, Xiangyang Ji, Yichen Wei. Fully Convolutional Instance-aware Semantic Segmentation. CVPR 2017.

Classes are ordered by the background, the first class, ..., and the L th class.

Parameters

- **x** (*Variable*) – 4D image variable.
- **scale** (*float*) – Amount of scaling applied to the raw image during preprocessing.

Returns

Returns tuple of five values listed below.

- **roi_ag_seg_scores**: Class-agnostic clipped mask scores for the proposed ROIs. Its shape is $(R', 2, RH, RW)$
- **ag_locs**: Class-agnostic offsets and scalings for the proposed RoIs. Its shape is $(R', 2, 4)$.
- **roi_cls_scores**: Class predictions for the proposed RoIs. Its shape is $(R', L+1)$.
- **rois**: RoIs proposed by RPN. Its shape is $(R', 4)$.
- **roi_indices**: Batch indices of RoIs. Its shape is $(R',)$.

Return type Variable, Variable, Variable, array, array

predict(*imgs*)

Segment object instances from images.

This method predicts instance-aware object regions for each image.

Parameters **imgs** (*iterable of numpy.ndarray*) – Arrays holding images of shape (B, C, H, W). All images are in CHW and RGB format and the range of their value is $[0, 255]$.

Returns

This method returns a tuple of three lists, `(masks, labels, scores)`.

- **masks**: A list of boolean arrays of shape (R, H, W), where R is the number of masks in a image. Each pixel holds value if it is inside the object inside or not.
- **labels** : A list of integer arrays of shape $(R,)$. Each value indicates the class of the masks. Values are in range $[0, L-1]$, where L is the number of the foreground classes.
- **scores** : A list of float arrays of shape $(R,)$. Each value indicates how confident the prediction is.

Return type tuple of lists

prepare(*img*)

Preprocess an image for feature extraction.

The length of the shorter edge is scaled to `self.min_size`. After the scaling, if the length of the longer edge is longer than `self.max_size`, the image is scaled to fit the longer edge to `self.max_size`.

After resizing the image, the image is subtracted by a mean image value `self.mean`.

Parameters **img** (*ndarray*) – An image. This is in CHW and RGB format. The range of its value is $[0, 255]$.

Returns A preprocessed image.

Return type ndarray

use_preset(*preset*)

Use the given preset during prediction.

This method changes values of `self.nms_thresh`, `self.score_thresh`, `self.mask_merge_thresh`, `self.binary_thresh`, `self.binary_thresh` and `self.min_drop_size`. These values are a threshold value used for non maximum suppression, a threshold value to discard low confidence proposals in *predict()*, a threshold value to merge mask in *predict()*, a threshold value to binalize segmentation scores in *predict()*, a limit number of predicted masks in one image and a threshold value to discard small bounding boxes respectively.

If the attributes need to be changed to something other than the values provided in the presets, please modify them by directly accessing the public attributes.

Parameters **preset** (*{'visualize', 'evaluate'}*) – A string to determine the preset to use.

FCISResNet101Head

class chainercv.experimental.links.model.fcis.**FCISResNet101Head**(*n_class*, *roi_size*, *group_size*, *spatial_scale*, *loc_normalize_mean*, *loc_normalize_std*, *iter2*, *initialW=None*)

FCIS Head for ResNet101 based implementation.

This class is used as a head for FCIS. This outputs class-agnostice segmentation scores, class-agnostic localizations and classification based on feature maps in the given RoIs.

Parameters

- **n_class** (*int*) – The number of classes possibly including the background.
- **roi_size** (*int*) – Height and width of the feature maps after Position Sensitive RoI pooling.
- **group_size** (*int*) – Group height and width for Position Sensitive ROI pooling.
- **spatial_scale** (*float*) – Scale of the roi is resized.
- **loc_normalize_mean** (*tuple of four floats*) – Mean values of localization estimates.
- **loc_normalize_std** (*tupler of four floats*) – Standard deviation of localization estimates.
- **iter2** (*bool*) – if the value is set `True`, Position Sensitive ROI pooling is executed twice. In the second time, Position Sensitive ROI pooling uses improved ROIs by the localization parameters calculated in the first time.
- **initialW** (*callable*) – Initializer for the layers.

mask_voting

chainercv.experimental.links.model.fcis.**mask_voting**(*seg_prob*, *bbox*, *cls_prob*, *size*, *score_thresh*, *nms_thresh*, *mask_merge_thresh*, *binary_thresh*, *limit=100*, *bg_label=0*)

Refine mask probabilities by merging multiple masks.

First, this function discard invalid masks with non maximum suppression. Then, it merges masks with weight calculated from class probabilities and iou. This function improves the mask qualities by merging overlapped masks predicted as the same object class.

Here are notations used. * R is the total number of RoIs produced in one image. * L is the number of classes excluding the background. * RH is the height of pooled image. * RW is the height of pooled image.

Parameters

- **seg_prob** (*array*) – A mask probability array whose shape is (R, RH, RW).
- **bbox** (*array*) – A bounding box array whose shape is $(R, 4)$.
- **cls_prob** (*array*) – A class probability array whose shape is $(R, L + 1)$.
- **size** (*tuple of int*) – Original image size.
- **score_thresh** (*float*) – A threshold value of the class score.
- **nms_thresh** (*float*) – A threshold value of non maximum suppression.
- **mask_merge_thresh** (*float*) – A threshold value of the bounding box iou for mask merging.
- **binary_thresh** (*float*) – A threshold value of mask score for mask merging.
- **limit** (*int*) – The maximum number of outputs.
- **bg_label** (*int*) – The id of the background label.

Returns

- **v_seg_prob**: Merged mask probability. Its shapes is (N, RH, RW).
- **v_bbox**: Bounding boxes for the merged masks. Its shape is $(N, 4)$.
- **v_label**: Class labels for the merged masks. Its shape is $(N,)$.
- **v_score**: Class probabilities for the merged masks. Its shape is $(N,)$.

Return type array, array, array, array

ResNet101Extractor

class chainercv.experimental.links.model.fcis.**ResNet101Extractor**(*initialW=None*)

ResNet101 Extractor for FCIS ResNet101 implementation.

This class is used as an extractor for FCISResNet101. This outputs feature maps. Dilated convolution is used in the C5 stage.

Parameters **initialW** – Initializer for ResNet101 extractor.

Train-only Utility

FCISTrainChain

class chainercv.experimental.links.model.fcis.**FCISTrainChain**(*fcis*, *rpn_sigma=3.0*, *roi_sigma=1.0*, *anchor_target_creator=<chainercv.links.mode object>*, *proposal_target_creator=<chainercv.experime object>*)

Calculate losses for FCIS and report them.

This is used to train FCIS in the joint training scheme[2].

The losses include:

- `rpn_loc_loss`: The localization loss for Region Proposal Network (RPN).
- `rpn_cls_loss`: The classification loss for RPN.
- `roi_loc_loss`: The localization loss for the head module.
- `roi_cls_loss`: The classification loss for the head module.
- `roi_mask_loss`: The mask loss for the head module.

Parameters

- **fcis** (*FCIS*) – A FCIS model for training.
- **rpn_sigma** (*float*) – Sigma parameter for the localization loss of Region Proposal Network (RPN). The default value is 3, which is the value used in[2].
- **roi_sigma** (*float*) – Sigma paramter for the localization loss of the head. The default value is 1, which is the value used in[2].
- **anchor_target_creator** – An instantiation of *AnchorTargetCreator*.
- **proposal_target_creator** – An instantiation of *ProposalTargetCreator*.

__call__(*imgs*, *masks*, *labels*, *bboxes*, *scale*)

Forward FCIS and calculate losses.

Here are notations used.

- N is the batch size.
- R is the number of bounding boxes per image.
- H is the image height.
- W is the image width.

Currently, only $N = 1$ is supported.

Parameters

- **imgs** (*Variable*) – A variable with a batch of images.
- **masks** (*Variable*) – A batch of masks. Its shape is (N, R, H, W).

[2] Yi Li, Haozhi Qi, Jifeng Dai, Xiangyang Ji, Yichen Wei. Fully Convolutional Instance-aware Semantic Segmentation. CVPR 2017.

- **labels** (*Variable*) – A batch of labels. Its shape is (N, R). The background is excluded from the definition, which means that the range of the value is $[0, L - 1]$. L is the number of foreground classes.
- **bboxes** (*Variable*) – A batch of bounding boxes. Its shape is $(N, R, 4)$.
- **scale** (*float or Variable*) – Amount of scaling applied to the raw image during preprocessing.

Returns Scalar loss variable. This is the sum of losses for Region Proposal Network and the head module.

Return type chainer.Variable

ProposalTargetCreator

class chainercv.experimental.links.model.fcis.**ProposalTargetCreator**(*n_sample=128*, *pos_ratio=0.25*, *pos_iou_thresh=0.5*, *neg_iou_thresh_hi=0.5*, *neg_iou_thresh_lo=0.1*, *binary_thresh=0.4*)

Assign ground truth classes, bounding boxes and masks to given RoIs.

The *__call__()* of this class generates training targets for each object proposal. This is used to train FCIS[3].

Parameters

- **n_sample** (*int*) – The number of sampled regions.
- **pos_ratio** (*float*) – Fraction of regions that is labeled as a foreground.
- **pos_iou_thresh** (*float*) – IoU threshold for a RoI to be considered as a foreground.
- **neg_iou_thresh_hi** (*float*) – RoI is considered to be the background if IoU is in [neg_iou_thresh_hi, neg_iou_thresh_hi).
- **neg_iou_thresh_lo** (*float*) – See above.
- **binary_thresh** (*float*) – Threshold for resized mask.

__call__(*roi*, *mask*, *label*, *bbox*, *loc_normalize_mean=(0.0, 0.0, 0.0, 0.0)*, *loc_normalize_std=(0.2, 0.2, 0.5, 0.5)*, *mask_size=(21, 21)*)

Assigns ground truth to sampled proposals.

This function samples total of `self.n_sample` RoIs from the combination of `roi`, `mask`, `label` and :obj: *bbox*. The RoIs are assigned with the ground truth class labels as well as bounding box offsets and scales to match the ground truth bounding boxes. As many as `pos_ratio * self.n_sample` RoIs are sampled as foregrounds.

Offsets and scales of bounding boxes are calculated using *chainercv.links.model.faster_rcnn.bbox2loc()*. Also, types of input arrays and output arrays are same.

Here are notations.

- S is the total number of sampled RoIs, which equals `self.n_sample`.
- L is number of object classes possibly including the background.
- H is the image height.

[3] Yi Li, Haozhi Qi, Jifeng Dai, Xiangyang Ji, Yichen Wei. Fully Convolutional Instance-aware Semantic Segmentation. CVPR 2017.

- W is the image width.
- RH is the mask height.
- RW is the mask width.

Parameters

- **roi** (*array*) – Region of Interests (RoIs) from which we sample. Its shape is $(R, 4)$
- **mask** (*array*) – The coordinates of ground truth masks. Its shape is (R', H, W).
- **label** (*array*) – Ground truth bounding box labels. Its shape is $(R',)$. Its range is $[0, L-1]$, where L is the number of foreground classes.
- **bbox** (*array*) – The coordinates of ground truth bounding boxes. Its shape is $(R', 4)$.
- **loc_normalize_mean** (*tuple of four floats*) – Mean values to normalize coordinates of bounding boxes.
- **loc_normalize_std** (*tuple of four floats*) – Standard deviation of the coordinates of bounding boxes.
- **mask_size** (*tuple of int or int*) – Generated mask size, which is equal to (RH, RW).

Returns

- **sample_roi**: Regions of interests that are sampled. Its shape is $(S, 4)$.
- **gt_roi_mask**: Masks assigned to sampled RoIs. Its shape is (S, RH, RW).
- **gt_roi_label**: Labels assigned to sampled RoIs. Its shape is $(S,)$. Its range is $[0, L]$. The label with value 0 is the background.
- **gt_roi_loc**: Offsets and scales to match the sampled RoIs to the ground truth bounding boxes. Its shape is $(S, 4)$.

Return type (array, array, array, array)

3.5 Extensions

3.5.1 Evaluator

DetectionCOCOEvaluator

class chainercv.extensions.**DetectionCOCOEvaluator**(*iterator*, *target*, *label_names=None*)

An extension that evaluates a detection model by MS COCO metric.

This extension iterates over an iterator and evaluates the prediction results. The results consist of average precisions (APs) and average recalls (ARs) as well as the mean of each (mean average precision and mean average recall). This extension reports the following values with keys. Please note that if `label_names` is not specified, only the mAPs and mARs are reported.

The underlying dataset of the iterator is assumed to return `img, bbox, label` or `img, bbox, label, area, crowded`.

key	description
ap/iou=0.50:0.95/area=all/max_dets=100/<label_names[l]>	[1]
ap/iou=0.50/area=all/max_dets=100/<label_names[l]>	[1]
ap/iou=0.75/area=all/max_dets=100/<label_names[l]>	[1]
ap/iou=0.50:0.95/area=small/max_dets=100/<label_names[l]>	[1][5]
ap/iou=0.50:0.95/area=medium/max_dets=100/<label_names[l]>	[1][5]
ap/iou=0.50:0.95/area=large/max_dets=100/<label_names[l]>	[1][5]
ar/iou=0.50:0.95/area=all/max_dets=1/<label_names[l]>	[2]
ar/iou=0.50/area=all/max_dets=10/<label_names[l]>	[2]
ar/iou=0.75/area=all/max_dets=100/<label_names[l]>	[2]
ar/iou=0.50:0.95/area=small/max_dets=100/<label_names[l]>	[2][5]
ar/iou=0.50:0.95/area=medium/max_dets=100/<label_names[l]>	[2][5]
ar/iou=0.50:0.95/area=large/max_dets=100/<label_names[l]>	[2][5]
map/iou=0.50:0.95/area=all/max_dets=100	[3]
map/iou=0.50/area=all/max_dets=100	[3]
map/iou=0.75/area=all/max_dets=100	[3]
map/iou=0.50:0.95/area=small/max_dets=100	[3][5]
map/iou=0.50:0.95/area=medium/max_dets=100	[3][5]
map/iou=0.50:0.95/area=large/max_dets=100	[3][5]
ar/iou=0.50:0.95/area=all/max_dets=1	[4]
ar/iou=0.50/area=all/max_dets=10	[4]
ar/iou=0.75/area=all/max_dets=100	[4]
ar/iou=0.50:0.95/area=small/max_dets=100	[4][5]
ar/iou=0.50:0.95/area=medium/max_dets=100	[4][5]
ar/iou=0.50:0.95/area=large/max_dets=100	[4][5]

Parameters

- **iterator** (*chainer.Iterator*) – An iterator. Each sample should be following tuple `img, bbox, label, area, crowded`.
- **target** (*chainer.Link*) – A detection link. This link must have `predict()` method that takes a list of images and returns `bboxes`, `labels` and `scores`.
- **label_names** (*iterable of strings*) – An iterable of names of classes. If this value is specified, average precision and average recalls for each class are reported.

DetectionVOCEvaluator

class chainercv.extensions.**DetectionVOCEvaluator**(*iterator*, *target*, *use_07_metric=False*, *label_names=None*)

An extension that evaluates a detection model by PASCAL VOC metric.

This extension iterates over an iterator and evaluates the prediction results by average precisions (APs) and mean of them (mean Average Precision, mAP). This extension reports the following values with keys. Please note that `'ap/<label_names[l]>'` is reported only if `label_names` is specified.

[1] Average precision for class `label_names[l]`, where l is the index of the class. If class l does not exist in either `pred_labels` or `gt_labels`, the corresponding value is set to `numpy.nan`.

[5] Skip if `gt_areas` is `None`.

[2] Average recall for class `label_names[l]`, where l is the index of the class. If class l does not exist in either `pred_labels` or `gt_labels`, the corresponding value is set to `numpy.nan`.

[3] The average of average precisions over classes.

[4] The average of average recalls over classes.

- `'map'`: Mean of average precisions (mAP).
- `'ap/<label_names[l]>'`: Average precision for class `label_names[l]`, where l is the index of the class. For example, this evaluator reports `'ap/aeroplane'`, `'ap/bicycle'`, etc. if `label_names` is `voc_bbox_label_names`. If there is no bounding box assigned to class `label_names[l]` in either ground truth or prediction, it reports `numpy.nan` as its average precision. In this case, mAP is computed without this class.

Parameters

- **iterator** (*chainer.Iterator*) – An iterator. Each sample should be following tuple `img, bbox, label` or `img, bbox, label, difficult`. `img` is an image, `bbox` is coordinates of bounding boxes, `label` is labels of the bounding boxes and `difficult` is whether the bounding boxes are difficult or not. If `difficult` is returned, difficult ground truth will be ignored from evaluation.
- **target** (*chainer.Link*) – A detection link. This link must have `predict()` method that takes a list of images and returns `bboxes`, `labels` and `scores`.
- **use_07_metric** (*bool*) – Whether to use PASCAL VOC 2007 evaluation metric for calculating average precision. The default value is `False`.
- **label_names** (*iterable of strings*) – An iterable of names of classes. If this value is specified, average precision for each class is also reported with the key `'ap/<label_names[l]>'`.

InstanceSegmentationCOCOEvaluator

class chainercv.extensions.**InstanceSegmentationCOCOEvaluator**(*iterator*, *target*, *label_names=None*)

An extension that evaluates a instance segmentation model by MS COCO metric.

This extension iterates over an iterator and evaluates the prediction results. The results consist of average precisions (APs) and average recalls (ARs) as well as the mean of each (mean average precision and mean average recall). This extension reports the following values with keys. Please note that if `label_names` is not specified, only the mAPs and mARs are reported.

The underlying dataset of the iterator is assumed to return `img, mask, label` or `img, mask, label, area, crowded`.

key	description
ap/iou=0.50:0.95/area=all/max_dets=100/<label_names[l]>	[6]
ap/iou=0.50/area=all/max_dets=100/<label_names[l]>	[6]
ap/iou=0.75/area=all/max_dets=100/<label_names[l]>	[6]
ap/iou=0.50:0.95/area=small/max_dets=100/<label_names[l]>	[6][10]
ap/iou=0.50:0.95/area=medium/max_dets=100/<label_names[l]>	[6][10]
ap/iou=0.50:0.95/area=large/max_dets=100/<label_names[l]>	[6][10]
ar/iou=0.50:0.95/area=all/max_dets=1/<label_names[l]>	[7]
ar/iou=0.50/area=all/max_dets=10/<label_names[l]>	[7]
ar/iou=0.75/area=all/max_dets=100/<label_names[l]>	[7]
ar/iou=0.50:0.95/area=small/max_dets=100/<label_names[l]>	[7][10]
ar/iou=0.50:0.95/area=medium/max_dets=100/<label_names[l]>	[7][10]
ar/iou=0.50:0.95/area=large/max_dets=100/<label_names[l]>	[7][10]
map/iou=0.50:0.95/area=all/max_dets=100	[8]
map/iou=0.50/area=all/max_dets=100	[8]
map/iou=0.75/area=all/max_dets=100	[8]
map/iou=0.50:0.95/area=small/max_dets=100	[8][10]
map/iou=0.50:0.95/area=medium/max_dets=100	[8][10]
map/iou=0.50:0.95/area=large/max_dets=100	[8][10]
ar/iou=0.50:0.95/area=all/max_dets=1	[9]
ar/iou=0.50/area=all/max_dets=10	[9]
ar/iou=0.75/area=all/max_dets=100	[9]
ar/iou=0.50:0.95/area=small/max_dets=100	[9][10]
ar/iou=0.50:0.95/area=medium/max_dets=100	[9][10]
ar/iou=0.50:0.95/area=large/max_dets=100	[9][10]

Parameters

- **iterator** (*chainer.Iterator*) – An iterator. Each sample should be following tuple `img, mask, label, area, crowded`.
- **target** (*chainer.Link*) – A detection link. This link must have `predict()` method that takes a list of images and returns `masks`, `labels` and `scores`.
- **label_names** (*iterable of strings*) – An iterable of names of classes. If this value is specified, average precision and average recalls for each class are reported.

InstanceSegmentationVOCEvaluator

class chainercv.extensions.**InstanceSegmentationVOCEvaluator**(*iterator*, *target*, *iou_thresh=0.5*, *use_07_metric=False*, *label_names=None*)

An evaluation extension of instance-segmentation by PASCAL VOC metric.

This extension iterates over an iterator and evaluates the prediction results by average precisions (APs) and mean of them (mean Average Precision, mAP). This extension reports the following values with keys. Please note that

[6] Average precision for class `label_names[l]`, where l is the index of the class. If class l does not exist in either `pred_labels` or `gt_labels`, the corresponding value is set to `numpy.nan`.

[10] Skip if `gt_areas` is `None`.

[7] Average recall for class `label_names[l]`, where l is the index of the class. If class l does not exist in either `pred_labels` or `gt_labels`, the corresponding value is set to `numpy.nan`.

[8] The average of average precisions over classes.

[9] The average of average recalls over classes.

'ap/<label_names[l]>' is reported only if label_names is specified.

- 'map': Mean of average precisions (mAP).
- 'ap/<label_names[l]>': Average precision for class label_names[l], where l is the index of the class. For example, this evaluator reports 'ap/aeroplane', 'ap/bicycle', etc. if label_names is sbd_instance_segmentation_label_names. If there is no bounding box assigned to class label_names[l] in either ground truth or prediction, it reports numpy.nan as its average precision. In this case, mAP is computed without this class.

Parameters

- **iterator** (*chainer.Iterator*) – An iterator. Each sample should be following tuple img, bbox, label or img, bbox, label, difficult. img is an image, bbox is coordinates of bounding boxes, label is labels of the bounding boxes and difficult is whether the bounding boxes are difficult or not. If difficult is returned, difficult ground truth will be ignored from evaluation.
- **target** (*chainer.Link*) – An instance-segmentation link. This link must have predict() method that takes a list of images and returns bboxes, labels and scores.
- **iou_thresh** (*float*) – Intersection over Union (IoU) threshold for calulating average precision. The default value is 0.5.
- **use_07_metric** (*bool*) – Whether to use PASCAL VOC 2007 evaluation metric for calculating average precision. The default value is False.
- **label_names** (*iterable of strings*) – An iterable of names of classes. If this value is specified, average precision for each class is also reported with the key 'ap/<label_names[l]>'.

SemanticSegmentationEvaluator

class chainercv.extensions.**SemanticSegmentationEvaluator**(*iterator*, *target*, *label_names=None*)

An extension that evaluates a semantic segmentation model.

This extension iterates over an iterator and evaluates the prediction results of the model by common evaluation metrics for semantic segmentation. This extension reports values with keys below. Please note that 'iou/<label_names[l]>' and 'class_accuracy/<label_names[l]>' are reported only if label_names is specified.

- 'miou': Mean of IoUs (mIoU).
- 'iou/<label_names[l]>': IoU for class label_names[l], where l is the index of the class. For example, if label_names is camvid_label_names, this evaluator reports 'iou/Sky', 'ap/Building', etc.
- 'mean_class_accuracy': Mean of class accuracies.
- 'class_accuracy/<label_names[l]>': Class accuracy for class label_names[l], where l is the index of the class.
- 'pixel_accuracy': Pixel accuracy.

If there is no label assigned to class label_names[l] in the ground truth, values corresponding to keys 'iou/<label_names[l]>' and 'class_accuracy/<label_names[l]>' are numpy.nan. In that case, the means of them are calculated by excluding them from calculation.

For details on the evaluation metrics, please see the documentation for *chainercv.evaluations.eval_semantic_segmentation()*.

See also:

chainercv.evaluations.eval_semantic_segmentation().

Parameters

- **iterator** (*chainer.Iterator*) – An iterator. Each sample should be following tuple `img, label`. `img` is an image, `label` is pixel-wise label.
- **target** (*chainer.Link*) – A semantic segmentation link. This link should have `predict()` method that takes a list of images and returns `labels`.
- **label_names** (*iterable of strings*) – An iterable of names of classes. If this value is specified, IoU and class accuracy for each class are also reported with the keys `'iou/<label_names[l]>'` and `'class_accuracy/<label_names[l]>'`.

3.5.2 Visualization Report

DetectionVisReport

class chainercv.extensions.**DetectionVisReport**(*iterator*, *target*, *label_names=None*, *file-name='detection_iter={iteration}_idx={index}.jpg'*)

An extension that visualizes output of a detection model.

This extension visualizes the predicted bounding boxes together with the ground truth bounding boxes.

Internally, this extension takes examples from an iterator, predict bounding boxes from the images in the examples, and visualizes them using *chainercv.visualizations.vis_bbox()*. The process can be illustrated in the following code.

```
batch = next(iterator)
# Convert batch -> imgs, gt_bboxes, gt_labels
pred_bboxes, pred_labels, pred_scores = target.predict(imgs)
# Visualization code
for img, gt_bbox, gt_label, pred_bbox, pred_label, pred_score \
        in zip(imgs, gt_boxes, gt_labels,
               pred_bboxes, pred_labels, pred_scores):
    # the ground truth
    vis_bbox(img, gt_bbox, gt_label)
    # the prediction
    vis_bbox(img, pred_bbox, pred_label, pred_score)
```

Note: `gt_bbox` and `pred_bbox` are float arrays of shape $(R, 4)$, where R is the number of bounding boxes in the image. Each bounding box is organized by $(y_{min}, x_{min}, y_{max}, x_{max})$ in the second axis.

`gt_label` and `pred_label` are intenger arrays of shape $(R,)$. Each label indicates the class of the bounding box.

`pred_score` is a float array of shape $(R,)$. Each score indicates how confident the prediction is.

Parameters

- **iterator** – Iterator object that produces images and ground truth.
- **target** – Link object used for detection.

- **label_names** (*iterable of strings*) – Name of labels ordered according to label ids. If this is `None`, labels will be skipped.
- **filename** (*str*) – Basename for the saved image. It can contain two keywords, `'{iteration}'` and `'{index}'`. They are replaced with the iteration of the trainer and the index of the sample when this extension save an image. The default value is `'detection_iter={iteration}_idx={index}.jpg'`.

3.6 Functions

3.6.1 Spatial Pooling

psroi_pooling_2d

chainercv.functions.**psroi_pooling_2d**(*x*, *rois*, *roi_indices*, *out_c*, *out_h*, *out_w*, *spatial_scale*, *group_size*)

Position Sensitive Region of Interest (ROI) pooling function.

This function computes position sensitive average of input spatial patch with the given region of interests. Each ROI is splitted into $(group_size, group_size)$ regions, and position sensitive values in each region is computed.

Parameters

- **x** (*Variable*) – Input variable. The shape is expected to be 4 dimentional: (n: batch, c: channel, h, height, w: width).
- **rois** (*array*) – Input roi. The shape is expected to be $(R, 4)$, and each datum is set as below: (y_min, x_min, y_max, x_max). The dtype is `numpy.float32`.
- **roi_indices** (*array*) – Input roi indices. The shape is expected to be $(R,)$. The dtype is `numpy.int32`.
- **out_c** (*int*) – Channels of output image after pooled.
- **out_h** (*int*) – Height of output image after pooled.
- **out_w** (*int*) – Width of output image after pooled.
- **spatial_scale** (*float*) – Scale of the roi is resized.
- **group_size** (*int*) – Position sensitive group size.

Returns Output variable.

Return type Variable

See the original paper proposing PSROIPooling: R-FCN.

3.7 Links

3.7.1 Model

General Chain

General Chain

FeaturePredictor

class chainercv.links.**FeaturePredictor**(*extractor*, *crop_size*, *scale_size=None*, *crop='center'*, *mean=None*)

Wrapper that adds a prediction method to a feature extraction link.

The *predict()* takes three steps to make a prediction.

1. Preprocess input images
2. Forward the preprocessed images to the network
3. Average features in the case when more than one crops are extracted.

Example

```
>>> from chainercv.links import VGG16
>>> from chainercv.links import FeaturePredictor
>>> base_model = VGG16()
>>> model = FeaturePredictor(base_model, 224, 256)
>>> prob = model.predict([img])
# Predicting multiple features
>>> model.extractor.pick = ['conv5_3', 'fc7']
>>> conv5_3, fc7 = model.predict([img])
```

When `self.crop == 'center'`, *predict()* extracts features from the center crop of the input images. When `self.crop == '10'`, *predict()* extracts features from patches that are ten-cropped from the input images.

When extracting more than one crops from an image, the output of *predict()* returns the average of the features computed from the crops.

Parameters

- **extractor** – A feature extraction link. This is a callable chain that takes a batch of images and returns a variable or a tuple of variables.
- **crop_size** (*int or tuple*) – The height and the width of an image after cropping in preprocessing. If this is an integer, the image is cropped to $(crop_size, crop_size)$.
- **scale_size** (*int or tuple*) – If `scale_size` is `None`, neither scaling nor resizing is conducted during preprocessing. This is the default behavior. If this is an integer, an image is resized so that the length of the shorter edge is equal to `scale_size`. If this is a tuple `(height, width)`, the image is resized to $(height, width)$.
- **crop** (*{'center', '10'}*) – Determines the style of cropping.
- **mean** (*numpy.ndarray*) – A mean value. If this is `None`, `extractor.mean` is used as the mean value.

predict(*imgs*)

Predict features from images.

Given N input images, this method outputs a batched array with batchsize N.

Parameters **imgs** (*iterable of numpy.ndarray*) – Array-images. All images are in CHW format and the range of their value is $[0, 255]$.

Returns A batch of features or a tuple of them.

Return type numpy.ndarray or tuple of numpy.ndarray

PickableSequentialChain

class chainercv.links.**PickableSequentialChain**

A sequential chain that can pick intermediate layers.

Callable objects, such as `chainer.Link` and `chainer.Function`, can be registered to this chain with `init_scope()`. This chain keeps the order of registrations and `__call__()` executes callables in that order. A `chainer.Link` object in the sequence will be added as a child link of this link.

`__call__()` returns single or multiple layers that are picked up through a stream of computation. These layers can be specified by `pick`, which contains the names of the layers that are collected. When `pick` is a string, single layer is returned. When `pick` is an iterable of strings, a tuple of layers is returned. The order of the layers is the same as the order of the strings in `pick`. When `pick` is `None`, the last layer is returned.

Examples

```
>>> import chainer.functions as F
>>> import chainer.links as L
>>> model = PickableSequentialChain()
>>> with model.init_scope():
>>>     model.l1 = L.Linear(None, 1000)
>>>     model.l1_relu = F.relu
>>>     model.l2 = L.Linear(None, 1000)
>>>     model.l2_relu = F.relu
>>>     model.l3 = L.Linear(None, 10)
>>> # This is layer l3
>>> layer3 = model(x)
>>> # The layers to be collected can be changed.
>>> model.pick = ('l2_relu', 'l1_relu')
>>> # These are layers l2_relu and l1_relu.
>>> layer2, layer1 = model(x)
```

Parameters

- **pick** (*string or iterable of strings*) – Names of layers that are collected during the forward pass.
- **layer_names** (*iterable of strings*) – Names of layers that can be collected from this chain. The names are ordered in the order of computation.

copy (**args*, ***kargs*)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the documentation for the argument `mode` below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameter variables under the returned link object is re-initialized by calling their `initialize()` method, so that all the parameters may have different initial values from the original link. `copy` means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. `share` means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default `mode` is `share`.

Returns Copied link object.

Return type Link

remove_unused()
Delete all layers that are not needed for the forward pass.

Feature Extraction

Feature extraction links extract feature(s) from given images.

ResNet

Feature Extraction Link

ResNet

class chainercv.links.model.resnet.**ResNet**(*n_layer*, *n_class=None*, *pretrained_model=None*, *mean=None*, *initialW=None*, *fc_kwargs={}*, *arch='fb'*)

Base class for ResNet architecture.

This is a pickable sequential link. The network can choose output layers from set of all intermediate layers. The attribute `pick` is the names of the layers that are going to be picked by `__call__()`. The attribute `layer_names` is the names of all layers that can be picked.

Examples

```
>>> model = ResNet50()
# By default, __call__ returns a probability score (after Softmax).
>>> prob = model(imgs)
>>> model.pick = 'res5'
# This is layer res5
>>> res5 = model(imgs)
>>> model.pick = ['res5', 'fc6']
>>> # These are layers res5 and fc6.
>>> res5, fc6 = model(imgs)
```

See also:

`chainercv.links.model.PickableSequentialChain`

When `pretrained_model` is the path of a pre-trained chainer model serialized as a `npz` file in the constructor, this chain model automatically initializes all the parameters with it. When a string in the prespecified set is provided, a pretrained model is loaded from weights distributed on the Internet. The list of pretrained models supported are as follows:

- `imagenet`: Loads weights trained with ImageNet. When `arch=='he'`, the weights distributed at Model Zoo are used.

Parameters

- **n_layer** (*int*) – The number of layers.

- **n_class** (*int*) – The number of classes. If None, the default values are used. If a supported pretrained model is used, the number of classes used to train the pretrained model is used. Otherwise, the number of classes in ILSVRC 2012 dataset is used.
- **pretrained_model** (*string*) – The destination of the pre-trained chainer model serialized as a npz file. If this is one of the strings described above, it automatically loads weights stored under a directory $CHAINER_DATASET_ROOT/pfnet/chainercv/models/, where $CHAINER_DATASET_ROOT is set as $HOME/.chainer/dataset unless you specify another value by modifying the environment variable.
- **mean** (*numpy.ndarray*) – A mean value. If None, the default values are used. If a supported pretrained model is used, the mean value used to train the pretrained model is used. Otherwise, the mean value calculated from ILSVRC 2012 dataset is used.
- **initialW** (*callable*) – Initializer for the weights of convolution kernels.
- **fc_kwargs** (*dict*) – Keyword arguments passed to initialize the chainer.links.Linear.
- **arch** (*string*) – If fb, use Facebook ResNet architecture. When he, use the architecture presented by the original ResNet paper. This option changes where to apply strided convolution. The default value is fb.

ResNet50

class chainercv.links.model.resnet.**ResNet50**(*n_class=None*, *pretrained_model=None*, *mean=None*, *initialW=None*, *fc_kwargs={}*, *arch='fb'*)

ResNet-50 Network.

Please consult the documentation for *ResNet*.

See also:

chainercv.links.model.resnet.ResNet

ResNet101

class chainercv.links.model.resnet.**ResNet101**(*n_class=None*, *pretrained_model=None*, *mean=None*, *initialW=None*, *fc_kwargs={}*, *arch='fb'*)

ResNet-101 Network.

Please consult the documentation for *ResNet*.

See also:

chainercv.links.model.resnet.ResNet

ResNet152

class chainercv.links.model.resnet.**ResNet152**(*n_class=None*, *pretrained_model=None*, *mean=None*, *initialW=None*, *fc_kwargs={}*, *arch='fb'*)

ResNet-152 Network.

Please consult the documentation for *ResNet*.

See also:

chainercv.links.model.resnet.ResNet

Utility

Bottleneck

class chainercv.links.model.resnet.**Bottleneck**(*in_channels*, *mid_channels*, *out_channels*, *stride=1*, *dilate=1*, *groups=1*, *initialW=None*, *bn_kwargs={}*, *residual_conv=False*, *stride_first=False*, *add_seblock=False*)

A bottleneck layer.

Parameters

- **in_channels** (*int*) – The number of channels of the input array.
- **mid_channels** (*int*) – The number of channels of intermediate arrays.
- **out_channels** (*int*) – The number of channels of the output array.
- **stride** (*int or tuple of ints*) – Stride of filter application.
- **dilate** (*int or tuple of ints*) – Dilation factor of filter applications. `dilate=d` and `dilate=(d, d)` are equivalent.
- **groups** (*int*) – The number of groups to use grouped convolution in the second layer. The default is one, where grouped convolution is not used.
- **initialW** (*callable*) – Initial weight value used in the convolutional layers.
- **bn_kwargs** (*dict*) – Keyword arguments passed to initialize `chainer.links.BatchNormalization`.
- **residual_conv** (*bool*) – If `True`, apply a 1x1 convolution to the residual.
- **stride_first** (*bool*) – If `True`, apply strided convolution with the first convolution layer. Otherwise, apply strided convolution with the second convolution layer.
- **add_seblock** (*bool*) – If `True`, apply a squeeze-and-excitation block to each residual block.

ResBlock

class chainercv.links.model.resnet.**ResBlock**(*n_layer*, *in_channels*, *mid_channels*, *out_channels*, *stride*, *dilate=1*, *groups=1*, *initialW=None*, *bn_kwargs={}*, *stride_first=False*, *add_seblock=False*)

A building block for ResNets.

in –> Bottleneck with residual_conv –> Bottleneck * (n_layer - 1) –> out

Parameters

- **n_layer** (*int*) – The number of layers used in the building block.
- **in_channels** (*int*) – The number of channels of the input array.
- **mid_channels** (*int*) – The number of channels of intermediate arrays.

- **out_channels** (*int*) – The number of channels of the output array.
- **stride** (*int or tuple of ints*) – Stride of filter application.
- **dilate** (*int or tuple of ints*) – Dilation factor of filter applications. `dilate=d` and `dilate=(d, d)` are equivalent.
- **groups** (*int*) – The number of groups to use grouped convolution in the second layer of each bottleneck. The default is one, where grouped convolution is not used.
- **initialW** (*callable*) – Initial weight value used in the convolutional layers.
- **bn_kwargs** (*dict*) – Keyword arguments passed to initialize `chainer.links.BatchNormalization`.
- **stride_first** (*bool*) – This determines the behavior of the bottleneck with a shortcut. If `True`, apply strided convolution with the first convolution layer. Otherwise, apply strided convolution with the second convolution layer.
- **add_seblock** (*bool*) – If `True`, apply a squeeze-and-excitation block to each residual block.

SEResNet

Feature Extraction Link

SEResNet

class chainercv.links.model.senet.**SEResNet**(*n_layer*, *n_class=None*, *pretrained_model=None*, *mean=None*, *initialW=None*, *fc_kwargs={}*)

Base class for SE-ResNet architecture.

This architecture is based on ResNet. A squeeze-and-excitation block is applied at the end of each non-identity branch of residual block. Please refer to the original paper for a detailed description of network architecture.

Similar to *chainercv.links.model.resnet.ResNet*, ImageNet pretrained weights are downloaded when `pretrained_model` argument is `imagenet`, originally distributed at the Github repository by one of the paper authors.

See also:

chainercv.links.model.resnet.ResNet *chainercv.links.connection.SEBlock*

Parameters

- **n_layer** (*int*) – The number of layers.
- **n_class** (*int*) – The number of classes. If `None`, the default values are used. If a supported pretrained model is used, the number of classes used to train the pretrained model is used. Otherwise, the number of classes in ILSVRC 2012 dataset is used.
- **pretrained_model** (*string*) – The destination of the pre-trained chainer model serialized as a `npz` file. If this is one of the strings described above, it automatically loads weights stored under a directory `$CHAINER_DATASET_ROOT/pfnet/chainercv/models/`, where `$CHAINER_DATASET_ROOT` is set as `$HOME/.chainer/dataset` unless you specify another value by modifying the environment variable.

- **mean** (*numpy.ndarray*) – A mean value. If `None`, the default values are used. If a supported pretrained model is used, the mean value used to train the pretrained model is used. Otherwise, the mean value calculated from ILSVRC 2012 dataset is used.
- **initialW** (*callable*) – Initializer for the weights of convolution kernels.
- **fc_kwargs** (*dict*) – Keyword arguments passed to initialize the `chainer.links.Linear`.

SEResNet50

class chainercv.links.model.senet.**SEResNet50**(*n_class=None*, *pretrained_model=None*, *mean=None*, *initialW=None*, *fc_kwargs={}*)

SE-ResNet-50 Network.

Please consult the documentation for *SEResNet*.

See also:

chainercv.links.model.senet.SEResNet

SEResNet101

class chainercv.links.model.senet.**SEResNet101**(*n_class=None*, *pretrained_model=None*, *mean=None*, *initialW=None*, *fc_kwargs={}*)

SE-ResNet-101 Network.

Please consult the documentation for *SEResNet*.

See also:

chainercv.links.model.senet.SEResNet

SEResNet152

class chainercv.links.model.senet.**SEResNet152**(*n_class=None*, *pretrained_model=None*, *mean=None*, *initialW=None*, *fc_kwargs={}*)

SE-ResNet-152 Network.

Please consult the documentation for *SEResNet*.

See also:

chainercv.links.model.senet.SEResNet

SEResNeXt

class chainercv.links.model.senet.**SEResNeXt**(*n_layer*, *n_class=None*, *pretrained_model=None*, *mean=None*, *initialW=None*, *fc_kwargs={}*)

Base class for SE-ResNeXt architecture.

ResNeXt is a ResNet-based architecture, where grouped convolution is adopted to the second convolution layer of each bottleneck block. In addition, a squeeze-and-excitation block is applied at the end of each non-identity

branch of residual block. Please refer to Aggregated Residual Transformations for Deep Neural Networks and Squeeze-and-Excitation Networks for detailed description of network architecture.

Similar to *chainercv.links.model.resnet.ResNet*, ImageNet pretrained weights are downloaded when `pretrained_model` argument is `imagenet`, originally distributed at the Github repository by one of the paper authors of SENet.

See also:

chainercv.links.model.resnet.ResNet *chainercv.links.model.senet.SEResNet* *chainercv.links.connection.SEBlock*

Parameters

- **n_layer** (*int*) – The number of layers.
- **n_class** (*int*) – The number of classes. If `None`, the default values are used. If a supported pretrained model is used, the number of classes used to train the pretrained model is used. Otherwise, the number of classes in ILSVRC 2012 dataset is used.
- **pretrained_model** (*string*) – The destination of the pre-trained chainer model serialized as a `npz` file. If this is one of the strings described above, it automatically loads weights stored under a directory `$CHAINER_DATASET_ROOT/pfnet/chainercv/models/`, where `$CHAINER_DATASET_ROOT` is set as `$HOME/.chainer/dataset` unless you specify another value by modifying the environment variable.
- **mean** (*numpy.ndarray*) – A mean value. If `None`, the default values are used. If a supported pretrained model is used, the mean value used to train the pretrained model is used. Otherwise, the mean value calculated from ILSVRC 2012 dataset is used.
- **initialW** (*callable*) – Initializer for the weights of convolution kernels.
- **fc_kwargs** (*dict*) – Keyword arguments passed to initialize the `chainer.links.Linear`.

SEResNeXt50

class chainercv.links.model.senet.**SEResNeXt50**(*n_class=None*, *pretrained_model=None*, *mean=None*, *initialW=None*, *fc_kwargs={}*)

SE-ResNeXt-50 Network

Please consult the documentation for *SEResNeXt*.

See also:

chainercv.links.model.senet.SEResNeXt

SEResNeXt101

class chainercv.links.model.senet.**SEResNeXt101**(*n_class=None*, *pretrained_model=None*, *mean=None*, *initialW=None*, *fc_kwargs={}*)

SE-ResNeXt-101 Network

Please consult the documentation for *SEResNeXt*.

See also:

chainercv.links.model.senet.SEResNeXt

VGG

VGG16

class chainercv.links.model.vgg.**VGG16**(*n_class=None*, *pretrained_model=None*, *mean=None*, *initialW=None*, *initial_bias=None*)

VGG-16 Network.

This is a pickable sequential link. The network can choose output layers from set of all intermediate layers. The attribute `pick` is the names of the layers that are going to be picked by `__call__()`. The attribute `layer_names` is the names of all layers that can be picked.

Examples

```
>>> model = VGG16()
# By default, __call__ returns a probability score (after Softmax).
>>> prob = model(imgs)
>>> model.pick = 'conv5_3'
# This is layer conv5_3 (after ReLU).
>>> conv5_3 = model(imgs)
>>> model.pick = ['conv5_3', 'fc6']
>>> # These are layers conv5_3 (after ReLU) and fc6 (before ReLU).
>>> conv5_3, fc6 = model(imgs)
```

See also:

`chainercv.links.model.PickableSequentialChain`

When `pretrained_model` is the path of a pre-trained chainer model serialized as a `npz` file in the constructor, this chain model automatically initializes all the parameters with it. When a string in the prespecified set is provided, a pretrained model is loaded from weights distributed on the Internet. The list of pretrained models supported are as follows:

- `imagenet`: Loads weights trained with ImageNet and distributed at Model Zoo.

Parameters

- **n_class** (*int*) – The number of classes. If `None`, the default values are used. If a supported pretrained model is used, the number of classes used to train the pretrained model is used. Otherwise, the number of classes in ILSVRC 2012 dataset is used.
- **pretrained_model** (*string*) – The destination of the pre-trained chainer model serialized as a `npz` file. If this is one of the strings described above, it automatically loads weights stored under a directory `$CHAINER_DATASET_ROOT/pfnet/chainercv/models/`, where `$CHAINER_DATASET_ROOT` is set as `$HOME/.chainer/dataset` unless you specify another value by modifying the environment variable.
- **mean** (*numpy.ndarray*) – A mean value. If `None`, the default values are used. If a supported pretrained model is used, the mean value used to train the pretrained model is used. Otherwise, the mean value calculated from ILSVRC 2012 dataset is used.
- **initialW** (*callable*) – Initializer for the weights.
- **initial_bias** (*callable*) – Initializer for the biases.

Detection

Detection links share a common method `predict()` to detect objects in images. For more details, please read `FasterRCNN.predict()`.

Faster R-CNN

Detection Link

FasterRCNNVGG16

class chainercv.links.model.faster_rcnn.**FasterRCNNVGG16**(*n_fg_class=None, pretrained_model=None, min_size=600, max_size=1000, ratios=[0.5, 1, 2], anchor_scales=[8, 16, 32], vgg_initialW=None, rpn_initialW=None, loc_initialW=None, score_initialW=None, proposal_creator_params={}*)

Faster R-CNN based on VGG-16.

When you specify the path of a pre-trained chainer model serialized as a `npz` file in the constructor, this chain model automatically initializes all the parameters with it. When a string in prespecified set is provided, a pretrained model is loaded from weights distributed on the Internet. The list of pretrained models supported are as follows:

- `voc07`: Loads weights trained with the trainval split of PASCAL VOC2007 Detection Dataset.
- `imagenet`: Loads weights trained with ImageNet Classfication task for the feature extractor and the head modules. Weights that do not have a corresponding layer in VGG-16 will be randomly initialized.

For descriptions on the interface of this model, please refer to *FasterRCNN*.

FasterRCNNVGG16 supports finer control on random initializations of weights by arguments `vgg_initialW`, `rpn_initialW`, `loc_initialW` and `score_initialW`. It accepts a callable that takes an array and edits its values. If `None` is passed as an initializer, the default initializer is used.

Parameters

- **n_fg_class** (*int*) – The number of classes excluding the background.
- **pretrained_model** (*string*) – The destination of the pre-trained chainer model serialized as a `npz` file. If this is one of the strings described above, it automatically loads weights stored under a directory `$CHAINER_DATASET_ROOT/pfnet/chainercv/models/`, where `$CHAINER_DATASET_ROOT` is set as `$HOME/.chainer/dataset` unless you specify another value by modifying the environment variable.
- **min_size** (*int*) – A preprocessing paramter for `prepare()`.
- **max_size** (*int*) – A preprocessing paramter for `prepare()`.
- **ratios** (*list of floats*) – This is ratios of width to height of the anchors.

- **anchor_scales** (*list of numbers*) – This is areas of anchors. Those areas will be the product of the square of an element in `anchor_scales` and the original area of the reference window.
- **vgg_initialW** (*callable*) – Initializer for the layers corresponding to the VGG-16 layers.
- **rpn_initialW** (*callable*) – Initializer for Region Proposal Network layers.
- **loc_initialW** (*callable*) – Initializer for the localization head.
- **score_initialW** (*callable*) – Initializer for the score head.
- **proposal_creator_params** (*dict*) – Key valued paramters for *ProposalCreator*.

Utility

bbox2loc

chainercv.links.model.faster_rcnn.**bbox2loc**(*src_bbox*, *dst_bbox*)

Encodes the source and the destination bounding boxes to "loc".

Given bounding boxes, this function computes offsets and scales to match the source bounding boxes to the target bounding boxes. Mathematcially, given a bounding box whose center is $(y, x) = p_y, p_x$ and size p_h, p_w and the target bounding box whose center is g_y, g_x and size g_h, g_w, the offsets and scales t_y, t_x, t_h, t_w can be computed by the following formulas.

- $t_y = \frac{(g_y - p_y)}{p_h}$
- $t_x = \frac{(g_x - p_x)}{p_w}$
- $t_h = \log(\frac{g_h}{p_h})$
- $t_w = \log(\frac{g_w}{p_w})$

The output is same type as the type of the inputs. The encoding formulas are used in works such as R-CNN[1].

Parameters

- **src_bbox** (*array*) – An image coordinate array whose shape is $(R, 4)$. R is the number of bounding boxes. These coordinates are $p_{ymin}, p_{xmin}, p_{ymax}, p_{xmax}$.
- **dst_bbox** (*array*) – An image coordinate array whose shape is $(R, 4)$. These coordinates are $g_{ymin}, g_{xmin}, g_{ymax}, g_{xmax}$.

Returns Bounding box offsets and scales from `src_bbox` to `dst_bbox`. This has shape $(R, 4)$. The second axis contains four values t_y, t_x, t_h, t_w.

Return type array

[1] Ross Girshick, Jeff Donahue, Trevor Darrell, Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. CVPR 2014.

FasterRCNN

class chainercv.links.model.faster_rcnn.**FasterRCNN**(*extractor*, *rpn*, *head*, *mean*, *min_size=600*, *max_size=1000*, *loc_normalize_mean=(0.0, 0.0, 0.0, 0.0)*, *loc_normalize_std=(0.1, 0.1, 0.2, 0.2)*)

Base class for Faster R-CNN.

This is a base class for Faster R-CNN links supporting object detection API[2]. The following three stages constitute Faster R-CNN.

1. **Feature extraction**: Images are taken and their feature maps are calculated.
2. **Region Proposal Networks**: Given the feature maps calculated in the previous stage, produce set of RoIs around objects.
3. **Localization and Classification Heads**: Using feature maps that belong to the proposed RoIs, classify the categories of the objects in the RoIs and improve localizations.

Each stage is carried out by one of the callable `chainer.Chain` objects `feature`, `rpn` and `head`.

There are two functions *predict()* and *__call__()* to conduct object detection. *predict()* takes images and returns bounding boxes that are converted to image coordinates. This will be useful for a scenario when Faster R-CNN is treated as a black box function, for instance. *__call__()* is provided for a scnerario when intermediate outputs are needed, for instance, for training and debugging.

Links that support obejct detection API have method *predict()* with the same interface. Please refer to *predict()* for further details.

Parameters

- **extractor** (*callable Chain*) – A callable that takes a BCHW image array and returns feature maps.
- **rpn** (*callable Chain*) – A callable that has the same interface as *RegionProposalNetwork*. Please refer to the documentation found there.
- **head** (*callable Chain*) – A callable that takes a BCHW array, RoIs and batch indices for RoIs. This returns class dependent localization paramters and class scores.
- **mean** (*numpy.ndarray*) – A value to be subtracted from an image in *prepare()*.
- **min_size** (*int*) – A preprocessing paramter for *prepare()*. Please refer to a docstring found for *prepare()*.
- **max_size** (*int*) – A preprocessing paramter for *prepare()*.
- **loc_normalize_mean** (*tuple of four floats*) – Mean values of localization estimates.
- **loc_normalize_std** (*tupler of four floats*) – Standard deviation of localization estimates.

__call__(*x*, *scales=None*)

Forward Faster R-CNN.

Scaling paramter `scales` is used by RPN to determine the threshold to select small objects, which are going to be rejected irrespective of their confidence scores.

Here are notations used.

[2] Shaoqing Ren, Kaiming He, Ross Girshick, Jian Sun. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. NIPS 2015.

- N is the number of batch size
- R' is the total number of RoIs produced across batches. Given R_i proposed RoIs from the i th image, $R' = \sum_{i=1}^{N} R_i$.
- L is the number of classes excluding the background.

Classes are ordered by the background, the first class, ..., and the L th class.

Parameters

- **x** (*Variable*) – 4D image variable.
- **scales** (*tuple of floats*) – Amount of scaling applied to each input image during preprocessing.

Returns

Returns tuple of four values listed below.

- **roi_cls_locs**: Offsets and scalings for the proposed RoIs. Its shape is $(R', (L+1) \times 4)$.
- **roi_scores**: Class predictions for the proposed RoIs. Its shape is $(R', L+1)$.
- **rois**: RoIs proposed by RPN. Its shape is $(R', 4)$.
- **roi_indices**: Batch indices of RoIs. Its shape is $(R',)$.

Return type Variable, Variable, array, array

predict(*imgs*)

Detect objects from images.

This method predicts objects for each image.

Parameters **imgs** (*iterable of numpy.ndarray*) – Arrays holding images. All images are in CHW and RGB format and the range of their value is $[0, 255]$.

Returns

This method returns a tuple of three lists, `(bboxes, labels, scores)`.

- **bboxes**: A list of float arrays of shape $(R, 4)$, where R is the number of bounding boxes in a image. Each bounding box is organized by $(y_{min}, x_{min}, y_{max}, x_{max})$ in the second axis.
- **labels** : A list of integer arrays of shape $(R,)$. Each value indicates the class of the bounding box. Values are in range $[0, L-1]$, where L is the number of the foreground classes.
- **scores** : A list of float arrays of shape $(R,)$. Each value indicates how confident the prediction is.

Return type tuple of lists

prepare(*img*)

Preprocess an image for feature extraction.

The length of the shorter edge is scaled to `self.min_size`. After the scaling, if the length of the longer edge is longer than `self.max_size`, the image is scaled to fit the longer edge to `self.max_size`.

After resizing the image, the image is subtracted by a mean image value `self.mean`.

Parameters **img** (*ndarray*) – An image. This is in CHW and RGB format. The range of its value is $[0, 255]$.

Returns A preprocessed image.

Return type ndarray

use_preset(*preset*)

Use the given preset during prediction.

This method changes values of `self.nms_thresh` and `self.score_thresh`. These values are a threshold value used for non maximum suppression and a threshold value to discard low confidence proposals in *predict()*, respectively.

If the attributes need to be changed to something other than the values provided in the presets, please modify them by directly accessing the public attributes.

Parameters **preset** (*{'visualize', 'evaluate'}*) – A string to determine the preset to use.

generate_anchor_base

chainercv.links.model.faster_rcnn.**generate_anchor_base**(*base_size=16, ratios=[0.5, 1, 2], anchor_scales=[8, 16, 32]*)

Generate anchor base windows by enumerating aspect ratio and scales.

Generate anchors that are scaled and modified to the given aspect ratios. Area of a scaled anchor is preserved when modifying to the given aspect ratio.

`R = len(ratios) * len(anchor_scales)` anchors are generated by this function. The `i * len(anchor_scales) + j` th anchor corresponds to an anchor generated by `ratios[i]` and `anchor_scales[j]`.

For example, if the scale is 8 and the ratio is 0.25, the width and the height of the base window will be stretched by 8. For modifying the anchor to the given aspect ratio, the height is halved and the width is doubled.

Parameters

- **base_size** (*number*) – The width and the height of the reference window.
- **ratios** (*list of floats*) – This is ratios of width to height of the anchors.
- **anchor_scales** (*list of numbers*) – This is areas of anchors. Those areas will be the product of the square of an element in `anchor_scales` and the original area of the reference window.

Returns An array of shape $(R, 4)$. Each element is a set of coordinates of a bounding box. The second axis corresponds to $(y_{min}, x_{min}, y_{max}, x_{max})$ of a bounding box.

Return type ndarray

loc2bbox

chainercv.links.model.faster_rcnn.**loc2bbox**(*src_bbox*, *loc*)

Decode bounding boxes from bounding box offsets and scales.

Given bounding box offsets and scales computed by *bbox2loc()*, this function decodes the representation to coordinates in 2D image coordinates.

Given scales and offsets t_y, t_x, t_h, t_w and a bounding box whose center is $(y, x) = p_y, p_x$ and size p_h, p_w, the decoded bounding box's center $\hat{g}_y$, $\hat{g}_x$ and size $\hat{g}_h$, $\hat{g}_w$ are calculated by the following formulas.

- $\hat{g}_y = p_h t_y + p_y$
- $\hat{g}_x = p_w t_x + p_x$

- $\hat{g}_h = p_h \exp(t_h)$
- $\hat{g}_w = p_w \exp(t_w)$

The decoding formulas are used in works such as R-CNN[3].

The output is same type as the type of the inputs.

Parameters

- **src_bbox** (*array*) – A coordinates of bounding boxes. Its shape is $(R, 4)$. These coordinates are $p_{ymin}, p_{xmin}, p_{ymax}, p_{xmax}$.
- **loc** (*array*) – An array with offsets and scales. The shapes of `src_bbox` and `loc` should be same. This contains values t_y, t_x, t_h, t_w.

Returns Decoded bounding box coordinates. Its shape is $(R, 4)$. The second axis contains four values $\hat{g}_{ymin}, \hat{g}_{xmin}, \hat{g}_{ymax}, \hat{g}_{xmax}$.

Return type array

ProposalCreator

class chainercv.links.model.faster_rcnn.**ProposalCreator**(*nms_thresh=0.7*, *n_train_pre_nms=12000*, *n_train_post_nms=2000*, *n_test_pre_nms=6000*, *n_test_post_nms=300*, *force_cpu_nms=False*, *min_size=16*)

Proposal regions are generated by calling this object.

The *__call__()* of this object outputs object detection proposals by applying estimated bounding box offsets to a set of anchors.

This class takes parameters to control number of bounding boxes to pass to NMS and keep after NMS. If the paramters are negative, it uses all the bounding boxes supplied or keep all the bounding boxes returned by NMS.

This class is used for Region Proposal Networks introduced in Faster R-CNN[4].

Parameters

- **nms_thresh** (*float*) – Threshold value used when calling NMS.
- **n_train_pre_nms** (*int*) – Number of top scored bounding boxes to keep before passing to NMS in train mode.
- **n_train_post_nms** (*int*) – Number of top scored bounding boxes to keep after passing to NMS in train mode.
- **n_test_pre_nms** (*int*) – Number of top scored bounding boxes to keep before passing to NMS in test mode.
- **n_test_post_nms** (*int*) – Number of top scored bounding boxes to keep after passing to NMS in test mode.
- **force_cpu_nms** (*bool*) – If this is `True`, always use NMS in CPU mode. If `False`, the NMS mode is selected based on the type of inputs.

[3] Ross Girshick, Jeff Donahue, Trevor Darrell, Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. CVPR 2014.

[4] Shaoqing Ren, Kaiming He, Ross Girshick, Jian Sun. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. NIPS 2015.

- **min_size** (*int*) – A paramter to determine the threshold on discarding bounding boxes based on their sizes.

__call__(*loc*, *score*, *anchor*, *img_size*, *scale=1.0*)

Propose RoIs.

Inputs `loc, score, anchor` refer to the same anchor when indexed by the same index.

On notations, R is the total number of anchors. This is equal to product of the height and the width of an image and the number of anchor bases per pixel.

Type of the output is same as the inputs.

Parameters

- **loc** (*array*) – Predicted offsets and scaling to anchors. Its shape is $(R, 4)$.
- **score** (*array*) – Predicted foreground probability for anchors. Its shape is $(R,)$.
- **anchor** (*array*) – Coordinates of anchors. Its shape is $(R, 4)$.
- **img_size** (*tuple of ints*) – A tuple `height, width`, which contains image size after scaling.
- **scale** (*float*) – The scaling factor used to scale an image after reading it from a file.

Returns An array of coordinates of proposal boxes. Its shape is $(S, 4)$. S is less than `self.n_test_post_nms` in test time and less than `self.n_train_post_nms` in train time. S depends on the size of the predicted bounding boxes and the number of bounding boxes discarded by NMS.

Return type array

RegionProposalNetwork

class chainercv.links.model.faster_rcnn.**RegionProposalNetwork**(*in_channels=512, mid_channels=512, ratios=[0.5, 1, 2], anchor_scales=[8, 16, 32], feat_stride=16, initialW=None, proposal_creator_params={}*)

Region Proposal Network introduced in Faster R-CNN.

This is Region Proposal Network introduced in Faster R-CNN[5]. This takes features extracted from images and propose class agnostic bounding boxes around "objects".

Parameters

- **in_channels** (*int*) – The channel size of input.
- **mid_channels** (*int*) – The channel size of the intermediate tensor.
- **ratios** (*list of floats*) – This is ratios of width to height of the anchors.

[5] Shaoqing Ren, Kaiming He, Ross Girshick, Jian Sun. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. NIPS 2015.

- **anchor_scales** (*list of numbers*) – This is areas of anchors. Those areas will be the product of the square of an element in `anchor_scales` and the original area of the reference window.
- **feat_stride** (*int*) – Stride size after extracting features from an image.
- **initialW** (*callable*) – Initial weight value. If `None` then this function uses Gaussian distribution scaled by 0.1 to initialize weight. May also be a callable that takes an array and edits its values.
- **proposal_creator_params** (*dict*) – Key valued paramters for *ProposalCreator*.

See also:

ProposalCreator

__call__(*x*, *img_size*, *scales=None*)

Forward Region Proposal Network.

Here are notations.

- N is batch size.
- C channel size of the input.
- H and W are height and witdh of the input feature.
- A is number of anchors assigned to each pixel.

Parameters

- **x** (*Variable*) – The Features extracted from images. Its shape is (N, C, H, W).
- **img_size** (*tuple of ints*) – A tuple `height, width`, which contains image size after scaling.
- **scales** (*tuple of floats*) – The amount of scaling done to each input image during preprocessing.

Returns

This is a tuple of five following values.

- **rpn_locs**: Predicted bounding box offsets and scales for anchors. Its shape is $(N, HWA, 4)$.
- **rpn_scores**: Predicted foreground scores for anchors. Its shape is $(N, HWA, 2)$.
- **rois**: A bounding box array containing coordinates of proposal boxes. This is a concatenation of bounding box arrays from multiple images in the batch. Its shape is $(R', 4)$. Given R_i predicted bounding boxes from the i th image, $R' = \sum_{i=1}^{N} R_i$.
- **roi_indices**: An array containing indices of images to which RoIs correspond to. Its shape is $(R',)$.
- **anchor**: Coordinates of enumerated shifted anchors. Its shape is $(HWA, 4)$.

Return type (Variable, Variable, array, array, array)

VGG16RoIHead

class chainercv.links.model.faster_rcnn.**VGG16RoIHead**(*n_class*, *roi_size*, *spatial_scale*, *vgg_initialW=None*, *loc_initialW=None*, *score_initialW=None*)

Faster R-CNN Head for VGG-16 based implementation.

This class is used as a head for Faster R-CNN. This outputs class-wise localizations and classification based on feature maps in the given RoIs.

Parameters

- **n_class** (*int*) – The number of classes possibly including the background.
- **roi_size** (*int*) – Height and width of the feature maps after RoI-pooling.
- **spatial_scale** (*float*) – Scale of the roi is resized.
- **vgg_initialW** (*callable*) – Initializer for the layers corresponding to the VGG-16 layers.
- **loc_initialW** (*callable*) – Initializer for the localization head.
- **score_initialW** (*callable*) – Initializer for the score head.

Train-only Utility

AnchorTargetCreator

class chainercv.links.model.faster_rcnn.**AnchorTargetCreator**(*n_sample=256*, *pos_iou_thresh=0.7*, *neg_iou_thresh=0.3*, *pos_ratio=0.5*)

Assign the ground truth bounding boxes to anchors.

Assigns the ground truth bounding boxes to anchors for training Region Proposal Networks introduced in Faster R-CNN[6].

Offsets and scales to match anchors to the ground truth are calculated using the encoding scheme of *bbox2loc()*.

Parameters

- **n_sample** (*int*) – The number of regions to produce.
- **pos_iou_thresh** (*float*) – Anchors with IoU above this threshold will be assigned as positive.
- **neg_iou_thresh** (*float*) – Anchors with IoU below this threshold will be assigned as negative.
- **pos_ratio** (*float*) – Ratio of positive regions in the sampled regions.

__call__(*bbox*, *anchor*, *img_size*)

Assign ground truth supervision to sampled subset of anchors.

Types of input arrays and output arrays are same.

[6] Shaoqing Ren, Kaiming He, Ross Girshick, Jian Sun. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. NIPS 2015.

Here are notations.

- S is the number of anchors.
- R is the number of bounding boxes.

Parameters

- **bbox** (*array*) – Coordinates of bounding boxes. Its shape is $(R, 4)$.
- **anchor** (*array*) – Coordinates of anchors. Its shape is $(S, 4)$.
- **img_size** (*tuple of ints*) – A tuple `H, W`, which is a tuple of height and width of an image.

Returns

- **loc**: Offsets and scales to match the anchors to the ground truth bounding boxes. Its shape is $(S, 4)$.
- **label**: Labels of anchors with values `(1=positive, 0=negative, -1=ignore)`. Its shape is $(S,)$.

Return type (array, array)

FasterRCNNTrainChain

class chainercv.links.model.faster_rcnn.**FasterRCNNTrainChain**(*faster_rcnn*, *rpn_sigma=3.0*, *roi_sigma=1.0*, *anchor_target_creator=<chainercv.links.mode object>*, *proposal_target_creator=<chainercv.links.mod object>*)

Calculate losses for Faster R-CNN and report them.

This is used to train Faster R-CNN in the joint training scheme[7].

The losses include:

- `rpn_loc_loss`: The localization loss for Region Proposal Network (RPN).
- `rpn_cls_loss`: The classification loss for RPN.
- `roi_loc_loss`: The localization loss for the head module.
- `roi_cls_loss`: The classification loss for the head module.

Parameters

- **faster_rcnn** (*FasterRCNN*) – A Faster R-CNN model that is going to be trained.
- **rpn_sigma** (*float*) – Sigma parameter for the localization loss of Region Proposal Network (RPN). The default value is 3, which is the value used in[7].
- **roi_sigma** (*float*) – Sigma paramter for the localization loss of the head. The default value is 1, which is the value used in[7].
- **anchor_target_creator** – An instantiation of *AnchorTargetCreator*.
- **proposal_target_creator** – An instantiation of *ProposalTargetCreator*.

[7] Shaoqing Ren, Kaiming He, Ross Girshick, Jian Sun. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. NIPS 2015.

__call__(*imgs*, *bboxes*, *labels*, *scales*)
Forward Faster R-CNN and calculate losses.

Here are notations used.

- N is the batch size.
- R is the number of bounding boxes per image.

Currently, only $N = 1$ is supported.

Parameters

- **imgs** (*Variable*) – A variable with a batch of images.
- **bboxes** (*Variable*) – A batch of bounding boxes. Its shape is $(N, R, 4)$.
- **labels** (*Variable*) – A batch of labels. Its shape is (N, R). The background is excluded from the definition, which means that the range of the value is $[0, L - 1]$. L is the number of foreground classes.
- **scales** (*Variable*) – Amount of scaling applied to each input image during preprocessing.

Returns Scalar loss variable. This is the sum of losses for Region Proposal Network and the head module.

Return type chainer.Variable

ProposalTargetCreator

class chainercv.links.model.faster_rcnn.**ProposalTargetCreator**(*n_sample=128*, *pos_ratio=0.25*, *pos_iou_thresh=0.5*, *neg_iou_thresh_hi=0.5*, *neg_iou_thresh_lo=0.0*)
Assign ground truth bounding boxes to given RoIs.

The *__call__()* of this class generates training targets for each object proposal. This is used to train Faster RCNN[8].

Parameters

- **n_sample** (*int*) – The number of sampled regions.
- **pos_ratio** (*float*) – Fraction of regions that is labeled as a foreground.
- **pos_iou_thresh** (*float*) – IoU threshold for a RoI to be considered as a foreground.
- **neg_iou_thresh_hi** (*float*) – RoI is considered to be the background if IoU is in [neg_iou_thresh_lo, neg_iou_thresh_hi).
- **neg_iou_thresh_lo** (*float*) – See above.

__call__(*roi*, *bbox*, *label*, *loc_normalize_mean=(0.0, 0.0, 0.0, 0.0)*, *loc_normalize_std=(0.1, 0.1, 0.2, 0.2)*)
Assigns ground truth to sampled proposals.

This function samples total of self.n_sample RoIs from the combination of roi and bbox. The RoIs are assigned with the ground truth class labels as well as bounding box offsets and scales to match

[8] Shaoqing Ren, Kaiming He, Ross Girshick, Jian Sun. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. NIPS 2015.

the ground truth bounding boxes. As many as `pos_ratio * self.n_sample` RoIs are sampled as foregrounds.

Offsets and scales of bounding boxes are calculated using *chainercv.links.model.faster_rcnn.bbox2loc()*. Also, types of input arrays and output arrays are same.

Here are notations.

- S is the total number of sampled RoIs, which equals `self.n_sample`.
- L is number of object classes possibly including the background.

Parameters

- **roi** (*array*) – Region of Interests (RoIs) from which we sample. Its shape is $(R, 4)$
- **bbox** (*array*) – The coordinates of ground truth bounding boxes. Its shape is $(R', 4)$.
- **label** (*array*) – Ground truth bounding box labels. Its shape is $(R',)$. Its range is $[0, L-1]$, where L is the number of foreground classes.
- **loc_normalize_mean** (*tuple of four floats*) – Mean values to normalize coordinates of bounding boxes.
- **loc_normalize_std** (*tuple of four floats*) – Standard deviation of the coordinates of bounding boxes.

Returns

- **sample_roi**: Regions of interests that are sampled. Its shape is $(S, 4)$.
- **gt_roi_loc**: Offsets and scales to match the sampled RoIs to the ground truth bounding boxes. Its shape is $(S, 4)$.
- **gt_roi_label**: Labels assigned to sampled RoIs. Its shape is $(S,)$. Its range is $[0, L]$. The label with value 0 is the background.

Return type (array, array, array)

FPN (Feature Pyramid Networks)

Detection Links

FasterRCNNFPNResnet50

class chainercv.links.model.fpn.**FasterRCNNFPNResNet50**(*n_fg_class=None*, *pretrained_model=None*)

Feature Pyramid Networks with ResNet-50.

This is a model of Feature Pyramid Networks[1]. This model uses `ResNet50` as its base feature extractor.

Parameters

- **n_fg_class** (*int*) – The number of classes excluding the background.
- **pretrained_model** (*string*) – The weight file to be loaded. This can take `'coco'`, *filepath* or `None`. The default value is `None`.
  - `'coco'`: Load weights trained on train split of MS COCO 2017. The weight file is downloaded and cached automatically. `n_fg_class` must be `80` or `None`.

[1] Tsung-Yi Lin et al. Feature Pyramid Networks for Object Detection. CVPR 2017

- `'imagenet'`: Load weights of ResNet-50 trained on ImageNet. The weight file is downloaded and cached automatically. This option initializes weights partially and the rests are initialized randomly. In this case, `n_fg_class` can be set to any number.
- *filepath*: A path of npz file. In this case, `n_fg_class` must be specified properly.
- `None`: Do not load weights.

FasterRCNNFPNResnet101

class chainercv.links.model.fpn.**FasterRCNNFPNResNet101**(*n_fg_class=None*, *pretrained_model=None*)

Feature Pyramid Networks with ResNet-101.

This is a model of Feature Pyramid Networks[2]. This model uses `ResNet101` as its base feature extractor.

Parameters

- **n_fg_class** (*int*) – The number of classes excluding the background.
- **pretrained_model** (*string*) – The weight file to be loaded. This can take `'coco'`, *filepath* or `None`. The default value is `None`.
  - `'coco'`: Load weights trained on train split of MS COCO 2017. The weight file is downloaded and cached automatically. `n_fg_class` must be `80` or `None`.
  - `'imagenet'`: Load weights of ResNet-101 trained on ImageNet. The weight file is downloaded and cached automatically. This option initializes weights partially and the rests are initialized randomly. In this case, `n_fg_class` can be set to any number.
  - *filepath*: A path of npz file. In this case, `n_fg_class` must be specified properly.
  - `None`: Do not load weights.

Utility

FasterRCNN

class chainercv.links.model.fpn.**FasterRCNN**(*extractor*, *rpn*, *head*)

Base class of Feature Pyramid Networks.

This is a base class of Feature Pyramid Networks[3].

Parameters

- **extractor** (*Link*) – A link that extracts feature maps. This link must have `scales`, `mean` and `__call__()`.
- **rpn** (*Link*) – A link that has the same interface as *RPN*. Please refer to the documentation found there.
- **head** (*Link*) – A link that has the same interface as *Head*. Please refer to the documentation found there.
- **nms_thresh** (*float*) – The threshold value for *non_maximum_suppression()*. The default value is `0.45`. This value can be changed directly or by using *use_preset()*.

[2] Tsung-Yi Lin et al. Feature Pyramid Networks for Object Detection. CVPR 2017
[3] Tsung-Yi Lin et al. Feature Pyramid Networks for Object Detection. CVPR 2017

- **score_thresh** (*float*) – The threshold value for confidence score. If a bounding box whose confidence score is lower than this value, the bounding box will be suppressed. The default value is `0.6`. This value can be changed directly or by using *use_preset()*.

predict(*imgs*)
Detect objects from images.

This method predicts objects for each image.

Parameters **imgs** (*iterable of numpy.ndarray*) – Arrays holding images. All images are in CHW and RGB format and the range of their value is $[0, 255]$.

Returns

This method returns a tuple of three lists, `(bboxes, labels, scores)`.

- **bboxes**: A list of float arrays of shape $(R, 4)$, where R is the number of bounding boxes in a image. Each bounding box is organized by $(y_{min}, x_{min}, y_{max}, x_{max})$ in the second axis.
- **labels** : A list of integer arrays of shape $(R,)$. Each value indicates the class of the bounding box. Values are in range $[0, L-1]$, where L is the number of the foreground classes.
- **scores** : A list of float arrays of shape $(R,)$. Each value indicates how confident the prediction is.

Return type tuple of lists

prepare(*imgs*)
Preprocess images.

Parameters **imgs** (*iterable of numpy.ndarray*) – Arrays holding images. All images are in CHW and RGB format and the range of their value is $[0, 255]$.

Returns preprocessed images and scales that were caluclated in prepocessing.

Return type Two arrays

use_preset(*preset*)
Use the given preset during prediction.

This method changes values of `nms_thresh` and `score_thresh`. These values are a threshold value used for non maximum suppression and a threshold value to discard low confidence proposals in *predict()*, respectively.

If the attributes need to be changed to something other than the values provided in the presets, please modify them by directly accessing the public attributes.

Parameters **preset** (*{'visualize', 'evaluate'}*) – A string to determine the preset to use.

FPN

class chainercv.links.model.fpn.**FPN**(*base*, *n_base_output*, *scales*)
An extractor class of Feature Pyramid Networks.

This class wraps a feature extractor and provides multi-scale features.

Parameters

- **base** (*Link*) – A base feature extractor. It should have `__call__()` and `mean`. `__call__()` should take a batch of images and return feature maps of them. The size of the $k+1$-th feature map should be the half as that of the k-th feature map.
- **n_base_output** (*int*) – The number of feature maps that `base` returns.
- **scales** (*tuple of floats*) – The scales of feature maps.

Head

class chainercv.links.model.fpn.**Head**(*n_class*, *scales*)
Head network of Feature Pyramid Networks.

Parameters

- **n_class** (*int*) – The number of classes including background.
- **scales** (*tuple of floats*) – The scales of feature maps.

__call__(*hs*, *rois*, *roi_indices*)
Calculates RoIs.

Parameters

- **hs** (*iterable of array*) – An iterable of feature maps.
- **rois** (*list of arrays*) – A list of arrays of shape: math: *(R_l, 4)*, where: math: *R_l* is the number of RoIs in the: math: *l*- th feature map.
- **roi_indices** (*list of arrays*) – A list of arrays of shape $(R_l,)$.

Returns

`locs` and `confs`.

- **locs**: An arrays whose shape is $(R, n_class, 4)$, where R is the total number of RoIs in the batch.

” **confs**: A list of array whose shape is (R, n_class).

Return type tuple of two arrays

decode(*rois*, *roi_indices*, *locs*, *confs*, *scales*, *sizes*, *nms_thresh*, *score_thresh*)
Decodes back to coordinates of RoIs.

This method decodes `locs` and `confs` returned by a FPN network back to `bboxes`, `labels` and `scores`.

Parameters

- **rois** (*iterable of arrays*) – An iterable of arrays of shape $(R_l, 4)$, where R_l is the number of RoIs in the l-th feature map.
- **roi_indices** (*iterable of arrays*) – An iterable of arrays of shape $(R_l,)$.
- **locs** (*array*) – An array whose shape is $(R, n_class, 4)$, where R is the total number of RoIs in the given batch.
- **confs** (*array*) – An array whose shape is (R, n_class).
- **scales** (*list of floats*) – A list of floats returned by `prepare()`
- **sizes** (*list of tuples of two ints*) – A list of (H_n, W_n), where H_n and W_n are height and width of the n-th image.

- **nms_thresh** (*float*) – The threshold value for *non_maximum_suppression()*.
- **score_thresh** (*float*) – The threshold value for confidence score.

Returns bboxes, labels and scores.

Return type

tuple of three list of arrays

- **bboxes**: A list of float arrays of shape $(R'_n, 4)$, where R'_n is the number of bounding boxes in the n-th image. Each bounding box is organized by $(y_{min}, x_{min}, y_{max}, x_{max})$ in the second axis.
- **labels** : A list of integer arrays of shape $(R'_n,)$. Each value indicates the class of the bounding box. Values are in range $[0, L - 1]$, where L is the number of the foreground classes.
- **scores** : A list of float arrays of shape $(R'_n,)$. Each value indicates how confident the prediction is.

distribute(*rois*, *roi_indices*)

Assigns Rois to feature maps according to their size.

Parameters

- **rois** (*array*) – An array of shape $(R, 4)$, where R is the total number of RoIs in the given batch.
- **roi_indices** (*array*) – An array of shape $(R,)$.

Returns

rois and roi_indices.

- **rois**: A list of arrays of shape $(R_l, 4)$, where R_l is the number of RoIs in the l-th feature map.
- **roi_indices** : A list of arrays of shape $(R_l,)$.

Return type tuple of two lists

RPN

class chainercv.links.model.fpn.**RPN**(*scales*)

Region Proposal Network of Feature Pyramid Networks.

Parameters **scales** (*tuple of floats*) – The scales of feature maps.

__call__(*hs*)

Calculates RoIs.

Parameters **hs** (*iterable of array*) – An iterable of feature maps.

Returns

locs and confs.

- **locs**: A list of arrays whose shape is $(N, K_l, 4)$, where N is the size of batch and K_l is the number of the anchor boxes of the l-th level.

" **confs**: A list of array whose shape is (N, K_l).

Return type tuple of two arrays

anchors (*sizes*)
Calculates anchor boxes.

Parameters **sizes** (*iterable of tuples of two ints*) – An iterable of (H_l, W_l), where H_l and W_l are height and width of the l-th feature map.

Returns The shape of the l-th array is $(H_l * W_l * A, 4)$, where A is the number of anchor ratios.

Return type list of arrays

decode (*locs*, *confs*, *anchors*, *in_shape*)
Decodes back to coordinates of RoIs.

This method decodes `locs` and `confs` returned by a FPN network back to `rois` and `roi_indices`.

Parameters

- **locs** (*list of arrays*) – A list of arrays whose shape is $(N, K_l, 4)$, where N is the size of batch and N_l is the number of the anchor boxes of the l-th level.
- **confs** (*list of arrays*) – A list of array whose shape is (N, K_l).
- **anchors** (*list of arrays*) – Anchor boxes returned by *anchors()*.
- **in_shape** (*tuple of ints*) – The shape of input of array the feature extractor.

Returns

`rois` and `roi_indices`.

- **rois**: An array of shape $(R, 4)$, where R is the total number of RoIs in the given batch.
- **roi_indices** : An array of shape $(R,)$.

Return type tuple of two arrays

Train-only Utility

head_loss_pre

chainercv.links.model.fpn.**head_loss_pre** (*rois*, *roi_indices*, *std*, *bboxes*, *labels*)
Loss function for Head (pre).

This function processes RoIs for *head_loss_post()*.

Parameters

- **rois** (*iterable of arrays*) – An iterable of arrays of shape $(R_l, 4)$, where R_l is the number of RoIs in the l-th feature map.
- **roi_indices** (*iterable of arrays*) – An iterable of arrays of shape $(R_l,)$.
- **std** (*tuple of floats*) – Two coefficients used for encoding bounding boxes.
- **bboxes** (*list of arrays*) – A list of arrays whose shape is $(R_n, 4)$, where R_n is the number of ground truth bounding boxes.
- **labels** – A list of arrays whose shape is $(R_n,)$.

head_loss_post

chainercv.links.model.fpn.**head_loss_post**(*locs*, *confs*, *roi_indices*, *gt_locs*, *gt_labels*, *batch-size*)

Loss function for Head (post).

Parameters

- **locs** (*array*) – An array whose shape is $(R, n_class, 4)$, where R is the total number of RoIs in the given batch.
- **confs** (*array*) – An iterable of arrays whose shape is (R, n_class).
- **roi_indices** (*list of arrays*) – A list of arrays returned by `head_locs_pre()`.
- **gt_locs** (*list of arrays*) – A list of arrays returned by `head_locs_pre()`.
- **gt_labels** (*list of arrays*) – A list of arrays returned by `head_locs_pre()`.
- **batchsize** (*int*) – The size of batch.

Returns `loc_loss` and `conf_loss`.

Return type tuple of two variables

rpn_loss

chainercv.links.model.fpn.**rpn_loss**(*locs*, *confs*, *anchors*, *sizes*, *bboxes*)

Loss function for RPN.

Parameters

- **locs** (*iterable of arrays*) – An iterable of arrays whose shape is $(N, K_l, 4)$, where K_l is the number of the anchor boxes of the l-th level.
- **confs** (*iterable of arrays*) – An iterable of arrays whose shape is (N, K_l).
- **anchors** (*list of arrays*) – A list of arrays returned by `anchors()`.
- **sizes** (*list of tuples of two ints*) – A list of (H_n, W_n), where H_n and W_n are height and width of the n-th image.
- **bboxes** (*list of arrays*) – A list of arrays whose shape is $(R_n, 4)$, where R_n is the number of ground truth bounding boxes.

Returns `loc_loss` and `conf_loss`.

Return type tuple of two variables

SSD (Single Shot Multibox Detector)

Detection Links

SSD300

class chainercv.links.model.ssd.**SSD300**(*n_fg_class=None*, *pretrained_model=None*)

Single Shot Multibox Detector with 300x300 inputs.

This is a model of Single Shot Multibox Detector[1]. This model uses *VGG16Extractor300* as its feature extractor.

Parameters

- **n_fg_class** (*int*) – The number of classes excluding the background.
- **pretrained_model** (*string*) – The weight file to be loaded. This can take `'voc0712'`, *filepath* or `None`. The default value is `None`.
  - `'voc0712'`: Load weights trained on trainval split of PASCAL VOC 2007 and 2012. The weight file is downloaded and cached automatically. `n_fg_class` must be `20` or `None`. These weights were converted from the Caffe model provided by the original implementation. The conversion code is *chainercv/examples/ssd/caffe2npz.py*.
  - `'imagenet'`: Load weights of VGG-16 trained on ImageNet. The weight file is downloaded and cached automatically. This option initializes weights partially and the rests are initialized randomly. In this case, `n_fg_class` can be set to any number.
  - *filepath*: A path of npz file. In this case, `n_fg_class` must be specified properly.
  - `None`: Do not load weights.

SSD512

class chainercv.links.model.ssd.**SSD512**(*n_fg_class=None*, *pretrained_model=None*)

Single Shot Multibox Detector with 512x512 inputs.

This is a model of Single Shot Multibox Detector[2]. This model uses *VGG16Extractor512* as its feature extractor.

Parameters

- **n_fg_class** (*int*) – The number of classes excluding the background.
- **pretrained_model** (*string*) – The weight file to be loaded. This can take `'voc0712'`, *filepath* or `None`. The default value is `None`.
  - `'voc0712'`: Load weights trained on trainval split of PASCAL VOC 2007 and 2012. The weight file is downloaded and cached automatically. `n_fg_class` must be `20` or `None`. These weights were converted from the Caffe model provided by the original implementation. The conversion code is *chainercv/examples/ssd/caffe2npz.py*.
  - `'imagenet'`: Load weights of VGG-16 trained on ImageNet. The weight file is downloaded and cached automatically. This option initializes weights partially and the rests are initialized randomly. In this case, `n_fg_class` can be set to any number.
  - *filepath*: A path of npz file. In this case, `n_fg_class` must be specified properly.
  - `None`: Do not load weights.

Utility

[1] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, Alexander C. Berg. SSD: Single Shot MultiBox Detector. ECCV 2016.

[2] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, Alexander C. Berg. SSD: Single Shot MultiBox Detector. ECCV 2016.

Multibox

class chainercv.links.model.ssd.**Multibox**(*n_class*, *aspect_ratios*, *initialW=None*, *initial_bias=None*)

Multibox head of Single Shot Multibox Detector.

This is a head part of Single Shot Multibox Detector[3]. This link computes `mb_locs` and `mb_confs` from feature maps. `mb_locs` contains information of the coordinates of bounding boxes and `mb_confs` contains confidence scores of each classes.

Parameters

- **n_class** (*int*) – The number of classes possibly including the background.
- **aspect_ratios** (*iterable of tuple or int*) – The aspect ratios of default bounding boxes for each feature map.
- **initialW** – An initializer used in `chainer.links.Convolution2d.__init__()`. The default value is `chainer.initializers.LeCunUniform`.
- **initial_bias** – An initializer used in `chainer.links.Convolution2d.__init__()`. The default value is `chainer.initializers.Zero`.

__call__(*xs*)

Compute loc and conf from feature maps

This method computes `mb_locs` and `mb_confs` from given feature maps.

Parameters **xs** (*iterable of chainer.Variable*) – An iterable of feature maps. The number of feature maps must be same as the number of `aspect_ratios`.

Returns

This method returns two `chainer.Variable`: `mb_locs` and `mb_confs`.

- **mb_locs**: A variable of float arrays of shape $(B, K, 4)$, where B is the number of samples in the batch and K is the number of default bounding boxes.
- **mb_confs**: A variable of float arrays of shape $(B, K, n_fg_class + 1)$.

Return type tuple of chainer.Variable

MultiboxCoder

class chainercv.links.model.ssd.**MultiboxCoder**(*grids*, *aspect_ratios*, *steps*, *sizes*, *variance*)

A helper class to encode/decode bounding boxes.

This class encodes `(bbox, label)` to `(mb_loc, mb_label)` and decodes `(mb_loc, mb_conf)` to `(bbox, label, score)`. These encoding/decoding are used in Single Shot Multibox Detector[4].

- `mb_loc`: An array representing offsets and scales from the default bounding boxes. Its shape is $(K, 4)$, where K is the number of the default bounding boxes. The second axis is composed by $(\Delta y, \Delta x, \Delta h, \Delta w)$. These values are computed by the following formulas.
  - $\Delta y = (b_y - m_y)/(m_h * v_0)$
  - $\Delta x = (b_x - m_x)/(m_w * v_0)$

[3] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, Alexander C. Berg. SSD: Single Shot MultiBox Detector. ECCV 2016.

[4] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, Alexander C. Berg. SSD: Single Shot MultiBox Detector. ECCV 2016.

- $\Delta h = log(b_h/m_h)/v_1$
- $\Delta w = log(b_w/m_w)/v_1$

(m_y, m_x) and (m_h, m_w) are center coodinates and size of a default bounding box. (b_y, b_x) and (b_h, b_w) are center coodinates and size of a given bounding boxes that is assined to the default bounding box. (v_0, v_1) are coefficients that can be set by argument `variance`.

- `mb_label`: An array representing classes of ground truth bounding boxes. Its shape is $(K,)$.
- `mb_conf`: An array representing classes of predicted bounding boxes. Its shape is $(K, n_fg_class + 1)$.

Parameters

- **grids** (*iterable of ints*) – An iterable of integers. Each integer indicates the size of a feature map.
- **aspect_ratios** (*iterable of tuples of ints*) – An iterable of tuples of integers used to compute the default bounding boxes. Each tuple indicates the aspect ratios of the default bounding boxes at each feature maps. The length of this iterable should be `len(grids)`.
- **steps** (*iterable of floats*) – The step size for each feature map. The length of this iterable should be `len(grids)`.
- **sizes** (*iterable of floats*) – The base size of default bounding boxes for each feature map. The length of this iterable should be `len(grids) + 1`.
- **variance** (*tuple of floats*) – Two coefficients for encoding/decoding the locations of bounding boxes. The first value is used to encode/decode coordinates of the centers. The second value is used to encode/decode the sizes of bounding boxes.

decode (*mb_loc*, *mb_conf*, *nms_thresh=0.45*, *score_thresh=0.6*)

Decodes back to coordinates and classes of bounding boxes.

This method decodes `mb_loc` and `mb_conf` returned by a SSD network back to `bbox`, `label` and `score`.

Parameters

- **mb_loc** (*array*) – A float array whose shape is $(K, 4)$, K is the number of default bounding boxes.
- **mb_conf** (*array*) – A float array whose shape is $(K, n_fg_class + 1)$.
- **nms_thresh** (*float*) – The threshold value for *non_maximum_suppression()*. The default value is `0.45`.
- **score_thresh** (*float*) – The threshold value for confidence score. If a bounding box whose confidence score is lower than this value, the bounding box will be suppressed. The default value is `0.6`.

Returns

This method returns a tuple of three arrays, `(bbox, label, score)`.

- **bbox**: A float array of shape $(R, 4)$, where R is the number of bounding boxes in a image. Each bounding box is organized by $(y_{min}, x_{min}, y_{max}, x_{max})$ in the second axis.
- **label** : An integer array of shape $(R,)$. Each value indicates the class of the bounding box.
- **score** : A float array of shape $(R,)$. Each value indicates how confident the prediction is.

Return type tuple of three arrays

encode(*bbox*, *label*, *iou_thresh=0.5*)
Encodes coordinates and classes of bounding boxes.

This method encodes `bbox` and `label` to `mb_loc` and `mb_label`, which are used to compute multibox loss.

Parameters

- **bbox** (*array*) – A float array of shape $(R, 4)$, where R is the number of bounding boxes in an image. Each bounding box is organized by $(y_{min}, x_{min}, y_{max}, x_{max})$ in the second axis.
- **label** (*array*) – An integer array of shape $(R,)$. Each value indicates the class of the bounding box.
- **iou_thresh** (*float*) – The threshold value to determine a default bounding box is assigned to a ground truth or not. The default value is `0.5`.

Returns

This method returns a tuple of two arrays, `(mb_loc, mb_label)`.

- **mb_loc**: A float array of shape $(K, 4)$, where K is the number of default bounding boxes.
- **mb_label**: An integer array of shape $(K,)$.

Return type tuple of two arrays

Normalize

class chainercv.links.model.ssd.**Normalize**(*n_channel*, *initial=0*, *eps=1e-05*)
Learnable L2 normalization[5].

This link normalizes input along the channel axis and scales it. The scale factors are trained channel-wise.

Parameters

- **n_channel** (*int*) – The number of channels.
- **initial** – A value to initialize the scale factors. It is pased to `chainer.initializers._get_initializer()`. The default value is 0.
- **eps** (*float*) – A small value to avoid zero-division. The default value is $1e-5$.

__call__(*x*)
Normalize input and scale it.

Parameters **x** (*chainer.Variable*) – A variable holding 4-dimensional array. Its `dtype` is `numpy.float32`.

Returns The shape and `dtype` are same as those of input.

Return type chainer.Variable

SSD

class chainercv.links.model.ssd.**SSD**(*extractor*, *multibox*, *steps*, *sizes*, *variance=(0.1, 0.2)*, *mean=0*)
Base class of Single Shot Multibox Detector.

[5] Wei Liu, Andrew Rabinovich, Alexander C. Berg. ParseNet: Looking Wider to See Better. ICLR 2016.

This is a base class of Single Shot Multibox Detector[6].

Parameters

- **extractor** – A link which extracts feature maps. This link must have `insize`, `grids` and *`__call__()`*.
  - `insize`: An integer which indicates the size of input images. Images are resized to this size before feature extraction.
  - `grids`: An iterable of integer. Each integer indicates the size of feature map. This value is used by `MultiBboxCoder`.
  - `__call__()`: A method which computes feature maps. It must take a batched images and return batched feature maps.
- **multibox** – A link which computes `mb_locs` and `mb_confs` from feature maps. This link must have `n_class`, `aspect_ratios` and *`__call__()`*.
  - `n_class`: An integer which indicates the number of classes. This value should include the background class.
  - `aspect_ratios`: An iterable of tuple of integer. Each tuple indicates the aspect ratios of default bounding boxes at each feature maps. This value is used by *`MultiboxCoder`*.
  - *`__call__()`*: A method which computes `mb_locs` and `mb_confs`. It must take a batched feature maps and return `mb_locs` and `mb_confs`.
- **steps** (*iterable of float*) – The step size for each feature map. This value is used by *`MultiboxCoder`*.
- **sizes** (*iterable of float*) – The base size of default bounding boxes for each feature map. This value is used by *`MultiboxCoder`*.
- **variance** (*tuple of floats*) – Two coefficients for decoding the locations of bounding boxes. This value is used by *`MultiboxCoder`*. The default value is `(0.1, 0.2)`.
- **nms_thresh** (*float*) – The threshold value for *`non_maximum_suppression()`*. The default value is `0.45`. This value can be changed directly or by using *`use_preset()`*.
- **score_thresh** (*float*) – The threshold value for confidence score. If a bounding box whose confidence score is lower than this value, the bounding box will be suppressed. The default value is `0.6`. This value can be changed directly or by using *`use_preset()`*.

__call__(*x*)

Compute localization and classification from a batch of images.

This method computes two variables, `mb_locs` and `mb_confs`. `self.coder.decode()` converts these variables to bounding box coordinates and confidence scores. These variables are also used in training SSD.

Parameters **x** (*chainer.Variable*) – A variable holding a batch of images. The images are preprocessed by `_prepare()`.

Returns

This method returns two variables, `mb_locs` and `mb_confs`.

- **mb_locs**: A variable of float arrays of shape $(B, K, 4)$, where B is the number of samples in the batch and K is the number of default bounding boxes.

[6] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, Alexander C. Berg. SSD: Single Shot MultiBox Detector. ECCV 2016.

- **mb_confs**: A variable of float arrays of shape $(B, K, n_fg_class + 1)$.

Return type tuple of chainer.Variable

predict(*imgs*)
Detect objects from images.

This method predicts objects for each image.

Parameters **imgs** (*iterable of numpy.ndarray*) – Arrays holding images. All images are in CHW and RGB format and the range of their value is $[0, 255]$.

Returns

This method returns a tuple of three lists, `(bboxes, labels, scores)`.

- **bboxes**: A list of float arrays of shape $(R, 4)$, where R is the number of bounding boxes in a image. Each bounding box is organized by $(y_{min}, x_{min}, y_{max}, x_{max})$ in the second axis.
- **labels** : A list of integer arrays of shape $(R,)$. Each value indicates the class of the bounding box. Values are in range $[0, L - 1]$, where L is the number of the foreground classes.
- **scores** : A list of float arrays of shape $(R,)$. Each value indicates how confident the prediction is.

Return type tuple of lists

to_cpu()
Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu(*device=None*)
Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

use_preset(*preset*)
Use the given preset during prediction.

This method changes values of `nms_thresh` and `score_thresh`. These values are a threshold value used for non maximum suppression and a threshold value to discard low confidence proposals in *predict()*, respectively.

If the attributes need to be changed to something other than the values provided in the presets, please modify them by directly accessing the public attributes.

Parameters **preset** (*{'visualize', 'evaluate'}*) – A string to determine the preset to use.

VGG16

class chainercv.links.model.ssd.**VGG16**
An extended VGG-16 model for SSD300 and SSD512.

This is an extended VGG-16 model proposed in[7]. The differences from original VGG-16[8] are shown below.

- `conv5_1`, `conv5_2` and `conv5_3` are changed from `Convolution2d` to `DilatedConvolution2d`.
- *Normalize* is inserted after `conv4_3`.
- The parameters of max pooling after `conv5_3` are changed.
- `fc6` and `fc7` are converted to `conv6` and `conv7`.

__call__(*...*) <==> *x*(*...*)

VGG16Extractor300

class chainercv.links.model.ssd.**VGG16Extractor300**
A VGG-16 based feature extractor for SSD300.

This is a feature extractor for *SSD300*. This extractor is based on *VGG16*.

__call__(*x*)
Compute feature maps from a batch of images.

This method extracts feature maps from `conv4_3`, `conv7`, `conv8_2`, `conv9_2`, `conv10_2`, and `conv11_2`.

Parameters **x** (*ndarray*) – An array holding a batch of images. The images should be resized to 300×300.

Returns Each variable contains a feature map.

Return type list of Variable

VGG16Extractor512

class chainercv.links.model.ssd.**VGG16Extractor512**
A VGG-16 based feature extractor for SSD512.

This is a feature extractor for *SSD512*. This extractor is based on *VGG16*.

__call__(*x*)
Compute feature maps from a batch of images.

This method extracts feature maps from `conv4_3`, `conv7`, `conv8_2`, `conv9_2`, `conv10_2`, `conv11_2`, and `conv12_2`.

Parameters **x** (*ndarray*) – An array holding a batch of images. The images should be resized to 512×512.

Returns Each variable contains a feature map.

Return type list of Variable

[7] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, Alexander C. Berg. SSD: Single Shot MultiBox Detector. ECCV 2016.
[8] Karen Simonyan, Andrew Zisserman. Very Deep Convolutional Networks for Large-Scale Image Recognition. ICLR 2015.

Train-only Utility

GradientScaling

class chainercv.links.model.ssd.**GradientScaling**(*rate*)

Optimizer/UpdateRule hook function for scaling gradient.

This hook function scales gradient by a constant value.

Parameters **rate** (*float*) – Coefficient for scaling.

Variables **rate** (*float*) – Coefficient for scaling.

multibox_loss

chainercv.links.model.ssd.**multibox_loss**(*mb_locs*, *mb_confs*, *gt_mb_locs*, *gt_mb_labels*, *k*, *comm=None*)

Computes multibox losses.

This is a loss function used in[9]. This function returns `loc_loss` and `conf_loss`. `loc_loss` is a loss for localization and `conf_loss` is a loss for classification. The formulas of these losses can be found in the equation (2) and (3) in the original paper.

Parameters

- **mb_locs** (*chainer.Variable or array*) – The offsets and scales for predicted bounding boxes. Its shape is $(B, K, 4)$, where B is the number of samples in the batch and K is the number of default bounding boxes.
- **mb_confs** (*chainer.Variable or array*) – The classes of predicted bounding boxes. Its shape is (B, K, n_class). This function assumes the first class is background (negative).
- **gt_mb_locs** (*chainer.Variable or array*) – The offsets and scales for ground truth bounding boxes. Its shape is $(B, K, 4)$.
- **gt_mb_labels** (*chainer.Variable or array*) – The classes of ground truth bounding boxes. Its shape is (B, K).
- **k** (*float*) – A coefficient which is used for hard negative mining. This value determines the ratio between the number of positives and that of mined negatives. The value used in the original paper is `3`.
- **comm** (*CommunicatorBase*) – A ChainerMN communicator. If it is specified, the number of positive examples is computed among all GPUs.

Returns This function returns two `chainer.Variable`: `loc_loss` and `conf_loss`.

Return type tuple of chainer.Variable

[9] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, Alexander C. Berg. SSD: Single Shot MultiBox Detector. ECCV 2016.

random_crop_with_bbox_constraints

chainercv.links.model.ssd.**random_crop_with_bbox_constraints**(*img*, *bbox*, *min_scale=0.3*, *max_scale=1*, *max_aspect_ratio=2*, *constraints=None*, *max_trial=50*, *return_param=False*)

Crop an image randomly with bounding box constraints.

This data augmentation is used in training of Single Shot Multibox Detector[10]. More details can be found in data augmentation section of the original paper.

Parameters

- **img** (*ndarray*) – An image array to be cropped. This is in CHW format.
- **bbox** (*ndarray*) – Bounding boxes used for constraints. The shape is $(R, 4)$. R is the number of bounding boxes.
- **min_scale** (*float*) – The minimum ratio between a cropped region and the original image. The default value is `0.3`.
- **max_scale** (*float*) – The maximum ratio between a cropped region and the original image. The default value is `1`.
- **max_aspect_ratio** (*float*) – The maximum aspect ratio of cropped region. The default value is `2`.
- **constaraints** (*iterable of tuples*) – An iterable of constraints. Each constraint should be `(min_iou, max_iou)` format. If you set `min_iou` or `max_iou` to `None`, it means not limited. If this argument is not specified, `((0.1, None), (0.3, None), (0.5, None), (0.7, None), (0.9, None), (None, 1))` will be used.
- **max_trial** (*int*) – The maximum number of trials to be conducted for each constraint. If this function can not find any region that satisfies the constraint in max_trial trials, this function skips the constraint. The default value is `50`.
- **return_param** (*bool*) – If `True`, this function returns information of intermediate values.

Returns

If `return_param = False`, returns an array `img` that is cropped from the input array.

If `return_param = True`, returns a tuple whose elements are `img, param`. `param` is a dictionary of intermediate parameters whose contents are listed below with key, value-type and the description of the value.

- **constraint** (*tuple*): The chosen constraint.
- **y_slice** (*slice*): A slice in vertical direction used to crop the input image.
- **x_slice** (*slice*): A slice in horizontal direction used to crop the input image.

Return type ndarray or (ndarray, dict)

[10] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, Alexander C. Berg. SSD: Single Shot MultiBox Detector. ECCV 2016.

random_distort

chainercv.links.model.ssd.**random_distort**(*img*, *brightness_delta=32*, *contrast_low=0.5*, *contrast_high=1.5*, *saturation_low=0.5*, *saturation_high=1.5*, *hue_delta=18*)

A color related data augmentation used in SSD.

This function is a combination of four augmentation methods: brightness, contrast, saturation and hue.

- brightness: Adding a random offset to the intensity of the image.
- contrast: Multiplying the intensity of the image by a random scale.
- saturation: Multiplying the saturation of the image by a random scale.
- hue: Adding a random offset to the hue of the image randomly.

This data augmentation is used in training of Single Shot Multibox Detector[11].

Note that this function requires `cv2`.

Parameters

- **img** (*ndarray*) – An image array to be augmented. This is in CHW and RGB format.
- **brightness_delta** (*float*) – The offset for saturation will be drawn from $[-brightness_delta, brightness_delta]$. The default value is `32`.
- **contrast_low** (*float*) – The scale for contrast will be drawn from $[contrast_low, contrast_high]$. The default value is `0.5`.
- **contrast_high** (*float*) – See `contrast_low`. The default value is `1.5`.
- **saturation_low** (*float*) – The scale for saturation will be drawn from $[saturation_low, saturation_high]$. The default value is `0.5`.
- **saturation_high** (*float*) – See `saturation_low`. The default value is `1.5`.
- **hue_delta** (*float*) – The offset for hue will be drawn from $[-hue_delta, hue_delta]$. The default value is `18`.

Returns An image in CHW and RGB format.

resize_with_random_interpolation

chainercv.links.model.ssd.**resize_with_random_interpolation**(*img*, *size*, *return_param=False*)

Resize an image with a randomly selected interpolation method.

This function is similar to *chainercv.transforms.resize()*, but this chooses the interpolation method randomly.

This data augmentation is used in training of Single Shot Multibox Detector[12].

Note that this function requires `cv2`.

Parameters

- **img** (*ndarray*) – An array to be transformed. This is in CHW format and the type should be `numpy.float32`.

[11] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, Alexander C. Berg. SSD: Single Shot MultiBox Detector. ECCV 2016.

[12] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, Alexander C. Berg. SSD: Single Shot MultiBox Detector. ECCV 2016.

- **size** (*tuple*) – This is a tuple of length 2. Its elements are ordered as (height, width).
- **return_param** (*bool*) – Returns information of interpolation.

Returns

If `return_param = False`, returns an array `img` that is the result of rotation.

If `return_param = True`, returns a tuple whose elements are `img, param`. `param` is a dictionary of intermediate parameters whose contents are listed below with key, value-type and the description of the value.

- **interpolatation**: The chosen interpolation method.

Return type ndarray or (ndarray, dict)

YOLO

Detection Links

YOLOv2

class chainercv.links.model.yolo.**YOLOv2**(*n_fg_class=None*, *pretrained_model=None*)

YOLOv2.

This is a model of YOLOv2[1]. This model uses *Darknet19Extractor* as its feature extractor.

Parameters

- **n_fg_class** (*int*) – The number of classes excluding the background.
- **pretrained_model** (*string*) – The weight file to be loaded. This can take `'voc0712'`, *filepath* or `None`. The default value is `None`.
  - `'voc0712'`: Load weights trained on trainval split of PASCAL VOC 2007 and 2012. The weight file is downloaded and cached automatically. `n_fg_class` must be `20` or `None`. These weights were converted from the darknet model provided by the original implementation. The conversion code is *chainercv/examples/yolo/darknet2npz.py*.
  - *filepath*: A path of npz file. In this case, `n_fg_class` must be specified properly.
  - `None`: Do not load weights.

YOLOv3

class chainercv.links.model.yolo.**YOLOv3**(*n_fg_class=None*, *pretrained_model=None*)

YOLOv3.

This is a model of YOLOv3[2]. This model uses *Darknet53Extractor* as its feature extractor.

Parameters

- **n_fg_class** (*int*) – The number of classes excluding the background.
- **pretrained_model** (*string*) – The weight file to be loaded. This can take `'voc0712'`, *filepath* or `None`. The default value is `None`.

[1] Joseph Redmon, Ali Farhadi. YOLO9000: Better, Faster, Stronger. CVPR 2017.
[2] Joseph Redmon, Ali Farhadi. YOLOv3: An Incremental Improvement. arXiv 2018.

- `'voc0712'`: Load weights trained on trainval split of PASCAL VOC 2007 and 2012. The weight file is downloaded and cached automatically. `n_fg_class` must be `20` or `None`. These weights were converted from the darknet model. The conversion code is *chainercv/examples/yolo/darknet2npz.py*.
- *filepath*: A path of npz file. In this case, `n_fg_class` must be specified properly.
- `None`: Do not load weights.

to_cpu()
Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu(*device=None*)
Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

Utility

ResidualBlock

class chainercv.links.model.yolo.**ResidualBlock**(**links*)
ChainList with a residual connection.

__call__(*...*) <==> *x(...)*

Darknet19Extractor

class chainercv.links.model.yolo.**Darknet19Extractor**
A Darknet19 based feature extractor for YOLOv2.

This is a feature extractor for *YOLOv2*

__call__(*x*)
Compute a feature map from a batch of images.

Parameters **x** (*ndarray*) – An array holding a batch of images. The images should be resized to 416×416.

Returns

Return type Variable

Darknet53Extractor

class chainercv.links.model.yolo.**Darknet53Extractor**
A Darknet53 based feature extractor for YOLOv3.

This is a feature extractor for *YOLOv3*

__call__(*x*)
Compute feature maps from a batch of images.

This method extracts feature maps from 3 layers.

Parameters **x** (*ndarray*) – An array holding a batch of images. The images should be resized to 416×416.

Returns Each variable contains a feature map.

Return type list of Variable

YOLOBase

class chainercv.links.model.yolo.**YOLOBase**(***links*)
Base class for YOLOv2 and YOLOv3.

A subclass of this class should have `extractor`, `__call__()`, and `_decode()`.

predict(*imgs*)
Detect objects from images.

This method predicts objects for each image.

Parameters **imgs** (*iterable of numpy.ndarray*) – Arrays holding images. All images are in CHW and RGB format and the range of their value is $[0, 255]$.

Returns

This method returns a tuple of three lists, `(bboxes, labels, scores)`.

- **bboxes**: A list of float arrays of shape $(R, 4)$, where R is the number of bounding boxes in a image. Each bounding box is organized by $(y_{min}, x_{min}, y_{max}, x_{max})$ in the second axis.
- **labels** : A list of integer arrays of shape $(R,)$. Each value indicates the class of the bounding box. Values are in range $[0, L - 1]$, where L is the number of the foreground classes.
- **scores** : A list of float arrays of shape $(R,)$. Each value indicates how confident the prediction is.

Return type tuple of lists

use_preset(*preset*)
Use the given preset during prediction.

This method changes values of `nms_thresh` and `score_thresh`. These values are a threshold value used for non maximum suppression and a threshold value to discard low confidence proposals in *predict()*, respectively.

If the attributes need to be changed to something other than the values provided in the presets, please modify them by directly accessing the public attributes.

Parameters **preset** (*{'visualize', 'evaluate'}*) – A string to determine the preset to use.

YOLOv2Base

class chainercv.links.model.yolo.**YOLOv2Base**(*n_fg_class=None*, *pretrained_model=None*)
Base class for YOLOv2 and YOLOv2Tiny.

A subclass of this class should have `_extractor`, `_models`, and `_anchors`.

__call__(*x*)
Compute localization, objectness, and classification from a batch of images.

This method computes three variables, `locs`, `objs`, and `confs`. `self._decode()` converts these variables to bounding box coordinates and confidence scores. These variables are also used in training YOLOv2.

Parameters **x** (*chainer.Variable*) – A variable holding a batch of images.

Returns

This method returns three variables, `locs`, `objs`, and `confs`.

- **locs**: A variable of float arrays of shape $(B, K, 4)$, where B is the number of samples in the batch and K is the number of default bounding boxes.
- **objs**: A variable of float arrays of shape (B, K).
- **confs**: A variable of float arrays of shape (B, K, n_fg_class).

Return type tuple of chainer.Variable

to_cpu()
Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu(*device=None*)
Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

Semantic Segmentation

Semantic segmentation links share a common method `predict()` to conduct semantic segmentation of images. For more details, please read *SegNetBasic.predict()*.

SegNet

Semantic Segmentation Link

SegNetBasic

class chainercv.links.model.segnet.**SegNetBasic**(*n_class=None*, *pretrained_model=None*, *initialW=None*)
SegNet Basic for semantic segmentation.

This is a SegNet[1] model for semantic segmenation. This is based on SegNetBasic model that is found here.

[1] Vijay Badrinarayanan, Alex Kendall and Roberto Cipolla "SegNet: A Deep Convolutional Encoder-Decoder Architecture for Image Segmentation." PAMI, 2017

When you specify the path of a pretrained chainer model serialized as a `npz` file in the constructor, this chain model automatically initializes all the parameters with it. When a string in prespecified set is provided, a pretrained model is loaded from weights distributed on the Internet. The list of pretrained models supported are as follows:

- `camvid`: Loads weights trained with the train split of CamVid dataset.

Parameters

- **n_class** (*int*) – The number of classes. If `None`, it can be infered if `pretrained_model` is given.
- **pretrained_model** (*string*) – The destination of the pretrained chainer model serialized as a `npz` file. If this is one of the strings described above, it automatically loads weights stored under a directory `$CHAINER_DATASET_ROOT/pfnet/chainercv/models/`, where `$CHAINER_DATASET_ROOT` is set as `$HOME/.chainer/dataset` unless you specify another value by modifying the environment variable.
- **initialW** (*callable*) – Initializer for convolution layers.

__call__(*x*)

Compute an image-wise score from a batch of images

Parameters **x** (*chainer.Variable*) – A variable with 4D image array.

Returns An image-wise score. Its channel size is `self.n_class`.

Return type chainer.Variable

predict(*imgs*)

Conduct semantic segmentations from images.

Parameters **imgs** (*iterable of numpy.ndarray*) – Arrays holding images. All images are in CHW and RGB format and the range of their values are $[0, 255]$.

Returns List of integer labels predicted from each image in the input list.

Return type list of numpy.ndarray

Classifiers

Classifier

PixelwiseSoftmaxClassifier

class chainercv.links.**PixelwiseSoftmaxClassifier**(*predictor*, *ignore_label=-1*, *class_weight=None*)

A pixel-wise classifier.

It computes the loss based on a given input/label pair for semantic segmentation.

Parameters

- **predictor** (*Link*) – Predictor network.
- **ignore_label** (*int*) – A class id that is going to be ignored in evaluation. The default value is -1.
- **class_weight** (*array*) – An array that contains constant weights that will be multiplied with the loss values along with the channel dimension. This will be used in `chainer.functions.softmax_cross_entropy()`.

__call__(*x*, *t*)
Computes the loss value for an image and label pair.

Parameters

- **x** (*Variable*) – A variable with a batch of images.
- **t** (*Variable*) – A variable with the ground truth image-wise label.

Returns Loss value.

Return type Variable

to_cpu()
Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu(*device=None*)
Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

3.7.2 Connection

Connection

Conv2DActiv

class chainercv.links.connection.**Conv2DActiv**(*in_channels*, *out_channels*, *ksize=None*, *stride=1*, *pad=0*, *dilate=1*, *nobias=False*, *initialW=None*, *initial_bias=None*, *activ=<function relu>*)
Convolution2D –> Activation

This is a chain that does two-dimensional convolution and applies an activation.

The arguments are the same as those of `chainer.links.Convolution2D` except for `activ`.

Example

There are several ways to initialize a *Conv2DActiv*.

1. Give the first three arguments explicitly:

```
>>> l = Conv2DActiv(5, 10, 3)
```

2. Omit `in_channels` or fill it with `None`:

In these ways, attributes are initialized at runtime based on the channel size of the input.

```
>>> l = Conv2DActiv(10, 3)
>>> l = Conv2DActiv(None, 10, 3)
```

Parameters

- **in_channels** (*int or None*) – Number of channels of input arrays. If `None`, parameter initialization will be deferred until the first forward data pass at which time the size will be determined.
- **out_channels** (*int*) – Number of channels of output arrays.
- **ksize** (*int or tuple of ints*) – Size of filters (a.k.a. kernels). `ksize=k` and `ksize=(k, k)` are equivalent.
- **stride** (*int or tuple of ints*) – Stride of filter applications. `stride=s` and `stride=(s, s)` are equivalent.
- **pad** (*int or tuple of ints*) – Spatial padding width for input arrays. `pad=p` and `pad=(p, p)` are equivalent.
- **dilate** (*int or tuple of ints*) – Dilation factor of filter applications. `dilate=d` and `dilate=(d, d)` are equivalent.
- **nobias** (*bool*) – If `True`, then this link does not use the bias term.
- **initialW** (*callable*) – Initial weight value. If `None`, the default initializer is used. May also be a callable that takes `numpy.ndarray` or `cupy.ndarray` and edits its value.
- **initial_bias** (*callable*) – Initial bias value. If `None`, the bias is set to 0. May also be a callable that takes `numpy.ndarray` or `cupy.ndarray` and edits its value.
- **activ** (*callable*) – An activation function. The default value is `chainer.functions.relu()`. If this is `None`, no activation is applied (i.e. the activation is the identity function).

Conv2DBNActiv

class chainercv.links.connection.**Conv2DBNActiv**(*in_channels*, *out_channels*, *ksize=None*, *stride=1*, *pad=0*, *dilate=1*, *groups=1*, *nobias=True*, *initialW=None*, *initial_bias=None*, *activ=<function relu>*, *bn_kwargs={}*)

Convolution2D –> Batch Normalization –> Activation

This is a chain that sequentially applies a two-dimensional convolution, a batch normalization and an activation.

The arguments are the same as that of `chainer.links.Convolution2D` except for `activ` and `bn_kwargs`. `bn_kwargs` can include `comm` key and a communicator of ChainerMN as the value to use `chainermn.links.MultiNodeBatchNormalization`. If `comm` is not included in `bn_kwargs`, `chainer.links.BatchNormalization` link from Chainer is used. Note that the default value for the `nobias` is changed to `True`.

Example

There are several ways to initialize a *Conv2DBNActiv*.

1. Give the first three arguments explicitly:

```
>>> l = Conv2DBNActiv(5, 10, 3)
```

2. Omit `in_channels` or fill it with `None`:

   In these ways, attributes are initialized at runtime based on the channel size of the input.

```
>>> l = Conv2DBNActiv(10, 3)
>>> l = Conv2DBNActiv(None, 10, 3)
```

Parameters

- **in_channels** (*int or None*) – Number of channels of input arrays. If `None`, parameter initialization will be deferred until the first forward data pass at which time the size will be determined.
- **out_channels** (*int*) – Number of channels of output arrays.
- **ksize** (*int or tuple of ints*) – Size of filters (a.k.a. kernels). `ksize=k` and `ksize=(k, k)` are equivalent.
- **stride** (*int or tuple of ints*) – Stride of filter applications. `stride=s` and `stride=(s, s)` are equivalent.
- **pad** (*int or tuple of ints*) – Spatial padding width for input arrays. `pad=p` and `pad=(p, p)` are equivalent.
- **dilate** (*int or tuple of ints*) – Dilation factor of filter applications. `dilate=d` and `dilate=(d, d)` are equivalent.
- **groups** (*int*) – The number of groups to use grouped convolution. The default is one, where grouped convolution is not used.
- **nobias** (*bool*) – If `True`, then this link does not use the bias term.
- **initialW** (*callable*) – Initial weight value. If `None`, the default initializer is used. May also be a callable that takes `numpy.ndarray` or `cupy.ndarray` and edits its value.
- **initial_bias** (*callable*) – Initial bias value. If `None`, the bias is set to 0. May also be a callable that takes `numpy.ndarray` or `cupy.ndarray` and edits its value.
- **activ** (*callable*) – An activation function. The default value is `chainer.functions.relu()`. If this is `None`, no activation is applied (i.e. the activation is the identity function).
- **bn_kwargs** (*dict*) – Keyword arguments passed to initialize `chainer.links.BatchNormalization`. If a ChainerMN communicator (`CommunicatorBase`) is given with the key `comm`, `MultiNodeBatchNormalization` will be used for the batch normalization. Otherwise, `BatchNormalization` will be used.

SEBlock

class chainercv.links.connection.**SEBlock**(*n_channel*, *ratio=16*)

A squeeze-and-excitation block.

This block is part of squeeze-and-excitation networks. Channel-wise multiplication weights are inferred from and applied to input feature map. Please refer to the original paper for more details.

See also:

chainercv.links.model.senet.SEResNet

Parameters

- **n_channel** (*int*) – The number of channels of the input and output array.
- **ratio** (*int*) – Reduction ratio of `n_channel` to the number of hidden layer units.

3.8 Transforms

3.8.1 Image

center_crop

`chainercv.transforms.`**center_crop**(*img*, *size*, *return_param=False*, *copy=False*)

Center crop an image by *size*.

An image is cropped to `size`. The center of the output image and the center of the input image are same.

Parameters

- **img** (*ndarray*) – An image array to be cropped. This is in CHW format.
- **size** (*tuple*) – The size of output image after cropping. This value is $(height, width)$.
- **return_param** (*bool*) – If `True`, this function returns information of slices.
- **copy** (*bool*) – If `False`, a view of `img` is returned.

Returns

If `return_param = False`, returns an array `out_img` that is cropped from the input array.

If `return_param = True`, returns a tuple whose elements are `out_img, param`. `param` is a dictionary of intermediate parameters whose contents are listed below with key, value-type and the description of the value.

- **y_slice** (*slice*): A slice used to crop the input image. The relation below holds together with `x_slice`.
- **x_slice** (*slice*): Similar to `y_slice`.

```
out_img = img[:, y_slice, x_slice]
```

Return type ndarray or (ndarray, dict)

flip

`chainercv.transforms.`**flip**(*img*, *y_flip=False*, *x_flip=False*, *copy=False*)

Flip an image in vertical or horizontal direction as specified.

Parameters

- **img** (*ndarray*) – An array that gets flipped. This is in CHW format.
- **y_flip** (*bool*) – Flip in vertical direction.
- **x_flip** (*bool*) – Flip in horizontal direction.
- **copy** (*bool*) – If False, a view of `img` will be returned.

Returns Transformed `img` in CHW format.

pca_lighting

chainercv.transforms.**pca_lighting**(*img*, *sigma*, *eigen_value=None*, *eigen_vector=None*)
AlexNet style color augmentation

This method adds a noise vector drawn from a Gaussian. The direction of the Gaussian is same as that of the principal components of the dataset.

This method is used in training of AlexNet[1].

Parameters

- **img** (*ndarray*) – An image array to be augmented. This is in CHW and RGB format.
- **sigma** (*float*) – Standard deviation of the Gaussian. In the original paper, this value is 10% of the range of intensity (25.5 if the range is $[0, 255]$).
- **eigen_value** (*ndarray*) – An array of eigen values. The shape has to be $(3,)$. If it is not specified, the values computed from ImageNet are used.
- **eigen_vector** (*ndarray*) – An array of eigen vectors. The shape has to be $(3, 3)$. If it is not specified, the vectors computed from ImageNet are used.

Returns An image in CHW format.

random_crop

chainercv.transforms.**random_crop**(*img*, *size*, *return_param=False*, *copy=False*)
Crop array randomly into *size*.

The input image is cropped by a randomly selected region whose shape is `size`.

Parameters

- **img** (*ndarray*) – An image array to be cropped. This is in CHW format.
- **size** (*tuple*) – The size of output image after cropping. This value is $(height, width)$.
- **return_param** (*bool*) – If `True`, this function returns information of slices.
- **copy** (*bool*) – If `False`, a view of `img` is returned.

Returns

If `return_param = False`, returns an array `out_img` that is cropped from the input array.

If `return_param = True`, returns a tuple whose elements are `out_img, param`. `param` is a dictionary of intermediate parameters whose contents are listed below with key, value-type and the description of the value.

- **y_slice** (*slice*): A slice used to crop the input image. The relation below holds together with `x_slice`.
- **x_slice** (*slice*): Similar to `x_slice`.

```
out_img = img[:, y_slice, x_slice]
```

Return type ndarray or (ndarray, dict)

[1] Alex Krizhevsky, Ilya Sutskever, Geoffrey E. Hinton. ImageNet Classification with Deep Convolutional Neural Networks. NIPS 2012.

random_expand

chainercv.transforms.**random_expand**(*img*, *max_ratio=4*, *fill=0*, *return_param=False*)

Expand an image randomly.

This method randomly place the input image on a larger canvas. The size of the canvas is (rH, rW), where (H, W) is the size of the input image and r is a random ratio drawn from $[1, max_ratio]$. The canvas is filled by a value `fill` except for the region where the original image is placed.

This data augmentation trick is used to create "zoom out" effect[2].

Parameters

- **img** (*ndarray*) – An image array to be augmented. This is in CHW format.
- **max_ratio** (*float*) – The maximum ratio of expansion. In the original paper, this value is 4.
- **fill** (*float, tuple or ndarray*) – The value of padded pixels. In the original paper, this value is the mean of ImageNet. If it is `numpy.ndarray`, its shape should be $(C, 1, 1)$, where C is the number of channels of `img`.
- **return_param** (*bool*) – Returns random parameters.

Returns

If `return_param = False`, returns an array `out_img` that is the result of expansion.

If `return_param = True`, returns a tuple whose elements are `out_img, param`. `param` is a dictionary of intermediate parameters whose contents are listed below with key, value-type and the description of the value.

- **ratio** (*float*): The sampled value used to make the canvas.
- **y_offset** (*int*): The y coordinate of the top left corner of the image after placing on the canvas.
- **x_offset** (*int*): The x coordinate of the top left corner of the image after placing on the canvas.

Return type ndarray or (ndarray, dict)

random_flip

chainercv.transforms.**random_flip**(*img*, *y_random=False*, *x_random=False*, *return_param=False*, *copy=False*)

Randomly flip an image in vertical or horizontal direction.

Parameters

- **img** (*ndarray*) – An array that gets flipped. This is in CHW format.
- **y_random** (*bool*) – Randomly flip in vertical direction.
- **x_random** (*bool*) – Randomly flip in horizontal direction.
- **return_param** (*bool*) – Returns information of flip.
- **copy** (*bool*) – If False, a view of `img` will be returned.

[2] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, Alexander C. Berg. SSD: Single Shot MultiBox Detector. ECCV 2016.

Returns

If `return_param = False`, returns an array `out_img` that is the result of flipping.

If `return_param = True`, returns a tuple whose elements are `out_img, param`. `param` is a dictionary of intermediate parameters whose contents are listed below with key, value-type and the description of the value.

- **y_flip** (*bool*): Whether the image was flipped in the vertical direction or not.
- **x_flip** (*bool*): Whether the image was flipped in the horizontal direction or not.

Return type ndarray or (ndarray, dict)

random_rotate

chainercv.transforms.**random_rotate**(*img*, *return_param=False*)

Randomly rotate images by 90, 180, 270 or 360 degrees.

Parameters

- **img** (*ndarray*) – An arrays that get flipped. This is in CHW format.
- **return_param** (*bool*) – Returns information of rotation.

Returns

If `return_param = False`, returns an array `out_img` that is the result of rotation.

If `return_param = True`, returns a tuple whose elements are `out_img, param`. `param` is a dictionary of intermediate parameters whose contents are listed below with key, value-type and the description of the value.

- **k** (*int*): The integer that represents the number of times the image is rotated by 90 degrees.

Return type ndarray or (ndarray, dict)

random_sized_crop

chainercv.transforms.**random_sized_crop**(*img*, *scale_ratio_range=(0.08, 1)*, *aspect_ratio_range=(0.75, 1.3333333333333333)*, *return_param=False*, *copy=False*)

Crop an image to random size and aspect ratio.

The size (H_{crop}, W_{crop}) and the left top coordinate (y_{start}, x_{start}) of the crop are calculated as follows:

- $H_{crop} = \lfloor \sqrt{s \times H \times W \times a} \rfloor$
- $W_{crop} = \lfloor \sqrt{s \times H \times W \div a} \rfloor$
- $y_{start} \sim Uniform\{0, H - H_{crop}\}$
- $x_{start} \sim Uniform\{0, W - W_{crop}\}$
- $s \sim Uniform(s_1, s_2)$
- $b \sim Uniform(a_1, a_2)$ and $a = b$ or $a = \frac{1}{b}$ in 50/50 probability.

Here, s_1, s_2 are the two floats in `scale_ratio_range` and a_1, a_2 are the two floats in `aspect_ratio_range`. Also, H and W are the height and the width of the image. Note that $s \approx \frac{H_{crop} \times W_{crop}}{H \times W}$ and $a \approx \frac{H_{crop}}{W_{crop}}$. The approximations come from flooring floats to integers.

Note: When it fails to sample a valid scale and aspect ratio for ten times, it picks values in a non-uniform way. If this happens, the selected scale ratio can be smaller than `scale_ratio_range[0]`.

Parameters

- **img** (*ndarray*) – An image array. This is in CHW format.
- **scale_ratio_range** (*tuple of two floats*) – Determines the distribution from which a scale ratio is sampled. The default values are selected so that the area of the crop is 8~100% of the original image. This is the default setting used to train ResNets in Torch style.
- **aspect_ratio_range** (*tuple of two floats*) – Determines the distribution from which an aspect ratio is sampled. The default values are $\frac{3}{4}$ and $\frac{4}{3}$, which are also the default setting to train ResNets in Torch style.
- **return_param** (*bool*) – Returns parameters if `True`.

Returns

If `return_param = False`, returns only the cropped image.

If `return_param = True`, returns a tuple of cropped image and `param`. `param` is a dictionary of intermediate parameters whose contents are listed below with key, value-type and the description of the value.

- **y_slice** (*slice*): A slice used to crop the input image. The relation below holds together with `x_slice`.
- **x_slice** (*slice*): Similar to `y_slice`.

```
out_img = img[:, y_slice, x_slice]
```

- **scale_ratio** (float): s in the description (see above).
- **aspect_ratio** (float): a in the description.

Return type ndarray or (ndarray, dict)

resize

chainercv.transforms.**resize**(*img*, *size*, *interpolation=2*)

Resize image to match the given shape.

The backend used by *resize()* is configured by `chainer.global_config.cv_resize_backend`. Two backends are supported: "cv2" and "PIL".

Parameters

- **img** (*ndarray*) – An array to be transformed. This is in CHW format and the type should be `numpy.float32`.
- **size** (*tuple*) – This is a tuple of length 2. Its elements are ordered as (height, width).
- **interpolation** (*int*) – Determines sampling strategy. This is one of `PIL.Image.NEAREST`, `PIL.Image.BILINEAR`, `PIL.Image.BICUBIC`, `PIL.Image.LANCZOS`. Bilinear interpolation is the default strategy.

Returns A resize array in CHW format.

Return type ndarray

resize_contain

chainercv.transforms.**resize_contain**(*img*, *size*, *fill=0*, *interpolation=2*, *return_param=False*)

Resize the image to fit in the given area while keeping aspect ratio.

If both the height and the width in `size` are larger than the height and the width of the `img`, the `img` is placed on the center with an appropriate padding to match `size`.

Otherwise, the input image is scaled to fit in a canvas whose size is `size` while preserving aspect ratio.

Parameters

- **img** (*ndarray*) – An array to be transformed. This is in CHW format.
- **size** (*tuple of two ints*) – A tuple of two elements: `height, width`. The size of the image after resizing.
- **fill** (*float, tuple or ndarray*) – The value of padded pixels. If it is `numpy.ndarray`, its shape should be $(C, 1, 1)$, where C is the number of channels of `img`.
- **interpolation** (*int*) – Determines sampling strategy. This is one of `PIL.Image.NEAREST`, `PIL.Image.BILINEAR`, `PIL.Image.BICUBIC`, `PIL.Image.LANCZOS`. Bilinear interpolation is the default strategy.
- **return_param** (*bool*) – Returns information of resizing and offsetting.

Returns

If `return_param = False`, returns an array `out_img` that is the result of resizing.

If `return_param = True`, returns a tuple whose elements are `out_img, param`. `param` is a dictionary of intermediate parameters whose contents are listed below with key, value-type and the description of the value.

- **y_offset** (*int*): The y coordinate of the top left corner of the image after placing on the canvas.
- **x_offset** (*int*): The x coordinate of the top left corner of the image after placing on the canvas.
- **scaled_size** (*tuple*): The size to which the image is scaled to before placing it on a canvas. This is a tuple of two elements: `height, width`.

Return type ndarray or (ndarray, dict)

rotate

chainercv.transforms.**rotate**(*img*, *angle*, *expand=True*, *fill=0*, *interpolation=2*)

Rotate images by degrees.

The backend used by *rotate()* is configured by `chainer.global_config.cv_rotate_backend`. Two backends are supported: “cv2” and “PIL”.

Parameters

- **img** (*ndarray*) – An arrays that get rotated. This is in CHW format.
- **angle** (*float*) – Counter clock-wise rotation angle (degree).
- **expand** (*bool*) – The output shaped is adapted or not. If `True`, the input image is contained complete in the output.

- **fill** (*float*) – The value used for pixels outside the boundaries.
- **interpolation** (*int*) – Determines sampling strategy. This is one of `PIL.Image.NEAREST`, `PIL.Image.BILINEAR`, `PIL.Image.BICUBIC`. Bilinear interpolation is the default strategy.

Returns returns an array `out_img` that is the result of rotation.

Return type ndarray

scale

`chainercv.transforms.`**scale**(*img*, *size*, *fit_short=True*, *interpolation=2*)

Rescales the input image to the given "size".

When `fit_short == True`, the input image will be resized so that the shorter edge will be scaled to length `size` after resizing. For example, if the height of the image is larger than its width, image will be resized to (size * height / width, size).

Otherwise, the input image will be resized so that the longer edge will be scaled to length `size` after resizing.

Parameters

- **img** (*ndarray*) – An image array to be scaled. This is in CHW format.
- **size** (*int*) – The length of the smaller edge.
- **fit_short** (*bool*) – Determines whether to match the length of the shorter edge or the longer edge to `size`.
- **interpolation** (*int*) – Determines sampling strategy. This is one of `PIL.Image.NEAREST`, `PIL.Image.BILINEAR`, `PIL.Image.BICUBIC`, `PIL.Image.LANCZOS`. Bilinear interpolation is the default strategy.

Returns A scaled image in CHW format.

Return type ndarray

ten_crop

`chainercv.transforms.`**ten_crop**(*img*, *size*)

Crop 10 regions from an array.

This method crops 10 regions. All regions will be in shape `size`. These regions consist of 1 center crop and 4 corner crops and horizontal flips of them.

The crops are ordered in this order.

- center crop
- top-left crop
- bottom-left crop
- top-right crop
- bottom-right crop
- center crop (flipped horizontally)
- top-left crop (flipped horizontally)
- bottom-left crop (flipped horizontally)

- top-right crop (flipped horizontally)
- bottom-right crop (flipped horizontally)

Parameters

- **img** (*ndarray*) – An image array to be cropped. This is in CHW format.
- **size** (*tuple*) – The size of output images after cropping. This value is $(height, width)$.

Returns The cropped arrays. The shape of tensor is $(10, C, H, W)$.

3.8.2 Bounding Box

crop_bbox

chainercv.transforms.**crop_bbox**(*bbox*, *y_slice=None*, *x_slice=None*, *allow_outside_center=True*, *return_param=False*)

Translate bounding boxes to fit within the cropped area of an image.

This method is mainly used together with image cropping. This method translates the coordinates of bounding boxes like *translate_bbox()*. In addition, this function truncates the bounding boxes to fit within the cropped area. If a bounding box does not overlap with the cropped area, this bounding box will be removed.

The bounding boxes are expected to be packed into a two dimensional tensor of shape $(R, 4)$, where R is the number of bounding boxes in the image. The second axis represents attributes of the bounding box. They are $(y_{min}, x_{min}, y_{max}, x_{max})$, where the four attributes are coordinates of the top left and the bottom right vertices.

Parameters

- **bbox** (*ndarray*) – Bounding boxes to be transformed. The shape is $(R, 4)$. R is the number of bounding boxes.
- **y_slice** (*slice*) – The slice of y axis.
- **x_slice** (*slice*) – The slice of x axis.
- **allow_outside_center** (*bool*) – If this argument is `False`, bounding boxes whose centers are outside of the cropped area are removed. The default value is `True`.
- **return_param** (*bool*) – If `True`, this function returns indices of kept bounding boxes.

Returns

If `return_param = False`, returns an array `bbox`.

If `return_param = True`, returns a tuple whose elements are `bbox, param`. `param` is a dictionary of intermediate parameters whose contents are listed below with key, value-type and the description of the value.

- **index** (*numpy.ndarray*): An array holding indices of used bounding boxes.
- **trancated_index** (*numpy.ndarray*): An array holding indices of truncated bounding boxes, with respect to **returned** `bbox`, rather than original `bbox`.

Return type ndarray or (ndarray, dict)

flip_bbox

chainercv.transforms.**flip_bbox**(*bbox*, *size*, *y_flip=False*, *x_flip=False*)

Flip bounding boxes accordingly.

The bounding boxes are expected to be packed into a two dimensional tensor of shape $(R, 4)$, where R is the number of bounding boxes in the image. The second axis represents attributes of the bounding box. They are $(y_{min}, x_{min}, y_{max}, x_{max})$, where the four attributes are coordinates of the top left and the bottom right vertices.

Parameters

- **bbox** (*ndarray*) – An array whose shape is $(R, 4)$. R is the number of bounding boxes.
- **size** (*tuple*) – A tuple of length 2. The height and the width of the image before resized.
- **y_flip** (*bool*) – Flip bounding box according to a vertical flip of an image.
- **x_flip** (*bool*) – Flip bounding box according to a horizontal flip of an image.

Returns Bounding boxes flipped according to the given flips.

Return type ndarray

resize_bbox

chainercv.transforms.**resize_bbox**(*bbox*, *in_size*, *out_size*)

Resize bounding boxes according to image resize.

The bounding boxes are expected to be packed into a two dimensional tensor of shape $(R, 4)$, where R is the number of bounding boxes in the image. The second axis represents attributes of the bounding box. They are $(y_{min}, x_{min}, y_{max}, x_{max})$, where the four attributes are coordinates of the top left and the bottom right vertices.

Parameters

- **bbox** (*ndarray*) – An array whose shape is $(R, 4)$. R is the number of bounding boxes.
- **in_size** (*tuple*) – A tuple of length 2. The height and the width of the image before resized.
- **out_size** (*tuple*) – A tuple of length 2. The height and the width of the image after resized.

Returns Bounding boxes rescaled according to the given image shapes.

Return type ndarray

rotate_bbox

chainercv.transforms.**rotate_bbox**(*bbox*, *angle*, *size*)

Rotate bounding boxes by degrees.

Parameters

- **bbox** (*ndarray*) – An array whose shape is $(R, 4)$. R is the number of bounding boxes.
- **angle** (*float*) – Counter clock-wise rotation angle (degree). image is rotated by 90 degrees.
- **size** (*tuple*) – A tuple of length 2. The height and the width of the image.

Returns Bounding boxes rescaled according to the given k.

Return type ndarray

translate_bbox

chainercv.transforms.**translate_bbox**(*bbox*, *y_offset=0*, *x_offset=0*)
Translate bounding boxes.

This method is mainly used together with image transforms, such as padding and cropping, which translates the left top point of the image from coordinate $(0, 0)$ to coordinate $(y, x) = (y_{offset}, x_{offset})$.

The bounding boxes are expected to be packed into a two dimensional tensor of shape $(R, 4)$, where R is the number of bounding boxes in the image. The second axis represents attributes of the bounding box. They are $(y_{min}, x_{min}, y_{max}, x_{max})$, where the four attributes are coordinates of the top left and the bottom right vertices.

Parameters

- **bbox** (*ndarray*) – Bounding boxes to be transformed. The shape is $(R, 4)$. R is the number of bounding boxes.
- **y_offset** (*int or float*) – The offset along y axis.
- **x_offset** (*int or float*) – The offset along x axis.

Returns Bounding boxes translated according to the given offsets.

Return type ndarray

3.8.3 Point

flip_point

chainercv.transforms.**flip_point**(*point*, *size*, *y_flip=False*, *x_flip=False*)
Modify points according to image flips.

Parameters

- **point** (*ndarray*) – Points in the image. The shape of this array is $(P, 2)$. P is the number of points in the image. The last dimension is composed of y and x coordinates of the points.
- **size** (*tuple*) – A tuple of length 2. The height and the width of the image, which is associated with the points.
- **y_flip** (*bool*) – Modify points according to a vertical flip of an image.
- **x_flip** (*bool*) – Modify keypoipoints according to a horizontal flip of an image.

Returns Points modified according to image flips.

Return type ndarray

resize_point

chainercv.transforms.**resize_point**(*point*, *in_size*, *out_size*)
Adapt point coordinates to the rescaled image space.

Parameters

- **point** (*ndarray*) – Points in the image. The shape of this array is $(P, 2)$. P is the number of points in the image. The last dimension is composed of y and x coordinates of the points.
- **in_size** (*tuple*) – A tuple of length 2. The height and the width of the image before resized.

- **out_size** (*tuple*) – A tuple of length 2. The height and the width of the image after resized.

Returns Points rescaled according to the given image shapes.

Return type ndarray

translate_point

chainercv.transforms.**translate_point**(*point*, *y_offset=0*, *x_offset=0*)

Translate points.

This method is mainly used together with image transforms, such as padding and cropping, which translates the top left point of the image to the coordinate $(y, x) = (y_{offset}, x_{offset})$.

Parameters

- **point** (*ndarray*) – Points in the image. The shape of this array is $(P, 2)$. P is the number of points in the image. The last dimension is composed of y and x coordinates of the points.
- **y_offset** (*int or float*) – The offset along y axis.
- **x_offset** (*int or float*) – The offset along x axis.

Returns Points modified translation of an image.

Return type ndarray

3.9 Visualizations

3.9.1 vis_bbox

chainercv.visualizations.**vis_bbox**(*img*, *bbox*, *label=None*, *score=None*, *label_names=None*, *instance_colors=None*, *alpha=1.0*, *linewidth=3.0*, *ax=None*)

Visualize bounding boxes inside image.

Example

```
>>> from chainercv.datasets import VOCBboxDataset
>>> from chainercv.datasets import voc_bbox_label_names
>>> from chainercv.visualizations import vis_bbox
>>> import matplotlib.pyplot as plt
>>> dataset = VOCBboxDataset()
>>> img, bbox, label = dataset[60]
>>> vis_bbox(img, bbox, label,
...          label_names=voc_bbox_label_names)
>>> plt.show()
```

This example visualizes by displaying the same colors for bounding boxes assigned to the same labels.

```
>>> from chainercv.datasets import VOCBboxDataset
>>> from chainercv.datasets import voc_bbox_label_names
>>> from chainercv.visualizations import vis_bbox
>>> from chainercv.visualizations.colormap import voc_colormap
```

(continues on next page)

(continued from previous page)

```
>>> import matplotlib.pyplot as plt
>>> dataset = VOCBboxDataset()
>>> img, bbox, label = dataset[61]
>>> colors = voc_colormap(label + 1)
>>> vis_bbox(img, bbox, label,
...          label_names=voc_bbox_label_names,
...          instance_colors=colors)
>>> plt.show()
```

Parameters

- **img** (*ndarray*) – An array of shape $(3, height, width)$. This is in RGB format and the range of its value is $[0, 255]$. If this is `None`, no image is displayed.
- **bbox** (*ndarray*) – An array of shape $(R, 4)$, where R is the number of bounding boxes in the image. Each element is organized by $(y_{min}, x_{min}, y_{max}, x_{max})$ in the second axis.
- **label** (*ndarray*) – An integer array of shape $(R,)$. The values correspond to id for label names stored in `label_names`. This is optional.
- **score** (*ndarray*) – A float array of shape $(R,)$. Each value indicates how confident the prediction is. This is optional.
- **label_names** (*iterable of strings*) – Name of labels ordered according to label ids. If this is `None`, labels will be skipped.
- **instance_colors** (*iterable of tuples*) – List of colors. Each color is RGB format and the range of its values is $[0, 255]$. The `i`-th element is the color used to visualize the `i`-th instance. If `instance_colors` is `None`, the red is used for all boxes.
- **alpha** (*float*) – The value which determines transparency of the bounding boxes. The range of this value is $[0, 1]$.
- **linewidth** (*float*) – The thickness of the edges of the bounding boxes.
- **ax** (*matplotlib.axes.Axis*) – The visualization is displayed on this axis. If this is `None` (default), a new axis is created.

Returns Returns the Axes object with the plot for further tweaking.

Return type Axes

3.9.2 vis_image

chainercv.visualizations.**vis_image**(*img*, *ax=None*)

Visualize a color image.

Parameters

- **img** (*ndarray*) – An array of shape $(3, height, width)$. This is in RGB format and the range of its value is $[0, 255]$. If this is `None`, no image is displayed.
- **ax** (*matplotlib.axes.Axis*) – The visualization is displayed on this axis. If this is `None` (default), a new axis is created.

Returns Returns the Axes object with the plot for further tweaking.

Return type Axes

3.9.3 vis_instance_segmentation

chainercv.visualizations.**vis_instance_segmentation**(*img*, *mask*, *label=None*, *score=None*, *label_names=None*, *instance_colors=None*, *alpha=0.7*, *ax=None*)

Visualize instance segmentation.

Example

This example visualizes an image and an instance segmentation.

```
>>> from chainercv.datasets import SBDInstanceSegmentationDataset
>>> from chainercv.datasets          ...      import sbd_instance_segmentation_
↪label_names
>>> from chainercv.visualizations import vis_instance_segmentation
>>> import matplotlib.pyplot as plt
>>> dataset = SBDInstanceSegmentationDataset()
>>> img, mask, label = dataset[0]
>>> vis_instance_segmentation(
...     img, mask, label,
...     label_names=sbd_instance_segmentation_label_names)
>>> plt.show()
```

This example visualizes an image, an instance segmentation and bounding boxes.

```
>>> from chainercv.datasets import SBDInstanceSegmentationDataset
>>> from chainercv.datasets          ...      import sbd_instance_segmentation_
↪label_names
>>> from chainercv.visualizations import vis_bbox
>>> from chainercv.visualizations import vis_instance_segmentation
>>> from chainercv.visualizations.colormap import voc_colormap
>>> from chainercv.utils import mask_to_bbox
>>> import matplotlib.pyplot as plt
>>> dataset = SBDInstanceSegmentationDataset()
>>> img, mask, label = dataset[0]
>>> bbox = mask_to_bbox(mask)
>>> colors = voc_colormap(list(range(1, len(mask) + 1)))
>>> ax = vis_bbox(img, bbox, label,
...     label_names=sbd_instance_segmentation_label_names,
...     instance_colors=colors, alpha=0.7, linewidth=0.5)
>>> vis_instance_segmentation(
...     None, mask, instance_colors=colors, alpha=0.7, ax=ax)
>>> plt.show()
```

Parameters

- **img** (*ndarray*) – An array of shape $(3, H, W)$. This is in RGB format and the range of its value is $[0, 255]$. If this is `None`, no image is displayed.
- **mask** (*ndarray*) – A bool array of shape :math‘(R, H, W)‘. If there is an object, the value of the pixel is `True`, and otherwise, it is `False`.
- **label** (*ndarray*) – An integer array of shape $(R,)$. The values correspond to id for label names stored in `label_names`.
- **score** (*ndarray*) – A float array of shape $(R,)$. Each value indicates how confident the prediction is. This is optional.

- **label_names** (*iterable of strings*) – Name of labels ordered according to label ids.
- **instance_colors** (*iterable of tuple*) – List of colors. Each color is RGB format and the range of its values is $[0, 255]$. The `i`-th element is the color used to visualize the `i`-th instance. If `instance_colors` is `None`, the default color map is used.
- **alpha** (*float*) – The value which determines transparency of the figure. The range of this value is $[0, 1]$. If this value is `0`, the figure will be completely transparent. The default value is `0.7`. This option is useful for overlaying the label on the source image.
- **ax** (*matplotlib.axes.Axis*) – The visualization is displayed on this axis. If this is `None` (default), a new axis is created.

Returns Returns `ax`. `ax` is an `matploblib.axes.Axes` with the plot.

Return type matploblib.axes.Axes

3.9.4 vis_point

chainercv.visualizations.**vis_point**(*img*, *point*, *mask=None*, *ax=None*)

Visualize points in an image.

Example

```
>>> import chainercv
>>> import matplotlib.pyplot as plt
>>> dataset = chainercv.datasets.CUBPointDataset()
>>> img, point, mask = dataset[0]
>>> chainercv.visualizations.vis_point(img, point, mask)
>>> plt.show()
```

Parameters

- **img** (*ndarray*) – An image of shape $(3, height, width)$. This is in RGB format and the range of its value is $[0, 255]$. This should be visualizable using `matplotlib.pyplot.imshow(img)`
- **point** (*ndarray*) – An array of point coordinates whose shape is $(P, 2)$, where P is the number of points. The second axis corresponds to y and x coordinates of the points.
- **mask** (*ndarray*) – A boolean array whose shape is $(P,)$. If i th element is `True`, the i th point is not displayed. If not specified, all points in `point` will be displayed.
- **ax** (*matplotlib.axes.Axes*) – If provided, plot on this axis.

Returns Returns the Axes object with the plot for further tweaking.

Return type Axes

3.9.5 vis_semantic_segmentation

chainercv.visualizations.**vis_semantic_segmentation**(*img*, *label*, *label_names=None*, *label_colors=None*, *ignore_label_color=(0, 0, 0)*, *alpha=1*, *all_label_names_in_legend=False*, *ax=None*)

Visualize a semantic segmentation.

Example

```
>>> from chainercv.datasets import VOCSemanticSegmentationDataset
>>> from chainercv.datasets          ...      import voc_semantic_segmentation_
↪label_colors
>>> from chainercv.datasets          ...      import voc_semantic_segmentation_
↪label_names
>>> from chainercv.visualizations import vis_semantic_segmentation
>>> import matplotlib.pyplot as plt
>>> dataset = VOCSemanticSegmentationDataset()
>>> img, label = dataset[60]
>>> ax, legend_handles = vis_semantic_segmentation(
...     img, label,
...     label_names=voc_semantic_segmentation_label_names,
...     label_colors=voc_semantic_segmentation_label_colors,
...     alpha=0.9)
>>> ax.legend(handles=legend_handles, bbox_to_anchor=(1, 1), loc=2)
>>> plt.show()
```

Parameters

- **img** (*ndarray*) – An array of shape $(3, height, width)$. This is in RGB format and the range of its value is $[0, 255]$. If this is `None`, no image is displayed.
- **label** (*ndarray*) – An integer array of shape $(height, width)$. The values correspond to id for label names stored in `label_names`.
- **label_names** (*iterable of strings*) – Name of labels ordered according to label ids.
- **label_colors** – (iterable of tuple): An iterable of colors for regular labels. Each color is RGB format and the range of its values is $[0, 255]$. If `colors` is `None`, the default color map is used.
- **ignore_label_color** (*tuple*) – Color for ignored label. This is RGB format and the range of its values is $[0, 255]$. The default value is `(0, 0, 0)`.
- **alpha** (*float*) – The value which determines transparency of the figure. The range of this value is $[0, 1]$. If this value is `0`, the figure will be completely transparent. The default value is `1`. This option is useful for overlaying the label on the source image.
- **all_label_names_in_legend** (*bool*) – Determines whether to include all label names in a legend. If this is `False`, the legend does not contain the names of unused labels. An unused label is defined as a label that does not appear in `label`. The default value is `False`.
- **ax** (*matplotlib.axes.Axis*) – The visualization is displayed on this axis. If this is `None` (default), a new axis is created.

Returns Returns `ax` and `legend_handles`. `ax` is an `matploblib.axes.Axes` with the plot. It can be used for further tweaking. `legend_handles` is a list of legends. It can be passed `matploblib.pyplot.legend()` to show a legend.

Return type matploblib.axes.Axes and list of matplotlib.patches.Patch

3.10 Utils

3.10.1 Bounding Box Utilities

bbox_iou

`chainercv.utils.`**`bbox_iou`**(*bbox_a*, *bbox_b*)

Calculate the Intersection of Unions (IoUs) between bounding boxes.

IoU is calculated as a ratio of area of the intersection and area of the union.

This function accepts both `numpy.ndarray` and `cupy.ndarray` as inputs. Please note that both `bbox_a` and `bbox_b` need to be same type. The output is same type as the type of the inputs.

Parameters

- **`bbox_a`** (*array*) – An array whose shape is $(N, 4)$. N is the number of bounding boxes. The dtype should be `numpy.float32`.
- **`bbox_b`** (*array*) – An array similar to `bbox_a`, whose shape is $(K, 4)$. The dtype should be `numpy.float32`.

Returns An array whose shape is (N, K). An element at index (n, k) contains IoUs between n th bounding box in `bbox_a` and k th bounding box in `bbox_b`.

Return type array

non_maximum_suppression

`chainercv.utils.`**`non_maximum_suppression`**(*bbox*, *thresh*, *score=None*, *limit=None*)

Suppress bounding boxes according to their IoUs.

This method checks each bounding box sequentially and selects the bounding box if the Intersection over Unions (IoUs) between the bounding box and the previously selected bounding boxes is less than `thresh`. This method is mainly used as postprocessing of object detection. The bounding boxes are selected from ones with higher scores. If `score` is not provided as an argument, the bounding box is ordered by its index in ascending order.

The bounding boxes are expected to be packed into a two dimensional tensor of shape $(R, 4)$, where R is the number of bounding boxes in the image. The second axis represents attributes of the bounding box. They are $(y_{min}, x_{min}, y_{max}, x_{max})$, where the four attributes are coordinates of the top left and the bottom right vertices.

`score` is a float array of shape $(R,)$. Each score indicates confidence of prediction.

This function accepts both `numpy.ndarray` and `cupy.ndarray` as an input. Please note that both `bbox` and `score` need to be the same type. The type of the output is the same as the input.

Parameters

- **`bbox`** (*array*) – Bounding boxes to be transformed. The shape is $(R, 4)$. R is the number of bounding boxes.
- **`thresh`** (*float*) – Threshold of IoUs.

- **score** (*array*) – An array of confidences whose shape is $(R,)$.
- **limit** (*int*) – The upper bound of the number of the output bounding boxes. If it is not specified, this method selects as many bounding boxes as possible.

Returns An array with indices of bounding boxes that are selected. They are sorted by the scores of bounding boxes in descending order. The shape of this array is $(K,)$ and its dtype is `numpy.int32`. Note that $K \leq R$.

Return type array

3.10.2 Download Utilities

cached_download

`chainercv.utils.`**`cached_download`**(*url*)

Downloads a file and caches it.

This is different from the original `cached_download()` in that the download progress is reported.

It downloads a file from the URL if there is no corresponding cache. After the download, this function stores a cache to the directory under the dataset root (see `set_dataset_root()`). If there is already a cache for the given URL, it just returns the path to the cache without downloading the same file.

Parameters **url** (*string*) – URL to download from.

Returns Path to the downloaded file.

Return type string

download_model

`chainercv.utils.`**`download_model`**(*url*)

Downloads a model file and puts it under model directory.

It downloads a file from the URL and puts it under model directory. For exmaple, if `url` is *http://example.com/subdir/model.npz*, the pretrained weights file will be saved to *$CHAINER_DATASET_ROOT/pfnet/chainercv/models/model.npz*. If there is already a file at the destination path, it just returns the path without downloading the same file.

Parameters **url** (*string*) – URL to download from.

Returns Path to the downloaded file.

Return type string

extractall

`chainercv.utils.`**`extractall`**(*file_path*, *destination*, *ext*)

Extracts an archive file.

This function extracts an archive file to a destination.

Parameters

- **file_path** (*string*) – The path of a file to be extracted.
- **destination** (*string*) – A directory path. The archive file will be extracted under this directory.

- **ext** (*string*) – An extension suffix of the archive file. This function supports `'.zip'`, `'.tar'`, `'.gz'` and `'.tgz'`.

3.10.3 Image Utilities

read_image

chainercv.utils.**read_image**(*path*, *dtype=<type 'numpy.float32'>*, *color=True*, *alpha=None*)

Read an image from a file.

This function reads an image from given file. The image is CHW format and the range of its value is $[0, 255]$. If `color = True`, the order of the channels is RGB.

The backend used by *read_image()* is configured by `chainer.global_config.cv_read_image_backend`. Two backends are supported: "cv2" and "PIL".

Parameters

- **path** (*string*) – A path of image file.
- **dtype** – The type of array. The default value is `float32`.
- **color** (*bool*) – This option determines the number of channels. If `True`, the number of channels is three. In this case, the order of the channels is RGB. This is the default behaviour. If `False`, this function returns a grayscale image.
- **alpha** (*None or {'ignore', 'blend_with_white', 'blend_with_black'}*) – Choose how RGBA images are handled. By default, an error is raised. Here are the other possible behaviors:
  - *'ignore'*: Ignore alpha channel.
  - *'blend_with_white'*: Blend RGB image multiplied by alpha on a white image.
  - *'blend_with_black'*: Blend RGB image multiplied by alpha on a black image.

Returns An image.

Return type ndarray

read_label

chainercv.utils.**read_label**(*path*, *dtype=<type 'numpy.int32'>*)

Read a label image from a file.

This function reads an label image from given file. If reading label doesn't work collectly, try *read_image()* with a parameter `color=True`.

Parameters

- **path** (*string*) – A path of image file.
- **dtype** – The type of array. The default value is `int32`.
- **color** (*bool*) – This option determines the number of channels. If `True`, the number of channels is three. In this case, the order of the channels is RGB. This is the default behaviour. If `False`, this function returns a grayscale image.

Returns An image.

Return type ndarray

tile_images

chainercv.utils.**tile_images**(*imgs*, *n_col*, *pad=2*, *fill=0*)
Make a tile of images

Parameters

- **imgs** (*numpy.ndarray*) – A batch of images whose shape is BCHW.
- **n_col** (*int*) – The number of columns in a tile.
- **pad** (*int or tuple of two ints*) – pad_y, pad_x. This is the amounts of padding in y and x directions. If this is an integer, the amounts of padding in the two directions are the same. The default value is 2.
- **fill** (*float, tuple or ndarray*) – The value of padded pixels. If it is numpy.ndarray, its shape should be $(C, 1, 1)$, where C is the number of channels of img.

Returns An image array in CHW format. The size of this image is $((H + pad_y) \times \lceil B/n_{n_{col}} \rceil, (W + pad_x) \times n_{col})$.

Return type ndarray

write_image

chainercv.utils.**write_image**(*img*, *path*)
Save an image to a file.

This function saves an image to given file. The image is in CHW format and the range of its value is $[0, 255]$.

Parameters

- **image** (*ndarray*) – An image to be saved.
- **path** (*string*) – The path of an image file.

3.10.4 Iterator Utilities

apply_to_iterator

chainercv.utils.**apply_to_iterator**(*func*, *iterator*, *n_input=1*, *hook=None*)
Apply a function/method to batches from an iterator.

This function applies a function/method to an iterator of batches.

It assumes that the iterator iterates over a collection of tuples that contain inputs to func(). Additionally, the tuples may contain values that are not used by func(). For convenience, we allow the iterator to iterate over a collection of inputs that are not tuple. Here is an illustration of the expected behavior of the iterator. This behaviour is the same as chainer.Iterator.

```
>>> batch = next(iterator)
>>> # batch: [in_val]
or
>>> # batch: [(in_val0, ..., in_val{n_input - 1})]
or
>>> # batch: [(in_val0, ..., in_val{n_input - 1}, rest_val0, ...)]
```

func() should take batch(es) of data and return batch(es) of computed values. Here is an illustration of the expected behavior of the function.

```
>>> out_vals = func([in_val0], ..., [in_val{n_input - 1}])
>>> # out_vals: [out_val]
or
>>> out_vals0, out_vals1, ... = func([in_val0], ..., [in_val{n_input - 1}])
>>> # out_vals0: [out_val0]
>>> # out_vals1: [out_val1]
```

With *apply_to_iterator()*, users can get iterator(s) of values returned by `func()`. It also returns iterator(s) of input values and values that are not used for computation.

```
>>> in_values, out_values, rest_values = apply_to_iterator(
>>>     func, iterator, n_input)
>>> # in_values: (iter of in_val0, ..., iter of in_val{n_input - 1})
>>> # out_values: (iter of out_val0, ...)
>>> # rest_values: (iter of rest_val0, ...)
```

Here is an exmple, which applies a pretrained Faster R-CNN to PASCAL VOC dataset.

```
>>> from chainer import iterators
>>>
>>> from chainercv.datasets import VOCBBoxDataset
>>> from chainercv.links import FasterRCNNVGG16
>>> from chainercv.utils import apply_to_iterator
>>>
>>> dataset = VOCBBoxDataset(year='2007', split='test')
>>> # next(iterator) -> [(img, gt_bbox, gt_label)]
>>> iterator = iterators.SerialIterator(
...     dataset, 2, repeat=False, shuffle=False)
>>>
>>> # model.predict([img]) -> ([pred_bbox], [pred_label], [pred_score])
>>> model = FasterRCNNVGG16(pretrained_model='voc07')
>>>
>>> in_values, out_values, rest_values = apply_to_iterator(
...     model.predict, iterator)
>>>
>>> # in_values contains one iterator
>>> imgs, = in_values
>>> # out_values contains three iterators
>>> pred_bboxes, pred_labels, pred_scores = out_values
>>> # rest_values contains two iterators
>>> gt_bboxes, gt_labels = rest_values
```

Parameters

- **func** – A callable that takes batch(es) of input data and returns computed data.
- **iterator** (*iterator*) – An iterator of batches. The first `n_input` elements in each sample are treated as input values. They are passed to `func`.
- **n_input** (*int*) – The number of input data. The default value is `1`.
- **hook** – A callable that is called after each iteration. `in_values`, `out_values`, and `rest_values` are passed as arguments. Note that these values do not contain data from the previous iterations.

Returns

This function returns three tuples of iterators: `in_values`, `out_values` and `rest_values`.

- `in_values`: A tuple of iterators. Each iterator returns a corresponding input value. For example, if `func()` takes `[in_val0], [in_val1]`, `next(in_values[0])` and `next(in_values[1])` will be `in_val0` and `in_val1`.
- `out_values`: A tuple of iterators. Each iterator returns a corresponding computed value. For example, if `func()` returns `([out_val0], [out_val1])`, `next(out_values[0])` and `next(out_values[1])` will be `out_val0` and `out_val1`.
- `rest_values`: A tuple of iterators. Each iterator returns a corresponding rest value. For example, if the `iterator` returns `[(in_val0, in_val1, rest_val0, rest_val1)]`, `next(rest_values[0])` and `next(rest_values[1])` will be `rest_val0` and `rest_val1`. If the input iterator does not give any rest values, this tuple will be empty.

Return type Three tuples of iterators

ProgressHook

class chainercv.utils.**ProgressHook**(*n_total=None*)

A hook class reporting the progress of iteration.

This is a hook class designed for `apply_prediction_to_iterator()`.

Parameters **n_total** (*int*) – The number of images. This argument is optional.

unzip

chainercv.utils.**unzip**(*iterable*)

Converts an iterable of tuples into a tuple of iterators.

This function converts an iterable of tuples into a tuple of iterators. This is an inverse function of `six.moves.zip()`.

```
>>> from chainercv.utils import unzip
>>> data = [(0, 'a'), (1, 'b'), (2, 'c'), (3, 'd'), (4, 'e')]
>>> int_iter, str_iter = unzip(data)
>>>
>>> next(int_iter)  # 0
>>> next(int_iter)  # 1
>>> next(int_iter)  # 2
>>>
>>> next(str_iter)  # 'a'
>>> next(str_iter)  # 'b'
>>> next(str_iter)  # 'c'
```

Parameters **iterable** (*iterable*) – An iterable of tuples. All tuples should have the same length.

Returns Each iterator corresponds to each element of input tuple. Note that each iterator stores values until they are popped. To reduce memory usage, it is recommended to delete unused iterators.

Return type tuple of iterators

3.10.5 Link Utilities

prepare_pretrained_model

chainercv.utils.**prepare_pretrained_model**(*param*, *pretrained_model*, *models*, *default={}*)

Select parameters based on the existence of pretrained model.

Parameters

- **param** (*dict*) – Map from the name of the parameter to values.
- **pretrained_model** (*string*) – Name of the pretrained weight, path to the pretrained weight or None.
- **models** (*dict*) – Map from the name of the pretrained weight to model, which is a dictionary containing the configuration used by the selected weight.

  model has four keys: param, overwritable, url and cv2.

  - **param** (*dict*): Parameters assigned to the pretrained weight.
  - **overwritable** (*set*): Names of parameters that are overwritable (i.e., param[key] != model['param'][key] is accepted).
  - **url** (*string*): Location of the pretrained weight.
  - **cv2** (*bool*): If True, a warning is raised if cv2 is not installed.

3.10.6 Mask Utilities

mask_iou

chainercv.utils.**mask_iou**(*mask_a*, *mask_b*)

Calculate the Intersection of Unions (IoUs) between masks.

IoU is calculated as a ratio of area of the intersection and area of the union.

This function accepts both numpy.ndarray and cupy.ndarray as inputs. Please note that both mask_a and mask_b need to be same type. The output is same type as the type of the inputs.

Parameters

- **mask_a** (*array*) – An array whose shape is (N, H, W). N is the number of masks. The dtype should be numpy.bool.
- **mask_b** (*array*) – An array similar to mask_a, whose shape is (K, H, W). The dtype should be numpy.bool.

Returns An array whose shape is (N, K). An element at index (n, k) contains IoUs between n th mask in mask_a and k th mask in mask_b.

Return type array

mask_to_bbox

chainercv.utils.**mask_to_bbox**(*mask*)

Compute the bounding boxes around the masked regions.

This function accepts both numpy.ndarray and cupy.ndarray as inputs.

Parameters **mask** (*array*) – An array whose shape is (R, H, W). R is the number of masks. The dtype should be `numpy.bool`.

Returns The bounding boxes around the masked regions. This is an array whose shape is $(R, 4)$. R is the number of bounding boxes. The dtype should be `numpy.float32`.

Return type array

3.10.7 Testing Utilities

assert_is_bbox

chainercv.utils.**assert_is_bbox**(*bbox*, *size=None*)

Checks if bounding boxes satisfy bounding box format.

This function checks if given bounding boxes satisfy bounding boxes format or not. If the bounding boxes do not satifiy the format, this function raises an `AssertionError`.

Parameters

- **bbox** (*ndarray*) – Bounding boxes to be checked.
- **size** (*tuple of ints*) – The size of an image. If this argument is specified, Each bounding box should be within the image.

assert_is_bbox_dataset

chainercv.utils.**assert_is_bbox_dataset**(*dataset*, *n_fg_class*, *n_example=None*)

Checks if a dataset satisfies the bounding box dataset API.

This function checks if a given dataset satisfies the bounding box dataset API or not. If the dataset does not satifiy the API, this function raises an `AssertionError`.

Parameters

- **dataset** – A dataset to be checked.
- **n_fg_class** (*int*) – The number of foreground classes.
- **n_example** (*int*) – The number of examples to be checked. If this argument is specified, this function picks examples ramdomly and checks them. Otherwise, this function checks all examples.

assert_is_detection_link

chainercv.utils.**assert_is_detection_link**(*link*, *n_fg_class*)

Checks if a link satisfies detection link APIs.

This function checks if a given link satisfies detection link APIs or not. If the link does not satifiy the APIs, this function raises an `AssertionError`.

Parameters

- **link** – A link to be checked.
- **n_fg_class** (*int*) – The number of foreground classes.

assert_is_image

chainercv.utils.**assert_is_image**(*img*, *color=True*, *check_range=True*)

Checks if an image satisfies image format.

This function checks if a given image satisfies image format or not. If the image does not satifiy the format, this function raises an AssertionError.

Parameters

- **img** (*ndarray*) – An image to be checked.
- **color** (*bool*) – A boolean that determines the expected channel size. If it is True, the number of channels should be 3. Otherwise, it should be 1. The default value is True.
- **check_range** (*bool*) – A boolean that determines whether the range of values are checked or not. If it is True, The values of image must be in $[0, 255]$. Otherwise, this function does not check the range. The default value is True.

assert_is_instance_segmentation_dataset

chainercv.utils.**assert_is_instance_segmentation_dataset**(*dataset*, *n_fg_class*, *n_example=None*)

Checks if a dataset satisfies instance segmentation dataset APIs.

This function checks if a given dataset satisfies instance segmentation dataset APIs or not. If the dataset does not satifiy the APIs, this function raises an AssertionError.

Parameters

- **dataset** – A dataset to be checked.
- **n_fg_class** (*int*) – The number of foreground classes.
- **n_example** (*int*) – The number of examples to be checked. If this argument is specified, this function picks examples ramdomly and checks them. Otherwise, this function checks all examples.

assert_is_label_dataset

chainercv.utils.**assert_is_label_dataset**(*dataset*, *n_class*, *n_example=None*, *color=True*)

Checks if a dataset satisfies the label dataset API.

This function checks if a given dataset satisfies the label dataset API or not. If the dataset does not satifiy the API, this function raises an AssertionError.

Parameters

- **dataset** – A dataset to be checked.
- **n_class** (*int*) – The number of classes.
- **n_example** (*int*) – The number of examples to be checked. If this argument is specified, this function picks examples ramdomly and checks them. Otherwise, this function checks all examples.
- **color** (*bool*) – A boolean that determines the expected channel size. If it is True, the number of channels should be 3. Otherwise, it should be 1. The default value is True.

assert_is_point

chainercv.utils.**assert_is_point**(*point*, *mask=None*, *size=None*)

Checks if points satisfy the format.

This function checks if given points satisfy the format and raises an `AssertionError` when the points violate the convention.

Parameters

- **point** (*ndarray*) – Points to be checked.
- **mask** (*ndarray*) – A mask of the points. If this is `None`, all points are regarded as valid.
- **size** (*tuple of ints*) – The size of an image. If this argument is specified, the coordinates of valid points are checked to be within the image.

assert_is_point_dataset

chainercv.utils.**assert_is_point_dataset**(*dataset*, *n_point=None*, *n_example=None*, *no_mask=False*)

Checks if a dataset satisfies the point dataset API.

This function checks if a given dataset satisfies the point dataset API or not. If the dataset does not satifiy the API, this function raises an `AssertionError`.

Parameters

- **dataset** – A dataset to be checked.
- **n_point** (*int*) – The number of expected points per image. If this is `None`, the number of points per image can be arbitrary.
- **n_example** (*int*) – The number of examples to be checked. If this argument is specified, this function picks examples ramdomly and checks them. Otherwise, this function checks all examples.
- **no_mask** (*bool*) – If `True`, we assume that point mask is always not contained. If `False`, point mask may or may not be contained.

assert_is_semantic_segmentation_dataset

chainercv.utils.**assert_is_semantic_segmentation_dataset**(*dataset*, *n_class*, *n_example=None*)

Checks if a dataset satisfies semantic segmentation dataset APIs.

This function checks if a given dataset satisfies semantic segmentation dataset APIs or not. If the dataset does not satifiy the APIs, this function raises an `AssertionError`.

Parameters

- **dataset** – A dataset to be checked.
- **n_class** (*int*) – The number of classes including background.
- **n_example** (*int*) – The number of examples to be checked. If this argument is specified, this function picks examples ramdomly and checks them. Otherwise, this function checks all examples.

assert_is_semantic_segmentation_link

chainercv.utils.**assert_is_semantic_segmentation_link**(*link*, *n_class*)
Checks if a link satisfies semantic segmentation link APIs.

This function checks if a given link satisfies semantic segmentation link APIs or not. If the link does not satifiy the APIs, this function raises an `AssertionError`.

Parameters

- **link** – A link to be checked.
- **n_class** (*int*) – The number of classes including background.

ConstantStubLink

class chainercv.utils.**ConstantStubLink**(*outputs*)
A chainer.Link that returns constant value(s).

This is a `chainer.Link` that returns constant `chainer.Variable` (s) when `__call__()` method is called.

Parameters **outputs** (*ndarray or tuple or ndarray*) – The value(s) of variable(s) returned by `__call__()`. If an array is specified, `__call__()` returns a `chainer.Variable`. Otherwise, it returns a tuple of `chainer.Variable`.

to_cpu()
Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu()
Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

generate_random_bbox

chainercv.utils.**generate_random_bbox**(*n*, *img_size*, *min_length*, *max_length*)
Generate valid bounding boxes with random position and shape.

Parameters

- **n** (*int*) – The number of bounding boxes.
- **img_size** (*tuple*) – A tuple of length 2. The height and the width of the image on which bounding boxes locate.
- **min_length** (*float*) – The minimum length of edges of bounding boxes.
- **max_length** (*float*) – The maximum length of edges of bounding boxes.

Returns Coordinates of bounding boxes. Its shape is $(R, 4)$. Here, R equals n. The second axis contains $y_{min}, x_{min}, y_{max}, x_{max}$, where $min_length \leq y_{max} - y_{min} < max_length$. and $min_length \leq x_{max} - x_{min} < max_length$

Return type numpy.ndarray

CHAPTER 4

Naming Conventions

Here are the notations used.

- B is the size of a batch.
- H is the height of an image.
- W is the width of an image.
- C is the number of channels.
- R is the total number of instances in an image.
- L is the number of classes.

4.1 Data objects

4.1.1 Images

- `imgs`: (B, C, H, W) or $[(C, H, W)]$
- `img`: (C, H, W)

Note: `image` is used for a name of a function or a class (e.g., *`chainercv.utils.write_image()`*).

4.1.2 Bounding boxes

- `bboxes`: $(B, R, 4)$ or $[(R, 4)]$
- `bbox`: $(R, 4)$
- `bb`: $(4,)$

4.1.3 Labels

name	classification	detection and instance segmentation	semantic segmentation	
`labels`	$(B,)$	(B, R) or $[(R,)]$	(B, H, W)	
`label`	$()$	$(R,)$	(H, W)	
`l`	r `lb`	–	$()$	–

4.1.4 Scores and probabilities

score represents an unbounded confidence value. On the other hand, probability is bounded in `[0, 1]` and sums to 1.

name	classification	detection and instance segmentation	semantic segmentation
`scores` or `probs`	(B, L)	(B, R, L) or $[(R, L)]$	(B, L, H, W)
`score` or `prob`	$(L,)$	(R, L)	(L, H, W)
`sc` or `pb`	–	$(L,)$	–

Note: Even for objects that satisfy the definition of probability, they can be named as `score`.

4.1.5 Instance segmentations

- `masks`: (B, R, H, W) or $[(R, H, W)]$
- `mask`: (R, H, W)
- `msk`: (H, W)

4.2 Attributing an additonal meaning to a basic data object

4.2.1 RoIs

- `rois`: $(R', 4)$, which consists of bounding boxes for multiple images. Assuming that there are B images each containing R_i bounding boxes, the formula $R' = \sum R_i$ is true.
- `roi_indices`: An array of shape $(R',)$ that contains batch indices of images to which bounding boxes correspond.
- `roi`: $(R, 4)$. This is RoIs for single image.

4.2.2 Attributes associated to RoIs

RoIs may have additional attributes, such as class scores and masks. These attributes are named by appending `roi_` (e.g., `scores`-like object is named as `roi_scores`).

- `roi_xs`: $(R',) + x_{shape}$
- `roi_x`: $(R,) + x_{shape}$

In the case of `scores` with shape $(L,)$, `roi_xs` would have shape (R', L).

Note: `roi_nouns = roi_noun = noun` when `batchsize=1`. Changing names interchangeably is fine.

4.2.3 Class-wise vs class-independent

`cls_nouns` is a multi-class version of `nouns`. For instance, `cls_locs` is $(B, R, L, 4)$ and `locs` is $(B, R, 4)$.

Note: `cls_probs` and `probs` can be used interchangeably in the case when there is no confusion.

4.2.4 Arbitrary input

`x` is a variable whose shape can be inferred from the context. It can be used only when there is no confusion on its shape. This is usually the case when naming an input to a neural network.

CHAPTER 5

License

5.1 Source Code

The source code of ChainerCV is licensed under MIT-License.

5.2 Pretrained Models

Pretrained models provided by ChainerCV are benefited from the following resources. See the following resources for the terms of use of a model with weights pretrained by any of such resources.

model	resource
ResNet50/101/152 (`imagenet`)	• ResNet50/101/152 (trained on ImageNet)
SEResNet50/101/152 (`imagenet`)	• SEResNet50/101/152 (trained on ImageNet)
SEResNeXt50/101 (`imagenet`)	• SEResNeXt50/101 (trained on ImageNet)
VGG16 (`imagenet`)	• VGG-16 (trained on ImageNet)
FasterRCNNVGG16 (`imagenet`)	• VGG-16 (trained on ImageNet)
FasterRCNNVGG16 (`voc07`/`voc0712`)	• VGG-16 (trained on ImageNet) • PASCAL VOC
SSD300/SSD512 (`imagenet`)	• VGG-16 (trained on ImageNet, FC reduced)
SSD300/SSD512 (`voc0712`)	• SSD300/SSD512 (trained on PASCAL VOC 2007 and 2012)
YOLOv2 (`voc0712`)	• Darknet19 (trained on ImageNet) • PASCAL VOC
YOLOv3 (`voc0712`)	• Darknet53 (trained on ImageNet) • PASCAL VOC
PSPNetResNet101 (`cityscapes`)	• PSPNet101 (trained on Cityscapes)
SegNetBasic (`camvid`)	• CamVid
FCISResNet101 (`sbd`)	• ResNet101 (trained on ImageNet) • SBD

CHAPTER 6

Indices and tables

- genindex
- modindex
- search

Bibliography

[Ren15] Shaoqing Ren, Kaiming He, Ross Girshick, Jian Sun. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. NIPS 2015.

[Liu16] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, Alexander C. Berg. SSD: Single Shot MultiBox Detector. ECCV 2016.

Python Module Index

C

chainercv, 19
chainercv.chainer_experimental, 19
chainercv.chainer_experimental.datasets.sliceable, 19
chainercv.chainer_experimental.training.extensions, 22
chainercv.datasets, 23
chainercv.evaluations, 34
chainercv.experimental.links.model.fcis, 47
chainercv.experimental.links.model.pspnet, 44
chainercv.experimental.links.model.yolo, 43
chainercv.extensions, 55
chainercv.functions, 61
chainercv.links, 61
chainercv.links.connection, 104
chainercv.links.model.faster_rcnn, 71
chainercv.links.model.fpn, 82
chainercv.links.model.resnet, 64
chainercv.links.model.segnet, 102
chainercv.links.model.senet, 67
chainercv.links.model.ssd, 88
chainercv.links.model.vgg, 70
chainercv.links.model.yolo, 99
chainercv.transforms, 107
chainercv.utils, 122
chainercv.visualizations, 117

Index

Symbols

__call__() (chainercv.experimental.links.model.fcis.FCIS method), 49
__call__() (chainercv.experimental.links.model.fcis.FCISTrainChain method), 53
__call__() (chainercv.experimental.links.model.fcis.ProposalTargetCreator method), 54
__call__() (chainercv.experimental.links.model.yolo.DarknetExtractor method), 44
__call__() (chainercv.links.PixelwiseSoftmaxClassifier method), 103
__call__() (chainercv.links.model.faster_rcnn.AnchorTargetCreator method), 79
__call__() (chainercv.links.model.faster_rcnn.FasterRCNN method), 73
__call__() (chainercv.links.model.faster_rcnn.FasterRCNNTrainChain method), 81
__call__() (chainercv.links.model.faster_rcnn.ProposalCreator method), 77
__call__() (chainercv.links.model.faster_rcnn.ProposalTargetCreator method), 81
__call__() (chainercv.links.model.faster_rcnn.RegionProposalNetwork method), 78
__call__() (chainercv.links.model.fpn.Head method), 85
__call__() (chainercv.links.model.fpn.RPN method), 86
__call__() (chainercv.links.model.segnet.SegNetBasic method), 103
__call__() (chainercv.links.model.ssd.Multibox method), 90
__call__() (chainercv.links.model.ssd.Normalize method), 92
__call__() (chainercv.links.model.ssd.SSD method), 93
__call__() (chainercv.links.model.ssd.VGG16 method), 95
__call__() (chainercv.links.model.ssd.VGG16Extractor300 method), 95
__call__() (chainercv.links.model.ssd.VGG16Extractor512 method), 95
__call__() (chainercv.links.model.yolo.Darknet19Extractor method), 100
__call__() (chainercv.links.model.yolo.Darknet53Extractor method), 100
__call__() (chainercv.links.model.yolo.ResidualBlock method), 100
__call__() (chainercv.links.model.yolo.YOLOv2Base method), 102

A

add_getter() (chainercv.chainer_experimental.datasets.sliceable.GetterDatas method), 20
ADE20KSemanticSegmentationDataset (class in chainercv.datasets), 26
ADE20KTestImageDataset (class in chainercv.datasets), 27
anchors() (chainercv.links.model.fpn.RPN method), 86
AnchorTargetCreator (class in chainercv.links.model.faster_rcnn), 79
apply_to_iterator() (in module chainercv.utils), 125
assert_is_bbox() (in module chainercv.utils), 129
assert_is_bbox_dataset() (in module chainercv.utils), 129
assert_is_detection_link() (in module chainercv.utils), 129
assert_is_image() (in module chainercv.utils), 130
assert_is_instance_segmentation_dataset() (in module chainercv.utils), 130
assert_is_label_dataset() (in module chainercv.utils), 130
assert_is_point() (in module chainercv.utils), 131
assert_is_point_dataset() (in module chainercv.utils), 131
assert_is_semantic_segmentation_dataset() (in module chainercv.utils), 131
assert_is_semantic_segmentation_link() (in module chainercv.utils), 132

B

bbox2loc() (in module chainercv.links.model.faster_rcnn), 72
bbox_iou() (in module chainercv.utils), 122
Bottleneck (class in chainercv.links.model.resnet), 66

C

cached_download() (in module chainercv.utils), 123
calc_detection_voc_ap() (in module chainercv.evaluations), 37
calc_detection_voc_prec_rec() (in module chainercv.evaluations), 37
calc_instance_segmentation_voc_prec_rec() (in module chainercv.evaluations), 41
calc_semantic_segmentation_confusion() (in module chainercv.evaluations), 42
calc_semantic_segmentation_iou() (in module chainercv.evaluations), 43
CamVidDataset (class in chainercv.datasets), 27
center_crop() (in module chainercv.transforms), 107
chainercv (module), 19
chainercv.chainer_experimental (module), 19
chainercv.chainer_experimental.datasets.sliceable (module), 19
chainercv.chainer_experimental.training.extensions (module), 22
chainercv.datasets (module), 23
chainercv.evaluations (module), 34
chainercv.experimental.links.model.fcis (module), 47
chainercv.experimental.links.model.pspnet (module), 44
chainercv.experimental.links.model.yolo (module), 43
chainercv.extensions (module), 55
chainercv.functions (module), 61
chainercv.links (module), 61, 103
chainercv.links.connection (module), 104
chainercv.links.model.faster_rcnn (module), 71
chainercv.links.model.fpn (module), 82
chainercv.links.model.resnet (module), 64
chainercv.links.model.segnet (module), 102
chainercv.links.model.senet (module), 67
chainercv.links.model.ssd (module), 88
chainercv.links.model.vgg (module), 70
chainercv.links.model.yolo (module), 99
chainercv.transforms (module), 107
chainercv.utils (module), 122
chainercv.visualizations (module), 117
CityscapesSemanticSegmentationDataset (class in chainercv.datasets), 28
CityscapesTestImageDataset (class in chainercv.datasets), 28
COCOBboxDataset (class in chainercv.datasets), 30
COCOInstanceSegmentationDataset (class in chainercv.datasets), 31
COCOSemanticSegmentationDataset (class in chainercv.datasets), 31
ConcatenatedDataset (class in chainercv.chainer_experimental.datasets.sliceable), 19
ConstantStubLink (class in chainercv.utils), 132
Conv2DActiv (class in chainercv.links.connection), 104
Conv2DBNActiv (class in chainercv.links.connection), 105
convolution_crop() (in module chainercv.experimental.links.model.pspnet), 45
copy() (chainercv.links.PickableSequentialChain method), 63
crop_bbox() (in module chainercv.transforms), 114
CUBLabelDataset (class in chainercv.datasets), 29
CUBPointDataset (class in chainercv.datasets), 29

D

Darknet19Extractor (class in chainercv.links.model.yolo), 100
Darknet53Extractor (class in chainercv.links.model.yolo), 100
DarknetExtractor (class in chainercv.experimental.links.model.yolo), 44
decode() (chainercv.links.model.fpn.Head method), 85
decode() (chainercv.links.model.fpn.RPN method), 87
decode() (chainercv.links.model.ssd.MultiboxCoder method), 91
DetectionCOCOEvaluator (class in chainercv.extensions), 55
DetectionVisReport (class in chainercv.extensions), 60
DetectionVOCEvaluator (class in chainercv.extensions), 56
directory_parsing_label_names() (in module chainercv.datasets), 24
DirectoryParsingLabelDataset (class in chainercv.datasets), 23
distribute() (chainercv.links.model.fpn.Head method), 86
download_model() (in module chainercv.utils), 123

E

encode() (chainercv.links.model.ssd.MultiboxCoder method), 91
eval_detection_coco() (in module chainercv.evaluations), 34
eval_detection_voc() (in module chainercv.evaluations), 36
eval_instance_segmentation_coco() (in module chainercv.evaluations), 38
eval_instance_segmentation_voc() (in module chainercv.evaluations), 40
eval_semantic_segmentation() (in module chainercv.evaluations), 41
extractall() (in module chainercv.utils), 123

F

FasterRCNN (class in chainercv.links.model.faster_rcnn), 73
FasterRCNN (class in chainercv.links.model.fpn), 83
FasterRCNNFPNResNet101 (class in chainercv.links.model.fpn), 83

FasterRCNNFPNResNet50 (class in chainercv.links.model.fpn), 82
FasterRCNNTrainChain (class in chainercv.links.model.faster_rcnn), 80
FasterRCNNVGG16 (class in chainercv.links.model.faster_rcnn), 71
FCIS (class in chainercv.experimental.links.model.fcis), 48
FCISResNet101 (class in chainercv.experimental.links.model.fcis), 47
FCISResNet101Head (class in chainercv.experimental.links.model.fcis), 51
FCISTrainChain (class in chainercv.experimental.links.model.fcis), 53
FeaturePredictor (class in chainercv.links), 62
flip() (in module chainercv.transforms), 107
flip_bbox() (in module chainercv.transforms), 114
flip_point() (in module chainercv.transforms), 116
FPN (class in chainercv.links.model.fpn), 84

G

generate_anchor_base() (in module chainercv.links.model.faster_rcnn), 75
generate_random_bbox() (in module chainercv.utils), 132
get_example_by_keys() (chainercv.chainer_experimental.datasets.sliceable.ConcatenatedDataset method), 20
get_example_by_keys() (chainercv.chainer_experimental.datasets.sliceable.GetterDataset method), 21
get_example_by_keys() (chainercv.chainer_experimental.datasets.sliceable.TupleDataset method), 21
GetterDataset (class in chainercv.chainer_experimental.datasets.sliceable), 20
GradientScaling (class in chainercv.links.model.ssd), 96

H

Head (class in chainercv.links.model.fpn), 85
head_loss_post() (in module chainercv.links.model.fpn), 88
head_loss_pre() (in module chainercv.links.model.fpn), 87

I

InstanceSegmentationCOCOEvaluator (class in chainercv.extensions), 57
InstanceSegmentationVOCEvaluator (class in chainercv.extensions), 58

L

loc2bbox() (in module chainercv.links.model.faster_rcnn), 75

M

make_shift() (in module chainercv.chainer_experimental.training.extensions), 22
mask_iou() (in module chainercv.utils), 128
mask_to_bbox() (in module chainercv.utils), 128
mask_voting() (in module chainercv.experimental.links.model.fcis), 52
MixUpSoftLabelDataset (class in chainercv.datasets), 24
Multibox (class in chainercv.links.model.ssd), 90
multibox_loss() (in module chainercv.links.model.ssd), 96
MultiboxCoder (class in chainercv.links.model.ssd), 90

N

non_maximum_suppression() (in module chainercv.utils), 122
Normalize (class in chainercv.links.model.ssd), 92

O

OnlineProductsDataset (class in chainercv.datasets), 32

P

pca_lighting() (in module chainercv.transforms), 108
PickableSequentialChain (class in chainercv.links), 63
PixelwiseSoftmaxClassifier (class in chainercv.links), 103
predict() (chainercv.experimental.links.model.fcis.FCIS method), 50
predict() (chainercv.experimental.links.model.pspnet.PSPNet method), 47
predict() (chainercv.links.FeaturePredictor method), 62
predict() (chainercv.links.model.faster_rcnn.FasterRCNN method), 74
predict() (chainercv.links.model.fpn.FasterRCNN method), 84
predict() (chainercv.links.model.segnet.SegNetBasic method), 103
predict() (chainercv.links.model.ssd.SSD method), 94
predict() (chainercv.links.model.yolo.YOLOBase method), 101
prepare() (chainercv.experimental.links.model.fcis.FCIS method), 50
prepare() (chainercv.links.model.faster_rcnn.FasterRCNN method), 74
prepare() (chainercv.links.model.fpn.FasterRCNN method), 84
prepare_pretrained_model() (in module chainercv.utils), 128
ProgressHook (class in chainercv.utils), 127
ProposalCreator (class in chainercv.links.model.faster_rcnn), 76
ProposalTargetCreator (class in chainercv.experimental.links.model.fcis), 54

ProposalTargetCreator (class in chainercv.links.model.faster_rcnn), 81
PSPNet (class in chainercv.experimental.links.model.pspnet), 46
PSPNetResNet101 (class in chainercv.experimental.links.model.pspnet), 44
PSPNetResNet50 (class in chainercv.experimental.links.model.pspnet), 45
psroi_pooling_2d() (in module chainercv.functions), 61

R

random_crop() (in module chainercv.transforms), 108
random_crop_with_bbox_constraints() (in module chainercv.links.model.ssd), 97
random_distort() (in module chainercv.links.model.ssd), 98
random_expand() (in module chainercv.transforms), 109
random_flip() (in module chainercv.transforms), 109
random_rotate() (in module chainercv.transforms), 110
random_sized_crop() (in module chainercv.transforms), 110
read_image() (in module chainercv.utils), 124
read_label() (in module chainercv.utils), 124
RegionProposalNetwork (class in chainercv.links.model.faster_rcnn), 77
remove_unused() (chainercv.links.PickableSequentialChain method), 64
ResBlock (class in chainercv.links.model.resnet), 66
ResidualBlock (class in chainercv.links.model.yolo), 100
resize() (in module chainercv.transforms), 111
resize_bbox() (in module chainercv.transforms), 115
resize_contain() (in module chainercv.transforms), 112
resize_point() (in module chainercv.transforms), 116
resize_with_random_interpolation() (in module chainercv.links.model.ssd), 98
ResNet (class in chainercv.links.model.resnet), 64
ResNet101 (class in chainercv.links.model.resnet), 65
ResNet101Extractor (class in chainercv.experimental.links.model.fcis), 52
ResNet152 (class in chainercv.links.model.resnet), 65
ResNet50 (class in chainercv.links.model.resnet), 65
rotate() (in module chainercv.transforms), 112
rotate_bbox() (in module chainercv.transforms), 115
RPN (class in chainercv.links.model.fpn), 86
rpn_loss() (in module chainercv.links.model.fpn), 88

S

SBDInstanceSegmentationDataset (class in chainercv.datasets), 34
scale() (in module chainercv.transforms), 113
SEBlock (class in chainercv.links.connection), 106
SegNetBasic (class in chainercv.links.model.segnet), 102
SemanticSegmentationEvaluator (class in chainercv.extensions), 59
SEResNet (class in chainercv.links.model.senet), 67
SEResNet101 (class in chainercv.links.model.senet), 68
SEResNet152 (class in chainercv.links.model.senet), 68
SEResNet50 (class in chainercv.links.model.senet), 68
SEResNeXt (class in chainercv.links.model.senet), 68
SEResNeXt101 (class in chainercv.links.model.senet), 69
SEResNeXt50 (class in chainercv.links.model.senet), 69
SiameseDataset (class in chainercv.datasets), 25
SSD (class in chainercv.links.model.ssd), 92
SSD300 (class in chainercv.links.model.ssd), 88
SSD512 (class in chainercv.links.model.ssd), 89

T

ten_crop() (in module chainercv.transforms), 113
tile_images() (in module chainercv.utils), 125
to_cpu() (chainercv.links.model.ssd.SSD method), 94
to_cpu() (chainercv.links.model.yolo.YOLOv2Base method), 102
to_cpu() (chainercv.links.model.yolo.YOLOv3 method), 100
to_cpu() (chainercv.links.PixelwiseSoftmaxClassifier method), 104
to_cpu() (chainercv.utils.ConstantStubLink method), 132
to_gpu() (chainercv.links.model.ssd.SSD method), 94
to_gpu() (chainercv.links.model.yolo.YOLOv2Base method), 102
to_gpu() (chainercv.links.model.yolo.YOLOv3 method), 100
to_gpu() (chainercv.links.PixelwiseSoftmaxClassifier method), 104
to_gpu() (chainercv.utils.ConstantStubLink method), 132
TransformDataset (class in chainercv.chainer_experimental.datasets.sliceable), 22
translate_bbox() (in module chainercv.transforms), 116
translate_point() (in module chainercv.transforms), 117
TupleDataset (class in chainercv.chainer_experimental.datasets.sliceable), 21

U

unzip() (in module chainercv.utils), 127
use_preset() (chainercv.experimental.links.model.fcis.FCIS method), 50
use_preset() (chainercv.links.model.faster_rcnn.FasterRCNN method), 75
use_preset() (chainercv.links.model.fpn.FasterRCNN method), 84
use_preset() (chainercv.links.model.ssd.SSD method), 94
use_preset() (chainercv.links.model.yolo.YOLOBase method), 101

V

VGG16 (class in chainercv.links.model.ssd), 95
VGG16 (class in chainercv.links.model.vgg), 70
VGG16Extractor300 (class in chainercv.links.model.ssd), 95
VGG16Extractor512 (class in chainercv.links.model.ssd), 95
VGG16RoIHead (class in chainercv.links.model.faster_rcnn), 79
vis_bbox() (in module chainercv.visualizations), 117
vis_image() (in module chainercv.visualizations), 118
vis_instance_segmentation() (in module chainercv.visualizations), 119
vis_point() (in module chainercv.visualizations), 120
vis_semantic_segmentation() (in module chainercv.visualizations), 121
VOCBboxDataset (class in chainercv.datasets), 32
VOCInstanceSegmentationDataset (class in chainercv.datasets), 33
VOCSemanticSegmentationDataset (class in chainercv.datasets), 33

W

write_image() (in module chainercv.utils), 125

Y

YOLOBase (class in chainercv.links.model.yolo), 101
YOLOv2 (class in chainercv.links.model.yolo), 99
YOLOv2Base (class in chainercv.links.model.yolo), 101
YOLOv2Tiny (class in chainercv.experimental.links.model.yolo), 43
YOLOv3 (class in chainercv.links.model.yolo), 99