
ASHPY Documentation

Release 0.1.3

ml@zuru.tech

Oct 07, 2019

CONTENTS:

1	Welcome To AshPy!	1
1.1	AshPy	1
2	Write The Docs!	11
2.1	The Whys	11
2.2	Documentation Architecture	12
2.3	Additional Materials	14
3	Getting Started	15
3.1	Datasets	15
3.2	Models	16
3.3	Creating a new Trainer	20
3.4	Complete Examples	20
4	API Reference	31
4.1	ashpy.ashtypes	31
4.2	ashpy.contexts	31
4.3	ashpy.keras	49
4.4	ashpy.layers	57
4.5	ashpy.losses	62
4.6	ashpy.metrics	110
4.7	ashpy.models	155
4.8	ashpy.modes	266
4.9	ashpy.trainers	267
5	AshPy Internals	297
5.1	Context	297
5.2	Executor	298
6	Dependencies Graph	301
6.1	ashpy.models	301
6.2	ashpy.trainers	301
6.3	ashpy.layers	302
6.4	ashpy.losses	302
6.5	ashpy.metrics	303
7	About	305
8	Indices and tables	307
	Python Module Index	309

WELCOME TO ASHPY!

Warning: AshPy is still a work in progress and may change substantially before the first proper release. The API is not mature enough to be considered stable, but we'll try to keep breaking changes to a minimum.

1.1 AshPy

AshPy is a TensorFlow 2.0 library for (**distributed**) training, evaluation, model selection, and fast prototyping. It is designed to ease the burden of setting up all the nuances of the architectures built to train complex custom deep learning models.

Quick Example | Features | Set Up | Usage | Dataset Output Format | Test

1.1.1 Quick Example

```
# define a distribution strategy
strategy = tf.distribute.MirroredStrategy()

# work inside the scope of the created strategy
with strategy.scope():

    # get the MNIST dataset
    train, validation = tf.keras.datasets.mnist.load_data()

    # process data if needed
    def process(images, labels):
```

(continues on next page)

(continued from previous page)

```
data_images = tf.data.Dataset.from_tensor_slices((images)).map(
    lambda x: tf.reshape(x, (28 * 28,))
)
data_images = data_images.map(
    lambda x: tf.image.convert_image_dtype(x, tf.float32)
)
data_labels = tf.data.Dataset.from_tensor_slices((labels))
dataset = tf.data.Dataset.zip((data_images, data_labels))
dataset = dataset.batch(1024 * 1)
return dataset

# apply the process function to the data
train, validation = (
    process(train[0], train[1]),
    process(validation[0], validation[1]),
)

# create the model
model = tf.keras.Sequential(
    [
        tf.keras.layers.Dense(10, activation=tf.nn.sigmoid),
        tf.keras.layers.Dense(10),
    ]
)

# define the optimizer
optimizer = tf.optimizers.Adam(1e-3)

# the loss is provided by the AshPy library
loss = ClassifierLoss(tf.losses.SparseCategoricalCrossentropy(from_logits=True))
logdir = "testlog"
epochs = 10

# the metrics are provided by the AshPy library
# and every metric with model_selection_operator != None performs
# model selection, saving the best model in a different folder per metric.
metrics = [
    ClassifierMetric(
        tf.metrics.Accuracy(), model_selection_operator=operator.gt
    ),
    ClassifierMetric(
        tf.metrics.BinaryAccuracy(), model_selection_operator=operator.gt
    ),
]

# define the AshPy trainer
trainer = ClassifierTrainer(
    model, optimizer, loss, epochs, metrics, logdir=logdir
)

# run the training process
trainer(train, validation)
```

1.1.2 Features

AshPy is a library designed to ease the burden of setting up all the nuances of the architectures built to train complex custom deep learning models. It provides both fully convolutional and fully connected models such as:

- autoencoder
- decoder
- encoder

and a fully convolutional:

- unet

Moreover, it provides already prepared trainers for a classifier model and GAN networks. In particular, in regards of the latter, it offers a basic GAN architecture with a Generator-Discriminator structure and an enhanced GAN architecture version made up of a Encoder-Generator-Discriminator structure.

AshPy it is developed around the concepts of *Executor*, *Context*, *Metric*, and *Strategies* that represents its foundations.

Executor An Executor is a class that helps to better generalize a training loop. With an Executor you can construct, for example, a custom loss function and put whatever computation you need inside it. You should define a `call` function inside your class and decorate it with `@Executor.reduce` header. Inside the `call` function you can take advantage of a context.

Context A Context is a useful class in which all the models, metrics, dataset and mode of your network are set. Passing the context around means that you can any time access to all what you need in order to performs any type of computation.

Metric A Metric is a class from which you can inherit to create your custom metric that can automatically keep track of the best performance of the model during training and, automatically save the best one doing what is called the *model selection*.

Strategies If you want to distribute your training across multiple GPUs, there is the `tf.distribute.Strategy` TensorFlow API with which you can distribute your models and training code with minimal code changes. AshPy implements this type of strategies internally and will check everything for you to apply the distribution strategy correctly. All you need to do is as simple as doing the following:

```
strategy = tf.distribute.MirroredStrategy()
with strategy.scope():

    generator = ConvGenerator(
        layer_spec_input_res=(7, 7),
        layer_spec_target_res=(28, 28),
        kernel_size=(5, 5),
        initial_filters=256,
        filters_cap=16,
        channels=1,
    )
    # rest of the code
    # with trainer definition and so on
```

i.e., create the strategy and put the rest of the code inside its scope.

In general AshPy aims to:

- Rapid model prototyping
- Enforcement of best practices & API consistency

- Remove duplicated and boilerplate code
- General usability by new project

NOTE: We invite you to read the full documentation on [the official website](#).

The following README aims to help you understand what you need to do to setup AshPy on your system and, with some examples, what you need to do to setup a complete training of your network. Moreover, it will explain some fundamental modules you need to understand to fully exploit the potential of the library.

1.1.3 Set up

Pip install

```
# Depending on GPU support you might want to install
# tensorflow-gpu or tensorflow
pip install tensorflow-gpu==2.0.0beta1
#pip install tensorflow==2.0.0beta1
pip install ashpy
```

Source install

Clone this repo, go inside the downloaded folder and install with:

```
pip install -e .
```

1.1.4 Usage

Let's quickly start with some examples.

Classifier

Let's say we want to train a classifier.

```
import operator
import tensorflow as tf
from ashpy.metrics import ClassifierMetric
from ashpy.trainers.classifier import ClassifierTrainer
from ashpy.losses.classifier import ClassifierLoss

def toy_dataset():
    inputs = tf.expand_dims(tf.range(1, 1000.0), -1)
    labels = tf.expand_dims([1 if tf.equal(tf.math.mod(tf.squeeze(i), 2), 0) else 0,
    ↪for i in inputs], -1)
    return tf.data.Dataset.from_tensor_slices((inputs, labels)).shuffle(10).batch(2)

model = tf.keras.Sequential([
    tf.keras.layers.Dense(10, activation=tf.nn.sigmoid),
    tf.keras.layers.Dense(2)
])

optimizer = tf.optimizers.Adam(1e-3)
```

(continues on next page)

(continued from previous page)

```

loss = ClassifierLoss(tf.losses.SparseCategoricalCrossentropy(from_logits=True))
logdir = "testlog"
epochs = 2

metrics = [
    ClassifierMetric(tf.metrics.Accuracy(), model_selection_operator=operator.gt),
    ClassifierMetric(tf.metrics.BinaryAccuracy(), model_selection_operator=operator.
↪gt),
]

trainer = ClassifierTrainer(model, optimizer, loss, epochs, metrics, logdir=logdir)

train, validation = toy_dataset(), toy_dataset()
trainer(train, validation)

```

Skipping the `toy_dataset()` function that creates a toy dataset, we'll give a look to the code step by step.

So, first of all we define a model and its optimizer. Here, the model is a very simple sequential Keras model defined as:

```

model = tf.keras.Sequential([
    tf.keras.layers.Dense(10, activation=tf.nn.sigmoid),
    tf.keras.layers.Dense(2)
])

optimizer = tf.optimizers.Adam(1e-3)

```

Then we define the loss:

```

loss = ClassifierLoss(tf.losses.SparseCategoricalCrossentropy(from_logits=True))

```

The `ClassifierLoss` loss defined above it is defined using an internal class called "Executor". The `Executor` is a class that let you define, alongside with a desired loss, the function that you want to use to "evaluate" that loss with all the needed parameters.

This works in conjunction with the following line (we will speak about the "metrics" and the other few definition lines in a minute):

```

trainer = ClassifierTrainer(model, optimizer, loss, epochs, metrics, logdir=logdir)

```

where a `ClassifierTrainer` is an object designed to run a specific training procedure adjusted, in this case, for a classifier.

The arguments of this function are the model, the optimizer, the loss, the number of epochs, the metrics and the logdir. We have already seen the definition of the model, the optimizer and of the loss. The definition of epochs, metrics and logdir happens here:

```

logdir = "testlog"
epochs = 2

metrics = [
    ClassifierMetric(tf.metrics.Accuracy(), model_selection_operator=operator.gt),
    ClassifierMetric(
        tf.metrics.BinaryAccuracy(), model_selection_operator=operator.gt),
]

```

What we need to underline here is the definition of the metrics because as you can see they are defined through the use of specific classes: `ClassifierMetric`. As for the `ClassifierTrainer`, the `ClassifierMetric` it

is a specified designed class for the Classifier. If you want to create a different metric you should inheriting from the Metric class provided by the Ash library. This kind of Metrics are useful because you can indicate a processing function to apply on predictions (e.g., `tf.argmax`) and an operator (e.g., `operator.gt` is the “greater than” operator) if you desire to activate the model selection during the training process based on that particular metric.

Finally, once the datasets has been set, you can start the training procedure calling the trainer object:

```
train, validation = toy_dataset(), toy_dataset()
trainer(train, validation)
```

1.1.5 GAN - Generative Adversarial Network

AshPy is equipped with two types of GAN network architectures:

- A plain GAN network with the classic structure Generator - Discriminator.
- A more elaborated GAN network architecture with the classic Generator - Discriminator structure plus an Encoder model (BiGAN like).

As for the previous classifier training example, let’s see for first a simple example of an entire “toy” code, regarding a simple plain GAN. At the end we will briefly touch upon the differences with the GAN network with the Encoder.

```
import operator
import tensorflow as tf
from ashpy.models.gans import ConvGenerator, ConvDiscriminator
from ashpy.metrics import InceptionScore
from ashpy.losses.gan import DiscriminatorMinMax, GeneratorBCE

generator = ConvGenerator(
    layer_spec_input_res=(7, 7),
    layer_spec_target_res=(28, 28),
    kernel_size=(5, 5),
    initial_filters=32,
    filters_cap=16,
    channels=1,
)

discriminator = ConvDiscriminator(
    layer_spec_input_res=(28, 28),
    layer_spec_target_res=(7, 7),
    kernel_size=(5, 5),
    initial_filters=16,
    filters_cap=32,
    output_shape=1,
)

# Losses
generator_bce = GeneratorBCE()
minmax = DiscriminatorMinMax()

# Real data
batch_size = 2
mnist_x, mnist_y = tf.zeros((100,28,28)), tf.zeros((100,))

# Trainer
epochs = 2
logdir = "testlog/adversarial"
```

(continues on next page)

(continued from previous page)

```

metrics = [
    InceptionScore(
        # Fake inception model
        ConvDiscriminator(
            layer_spec_input_res=(299, 299),
            layer_spec_target_res=(7, 7),
            kernel_size=(5, 5),
            initial_filters=16,
            filters_cap=32,
            output_shape=10,
        ),
        model_selection_operator=operator.gt,
        logdir=logdir,
    )
]

trainer = AdversarialTrainer(
    generator,
    discriminator,
    tf.optimizers.Adam(1e-4),
    tf.optimizers.Adam(1e-4),
    generator_bce,
    minmax,
    epochs,
    metrics,
    logdir,
)

# Dataset
noise_dataset = tf.data.Dataset.from_tensors(0).repeat().map(
    lambda _: tf.random.normal(shape=(100,), dtype=tf.float32, mean=0.0, stddev=1)
).batch(batch_size).prefetch(1)

# take only 2 samples to speed up tests
real_data = tf.data.Dataset.from_tensor_slices(
    (tf.expand_dims(mnist_x, -1), tf.expand_dims(mnist_y, -1))
).take(batch_size).batch(batch_size).prefetch(1)

# Add noise in the same dataset, just by mapping.
# The return type of the dataset must be: tuple(tuple(a,b), noise)
dataset = real_data.map(lambda x, y: ((x, y), tf.random.normal(shape=(batch_size, 100))))

trainer(dataset)

```

First we define the generator and discriminator of the GAN architecture:

```

generator = ConvGenerator(
    layer_spec_input_res=(7, 7),
    layer_spec_target_res=(28, 28),
    kernel_size=(5, 5),
    initial_filters=32,
    filters_cap=16,
    channels=1,
)

```

(continues on next page)

(continued from previous page)

```
discriminator = ConvDiscriminator(
    layer_spec_input_res=(28, 28),
    layer_spec_target_res=(7, 7),
    kernel_size=(5, 5),
    initial_filters=16,
    filters_cap=32,
    output_shape=1,
)
```

and then we define the losses:

```
# Losses
generator_bce = GeneratorBCE()
minmax = DiscriminatorMinMax()
```

where `GeneratorBCE()` and `DiscriminatorMinMax()` are the losses defined inheriting `Executor`. Again, as we have seen in the previous classifier example, you can customize this type (the ones inheriting from the `Executor`) of losses.

The metrics are defined as follow:

```
metrics = [
    InceptionScore(
        # Fake inception model
        ConvDiscriminator(
            layer_spec_input_res=(299, 299),
            layer_spec_target_res=(7, 7),
            kernel_size=(5, 5),
            initial_filters=16,
            filters_cap=32,
            output_shape=10,
        ),
        model_selection_operator=operator.gt,
        logdir=logdir,
    )
]
```

and in particular here we have the `InceptionScore` metric constructed on the fly with the `ConvDiscriminator` class provided by `AshPy`.

Finally, the actual trainer is constructed and then called:

```
trainer = AdversarialTrainer(
    generator,
    discriminator,
    tf.optimizers.Adam(1e-4),
    tf.optimizers.Adam(1e-4),
    generator_bce,
    minmax,
    epochs,
    metrics,
    logdir,
)
```

```
trainer(dataset)
```

The main difference with a GAN architecture with an Encoder is that we would have the encoder loss:

```
encoder_bce = EncoderBCE()
```

an encoder accuracy metric:

```
metrics = [EncodingAccuracy(classifier, model_selection_operator=operator.gt,
↳logdir=logdir)]
```

and an EncoderTrainer:

```
trainer = EncoderTrainer(
    generator,
    discriminator,
    encoder,
    tf.optimizers.Adam(1e-4),
    tf.optimizers.Adam(1e-5),
    tf.optimizers.Adam(1e-6),
    generator_bce,
    minmax,
    encoder_bce,
    epochs,
    metrics=metrics,
    logdir=logdir,
)
```

Note that the `EncoderTrainer` indicates a trainer of a GAN network with an Encoder and not a trainer of an Encoder itself.

1.1.6 Dataset Output Format

In order to standardize the GAN training, AshPy requires the input dataset to be in a common format. In particular, the dataset return type must always be in the format showed below, where the first element of the tuple is the discriminator input, and the second is the generator input.

```
tuple(tuple(a,b), noise)
```

Where `a` is the input sample, `b` is the label/condition (if any, otherwise fill it with 0), and `noise` is the latent vector of input.

To train Pix2Pix-like architecture, that have no `noise` as `ConvGenerator` input, just return the values in the format `(tuple(a,b), b)` since the condition is the generator output.

1.1.7 Test

In order to run the doctests first you need to install the `pytest-sphinx` package and `pytest-cov` package:

```
pip --no-cache-dir install --upgrade git+https://github.com/thisch/pytest-sphinx.git
↳pytest
pip install pytest-cov
```

Then you can run tests with coverage statistics with:

```
pytest -x -s -vvv --doctest-modules ashpy --cov=ashpy
```


WRITE THE DOCS!

Ash being a project built with a **Documentation Driven** approach means that a solid, automated documentation procedure is a mandatory requirement.

The core components of our systems are:

- Sphinx for the documentation generation
- reStructuredText as the markup language
- Google Style docstrings for in-code documentation
- `'vale'` and `vale-styles`
- Automatic internal deployment via GitLab Pages CI/CD integration

This document goal is threefold:

1. Explaining the Documentation Architecture, the steps taken to automate it and defending such choices
2. Serve as a future reference for other projects
3. Act as an example for the Guide format and a demo of Sphinx + reST superpowers
4. **Convince you of the need to always be on the lookout for errors even in a perfect** system.

2.1 The Whys

2.1.1 Why Sphinx?

Sphinx is the most used documentation framework for Python, developed for the Standard library itself it's now adopted by all the most known third party libraries. What makes Sphinx so great is the combination of extensibility via themes, extensions and what not, coupled with a plethora of builtin functionalities that make writing docs a breeze.:

An example from Sphinx Site:

- Output formats: HTML (including Windows HTML Help), LaTeX (for printable PDF versions), ePub, Texinfo, manual pages, plain text
- Extensive cross-references: semantic markup and automatic links for functions, classes, citations, glossary terms and similar pieces of information
- Hierarchical structure: easy definition of a document tree, with automatic links to siblings, parents and children
- Automatic indices: general index as well as a language-specific module indices
- Code handling: automatic highlighting using the Pygments highlighter

- Extensions: automatic testing of code snippets, inclusion of docstrings from Python modules (API docs), and more
- Contributed extensions: more than 50 extensions contributed by users in a second repository; most of them installable from PyPI

2.1.2 Why reST?

More than why reST, the real question is *Why not Markdown?*

While Markdown can be easier and slightly quicker to write, it does not offer the same level of fine grained control, necessary for an effort as complex as technical writing, without sacrificing portability.

Eric Holscher has an aptly named article: [Why You Shouldn't Use "Markdown" for Documentation](#), he is one of the greatest documentation advocate out there. Go and read his articles, they are beautiful.

2.1.3 Why Google Style for Docstrings?

Google Docstrings are to us the best way to organically combine code and documentation. Leveraging Napoleon, a Sphinx extension offering automatic documentation support for both Numpy and Google docstrings style, we can write easy to read docstrings and still be able to use `autodoc` and `autosummary` directives.

2.2 Documentation Architecture

2.2.1 Tutorials, Guides, Complex Examples

Any form of documentation which is not generated from the codebase should go here. Parent/Entry Point reStructuredText file, should be added in `docs/src` and then referenced in `index.rst`

2.2.2 API Reference

API reference contains the full API documentation automatically generated from our codebase. The only manual step required is adding the module you want to document to the `api.rst` located inside `docs/source`.

Automate all the docs!

Classes, Functions, Exceptions: Annotate them normally, they do not require anything else.

Autosummary & submodules with imports: A painful story

Exposing Python objects to their parent module by importing them in its `__init__.py` file, breaks the `autosummary` directives when combining it with the automatic generation of stub files. Currently there's no way of making `autosummary` aware of the imported objects thus if you desire to document that API piece you need to find a workaround.

Example

Suppose we have the following structure:

```
keras/
  |---> __init__.py
  |
  |---> models.py
```

And that these two file contains respectively:

- `__init__.py`

```
from .models import Model

__ALL__ = ["Model"]
```

- `models.py`

```
class Model:
    pass
```

Calling the `autosummary` directive (with the `toctree` option) on `keras` will not generate stub files for `keras.Model` causing it to not show in the Table of Contents of our API reference.

To circumvent this limitation it is ideal to insert some manual labour into the `keras` docstring.

- `__init__.py`

```
"""
Documentation example.

.. rubric:: Classes

.. autosummary:: Classes
   :toctree: _autosummary
   :nosignatures:

   keras.Model

.. rubric:: Submodules

.. autosummary:: keras.models
   :toctree: _autosummary
   :nosignatures:
   :template: autosummary/submodule.rst

   keras.models
"""
from .models import Model

__ALL__ = ["Model"]
```

This way `autosummary` will produce the proper API documentation. The same approach applies also when exposing functions, exceptions, and modules.

Note: used when annotating submodules.

2.2.3 Inheritance Diagrams

Inheritance Diagrams are drawn using `sphinx.ext.inheritance_diagram` and `sphinx.ext.graphviz`.

The `autosummary` template for classes has been modified in order to automatically generate an inheritance diagram just below the title.

An `Inheritance Diagrams` page is manually created in order to showcase all the diagrams in one single page. The page gives a quick overview of the relations between the classes of each module.

2.3 Additional Materials

- [Sphinx](#)
- [Google Python Style](#)
- [Google Developer Documentation Style Guide](#)
- [Napoleon - Example Google Style Python Docstrings](#)
- [Read the Docs Sphinx Theme](#)
- [Write the Docs](#)
- [reStructuredText Primer](#)
- [vale](#)
- [Eric Holscher](#)

GETTING STARTED

3.1 Datasets

AshPy supports `tf.data.Dataset` format.

We highly encourage you to use [Tensorflow Datasets](#) to manage and use your datasets in an handy way.

```
pip install tfds-nightly
```

3.1.1 Classification

In order to create a dataset for classification:

```
import tensorflow_datasets as tfds

from ashpy.trainers import ClassifierTrainer

def extract_fn(example):
    return example["image"], example["label"]

def main():
    ds_train, ds_validation = tfds.load(name="mnist", split=["train", "validation"])

    # build the input pipeline
    ds_train = ds_train.batch(BATCH_SIZE).prefetch(1)
    ds_train = ds_train.map(extract_fn)

    # same for validation
    ...

    # define model, loss, optimizer
    ...

    # define the classifier trainer
    trainer = ClassifierTrainer(model, optimizer, loss, epochs, metrics, _
↪ logdir=logdir)

    # train
    trainer.train(ds_train, ds_validation)
```

3.1.2 GANs

In order to create a datasets for a (Conditional) GANs:

```
import tensorflow_datasets as tfds

from ashpy.trainers import AdversarialTrainer

def extract_fn(example):
    # the ashpy input must be (real, condition), condition
    return (example["image"], example["label"], example["label"])

def main():
    ds_train = tfds.load(name="mnist", split="train")

    # build the input pipeline
    ds_train = ds_train.batch(BATCH_SIZE).prefetch(1)
    ds_train = ds_train.map(extract_fn)

    # define models, losses, optimizers
    ...

    # define the adversarial trainer
    trainer = AdversarialTrainer(generator,
                                  discriminator,
                                  generator_optimizer,
                                  discriminator_optimizer,
                                  generator_loss,
                                  discriminator_loss,
                                  epochs,
                                  metrics,
                                  logdir,
                                  )

    # train
    trainer.train(ds_train)
```

3.2 Models

AshPy supports [Keras](#) models as inputs. You can use an AshPy predefined model or you can implement your own model.

3.2.1 Using an AshPy model

```
import tensorflow_datasets as tfds

from ashpy.trainers import ClassifierTrainer
from ashpy.models import UNet

def main():

    # create the dataset and the input pipeline
```

(continues on next page)

(continued from previous page)

```

# define models, loss, optimizer
model = UNet(
    input_res,
    min_res,
    kernel_size,
    initial_filters,
    filters_cap,
    channels,
    use_dropout_encoder,
    use_dropout_decoder,
    dropout_prob,
    use_attention,
)

# define the classifier trainer
trainer = AdversarialTrainer(generator,
    discriminator,
    generator_optimizer,
    discriminator_optimizer,
    generator_loss,
    discriminator_loss,
    epochs,
    metrics,
    logdir,
)

# train
trainer.train(ds_train)

```

3.2.2 Creating a Model

It's very easy to create a simple model, since AshPy's models are Keras' models.

```

from ashpy.layers import Attention, InstanceNormalization

def downsample(
    filters,
    apply_normalization=True,
    attention=False,
    activation=tf.keras.layers.LeakyReLU(alpha=0.2),
    size=3,
):
    initializer = tf.random_normal_initializer(0.0, 0.02)

    result = tf.keras.Sequential()
    result.add(
        tf.keras.layers.Conv2D(
            filters,
            size,
            strides=2,
            padding="same",
            kernel_initializer=initializer,
            use_bias=not apply_normalization,
        )
    )

```

(continues on next page)

```

if apply_normalization:
    result.add(InstanceNormalization())

result.add(activation)

if attention:
    result.add(Attention(filters))

return result

def upsample(
    filters,
    apply_dropout=False,
    apply_normalization=True,
    attention=False,
    activation=tf.keras.layers.ReLU(),
    size=3,
):
    initializer = tf.random_normal_initializer(0.0, 0.02)

    result = tf.keras.Sequential()
    result.add(tf.keras.layers.UpSampling2D(size=(2, 2)))
    result.add(tf.keras.layers.ZeroPadding2D(padding=(1, 1)))

    result.add(
        tf.keras.layers.Conv2D(
            filters,
            size,
            strides=1,
            padding="valid",
            kernel_initializer=initializer,
            use_bias=False,
        )
    )

    if apply_dropout:
        result.add(tf.keras.layers.Dropout(0.5))

    if apply_normalization:
        result.add(Normalizer())

    result.add(activation)

    if attention:
        result.add(Attention(filters))

    return result

def Generator(attention, output_channels=3):
    down_stack = [
        downsample(32, apply_normalization=False), # 256
        downsample(32), # 128
        downsample(64, attention=attention), # 64
        downsample(64), # 32

```

(continues on next page)

(continued from previous page)

```

    downsample(64), # 16
    downsample(128), # 8
    downsample(128), # 4
    downsample(256), # 2
    downsample(512, apply_normalization=False), # 1
]

up_stack = [
    upsample(256, apply_dropout=True), # 2
    upsample(128, apply_dropout=True), # 4
    upsample(128, apply_dropout=True), # 8
    upsample(64), # 16
    upsample(64), # 32
    upsample(64, attention=attention), # 64
    upsample(32), # 128
    upsample(32), # 256
    upsample(32), # 512
]

inputs = tf.keras.layers.Input(shape=[None, None, 1])
x = inputs

# Downsampling through the model
skips = []
for down in down_stack:
    x = down(x)
    skips.append(x)

skips = reversed(skips[:-1])

# Upsampling and establishing the skip connections
for up, skip in zip(up_stack, skips):
    x = up(x)
    x = tf.keras.layers.Concatenate()([x, skip])

last = upsample(
    output_channels,
    activation=tf.keras.layers.Activation(tf.nn.tanh),
    apply_normalization=False,
)

x = last(x)

return tf.keras.Model(inputs=inputs, outputs=x)

```

In this way we have created a new model to be used inside AshPy.

3.2.3 Inheriting from `ashpy.models.Conv2DInterface`

The third possibility you have to create a new model is to inherit from the `ashpy.models.convolutional.interfaces.Conv2DInterface`.

This class offers the basic methods to implement in a simple way a new model.

3.3 Creating a new Trainer

AshPy has different generic trainers. Trainers implement the basic training loop together with distribution strategy management and logging. By now the only distribution strategy handled is the `tf.distribute.MirroredStrategy`.

3.4 Complete Examples

3.4.1 Classifier

```
1 # Copyright 2019 Zuru Tech HK Limited. All Rights Reserved.
2 #
3 # Licensed under the Apache License, Version 2.0 (the "License");
4 # you may not use this file except in compliance with the License.
5 # You may obtain a copy of the License at
6 #
7 # http://www.apache.org/licenses/LICENSE-2.0
8 #
9 # Unless required by applicable law or agreed to in writing, software
10 # distributed under the License is distributed on an "AS IS" BASIS,
11 # WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.
12 # See the License for the specific language governing permissions and
13 # limitations under the License.
14
15 """
16 Example of Multi-GPU classifier trainer
17 """
18
19 import operator
20
21 import tensorflow as tf
22
23 from ashpy.losses import ClassifierLoss
24 from ashpy.metrics import ClassifierMetric
25 from ashpy.trainers import ClassifierTrainer
26
27
28 def main():
29     """How to use ash to train a classifier, measure the
30     performance and perform model selection."""
31
32     strategy = tf.distribute.MirroredStrategy()
33     with strategy.scope():
34         train, validation = tf.keras.datasets.mnist.load_data()
35
36         def process(images, labels):
37             data_images = tf.data.Dataset.from_tensor_slices((images)).map(
38                 lambda x: tf.reshape(x, (28 * 28,))
39             )
40             data_images = data_images.map(
41                 lambda x: tf.image.convert_image_dtype(x, tf.float32)
42             )
```

(continues on next page)

(continued from previous page)

```

43     data_labels = tf.data.Dataset.from_tensor_slices((labels))
44     dataset = tf.data.Dataset.zip((data_images, data_labels))
45     dataset = dataset.batch(1024 * 1)
46     return dataset
47
48     train, validation = (
49         process(train[0], train[1]),
50         process(validation[0], validation[1]),
51     )
52
53     model = tf.keras.Sequential(
54         [
55             tf.keras.layers.Dense(10, activation=tf.nn.sigmoid),
56             tf.keras.layers.Dense(10),
57         ]
58     )
59     optimizer = tf.optimizers.Adam(1e-3)
60     loss = ClassifierLoss(tf.losses.SparseCategoricalCrossentropy(from_
↪ logits=True))
61     logdir = "testlog"
62     epochs = 10
63
64     metrics = [
65         ClassifierMetric(
66             tf.metrics.Accuracy(), model_selection_operator=operator.gt
67         ),
68         ClassifierMetric(
69             tf.metrics.BinaryAccuracy(), model_selection_operator=operator.gt
70         ),
71     ]
72
73     trainer = ClassifierTrainer(
74         model, optimizer, loss, epochs, metrics, logdir=logdir
75     )
76     trainer(train, validation)
77
78
79 if __name__ == "__main__":
80     main()

```

3.4.2 GANs

BiGAN

```

1  # Copyright 2019 Zuru Tech HK Limited. All Rights Reserved.
2  #
3  # Licensed under the Apache License, Version 2.0 (the "License");
4  # you may not use this file except in compliance with the License.
5  # You may obtain a copy of the License at
6  #
7  # http://www.apache.org/licenses/LICENSE-2.0
8  #
9  # Unless required by applicable law or agreed to in writing, software
10 # distributed under the License is distributed on an "AS IS" BASIS,

```

(continues on next page)

(continued from previous page)

```

11 # WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.
12 # See the License for the specific language governing permissions and
13 # limitations under the License.
14
15 """Bigan dummy implementation."""
16
17 import operator
18
19 import tensorflow as tf
20 from tensorflow import keras
21
22 from ashpy.losses import DiscriminatorMinMax, EncoderBCE, GeneratorBCE
23 from ashpy.metrics import EncodingAccuracy
24 from ashpy.trainers import EncoderTrainer
25
26
27 def main():
28     """Main train loop and models definition."""
29
30     def real_gen():
31         """generator of real values."""
32         for _ in tf.range(100):
33             yield ((10.0,), (0,))
34
35     num_classes = 1
36     latent_dim = 100
37
38     generator = keras.Sequential([keras.layers.Dense(1)])
39
40     left_input = tf.keras.layers.Input(shape=(1,))
41     left = tf.keras.layers.Dense(10, activation=tf.nn.elu)(left_input)
42
43     right_input = tf.keras.layers.Input(shape=(latent_dim,))
44     right = tf.keras.layers.Dense(10, activation=tf.nn.elu)(right_input)
45
46     net = tf.keras.layers.Concatenate()([left, right])
47     out = tf.keras.layers.Dense(1)(net)
48
49     discriminator = tf.keras.Model(inputs=[left_input, right_input], outputs=[out])
50
51     encoder = keras.Sequential([keras.layers.Dense(latent_dim)])
52     generator_bce = GeneratorBCE()
53     encoder_bce = EncoderBCE()
54     minmax = DiscriminatorMinMax()
55
56     epochs = 100
57     logdir = "log/adversarial/encoder"
58
59     # Fake pre-trained classifier
60     classifier = tf.keras.Sequential(
61         [tf.keras.layers.Dense(10), tf.keras.layers.Dense(num_classes)]
62     )
63
64     metrics = [
65         EncodingAccuracy(
66             classifier, model_selection_operator=operator.gt, logdir=logdir
67         )

```

(continues on next page)

(continued from previous page)

```

68 ]
69
70 trainer = EncoderTrainer(
71     generator,
72     discriminator,
73     encoder,
74     tf.optimizers.Adam(1e-4),
75     tf.optimizers.Adam(1e-5),
76     tf.optimizers.Adam(1e-6),
77     generator_bce,
78     minmax,
79     encoder_bce,
80     epochs,
81     metrics=metrics,
82     logdir=logdir,
83 )
84
85 batch_size = 10
86 discriminator_input = tf.data.Dataset.from_generator(
87     real_gen, (tf.float32, tf.int64), ((1), (1))
88 ).batch(batch_size)
89
90 dataset = discriminator_input.map(
91     lambda x, y: ((x, y), tf.random.normal(shape=(batch_size, latent_dim)))
92 )
93
94 trainer(dataset)
95
96
97 if __name__ == "__main__":
98     main()

```

MNIST

```

1  # Copyright 2019 Zuru Tech HK Limited. All Rights Reserved.
2  #
3  # Licensed under the Apache License, Version 2.0 (the "License");
4  # you may not use this file except in compliance with the License.
5  # You may obtain a copy of the License at
6  #
7  # http://www.apache.org/licenses/LICENSE-2.0
8  #
9  # Unless required by applicable law or agreed to in writing, software
10 # distributed under the License is distributed on an "AS IS" BASIS,
11 # WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.
12 # See the License for the specific language governing permissions and
13 # limitations under the License.
14
15 """Adversarial trainer example."""
16
17 import tensorflow as tf
18 from tensorflow import keras # pylint: disable=no-name-in-module
19
20 from ashpy.losses import DiscriminatorMinMax, GeneratorBCE
21 from ashpy.metrics import InceptionScore

```

(continues on next page)

(continued from previous page)

```
22 from ashpy.models.gans import ConvDiscriminator, ConvGenerator
23 from ashpy.trainers import AdversarialTrainer
24
25
26 def main():
27     """Main."""
28
29     strategy = tf.distribute.MirroredStrategy()
30     with strategy.scope():
31
32         generator = ConvGenerator(
33             layer_spec_input_res=(7, 7),
34             layer_spec_target_res=(28, 28),
35             kernel_size=(5, 5),
36             initial_filters=256,
37             filters_cap=16,
38             channels=1,
39         )
40
41         discriminator = ConvDiscriminator(
42             layer_spec_input_res=(28, 28),
43             layer_spec_target_res=(7, 7),
44             kernel_size=(5, 5),
45             initial_filters=32,
46             filters_cap=128,
47             output_shape=1,
48         )
49
50         # Losses
51         generator_bce = GeneratorBCE()
52         minmax = DiscriminatorMinMax()
53
54         # Trainer
55         logdir = "log/adversarial"
56
57         # InceptionScore: keep commented until the issues
58         # https://github.com/tensorflow/tensorflow/issues/28599
59         # https://github.com/tensorflow/hub/issues/295
60         # Haven't been solved and merged into tf2
61
62         metrics = [
63             # InceptionScore(
64             #     InceptionScore.get_or_train_inception(
65             #         mnist_dataset,
66             #         "mnist",
67             #         num_classes=10,
68             #         epochs=1,
69             #         fine_tuning=False,
70             #         logdir=logdir,
71             #     ),
72             #     model_selection_operator=operator.gt,
73             #     logdir=logdir,
74             # )
75         ]
76
77         epochs = 50
78         trainer = AdversarialTrainer(
```

(continues on next page)

(continued from previous page)

```

79     generator,
80     discriminator,
81     tf.optimizers.Adam(1e-4),
82     tf.optimizers.Adam(1e-4),
83     generator_bce,
84     minmax,
85     epochs,
86     metrics,
87     logdir,
88 )
89
90 batch_size = 512
91
92 # Real data
93 mnist_x, mnist_y = keras.datasets.mnist.load_data()[0]
94
95 def iterator():
96     """Iterator in order to do not load in memory all the dataset."""
97     for image, label in zip(mnist_x, mnist_y):
98         yield tf.image.convert_image_dtype(
99             tf.expand_dims(image, -1), tf.float32
100             ), tf.expand_dims(label, -1)
101
102 real_data = (
103     tf.data.Dataset.from_generator(
104         iterator, (tf.float32, tf.int64), ((28, 28, 1), (1,))
105     )
106     .batch(batch_size)
107     .prefetch(1)
108 )
109
110 # Add noise in the same dataset, just by mapping.
111 # The return type of the dataset must be: tuple(tuple(a,b), noise)
112 dataset = real_data.map(
113     lambda x, y: ((x, y), tf.random.normal(shape=(batch_size, 100)))
114 )
115
116 trainer(dataset)
117
118
119 if __name__ == "__main__":
120     main()

```

Facades (Pix2Pix)

```

1 # Copyright 2019 Zuru Tech HK Limited. All Rights Reserved.
2 #
3 # Licensed under the Apache License, Version 2.0 (the "License");
4 # you may not use this file except in compliance with the License.
5 # You may obtain a copy of the License at
6 #
7 # http://www.apache.org/licenses/LICENSE-2.0
8 #
9 # Unless required by applicable law or agreed to in writing, software
10 # distributed under the License is distributed on an "AS IS" BASIS,

```

(continues on next page)

(continued from previous page)

```

11 # WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.
12 # See the License for the specific language governing permissions and
13 # limitations under the License.
14
15 """
16 Pix2Pix on Facades Datasets dummy implementation.
17 Input Pipeline taken from: https://www.tensorflow.org/beta/tutorials/generative/
18 ↪pix2pix
19 """
20 import os
21
22 import tensorflow as tf
23
24 from ashpy import LogEvalMode
25 from ashpy.losses.gan import Pix2PixLoss, get_adversarial_loss_discriminator
26 from ashpy.models.convolutional.discriminators import PatchDiscriminator
27 from ashpy.models.convolutional.unet import FUNet
28 from ashpy.trainers.gan import AdversarialTrainer
29
30 _URL = "https://people.eecs.berkeley.edu/~tinghuiz/projects/pix2pix/datasets/facades.
31 ↪tar.gz"
32
33 PATH_TO_ZIP = tf.keras.utils.get_file("facades.tar.gz", origin=_URL, extract=True)
34 PATH = os.path.join(os.path.dirname(PATH_TO_ZIP), "facades/")
35
36 BUFFER_SIZE = 100
37 BATCH_SIZE = 1
38 IMG_WIDTH = 256
39 IMG_HEIGHT = 256
40
41 def load(image_file):
42     """Load the image from file path."""
43     image = tf.io.read_file(image_file)
44     image = tf.image.decode_jpeg(image)
45
46     width = tf.shape(image)[1]
47
48     width = width // 2
49     real_image = image[:, :width, :]
50     input_image = image[:, width:, :]
51
52     input_image = tf.cast(input_image, tf.float32)
53     real_image = tf.cast(real_image, tf.float32)
54
55     return input_image, real_image
56
57 def resize(input_image, real_image, height, width):
58     """Resize input_image and real_image to height x width."""
59     input_image = tf.image.resize(
60         input_image, [height, width], method=tf.image.ResizeMethod.NEAREST_NEIGHBOR
61     )
62     real_image = tf.image.resize(
63         real_image, [height, width], method=tf.image.ResizeMethod.NEAREST_NEIGHBOR
64     )
65

```

(continues on next page)

(continued from previous page)

```

66     return input_image, real_image
67
68
69 def random_crop(input_image, real_image):
70     """Random crop both input_image and real_image."""
71     stacked_image = tf.stack([input_image, real_image], axis=0)
72     cropped_image = tf.image.random_crop(
73         stacked_image, size=[2, IMG_HEIGHT, IMG_WIDTH, 3]
74     )
75
76     return cropped_image[0], cropped_image[1]
77
78
79 def normalize(input_image, real_image):
80     """
81     Normalize images in [-1, 1].
82     """
83     input_image = (input_image / 127.5) - 1
84     real_image = (real_image / 127.5) - 1
85
86     return input_image, real_image
87
88
89 def load_image_train(image_file):
90     """Load and process the image_file to be ready for the training."""
91     input_image, real_image = load(image_file)
92     input_image, real_image = random_jitter(input_image, real_image)
93     input_image, real_image = normalize(input_image, real_image)
94
95     return input_image, real_image
96
97
98 @tf.function
99 def random_jitter(input_image, real_image):
100     """Apply random jitter to both input_image and real_image."""
101     # resizing to 286 x 286 x 3
102     input_image, real_image = resize(input_image, real_image, 286, 286)
103
104     # randomly cropping to 256 x 256 x 3
105     input_image, real_image = random_crop(input_image, real_image)
106
107     if tf.random.uniform(()) > 0.5:
108         # random mirroring
109         input_image = tf.image.flip_left_right(input_image)
110         real_image = tf.image.flip_left_right(real_image)
111
112     return input_image, real_image
113
114
115 def main(
116     kernel_size=5,
117     learning_rate_d=2e-4,
118     learning_rate_g=2e-4,
119     g_input_res=IMG_WIDTH,
120     g_min_res=1,
121     g_initial_filters=64,
122     g_filters_cap=512,

```

(continues on next page)

(continued from previous page)

```

123     use_dropout_encoder=False,
124     use_dropout_decoder=True,
125     d_target_res=32,
126     d_initial_filters=64,
127     d_filters_cap=512,
128     use_dropout_discriminator=False,
129     dataset_name="facades",
130     resolution=256,
131     epochs=100_000,
132     dropout_prob=0.3,
133     l1_loss_weight=100,
134     gan_loss_weight=1,
135     use_attention_d=False,
136     use_attention_g=False,
137     channels=3,
138     gan_loss_type=1,
139 ):
140     """Main train loop and models definition."""
141
142     generator = FUNet(
143         input_res=g_input_res,
144         min_res=g_min_res,
145         kernel_size=kernel_size,
146         initial_filters=g_initial_filters,
147         filters_cap=g_filters_cap,
148         channels=channels, # color_to_label_tensor.shape[0],
149         use_dropout_encoder=use_dropout_encoder,
150         use_dropout_decoder=use_dropout_decoder,
151         dropout_prob=dropout_prob,
152         use_attention=use_attention_g,
153     )
154     discriminator = PatchDiscriminator(
155         input_res=resolution,
156         min_res=d_target_res,
157         initial_filters=d_initial_filters,
158         kernel_size=kernel_size,
159         filters_cap=d_filters_cap,
160         use_dropout=use_dropout_discriminator,
161         dropout_prob=dropout_prob,
162         use_attention=use_attention_d,
163     )
164
165     discriminator_loss = get_adversarial_loss_discriminator(gan_loss_type)()
166     generator_loss = Pix2PixLoss(
167         l1_loss_weight=l1_loss_weight,
168         adversarial_loss_weight=gan_loss_weight,
169         adversarial_loss_type=gan_loss_type,
170     )
171
172     metrics = []
173     logdir = f'{"log"}/{dataset_name}/run2'
174
175     if not os.path.exists(logdir):
176         os.makedirs(logdir)
177
178     trainer = AdversarialTrainer(
179         generator=generator,

```

(continues on next page)

(continued from previous page)

```
180     discriminator=discriminator,
181     generator_optimizer=tf.optimizers.Adam(learning_rate_g, beta_1=0.5),
182     discriminator_optimizer=tf.optimizers.Adam(learning_rate_d, beta_1=0.5),
183     generator_loss=generator_loss,
184     discriminator_loss=discriminator_loss,
185     epochs=epochs,
186     metrics=metrics,
187     logdir=logdir,
188     log_eval_mode=LogEvalMode.TEST,
189 )
190
191 train_dataset = tf.data.Dataset.list_files(PATH + "train/*.jpg")
192 train_dataset = train_dataset.shuffle(BUFFER_SIZE)
193 train_dataset = train_dataset.map(load_image_train)
194 train_dataset = train_dataset.batch(BATCH_SIZE)
195
196 train_dataset = train_dataset.map(lambda x, y: ((y, x), x))
197
198 trainer(
199     # generator_input,
200     train_dataset
201 )
202
203
204 if __name__ == "__main__":
205     main()
```


API REFERENCE

<i>ashpy.ashtypes</i>	Custom Type-Aliases.
<i>ashpy.contexts</i>	Contexts help gaining an easier control over the model selection and testing process of the models.
<i>ashpy.keras</i>	Custom extensions of standard Keras components.
<i>ashpy.layers</i>	Collection of layers.
<i>ashpy.losses</i>	Collection of Losses.
<i>ashpy.metrics</i>	Collection of Metrics.
<i>ashpy.models</i>	Collection of Models.
<i>ashpy.modes</i>	Various modalities used to configure certain ash behaviours.
<i>ashpy.trainers</i>	Trainers help reducing boilerplate code by bootstrapping models training.

4.1 ashpy.ashtypes

Custom Type-Aliases. Mostly here for brevity in documentation.

```
TScalar = typing.Union[int, float]
```

```
TWeight = typing.Union[TScalar, tf.Tensor, typing.Callable[... , TScalar]]
```

4.2 ashpy.contexts

Contexts help gaining an easier control over the model selection and testing process of the models.

Classes

<i>base_context.BaseContext</i>	<i>ashpy.contexts.BaseContext</i> provide an interface for all contexts.
<i>classifier.ClassifierContext</i>	<i>ashpy.ClassifierContext</i> provide the standard functions to test a classifier.
<i>gan.GANContext</i>	<i>ashpy.contexts.gan.GANContext</i> measure the specified metrics on the GAN.
<i>gan.GANEncoderContext</i>	<i>ashpy.contexts.gan.GANEncoderContext</i> measure the specified metrics on the GAN.

4.2.1 BaseContext

Inheritance Diagram



```

class ashpy.contexts.base_context.BaseContext (metrics=None,          dataset=None,
                                                log_eval_mode=<LogEvalMode.TEST:
0>,          global_step=<tf.Variable
'global_step:0' shape=() dtype=int64,
numpy=0>, ckpt=None)
  
```

Bases: object

ashpy.contexts.BaseContext provide an interface for all contexts.

Methods

<code>__init__</code> ([metrics, dataset, log_eval_mode, ...])	Initialize the Context.
<code>measure_metrics</code> ()	Measure the metrics.
<code>model_selection</code> ()	Use the metrics to perform model selection.

Attributes

<code>dataset</code>	Retrieve the dataset.
<code>global_step</code>	Retrieve the global_step.
<code>log_eval_mode</code>	Retrieve model(s) mode.
<code>metrics</code>	Retrieve the metrics.

```

__init__(metrics=None,          dataset=None,          log_eval_mode=<LogEvalMode.TEST: 0>,
          global_step=<tf.Variable 'global_step:0' shape=() dtype=int64, numpy=0>, ckpt=None)
Initialize the Context.
  
```

Parameters

- **metrics** (list of [`ashpy.metrics.metric.Metric`]) – List of `ashpy.metrics.metric.Metric` objects.
- **dataset** (`tf.data.Dataset`) – The dataset to use, that contains everything needed to use the model in this context.
- **log_eval_mode** (`ashpy.modes.LogEvalMode`) – Models' mode to use when evaluating and logging.
- **global_step** (`tf.Variable`) – Keeps track of the training steps.
- **ckpt** (`tf.train.Checkpoint`) – Checkpoint to use to keep track of models status.

Return type None

`_validate_metrics()`

Check if every metric is an `ashpy.metrics.Metric`.

property `dataset`

Retrieve the dataset.

Return type `DatasetV2`

Returns `tf.data.Dataset`.

property `global_step`

Retrieve the `global_step`.

Return type `Variable`

Returns `tf.Variable`.

property `log_eval_mode`

Retrieve model(s) mode.

Return type `LogEvalMode`

Returns `ashpy.modes.LogEvalMode`.

`measure_metrics()`

Measure the metrics.

Return type None

property `metrics`

Retrieve the metrics.

Return type `List[Metric]`

Returns list of [`ashpy.metrics.metric.Metric`].

`model_selection()`

Use the metrics to perform model selection.

Return type None

4.2.2 ClassifierContext

Inheritance Diagram



```
class ashpy.contexts.classifier.ClassifierContext (classifier_model=None, loss=None,
dataset=None, metrics=None,
log_eval_mode=<LogEvalMode.TEST:
0>, global_step=<tf.Variable
'global_step:0' shape=()
dtype=int64, numpy=0>,
ckpt=None)
```

Bases: *ashpy.contexts.base_context.BaseContext*

ashpy.ClassifierContext provide the standard functions to test a classifier.

Methods

<code>__init__</code> (<i>classifier_model, loss, dataset, ...</i>)	Instantiate the <i>ashpy.contexts.classifier.ClassifierContext</i> context.
---	---

Attributes

<code>classifier_model</code>	Retrieve the Model Object.
<code>dataset</code>	Retrieve the dataset.
<code>global_step</code>	Retrieve the <code>global_step</code> .
<code>log_eval_mode</code>	Retrieve model(s) mode.
<code>loss</code>	Retrieve the loss value.
<code>metrics</code>	Retrieve the metrics.

```
__init__ (classifier_model=None, loss=None, dataset=None, metrics=None,
log_eval_mode=<LogEvalMode.TEST: 0>, global_step=<tf.Variable 'global_step:0'
shape=() dtype=int64, numpy=0>, ckpt=None)
Instantiate the ashpy.contexts.classifier.ClassifierContext context.
```

Parameters

- **classifier_model** (*tf.keras.Model*) – A *tf.keras.Model* model.
- **loss** (*ashpy.losses.Executor*) – Loss function, format `f(y_true, y_pred)`.
- **dataset** (*tf.data.Dataset*) – The test dataset.
- **metrics** (list of [*ashpy.metrics.metric.Metric*]) – List of *ashpy.metrics.metric.Metric* with which to measure training and validation data performances.
- **log_eval_mode** (*ashpy.modes.LogEvalMode*) – Models' mode to use when evaluating and logging.
- **global_step** (*tf.Variable*) – *tf.Variable* that keeps track of the training steps.
- **ckpt** (*tf.train.Checkpoint*) – checkpoint to use to keep track of models status.

Return type `None`

property classifier_model

Retrieve the Model Object.

Return type `Model`

Returns `tf.keras.Model`.

property loss

Retrieve the loss value.

Return type Optional[*Executor*]

4.2.3 GANContext

Inheritance Diagram



```

class ashpy.contexts.gan.GANContext (dataset=None, generator_model=None, dis-
    criminator_model=None, generator_loss=None,
    discriminator_loss=None, metrics=None,
    log_eval_mode=<LogEvalMode.TRAIN: 1>,
    global_step=<tf.Variable 'global_step:0' shape=()
    dtype=int64, numpy=0>, ckpt=None)
  
```

Bases: *ashpy.contexts.base_context.BaseContext*

ashpy.contexts.gan.GANContext measure the specified metrics on the GAN.

Methods

<code>__init__</code> ([dataset, generator_model, ...])	Initialize the Context.
---	-------------------------

Attributes

<code>dataset</code>	Retrieve the dataset.
<code>discriminator_loss</code>	Retrieve the discriminator loss.
<code>discriminator_model</code>	Retrieve the discriminator model.
<code>generator_loss</code>	Retrieve the generator loss.
<code>generator_model</code>	Retrieve the generator model.
<code>global_step</code>	Retrieve the global_step.
<code>log_eval_mode</code>	Retrieve model(s) mode.
<code>metrics</code>	Retrieve the metrics.

```

__init__ (dataset=None, generator_model=None, discriminator_model=None, generator_loss=None,
    discriminator_loss=None, metrics=None, log_eval_mode=<LogEvalMode.TRAIN: 1>,
    global_step=<tf.Variable 'global_step:0' shape=() dtype=int64, numpy=0>, ckpt=None)
  Initialize the Context.
  
```

Parameters

- **dataset** (`tf.data.Dataset`) – Dataset of tuples. [0] true dataset, [1] generator input dataset.
- **generator_model** (`tf.keras.Model`) – The generator.
- **discriminator_model** (`tf.keras.Model`) – The discriminator.
- **generator_loss** (`ashpy.losses.Executor()`) – The generator loss.
- **discriminator_loss** (`ashpy.losses.Executor()`) – The discriminator loss.
- **metrics** (list of [`ashpy.metrics.metric.Metric`]) – All the metrics to be used to evaluate the model.
- **log_eval_mode** (`ashpy.modes.LogEvalMode`) – Models' mode to use when evaluating and logging.
- **global_step** (`tf.Variable`) – *tf.Variable* that keeps track of the training steps.
- **ckpt** (`tf.train.Checkpoint`) – checkpoint to use to keep track of models status.

Return type `None`

property discriminator_loss

Retrieve the discriminator loss.

Return type `Optional[Executor]`

property discriminator_model

Retrieve the discriminator model.

Return type `Model`

Returns `tf.keras.Model`.

property generator_loss

Retrieve the generator loss.

Return type `Optional[Executor]`

property generator_model

Retrieve the generator model.

Return type `Model`

Returns `tf.keras.Model`.

4.2.4 GANEncoderContext

Inheritance Diagram



```
class ashpy.contexts.gan.GANEncoderContext (dataset=None, generator_model=None,
discriminator_model=None, encoder_model=None, generator_loss=None,
discriminator_loss=None, encoder_loss=None, metrics=None,
log_eval_mode=<LogEvalMode.TRAIN: 1>,
global_step=<tf.Variable 'global_step:0'
shape=() dtype=int64, numpy=0>,
ckpt=None)
```

Bases: *ashpy.contexts.gan.GANContext*

ashpy.contexts.gan.GANEncoderContext measure the specified metrics on the GAN.

Methods

<code>__init__</code> ([dataset, generator_model, ...])	Initialize the Context.
---	-------------------------

Attributes

<code>dataset</code>	Retrieve the dataset.
<code>discriminator_loss</code>	Retrieve the discriminator loss.
<code>discriminator_model</code>	Retrieve the discriminator model.
<code>encoder_loss</code>	Retrieve the encoder loss.
<code>encoder_model</code>	Retrieve the encoder model.
<code>generator_loss</code>	Retrieve the generator loss.
<code>generator_model</code>	Retrieve the generator model.
<code>global_step</code>	Retrieve the <code>global_step</code> .
<code>log_eval_mode</code>	Retrieve model(s) mode.
<code>metrics</code>	Retrieve the metrics.

```
__init__(dataset=None, generator_model=None, discriminator_model=None,
encoder_model=None, generator_loss=None, discriminator_loss=None, encoder_loss=None,
metrics=None, log_eval_mode=<LogEvalMode.TRAIN: 1>, global_step=<tf.Variable
'global_step:0' shape=() dtype=int64, numpy=0>, ckpt=None)
Initialize the Context.
```

Parameters

- **dataset** (`tf.data.Dataset`) – Dataset of tuples. [0] true dataset, [1] generator input dataset.
- **generator_model** (`tf.keras.Model`) – The generator.
- **discriminator_model** (`tf.keras.Model`) – The discriminator.
- **encoder_model** (`tf.keras.Model`) – The encoder.
- **generator_loss** (`ashpy.losses.Executor()`) – The generator loss.
- **discriminator_loss** (`ashpy.losses.Executor()`) – The discriminator loss.
- **encoder_loss** (`ashpy.losses.Executor()`) – The encoder loss.
- **metrics** (list of [`ashpy.metrics.metric.Metric`]) – All the metrics to be used to evaluate the model.

- **log_eval_mode** (*ashpy.modes.LogEvalMode*) – Models’ mode to use when evaluating and logging.
- **global_step** (*tf.Variable*) – *tf.Variable* that keeps track of the training steps.
- **ckpt** (*tf.train.Checkpoint*) – checkpoint to use to keep track of models status.

Return type `None`

property encoder_loss

Retrieve the encoder loss.

Return type `Optional[Executor]`

property encoder_model

Retrieve the encoder model.

Return type `Model`

Returns `tf.keras.Model`.

Modules

<i>base_context</i>	Primitive Context Interface.
<i>classifier</i>	Classifier Context.
<i>gan</i>	GANContext measures the specified metrics on the GAN.

4.2.5 base_context

Primitive Context Interface.

Contexts are checkpointable (subclassed from *tf.train.Checkpoint*) collections of variable encapsulated in a Python Class as a way to seamlessly handle information transfer.

Classes

<i>BaseContext</i>	<code>ashpy.contexts.BaseContext</code> provide an interface for all contexts.
--------------------	--

BaseContext

Inheritance Diagram



```
class ashpy.contexts.base_context.BaseContext (metrics=None, dataset=None,
                                             log_eval_mode=<LogEvalMode.TEST:
                                             0>, global_step=<tf.Variable
                                             'global_step:0' shape=() dtype=int64,
                                             numpy=0>, ckpt=None)
```

Bases: `object`

`ashpy.contexts.BaseContext` provide an interface for all contexts.

Methods

<code>__init__([metrics, dataset, log_eval_mode, ...])</code>	Initialize the Context.
<code>measure_metrics()</code>	Measure the metrics.
<code>model_selection()</code>	Use the metrics to perform model selection.

Attributes

<code>dataset</code>	Retrieve the dataset.
<code>global_step</code>	Retrieve the global_step.
<code>log_eval_mode</code>	Retrieve model(s) mode.
<code>metrics</code>	Retrieve the metrics.

```
__init__ (metrics=None, dataset=None, log_eval_mode=<LogEvalMode.TEST: 0>,
          global_step=<tf.Variable 'global_step:0' shape=() dtype=int64, numpy=0>, ckpt=None)
Initialize the Context.
```

Parameters

- **metrics** (list of [`ashpy.metrics.metric.Metric`]) – List of `ashpy.metrics.metric.Metric` objects.
- **dataset** (`tf.data.Dataset`) – The dataset to use, that contains everything needed to use the model in this context.
- **log_eval_mode** (`ashpy.modes.LogEvalMode`) – Models' mode to use when evaluating and logging.
- **global_step** (`tf.Variable`) – Keeps track of the training steps.
- **ckpt** (`tf.train.Checkpoint`) – Checkpoint to use to keep track of models status.

Return type `None`

```
_validate_metrics ()
Check if every metric is an ashpy.metrics.Metric.
```

```
property dataset
Retrieve the dataset.
```

Return type `DatasetV2`

Returns `tf.data.Dataset`.

```
property global_step
Retrieve the global_step.
```

Return type Variable

Returns `tf.Variable`.

property `log_eval_mode`

Retrieve model(s) mode.

Return type `LogEvalMode`

Returns `ashpy.modes.LogEvalMode`.

measure_metrics ()

Measure the metrics.

Return type None

property `metrics`

Retrieve the metrics.

Return type `List[Metric]`

Returns list of [`ashpy.metrics.metric.Metric`].

model_selection ()

Use the metrics to perform model selection.

Return type None

```
class ashpy.contexts.base_context.BaseContext (metrics=None, dataset=None,
                                              log_eval_mode=<LogEvalMode.TEST:
                                              0>, global_step=<tf.Variable
                                              'global_step:0' shape=() dtype=int64,
                                              numpy=0>, ckpt=None)
```

Bases: `object`

`ashpy.contexts.BaseContext` provide an interface for all contexts.

```
__init__ (metrics=None, dataset=None, log_eval_mode=<LogEvalMode.TEST: 0>,
          global_step=<tf.Variable 'global_step:0' shape=() dtype=int64, numpy=0>, ckpt=None)
Initialize the Context.
```

Parameters

- **metrics** (list of [`ashpy.metrics.metric.Metric`]) – List of `ashpy.metrics.metric.Metric` objects.
- **dataset** (`tf.data.Dataset`) – The dataset to use, that contains everything needed to use the model in this context.
- **log_eval_mode** (`ashpy.modes.LogEvalMode`) – Models' mode to use when evaluating and logging.
- **global_step** (`tf.Variable`) – Keeps track of the training steps.
- **ckpt** (`tf.train.Checkpoint`) – Checkpoint to use to keep track of models status.

Return type None

_validate_metrics ()

Check if every metric is an `ashpy.metrics.Metric`.

property `dataset`

Retrieve the dataset.

Return type `DatasetV2`

Returns `tf.data.Dataset`.

property global_step

Retrieve the global_step.

Return type Variable**Returns** `tf.Variable`.**property log_eval_mode**

Retrieve model(s) mode.

Return type `LogEvalMode`**Returns** `ashpy.modes.LogEvalMode`.**measure_metrics ()**

Measure the metrics.

Return type None**property metrics**

Retrieve the metrics.

Return type `List[Metric]`**Returns** list of `[ashpy.metrics.metric.Metric]`.**model_selection ()**

Use the metrics to perform model selection.

Return type None

4.2.6 classifier

Classifier Context.

Classes

ClassifierContext`ashpy.ClassifierContext` provide the standard functions to test a classifier.

ClassifierContext

Inheritance Diagram



```
class ashpy.contexts.classifier.ClassifierContext (classifier_model=None, loss=None,
dataset=None, metrics=None,
log_eval_mode=<LogEvalMode.TEST:
0>, global_step=<tf.Variable
'global_step:0' shape=()
dtype=int64, numpy=0>,
ckpt=None)
```

Bases: *ashpy.contexts.base_context.BaseContext*

ashpy.ClassifierContext provide the standard functions to test a classifier.

Methods

<code>__init__</code> (<i>classifier_model, loss, dataset, ...</i>)	Instantiate the <i>ashpy.contexts.classifier.ClassifierContext</i> context.
---	---

Attributes

<code>classifier_model</code>	Retrieve the Model Object.
<code>dataset</code>	Retrieve the dataset.
<code>global_step</code>	Retrieve the <code>global_step</code> .
<code>log_eval_mode</code>	Retrieve model(s) mode.
<code>loss</code>	Retrieve the loss value.
<code>metrics</code>	Retrieve the metrics.

```
__init__ (classifier_model=None, loss=None, dataset=None, metrics=None,
log_eval_mode=<LogEvalMode.TEST: 0>, global_step=<tf.Variable 'global_step:0'
shape=() dtype=int64, numpy=0>, ckpt=None)
Instantiate the ashpy.contexts.classifier.ClassifierContext context.
```

Parameters

- **classifier_model** (*tf.keras.Model*) – A *tf.keras.Model* model.
- **loss** (*ashpy.losses.Executor*) – Loss function, format `f(y_true, y_pred)`.
- **dataset** (*tf.data.Dataset*) – The test dataset.
- **metrics** (list of [*ashpy.metrics.metric.Metric*]) – List of *ashpy.metrics.metric.Metric* with which to measure training and validation data performances.
- **log_eval_mode** (*ashpy.modes.LogEvalMode*) – Models' mode to use when evaluating and logging.
- **global_step** (*tf.Variable*) – *tf.Variable* that keeps track of the training steps.
- **ckpt** (*tf.train.Checkpoint*) – checkpoint to use to keep track of models status.

Return type `None`

property classifier_model

Retrieve the Model Object.

Return type `Model`

Returns `tf.keras.Model`.

property loss

Retrieve the loss value.

Return type Optional[*Executor*]

```
class ashpy.contexts.classifier.ClassifierContext (classifier_model=None, loss=None,
                                                dataset=None, metrics=None,
                                                log_eval_mode=<LogEvalMode.TEST:
0>, global_step=<tf.Variable
'global_step:0' shape=()
dtype=int64, numpy=0>,
                                                ckpt=None)
```

Bases: *ashpy.contexts.base_context.BaseContext**ashpy.ClassifierContext* provide the standard functions to test a classifier.

```
__init__ (classifier_model=None, loss=None, dataset=None, metrics=None,
          log_eval_mode=<LogEvalMode.TEST: 0>, global_step=<tf.Variable 'global_step:0'
          shape=() dtype=int64, numpy=0>, ckpt=None)
```

Instantiate the *ashpy.contexts.classifier.ClassifierContext* context.**Parameters**

- **classifier_model** (*tf.keras.Model*) – A *tf.keras.Model* model.
- **loss** (*ashpy.losses.Executor*) – Loss function, format $f(y_{\text{true}}, y_{\text{pred}})$.
- **dataset** (*tf.data.Dataset*) – The test dataset.
- **metrics** (list of [*ashpy.metrics.metric.Metric*]) – List of *ashpy.metrics.metric.Metric* with which to measure training and validation data performances.
- **log_eval_mode** (*ashpy.modes.LogEvalMode*) – Models' mode to use when evaluating and logging.
- **global_step** (*tf.Variable*) – *tf.Variable* that keeps track of the training steps.
- **ckpt** (*tf.train.Checkpoint*) – checkpoint to use to keep track of models status.

Return type None**property classifier_model**

Retrieve the Model Object.

Return type Model**Returns** *tf.keras.Model*.**property loss**

Retrieve the loss value.

Return type Optional[*Executor*]

4.2.7 gan

GANContext measures the specified metrics on the GAN.

Classes

<code>GANContext</code>	<code>ashpy.contexts.gan.GANContext</code> measure the specified metrics on the GAN.
<code>GANEncoderContext</code>	<code>ashpy.contexts.gan.GANEncoderContext</code> measure the specified metrics on the GAN.

GANContext

Inheritance Diagram



```

class ashpy.contexts.gan.GANContext (dataset=None, generator_model=None, dis-
                                     criminator_model=None, generator_loss=None,
                                     discriminator_loss=None, metrics=None,
                                     log_eval_mode=<LogEvalMode.TRAIN: 1>,
                                     global_step=<tf.Variable 'global_step:0' shape=()
                                     dtype=int64, numpy=0>, ckpt=None)
Bases: ashpy.contexts.base_context.BaseContext
ashpy.contexts.gan.GANContext measure the specified metrics on the GAN.
  
```

Methods

<code>__init__</code> ([dataset, generator_model, ...])	Initialize the Context.
---	-------------------------

Attributes

<code>dataset</code>	Retrieve the dataset.
<code>discriminator_loss</code>	Retrieve the discriminator loss.
<code>discriminator_model</code>	Retrieve the discriminator model.
<code>generator_loss</code>	Retrieve the generator loss.
<code>generator_model</code>	Retrieve the generator model.
<code>global_step</code>	Retrieve the global_step.
<code>log_eval_mode</code>	Retrieve model(s) mode.
<code>metrics</code>	Retrieve the metrics.

```

__init__(dataset=None, generator_model=None, discriminator_model=None, generator_loss=None,
          discriminator_loss=None, metrics=None, log_eval_mode=<LogEvalMode.TRAIN: 1>,
          global_step=<tf.Variable 'global_step:0' shape=() dtype=int64, numpy=0>, ckpt=None)
Initialize the Context.
  
```

Parameters

- **dataset** (`tf.data.Dataset`) – Dataset of tuples. [0] true dataset, [1] generator input dataset.
- **generator_model** (`tf.keras.Model`) – The generator.
- **discriminator_model** (`tf.keras.Model`) – The discriminator.
- **generator_loss** (`ashpy.losses.Executor()`) – The generator loss.
- **discriminator_loss** (`ashpy.losses.Executor()`) – The discriminator loss.
- **metrics** (list of [`ashpy.metrics.metric.Metric`]) – All the metrics to be used to evaluate the model.
- **log_eval_mode** (`ashpy.modes.LogEvalMode`) – Models' mode to use when evaluating and logging.
- **global_step** (`tf.Variable`) – *tf.Variable* that keeps track of the training steps.
- **ckpt** (`tf.train.Checkpoint`) – checkpoint to use to keep track of models status.

Return type `None`

property discriminator_loss
Retrieve the discriminator loss.

Return type `Optional[Executor]`

property discriminator_model
Retrieve the discriminator model.

Return type `Model`

Returns `tf.keras.Model`.

property generator_loss
Retrieve the generator loss.

Return type `Optional[Executor]`

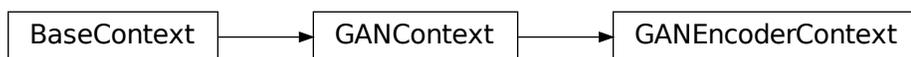
property generator_model
Retrieve the generator model.

Return type `Model`

Returns `tf.keras.Model`.

GANEncoderContext

Inheritance Diagram



```
class ashpy.contexts.gan.GANEncoderContext (dataset=None, generator_model=None,
discriminator_model=None, encoder_model=None, generator_loss=None,
discriminator_loss=None, encoder_loss=None, metrics=None,
log_eval_mode=<LogEvalMode.TRAIN: 1>,
global_step=<tf.Variable 'global_step:0'
shape=() dtype=int64, numpy=0>,
ckpt=None)
```

Bases: *ashpy.contexts.gan.GANContext*

ashpy.contexts.gan.GANEncoderContext measure the specified metrics on the GAN.

Methods

<code>__init__</code> ([dataset, generator_model, ...])	Initialize the Context.
---	-------------------------

Attributes

<code>dataset</code>	Retrieve the dataset.
<code>discriminator_loss</code>	Retrieve the discriminator loss.
<code>discriminator_model</code>	Retrieve the discriminator model.
<code>encoder_loss</code>	Retrieve the encoder loss.
<code>encoder_model</code>	Retrieve the encoder model.
<code>generator_loss</code>	Retrieve the generator loss.
<code>generator_model</code>	Retrieve the generator model.
<code>global_step</code>	Retrieve the <code>global_step</code> .
<code>log_eval_mode</code>	Retrieve model(s) mode.
<code>metrics</code>	Retrieve the metrics.

```
__init__ (dataset=None, generator_model=None, discriminator_model=None, en-
coder_model=None, generator_loss=None, discriminator_loss=None, encoder_loss=None,
metrics=None, log_eval_mode=<LogEvalMode.TRAIN: 1>, global_step=<tf.Variable
'global_step:0' shape=() dtype=int64, numpy=0>, ckpt=None)
Initialize the Context.
```

Parameters

- **dataset** (`tf.data.Dataset`) – Dataset of tuples. [0] true dataset, [1] generator input dataset.
- **generator_model** (`tf.keras.Model`) – The generator.
- **discriminator_model** (`tf.keras.Model`) – The discriminator.
- **encoder_model** (`tf.keras.Model`) – The encoder.
- **generator_loss** (`ashpy.losses.Executor()`) – The generator loss.
- **discriminator_loss** (`ashpy.losses.Executor()`) – The discriminator loss.
- **encoder_loss** (`ashpy.losses.Executor()`) – The encoder loss.
- **metrics** (list of [`ashpy.metrics.metric.Metric`]) – All the metrics to be used to evaluate the model.

- **log_eval_mode** (*ashpy.modes.LogEvalMode*) – Models’ mode to use when evaluating and logging.
- **global_step** (*tf.Variable*) – *tf.Variable* that keeps track of the training steps.
- **ckpt** (*tf.train.Checkpoint*) – checkpoint to use to keep track of models status.

Return type `None`

property encoder_loss

Retrieve the encoder loss.

Return type `Optional[Executor]`

property encoder_model

Retrieve the encoder model.

Return type `Model`

Returns `tf.keras.Model`.

```
class ashpy.contexts.gan.GANContext (dataset=None, generator_model=None, dis-
    criminator_model=None, generator_loss=None,
    discriminator_loss=None, metrics=None,
    log_eval_mode=<LogEvalMode.TRAIN: 1>,
    global_step=<tf.Variable 'global_step:0' shape=()
    dtype=int64, numpy=0>, ckpt=None)
```

Bases: *ashpy.contexts.base_context.BaseContext*

ashpy.contexts.gan.GANContext measure the specified metrics on the GAN.

```
__init__ (dataset=None, generator_model=None, discriminator_model=None, generator_loss=None,
    discriminator_loss=None, metrics=None, log_eval_mode=<LogEvalMode.TRAIN: 1>,
    global_step=<tf.Variable 'global_step:0' shape=() dtype=int64, numpy=0>, ckpt=None)
Initialize the Context.
```

Parameters

- **dataset** (*tf.data.Dataset*) – Dataset of tuples. [0] true dataset, [1] generator input dataset.
- **generator_model** (*tf.keras.Model*) – The generator.
- **discriminator_model** (*tf.keras.Model*) – The discriminator.
- **generator_loss** (*ashpy.losses.Executor()*) – The generator loss.
- **discriminator_loss** (*ashpy.losses.Executor()*) – The discriminator loss.
- **metrics** (list of [*ashpy.metrics.metric.Metric*]) – All the metrics to be used to evaluate the model.
- **log_eval_mode** (*ashpy.modes.LogEvalMode*) – Models’ mode to use when evaluating and logging.
- **global_step** (*tf.Variable*) – *tf.Variable* that keeps track of the training steps.
- **ckpt** (*tf.train.Checkpoint*) – checkpoint to use to keep track of models status.

Return type `None`

property discriminator_loss

Retrieve the discriminator loss.

Return type `Optional[Executor]`

property discriminator_model

Retrieve the discriminator model.

Return type Model

Returns `tf.keras.Model`.

property generator_loss

Retrieve the generator loss.

Return type Optional[*Executor*]

property generator_model

Retrieve the generator model.

Return type Model

Returns `tf.keras.Model`.

```
class ashpy.contexts.gan.GANEncoderContext (dataset=None, generator_model=None,
                                           discriminator_model=None, encoder_model=None,
                                           generator_loss=None, discriminator_loss=None,
                                           encoder_loss=None, metrics=None,
                                           log_eval_mode=<LogEvalMode.TRAIN: 1>,
                                           global_step=<tf.Variable 'global_step:0'
                                           shape=() dtype=int64, numpy=0>,
                                           ckpt=None)
```

Bases: `ashpy.contexts.gan.GANContext`

`ashpy.contexts.gan.GANEncoderContext` measure the specified metrics on the GAN.

```
__init__ (dataset=None, generator_model=None, discriminator_model=None, encoder_model=None,
          generator_loss=None, discriminator_loss=None, encoder_loss=None,
          metrics=None, log_eval_mode=<LogEvalMode.TRAIN: 1>, global_step=<tf.Variable
          'global_step:0' shape=() dtype=int64, numpy=0>, ckpt=None)
```

Initialize the Context.

Parameters

- **dataset** (`tf.data.Dataset`) – Dataset of tuples. [0] true dataset, [1] generator input dataset.
- **generator_model** (`tf.keras.Model`) – The generator.
- **discriminator_model** (`tf.keras.Model`) – The discriminator.
- **encoder_model** (`tf.keras.Model`) – The encoder.
- **generator_loss** (`ashpy.losses.Executor()`) – The generator loss.
- **discriminator_loss** (`ashpy.losses.Executor()`) – The discriminator loss.
- **encoder_loss** (`ashpy.losses.Executor()`) – The encoder loss.
- **metrics** (list of [`ashpy.metrics.metric.Metric`]) – All the metrics to be used to evaluate the model.
- **log_eval_mode** (`ashpy.modes.LogEvalMode`) – Models' mode to use when evaluating and logging.
- **global_step** (`tf.Variable`) – `tf.Variable` that keeps track of the training steps.
- **ckpt** (`tf.train.Checkpoint`) – checkpoint to use to keep track of models status.

Return type None

property encoder_loss

Retrieve the encoder loss.

Return type Optional[*Executor*]

property encoder_model

Retrieve the encoder model.

Return type Model

Returns `tf.keras.Model`.

4.3 ashpy.keras

Custom extensions of standard Keras components.

Modules

<i>losses</i>	Custom Keras losses, used by the AshPy executors.
---------------	---

4.3.1 losses

Custom Keras losses, used by the AshPy executors.

Classes

<i>DHingeLoss</i>	Discriminator Hinge Loss as Keras Metric.
<i>DLeastSquare</i>	Discriminator Least Square Loss as <code>tf.keras.losses.Loss</code> .
<i>DMinMax</i>	Implementation of MinMax Discriminator loss as <code>tf.keras.losses.Loss</code> .
<i>GHingeLoss</i>	Generator Hinge Loss as Keras Metric.
<i>L1</i>	L1 Loss implementation as <code>tf.keras.losses.Loss</code> .

DHingeLoss

Inheritance Diagram



class `ashpy.keras.losses.DHingeLoss`

Bases: `tensorflow.python.keras.losses.Loss`

Discriminator Hinge Loss as Keras Metric. See Geometric GAN¹ for more details.

The Discriminator Hinge loss is the hinge version of the adversarial loss. The Hinge loss is defined as:

$$L_{\text{hinge}} = \max(0, 1 - ty)$$

where y is the Discriminator output and t is the target class (+1 or -1 in the case of binary classification).

For the case of GANs:

$$L_{D_{\text{hinge}}} = -\mathbb{E}_{(x,y) \sim p_{\text{data}}}[\min(0, -1 + D(x, y))] - \mathbb{E}_{x \sim p_x, y \sim p_{\text{data}}}[\min(0, -1 - D(G(z), y))]$$

Methods

<code>__init__()</code>	Initialize the Loss.
<code>call(d_real, d_fake)</code>	Compute the hinge loss

Attributes

<code>reduction</code>	Return the current <i>reduction</i> for this type of loss.
------------------------	--

`__init__()`

Initialize the Loss.

Return type None

`call(d_real, d_fake)`

Compute the hinge loss

Return type Tensor

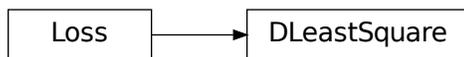
property `reduction`

Return the current *reduction* for this type of loss.

Return type `ReductionV2`

DLeastSquare

Inheritance Diagram



¹ Geometric GAN <https://arxiv.org/abs/1705.02894>

class `ashpy.keras.losses.DLeastSquare`

Bases: `tensorflow.python.keras.losses.Loss`

Discriminator Least Square Loss as `tf.keras.losses.Loss`.

Methods

<code>__init__()</code>	Least square Loss for Discriminator.
<code>call(d_real, d_fake)</code>	Compute the Least Square Loss.

Attributes

<code>reduction</code>	Return the reduction type for this loss.
------------------------	--

`__init__()`

Least square Loss for Discriminator.

Reference: Least Squares Generative Adversarial Networks¹.

Basically the Mean Squared Error between the discriminator output when evaluated in fake samples and 0 and the discriminator output when evaluated in real samples and 1: For the unconditioned case this is:

$$L_D = \frac{1}{2} E[(D(x) - 1)^2 + (0 - D(G(z)))^2]$$

where x are real samples and z is the latent vector.

For the conditioned case this is:

$$L_D = \frac{1}{2} E[(D(x, c) - 1)^2 + (0 - D(G(c), c))^2]$$

where c is the condition and x are real samples.

Return type `None`

call (`d_real`, `d_fake`)

Compute the Least Square Loss.

Parameters

- **d_real** (`tf.Tensor`) – Discriminator evaluated in real samples.
- **d_fake** (`tf.Tensor`) – Discriminator evaluated in fake samples.

Return type `Tensor`

Returns `tf.Tensor` – Loss.

property `reduction`

Return the reduction type for this loss.

Return type `ReductionV2`

Returns `tf.keras.losses.Reduction` – Reduction.

¹ Least Squares Generative Adversarial Networks <https://arxiv.org/abs/1611.04076>

DMinMax

Inheritance Diagram



class `ashpy.keras.losses.DMinMax` (*from_logits=True, label_smoothing=0.0*)

Bases: `tensorflow.python.keras.losses.Loss`

Implementation of MinMax Discriminator loss as `tf.keras.losses.Loss`.

$$L_D = -\frac{1}{2} E[\log(D(x)) + \log(1 - D(G(z)))]$$

Methods

<code>__init__</code> (<i>from_logits, label_smoothing</i>)	Initialize the loss.
<code>call</code> (<i>d_real, d_fake</i>)	Compute the MinMax Loss.

Attributes

<code>reduction</code>	Return the reduction type of this loss.
------------------------	---

`__init__` (*from_logits=True, label_smoothing=0.0*)

Initialize the loss.

Return type `None`

`call` (*d_real, d_fake*)

Compute the MinMax Loss.

Play the DiscriminatorMinMax game between the discriminator computed in real and the discriminator compute with fake inputs.

Parameters

- **d_real** (`tf.Tensor`) – Real data.
- **d_fake** (`tf.Tensor`) – Fake (generated) data.

Return type `Tensor`

Returns `tf.Tensor` – Output Tensor.

property `reduction`

Return the reduction type of this loss.

Return type `ReductionV2`

Returns `:py:classes:'tf.keras.losses.Reduction'` – Reduction.

GHingeLoss

Inheritance Diagram



class `ashpy.keras.losses.GHingeLoss`

Bases: `tensorflow.python.keras.losses.Loss`

Generator Hinge Loss as Keras Metric. See Geometric GAN¹ for more details.

The Generator Hinge loss is the hinge version of the adversarial loss. The Hinge loss is defined as:

$$L_{\text{hinge}} = \max(0, 1 - ty)$$

where y is the Discriminator output and t is the target class (+1 or -1 in the case of binary classification). The target class of the generated images is +1.

For the case of GANs

$$L_{G_{\text{hinge}}} = -\mathbb{E}_{(x \sim p_x, y \sim p_{data})}[\min(0, -1 + D(G(x), y))]$$

This can be simply approximated as:

$$L_{G_{\text{hinge}}} = -\mathbb{E}_{(x \sim p_x, y \sim p_{data})}[D(G(x), y)]$$

Methods

<code>__init__()</code>	Initialize the Loss.
<code>call(d_real, d_fake)</code>	Computes the hinge loss

Attributes

<code>reduction</code>	Return the current <i>reduction</i> for this type of loss.
------------------------	--

`__init__()`
Initialize the Loss.

Return type None

`call(d_real, d_fake)`

¹ Geometric GAN <https://arxiv.org/abs/1705.02894>

Computes the hinge loss

Return type Tensor

property reduction

Return the current *reduction* for this type of loss.

Return type ReductionV2

L1

Inheritance Diagram



class `ashpy.keras.losses.L1`

Bases: `tensorflow.python.keras.losses.Loss`

L1 Loss implementation as `tf.keras.losses.Loss`.

Methods

<code>__init__()</code>	Initialize the Loss.
<code>call(x, y)</code>	Compute the mean of the l1 between x and y.

Attributes

<code>reduction</code>	Return the current <i>reduction</i> for this type of loss.
------------------------	--

`__init__()`
Initialize the Loss.

Return type None

`call(x, y)`
Compute the mean of the l1 between x and y.

Return type Tensor

property reduction
Return the current *reduction* for this type of loss.

Return type ReductionV2

class `ashpy.keras.losses.DHingeLoss`

Bases: `tensorflow.python.keras.losses.Loss`

Discriminator Hinge Loss as Keras Metric. See Geometric GAN [1]_ for more details.

The Discriminator Hinge loss is the hinge version of the adversarial loss. The Hinge loss is defined as:

$$L_{\text{hinge}} = \max(0, 1 - ty)$$

where y is the Discriminator output and t is the target class (+1 or -1 in the case of binary classification).

For the case of GANs:

$$L_{D_{\text{hinge}}} = -\mathbb{E}_{(x,y) \sim p_{\text{data}}}[\min(0, -1 + D(x, y))] - \mathbb{E}_{x \sim p_x, y \sim p_{\text{data}}}[\min(0, -1 - D(G(z), y))]$$

__init__ ()

Initialize the Loss.

Return type None

call (*d_real*, *d_fake*)

Compute the hinge loss

Return type Tensor

property reduction

Return the current *reduction* for this type of loss.

Return type ReductionV2

class `ashpy.keras.losses.DLeastSquare`

Bases: `tensorflow.python.keras.losses.Loss`

Discriminator Least Square Loss as `tf.keras.losses.Loss`.

__init__ ()

Least square Loss for Discriminator.

Reference: Least Squares Generative Adversarial Networks [1]_.

Basically the Mean Squared Error between the discriminator output when evaluated in fake samples and 0 and the discriminator output when evaluated in real samples and 1: For the unconditioned case this is:

$$L_D = \frac{1}{2} E[(D(x) - 1)^2 + (0 - D(G(z)))^2]$$

where x are real samples and z is the latent vector.

For the conditioned case this is:

$$L_D = \frac{1}{2} E[(D(x, c) - 1)^2 + (0 - D(G(c), c))^2]$$

where c is the condition and x are real samples.

Return type None

call (*d_real*, *d_fake*)

Compute the Least Square Loss.

Parameters

- **d_real** (`tf.Tensor`) – Discriminator evaluated in real samples.
- **d_fake** (`tf.Tensor`) – Discriminator evaluated in fake samples.

Return type Tensor

Returns `tf.Tensor` – Loss.

property reduction

Return the reduction type for this loss.

Return type `ReductionV2`

Returns `tf.keras.losses.Reduction` – Reduction.

class `ashpy.keras.losses.DMinMax` (*from_logits=True, label_smoothing=0.0*)

Bases: `tensorflow.python.keras.losses.Loss`

Implementation of MinMax Discriminator loss as `tf.keras.losses.Loss`.

$$L_D = -\frac{1}{2}E[\log(D(x)) + \log(1 - D(G(z)))]$$

__init__ (*from_logits=True, label_smoothing=0.0*)

Initialize the loss.

Return type `None`

call (*d_real, d_fake*)

Compute the MinMax Loss.

Play the DiscriminatorMinMax game between the discriminator computed in real and the discriminator compute with fake inputs.

Parameters

- **d_real** (`tf.Tensor`) – Real data.
- **d_fake** (`tf.Tensor`) – Fake (generated) data.

Return type `Tensor`

Returns `tf.Tensor` – Output Tensor.

property reduction

Return the reduction type of this loss.

Return type `ReductionV2`

Returns **:py:classes:'tf.keras.losses.Reduction'** – Reduction.

class `ashpy.keras.losses.GHingeLoss`

Bases: `tensorflow.python.keras.losses.Loss`

Generator Hinge Loss as Keras Metric. See Geometric GAN [1]_ for more details.

The Generator Hinge loss is the hinge version of the adversarial loss. The Hinge loss is defined as:

$$L_{\text{hinge}} = \max(0, 1 - ty)$$

where y is the Discriminator output and t is the target class (+1 or -1 in the case of binary classification). The target class of the generated images is +1.

For the case of GANs

$$L_{G_{\text{hinge}}} = -\mathbb{E}_{(x \sim p_x, y \sim p_{data})}[\min(0, -1 + D(G(x), y))]$$

This can be simply approximated as:

$$L_{G_{\text{hinge}}} = -\mathbb{E}_{(x \sim p_x, y \sim p_{data})}[D(G(x), y)]$$

```

__init__()
    Initialize the Loss.

    Return type None

call(d_real, d_fake)
    Computes the hinge loss

    Return type Tensor

property reduction
    Return the current reduction for this type of loss.

    Return type ReductionV2

class ashpy.keras.losses.L1
    Bases: tensorflow.python.keras.losses.Loss

    L1 Loss implementation as tf.keras.losses.Loss.

    __init__()
        Initialize the Loss.

        Return type None

    call(x, y)
        Compute the mean of the l1 between x and y.

        Return type Tensor

    property reduction
        Return the current reduction for this type of loss.

        Return type ReductionV2

```

4.4 ashpy.layers

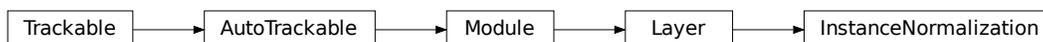
Collection of layers.

Layers

<code>instance_normalization.</code>	Instance Normalization Layer (used by Pix2Pix [1]_ and Pix2PixHD [2]_).
<code>InstanceNormalization</code>	
<code>attention.Attention</code>	Attention Layer from Self-Attention GAN [1]_.

4.4.1 InstanceNormalization

Inheritance Diagram



```
class ashpy.layers.instance_normalization.InstanceNormalization(eps=1e-06,
                                                             beta_initializer='zeros',
                                                             gamma_initializer='ones')
```

Bases: tensorflow.python.keras.engine.base_layer.Layer

Instance Normalization Layer (used by Pix2Pix¹ and Pix2PixHD²).

Basically it's a normalization done at instance level. The implementation follows the basic implementation of the Batch Normalization Layer.

- Direct Usage:

```
x = tf.ones((1, 10, 10, 64))

# instantiate attention layer as model.
normalization = InstanceNormalization()

# evaluate passing x.
output = normalization(x)

# the output shape is.
# the same as the input shape.
print(output.shape)
```

- Inside a Model:

```
def MyModel():
    inputs = tf.keras.layers.Input(shape=[None, None, 64])
    normalization = InstanceNormalization()
    return tf.keras.Model(inputs=inputs,
    ↪ outputs=normalization(inputs))

x = tf.ones((1, 10, 10, 64))
model = MyModel()
output = model(x)

print(output.shape)
```

```
(1, 10, 10, 64)
```

Methods

<code>__init__(<i>eps, beta_initializer, ...</i>)</code>	Initialize the layer.
<code>build(<i>input_shape</i>)</code>	Assemble the layer.
<code>call(<i>inputs[, training]</i>)</code>	Perform the computation.

Attributes

¹ Image-to-Image Translation with Conditional Adversarial Networks <https://arxiv.org/abs/1611.07004>

² High-Resolution Image Synthesis and Semantic Manipulation with Conditional GANs <https://arxiv.org/abs/1711.11585>

<code>activity_regularizer</code>	Optional regularizer function for the output of this layer.
<code>dtype</code>	
<code>dynamic</code>	
<code>inbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>input</code>	Retrieves the input tensor(s) of a layer.
<code>input_mask</code>	Retrieves the input mask tensor(s) of a layer.
<code>input_shape</code>	Retrieves the input shape(s) of a layer.
<code>input_spec</code>	
<code>losses</code>	Losses which are associated with this <i>Layer</i> .
<code>metrics</code>	
<code>name</code>	Returns the name of this module as passed or determined in the ctor.
<code>name_scope</code>	Returns a <i>tf.name_scope</i> instance for this class.
<code>non_trainable_variables</code>	
<code>non_trainable_weights</code>	
<code>outbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>output</code>	Retrieves the output tensor(s) of a layer.
<code>output_mask</code>	Retrieves the output mask tensor(s) of a layer.
<code>output_shape</code>	Retrieves the output shape(s) of a layer.
<code>submodules</code>	Sequence of all sub-modules.
<code>trainable</code>	
<code>trainable_variables</code>	Sequence of variables owned by this module and it's submodules.
<code>trainable_weights</code>	
<code>updates</code>	
<code>variables</code>	Returns the list of all layer variables/weights.
<code>weights</code>	Returns the list of all layer variables/weights.

`__init__` (*eps=1e-06*, *beta_initializer='zeros'*, *gamma_initializer='ones'*)
 Initialize the layer.

Parameters

- **eps** (*float*) – Variance_epsilon used by batch_norm layer.
- **beta_initializer** (*str*) – Initializer for the beta variable.
- **gamma_initializer** (*str*) – Initializer for the gamma variable.

Return type None

build (*input_shape*)
 Assemble the layer.

Parameters **input_shape** (*tuple of (int)*) – Specifies the shape of the input accepted by the layer.

Return type None

call (*inputs*, *training=False*)
 Perform the computation.

Parameters

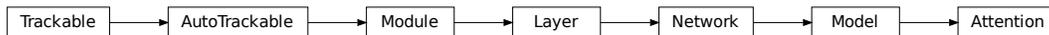
- **inputs** (`tf.Tensor`) – Inputs for the computation.
- **training** (`bool`) – Controls for training or evaluation mode.

Return type `Tensor`

Returns `tf.Tensor` – Output Tensor.

4.4.2 Attention

Inheritance Diagram



class `ashpy.layers.attention.Attention` (*filters*)

Bases: `tensorflow.python.keras.engine.training.Model`

Attention Layer from Self-Attention GAN¹.

First we extract features from the previous layer:

$$f(x) = W_f x$$

$$g(x) = W_g x$$

$$h(x) = W_h x$$

Then we calculate the importance matrix:

$$\beta_{j,i} = \frac{\exp(s_{i,j})}{\sum_{i=1}^N \exp(s_{ij})}$$

$\beta_{j,i}$ indicates the extent to which the model attends to the i^{th} location when synthethizing the j^{th} region.

Then we calculate the output of the attention layer $(o_1, \dots, o_N) \in \mathbb{R}^{C \times N}$:

$$o_j = \sum_{i=1}^N \beta_{j,i} h(x_i)$$

Finally we combine the (scaled) attention and the input to get the final output of the layer:

$$y_i = \gamma o_i + x_i$$

where γ is initialized as 0.

Examples

- Direct Usage:

¹ Self-Attention Generative Adversarial Networks <https://arxiv.org/abs/1805.08318>

```
x = tf.ones((1, 10, 10, 64))

# instantiate attention layer as model
attention = Attention(64)

# evaluate passing x
output = attention(x)

# the output shape is
# the same as the input shape
print(output.shape)
```

- Inside a Model:

```
def MyModel():
    inputs = tf.keras.layers.Input(shape=[None, None, 64])
    attention = Attention(64)
    return tf.keras.Model(inputs=inputs, outputs=attention(inputs))

x = tf.ones((1, 10, 10, 64))
model = MyModel()
output = model(x)

print(output.shape)
```

```
(1, 10, 10, 64)
```

Methods

<code>__init__(filters)</code>	Build the Attention Layer.
<code>call(inputs[, training])</code>	Perform the computation.

Attributes

<code>activity_regularizer</code>	Optional regularizer function for the output of this layer.
<code>dtype</code>	
<code>dynamic</code>	
<code>inbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>input</code>	Retrieves the input tensor(s) of a layer.
<code>input_mask</code>	Retrieves the input mask tensor(s) of a layer.
<code>input_shape</code>	Retrieves the input shape(s) of a layer.
<code>input_spec</code>	Gets the network's input specs.
<code>layers</code>	
<code>losses</code>	Losses which are associated with this <i>Layer</i> .
<code>metrics</code>	Returns the model's metrics added using <i>compile</i> , <i>add_metric</i> APIs.
<code>metrics_names</code>	Returns the model's display labels for all outputs.
<code>name</code>	Returns the name of this module as passed or determined in the ctor.

Continued on next page

Table 39 – continued from previous page

<code>name_scope</code>	Returns a <i>tf.name_scope</i> instance for this class.
<code>non_trainable_variables</code>	
<code>non_trainable_weights</code>	
<code>outbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>output</code>	Retrieves the output tensor(s) of a layer.
<code>output_mask</code>	Retrieves the output mask tensor(s) of a layer.
<code>output_shape</code>	Retrieves the output shape(s) of a layer.
<code>run_eagerly</code>	Settable attribute indicating whether the model should run eagerly.
<code>sample_weights</code>	
<code>state_updates</code>	Returns the <i>updates</i> from all layers that are stateful.
<code>stateful</code>	
<code>submodules</code>	Sequence of all sub-modules.
<code>trainable</code>	
<code>trainable_variables</code>	Sequence of variables owned by this module and it's submodules.
<code>trainable_weights</code>	
<code>updates</code>	
<code>variables</code>	Returns the list of all layer variables/weights.
<code>weights</code>	Returns the list of all layer variables/weights.

`__init__` (*filters*)

Build the Attention Layer.

Parameters `filters` (*int*) – Number of filters of the input tensor. It should be preferably a multiple of 8.

Return type None

call (*inputs, training=False*)

Perform the computation.

Parameters

- **inputs** (*tf.Tensor*) – Inputs for the computation.
- **training** (*bool*) – Controls for training or evaluation mode.

Return type Tensor

Returns *tf.Tensor* – Output Tensor.

4.5 ashpy.losses

Collection of Losses.

Executor

<code>executor.Executor</code>	Carry a function and the way of executing it.
<code>executor.SumExecutor</code>	The sum executor.

4.5.1 Executor

Inheritance Diagram



class `ashpy.losses.executor.Executor` (*fn=None*)

Bases: `object`

Carry a function and the way of executing it. Given a context.

Methods

<code>__init__(fn)</code>	Initialize the Executor.
<code>call(context, **kwargs)</code>	Execute the function, using the information provided by the context.
<code>reduce_loss(call_fn)</code>	Create a Decorator to reduce Losses.

Attributes

<code>fn</code>	Return the Keras loss function to execute.
<code>global_batch_size</code>	Global batch size comprises the batch size for each cpu.
<code>weight</code>	Return the loss weight.

__init__ (*fn=None*)
Initialize the Executor.

Parameters `fn` (`tf.keras.losses.Loss`) – A Keras Loss to execute.

Return type `None`

Returns `None`

abstract call (*context, **kwargs*)
Execute the function, using the information provided by the context.

Parameters `context` (`ashpy.contexts.BaseContext`) – The function execution Context.

Return type `Tensor`

Returns `tf.Tensor` – Output Tensor.

property fn
Return the Keras loss function to execute.

Return type `Loss`

Returns `tf.keras.losses.Loss` – Keras Loss.

property `global_batch_size`

Global batch size comprises the batch size for each cpu.

Calculated as `batch_size_for_replica*replica_numbers`.

Return type `int`

Returns `int` – Global Batch size value.

static `reduce_loss` (*call_fn*)

Create a Decorator to reduce Losses. Used to simplify things.

Apply a `reduce` `sum` operation to the loss and divide the result by the batch size.

Parameters `call_fn` (`typing.Callable`) – The executor call method.

Return type `Callable`

Returns `typing.Callable` – The decorated function.

property `weight`

Return the loss weight.

This weight is multiplied by the loss value. This is useful when working with multiples losses.

Return type `Callable[...float]`

Returns `typing.Callable` – Callable returning the weight (`float`).

4.5.2 SumExecutor

Inheritance Diagram



class `ashpy.losses.executor.SumExecutor` (*executors*)

Bases: `ashpy.losses.executor.Executor`

The sum executor. Executes the call of each fn and weights the losses.

Each Executor gets called (thus reducing its carried function), the results are then summed together.

Methods

`__init__`(*executors*)

Initialize the SumExecutor.

`call`(*args, **kwargs)

Evaluate and sum together the Executors.

Attributes

<code>executors</code>	Return the List of Executors.
<code>fn</code>	Return the Keras loss function to execute.
<code>global_batch_size</code>	Global batch size comprises the batch size for each cpu.
<code>weight</code>	Return the loss weight.

`__init__(executors)`

Initialize the SumExecutor.

Parameters `executors` (`list` of [`ashpy.executors.Executor`]) – Array of `ashpy.executors.Executor` to sum evaluate and sum together.

Return type `None`

Returns `None`

`call(*args, **kwargs)`

Evaluate and sum together the Executors.

Return type `Tensor`

Returns `:py:classes:'tf.Tensor'` – Output Tensor.

property `executors`

Return the List of Executors.

Return type `List[Executor]`

property `global_batch_size`

Global batch size comprises the batch size for each cpu.

Calculated as `batch_size_for_replica*replica_numbers`.

Return type `int`

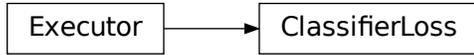
Returns `int` – Global Batch size value.

Classifier

<code>classifier.ClassifierLoss</code>	Classifier Loss Executor using the classifier model, instantiated with a fn.
--	--

4.5.3 ClassifierLoss

Inheritance Diagram



class `ashpy.losses.classifier.ClassifierLoss` (*fn*)

Bases: `ashpy.losses.executor.Executor`

Classifier Loss Executor using the classifier model, instantiated with a fn.

Methods

<code>__init__(fn)</code>	Initialize <i>ClassifierLoss</i> .
<code>call(*args, **kwargs)</code>	

Attributes

<code>fn</code>	Return the Keras loss function to execute.
<code>global_batch_size</code>	Global batch size comprises the batch size for each cpu.
<code>weight</code>	Return the loss weight.

`__init__(fn)`

Initialize *ClassifierLoss*.

Parameters `fn` (`tf.keras.losses.Loss`) – Classification Loss function, should take as input labels and prediction.

Return type None

Returns None

GAN

<code>gan.GANExecutor</code>	Executor for GANs.
<code>gan.AdversarialLossType</code>	Enumeration for Adversarial Losses.
<code>gan.GeneratorAdversarialLoss</code>	Base class for the adversarial loss of the generator.
<code>gan.DiscriminatorAdversarialLoss</code>	Base class for the adversarial loss of the discriminator.
<code>gan.GeneratorBCE</code>	The Binary CrossEntropy computed among the generator and the 1 label.
<code>gan.GeneratorLSGAN</code>	Least Square GAN Loss for generator.

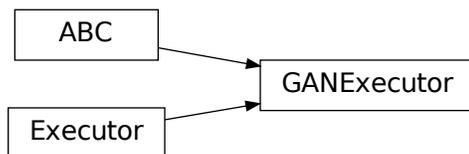
Continued on next page

Table 48 – continued from previous page

<code>gan.GeneratorL1</code>	L1 loss between the generator output and the target.
<code>gan.GeneratorHingeLoss</code>	Hinge loss for the Generator.
<code>gan.FeatureMatchingLoss</code>	Conditional GAN Feature matching loss.
<code>gan.CategoricalCrossEntropy</code>	Categorical Cross Entropy between generator output and target.
<code>gan.Pix2PixLoss</code>	Pix2Pix Loss.
<code>gan.Pix2PixLossSemantic</code>	Semantic Pix2Pix Loss.
<code>gan.EncoderBCE</code>	The Binary Cross Entropy computed among the encoder and the 0 label.
<code>gan.DiscriminatorMinMax</code>	The min-max game played by the discriminator.
<code>gan.DiscriminatorLSGAN</code>	Least square Loss for discriminator.
<code>gan.DiscriminatorHingeLoss</code>	Hinge loss for the Discriminator.
<code>gan.get_adversarial_loss_discriminator</code>	Return the correct loss for the Discriminator.
<code>gan.get_adversarial_loss_generator</code>	Return the correct loss for the Generator.

4.5.4 GANExecutor

Inheritance Diagram



class `ashpy.losses.gan.GANExecutor` (*fn=None*)
 Bases: `ashpy.losses.executor.Executor`, `abc.ABC`

Executor for GANs.

Implements the basic functions needed by the GAN losses.

Methods

<code>__init__</code> ([fn])	Initialize the GANExecutor.
<code>call</code> (context, **kwargs)	Execute the function, using the information provided by the context.
<code>get_discriminator_inputs</code> (context, ...)	Return the discriminator inputs.

Attributes

<code>fn</code>	Return the Keras loss function to execute.
<code>global_batch_size</code>	Global batch size comprises the batch size for each cpu.
<code>weight</code>	Return the loss weight.

`__init__` (*fn=None*)

Initialize the GANExecutor.

Parameters `fn` (`tf.keras.losses.Loss`) – A Keras Loss to execute.

Returns `None`

abstract call (*context, **kwargs*)

Execute the function, using the information provided by the context.

Parameters `context` (`ashpy.contexts.BaseContext`) – The function execution Context.

Returns `tf.Tensor` – Output Tensor.

static get_discriminator_inputs (*context, fake_or_real, condition, training*)

Return the discriminator inputs. If needed it uses the encoder.

The current implementation uses the number of inputs to determine whether the discriminator is conditioned or not.

Parameters

- **context** (`ashpy.contexts.gan.GANContext`) – Context for GAN models.
- **fake_or_real** (`tf.Tensor`) – Discriminator input tensor, it can be fake (generated) or real.
- **condition** (`tf.Tensor`) – Discriminator condition (it can also be generator noise).
- **training** (`bool`) – whether is training phase or not

Return type `Union[Tensor, List[Tensor]]`

Returns The discriminator inputs.

4.5.5 AdversarialLossType

Inheritance Diagram



class `ashpy.losses.gan.AdversarialLossType`

Bases: `enum.Enum`

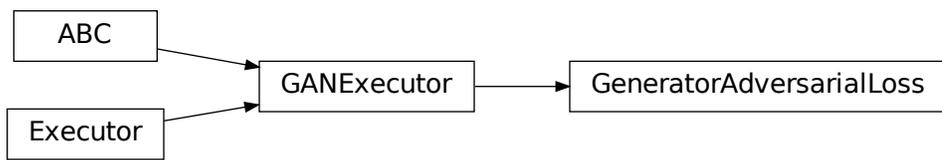
Enumeration for Adversarial Losses. Implemented: GAN and LSGAN.

Attributes

GAN
HINGE_LOSS
LSGAN

4.5.6 GeneratorAdversarialLoss

Inheritance Diagram



class `ashpy.losses.gan.GeneratorAdversarialLoss` (*loss_fn=None*)
 Bases: `ashpy.losses.gan.GANExecutor`
 Base class for the adversarial loss of the generator.

Methods

<code>__init__</code> ([loss_fn])	Initialize the Executor.
<code>call</code> (*args, **kwargs)	

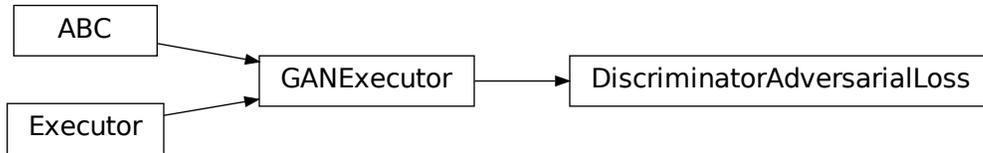
Attributes

<code>fn</code>	Return the Keras loss function to execute.
<code>global_batch_size</code>	Global batch size comprises the batch size for each cpu.
<code>weight</code>	Return the loss weight.

`__init__` (*loss_fn=None*)
 Initialize the Executor.
Parameters `loss_fn` (`tf.keras.losses.Loss`) – Keras Loss function to call passing (`tf.ones_like(d_fake_i), d_fake_i`).
Return type `None`

4.5.7 DiscriminatorAdversarialLoss

Inheritance Diagram



class `ashpy.losses.gan.DiscriminatorAdversarialLoss` (*loss_fn=None*)

Bases: `ashpy.losses.gan.GANExecutor`

Base class for the adversarial loss of the discriminator.

Methods

<code>__init__</code> (<i>loss_fn</i>)	Initialize the Executor.
<code>call</code> (*args, **kwargs)	

Attributes

<code>fn</code>	Return the Keras loss function to execute.
<code>global_batch_size</code>	Global batch size comprises the batch size for each cpu.
<code>weight</code>	Return the loss weight.

`__init__` (*loss_fn=None*)
Initialize the Executor.

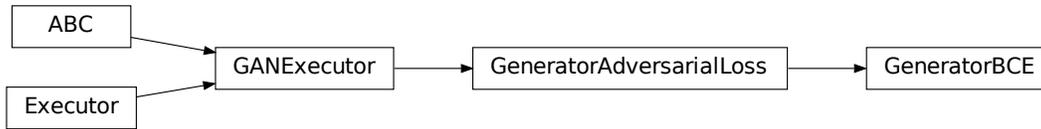
Parameters

- **loss_fn** (`tf.keras.losses.Loss`) – Loss function call passing
- **d_fake** (`((d_real,))`) –

Return type `None`

4.5.8 GeneratorBCE

Inheritance Diagram



class `ashpy.losses.gan.GeneratorBCE` (*from_logits=True*)
 Bases: `ashpy.losses.gan.GeneratorAdversarialLoss`

The Binary CrossEntropy computed among the generator and the 1 label.

$$L_G = E[\log(D(G(z)))]$$

Methods

<code>__init__</code> ([from_logits])	Initialize the BCE Loss for the Generator.
---------------------------------------	--

Attributes

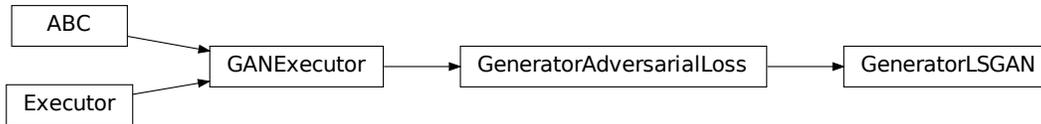
<code>fn</code>	Return the Keras loss function to execute.
<code>global_batch_size</code>	Global batch size comprises the batch size for each cpu.
<code>weight</code>	Return the loss weight.

`__init__` (*from_logits=True*)
 Initialize the BCE Loss for the Generator.

Return type None

4.5.9 GeneratorLSGAN

Inheritance Diagram



class `ashpy.losses.gan.GeneratorLSGAN`

Bases: `ashpy.losses.gan.GeneratorAdversarialLoss`

Least Square GAN Loss for generator.

Reference: <https://arxiv.org/abs/1611.04076>

Note: Basically the Mean Squared Error between the discriminator output when evaluated in fake and 1.

$$L_G = \frac{1}{2} E[(1 - D(G(z)))^2]$$

Methods

<code>__init__()</code>	Initialize the Least Square Loss for the Generator.
-------------------------	---

Attributes

<code>fn</code>	Return the Keras loss function to execute.
<code>global_batch_size</code>	Global batch size comprises the batch size for each cpu.
<code>weight</code>	Return the loss weight.

`__init__()`
Initialize the Least Square Loss for the Generator.

Return type `None`

4.5.10 GeneratorL1

Inheritance Diagram



class `ashpy.losses.gan.GeneratorL1`
 Bases: `ashpy.losses.gan.GANExecutor`
 L1 loss between the generator output and the target.

$$L_G = E\|x - G(z)\|_1$$

Where x is the target and $G(z)$ is generated image.

Methods

<code>__init__()</code>	Initialize the Executor.
<code>call(*args, **kwargs)</code>	

Attributes

<code>fn</code>	Return the Keras loss function to execute.
<code>global_batch_size</code>	Global batch size comprises the batch size for each cpu.
<code>weight</code>	Return the loss weight.

`__init__()`
 Initialize the Executor.
Return type None

4.5.11 GeneratorHingeLoss

Inheritance Diagram



class `ashpy.losses.gan.GeneratorHingeLoss`

Bases: `ashpy.losses.gan.GeneratorAdversarialLoss`

Hinge loss for the Generator. See Geometric GAN¹ for more details.

Methods

<code>__init__()</code>	Initialize the Least Square Loss for the Generator.
-------------------------	---

Attributes

<code>fn</code>	Return the Keras loss function to execute.
<code>global_batch_size</code>	Global batch size comprises the batch size for each cpu.
<code>weight</code>	Return the loss weight.

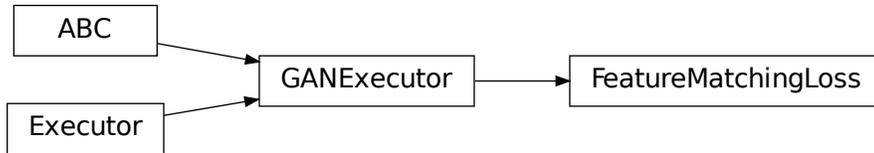
`__init__()`
Initialize the Least Square Loss for the Generator.

Return type `None`

4.5.12 FeatureMatchingLoss

¹ Geometric GAN <https://arxiv.org/abs/1705.02894>

Inheritance Diagram



class `ashpy.losses.gan.FeatureMatchingLoss`

Bases: `ashpy.losses.gan.GANExecutor`

Conditional GAN Feature matching loss.

The loss is computed for each example and it's the L1 (MAE) of the feature difference. Implementation of pix2pix HD: <https://github.com/NVIDIA/pix2pixHD>

$$FM = \sum_{i=0}^N \frac{1}{M_i} \|D_i(x, c) - D_i(G(c), c)\|_1$$

Where:

- D_i is the i -th layer of the discriminator
- N is the total number of layer of the discriminator
- M_i is the number of components for the i -th layer
- x is the target image
- c is the condition
- $G(c)$ is the generated image from the condition c
- $\| \cdot \|_1$ stands for norm 1.

This is for a single example: basically for each layer of the discriminator we compute the absolute error between the layer evaluated in real examples and in fake examples. Then we average along the batch. In the case where D_i is a multidimensional tensor we simply calculate the mean over the axis 1,2,3.

Methods

<code>__init__()</code>	Initialize the Executor.
<code>call(*args, **kwargs)</code>	

Attributes

<code>fn</code>	Return the Keras loss function to execute.
-----------------	--

Continued on next page

Table 65 – continued from previous page

<code>global_batch_size</code>	Global batch size comprises the batch size for each cpu.
<code>weight</code>	Return the loss weight.

`__init__()`
 Initialize the Executor.
Return type None

4.5.13 CategoricalCrossEntropy

Inheritance Diagram



class `ashpy.losses.gan.CategoricalCrossEntropy`

Bases: `ashpy.losses.executor.Executor`

Categorical Cross Entropy between generator output and target.

Useful when the output of the generator is a distribution over classes.

..note:: The target must be represented in one hot notation.

Methods

<code>__init__()</code>	Initialize the Categorical Cross Entropy Executor.
<code>call(*args, **kwargs)</code>	

Attributes

<code>fn</code>	Return the Keras loss function to execute.
<code>global_batch_size</code>	Global batch size comprises the batch size for each cpu.
<code>weight</code>	Return the loss weight.

`__init__()`
 Initialize the Categorical Cross Entropy Executor.
Return type None

4.5.14 Pix2PixLoss

Inheritance Diagram



```

class ashpy.losses.gan.Pix2PixLoss (l1_loss_weight=100.0, adversarial_loss_weight=1.0,
                                   feature_matching_weight=10.0, adversarial_loss_type=<AdversarialLossType.GAN: 1>,
                                   use_feature_matching_loss=False)
  
```

Bases: *ashpy.losses.executor.SumExecutor*

Pix2Pix Loss.

Weighted sum of *ashpy.losses.gan.GeneratorL1*, *ashpy.losses.gan.AdversarialLossG* and *ashpy.losses.gan.FeatureMatchingLoss*.

Used by Pix2Pix [1] and Pix2PixHD [2]

Methods

<code>__init__</code> ([l1_loss_weight, ...])	Initialize the loss.
---	----------------------

Attributes

<code>executors</code>	Return the List of Executors.
<code>fn</code>	Return the Keras loss function to execute.
<code>global_batch_size</code>	Global batch size comprises the batch size for each cpu.
<code>weight</code>	Return the loss weight.

```

__init__(l1_loss_weight=100.0, adversarial_loss_weight=1.0, feature_matching_weight=10.0, adversarial_loss_type=<AdversarialLossType.GAN: 1>, use_feature_matching_loss=False)
  
```

Initialize the loss.

Weighted sum of *ashpy.losses.gan.GeneratorL1*, *ashpy.losses.gan.AdversarialLossG* and *ashpy.losses.gan.FeatureMatchingLoss*.

Parameters

- **l1_loss_weight** (*ashpy.ashtypes.TWeight*) – Weight of L1 loss.
- **adversarial_loss_weight** (*ashpy.ashtypes.TWeight*) – Weight of adversarial loss.
- **feature_matching_weight** (*ashpy.ashtypes.TWeight*) – Weight of the fea-

ture matching loss.

- **adversarial_loss_type** (*ashpy.losses.gan.AdversarialLossType*) – Adversarial loss type (*ashpy.losses.gan.AdversarialLossType.GAN* or *ashpy.losses.gan.AdversarialLossType.LSGAN*).
- **use_feature_matching_loss** (*bool*) – if True use also uses *ashpy.losses.gan.FeatureMatchingLoss*.

Return type *None*

4.5.15 Pix2PixLossSemantic

Inheritance Diagram



```

class ashpy.losses.gan.Pix2PixLossSemantic (cross_entropy_weight=100.0,          ad-
                                             adversarial_loss_weight=1.0,          fea-
                                             ture_matching_weight=10.0,          adversar-
                                             ial_loss_type=<AdversarialLossType.GAN:
                                             1>, use_feature_matching_loss=False)
  
```

Bases: *ashpy.losses.executor.SumExecutor*

Semantic Pix2Pix Loss.

Weighted sum of *ashpy.losses.gan.CategoricalCrossEntropy*, *ashpy.losses.gan.AdversarialLossG* and *ashpy.losses.gan.FeatureMatchingLoss*.

Methods

<code>__init__</code> ([cross_entropy_weight, ...])	Initialize the Executor.
---	--------------------------

Attributes

<code>executors</code>	Return the List of Executors.
<code>fn</code>	Return the Keras loss function to execute.
<code>global_batch_size</code>	Global batch size comprises the batch size for each cpu.
<code>weight</code>	Return the loss weight.

```
__init__(cross_entropy_weight=100.0, adversarial_loss_weight=1.0, feature_matching_weight=10.0, adversarial_loss_type=<AdversarialLossType.GAN: 1>, use_feature_matching_loss=False)
```

Initialize the Executor.

Weighted sum of `ashpy.losses.gan.CategoricalCrossEntropy`, `ashpy.losses.gan.AdversarialLossG` and `ashpy.losses.gan.FeatureMatchingLoss`

Parameters

- **cross_entropy_weight** (`ashpy.ashtypes.TWeight`) – Weight of the categorical cross entropy loss.
- **adversarial_loss_weight** (`ashpy.ashtypes.TWeight`) – Weight of the adversarial loss.
- **feature_matching_weight** (`ashpy.ashtypes.TWeight`) – Weight of the feature matching loss.
- **adversarial_loss_type** (`ashpy.losses.gan.AdversarialLossType`) – type of adversarial loss, see `ashpy.losses.gan.AdversarialLossType`
- **use_feature_matching_loss** (`bool`) – whether to use feature matching loss or not

4.5.16 EncoderBCE

Inheritance Diagram



class `ashpy.losses.gan.EncoderBCE` (*from_logits=True*)

Bases: `ashpy.losses.executor.Executor`

The Binary Cross Entropy computed among the encoder and the 0 label.

Methods

<code>__init__</code> ([from_logits])	Initialize the Executor.
<code>call</code> (*args, **kwargs)	

Attributes

<code>fn</code>	Return the Keras loss function to execute.
-----------------	--

Continued on next page

Table 73 – continued from previous page

<code>global_batch_size</code>	Global batch size comprises the batch size for each cpu.
<code>weight</code>	Return the loss weight.

`__init__` (*from_logits=True*)

Initialize the Executor.

Return type None

4.5.17 DiscriminatorMinMax

Inheritance Diagram



class `ashpy.losses.gan.DiscriminatorMinMax` (*from_logits=True, label_smoothing=0.0*)

Bases: `ashpy.losses.gan.DiscriminatorAdversarialLoss`

The min-max game played by the discriminator.

$$L_D = -\frac{1}{2}E[\log(D(x)) + \log(1 - D(G(z)))]$$

Methods

<code>__init__</code> ([<i>from_logits</i> , <i>label_smoothing</i>])	Initialize Loss.
---	------------------

Attributes

<code>fn</code>	Return the Keras loss function to execute.
<code>global_batch_size</code>	Global batch size comprises the batch size for each cpu.
<code>weight</code>	Return the loss weight.

`__init__` (*from_logits=True, label_smoothing=0.0*)

Initialize Loss.

4.5.18 DiscriminatorLSGAN

Inheritance Diagram



class `ashpy.losses.gan.DiscriminatorLSGAN`

Bases: `ashpy.losses.gan.DiscriminatorAdversarialLoss`

Least square Loss for discriminator.

Reference: Least Squares Generative Adversarial Networks¹ .

Basically the Mean Squared Error between the discriminator output when evaluated in fake samples and 0 and the discriminator output when evaluated in real samples and 1: For the unconditioned case this is:

$$L_D = \frac{1}{2}E[(D(x) - 1)^2 + (0 - D(G(z)))^2]$$

where x are real samples and z is the latent vector.

For the conditioned case this is:

$$L_D = \frac{1}{2}E[(D(x, c) - 1)^2 + (0 - D(G(c), c))^2]$$

where c is the condition and x are real samples.

Methods

<code>__init__()</code>	Initialize loss.
-------------------------	------------------

Attributes

<code>fn</code>	Return the Keras loss function to execute.
<code>global_batch_size</code>	Global batch size comprises the batch size for each cpu.
<code>weight</code>	Return the loss weight.

`__init__()`
Initialize loss.

Return type None

¹ Least Squares Generative Adversarial Networks <https://arxiv.org/abs/1611.04076>

4.5.19 DiscriminatorHingeLoss

Inheritance Diagram



class `ashpy.losses.gan.DiscriminatorHingeLoss`

Bases: `ashpy.losses.gan.DiscriminatorAdversarialLoss`

Hinge loss for the Discriminator.

See Geometric GAN¹ for more details.

Methods

<code>__init__()</code>	Initialize the Least Square Loss for the Generator.
-------------------------	---

Attributes

<code>fn</code>	Return the Keras loss function to execute.
<code>global_batch_size</code>	Global batch size comprises the batch size for each cpu.
<code>weight</code>	Return the loss weight.

`__init__()`
Initialize the Least Square Loss for the Generator.

Return type None

4.5.20 `ashpy.losses.gan.get_adversarial_loss_discriminator`

`ashpy.losses.gan.get_adversarial_loss_discriminator` (*adversarial_loss_type*=<`AdversarialLossType.GAN`:
`1`>)

Return the correct loss for the Discriminator.

Parameters `adversarial_loss_type` (`ashpy.losses.gan.AdversarialLossType`) – Type of loss (`ashpy.losses.gan.AdversarialLossType.GAN` or `ashpy.losses.gan.AdversarialLossType.LSGAN`)

Return type `Type[Executor]`

¹ Geometric GAN <https://arxiv.org/abs/1705.02894>

Returns The correct (*ashpy.losses.executor.Executor*) (to be instantiated).

4.5.21 ashpy.losses.gan.get_adversarial_loss_generator

`ashpy.losses.gan.get_adversarial_loss_generator` (*adversarial_loss_type*=<*AdversarialLossType.GAN*:
I>)

Return the correct loss for the Generator.

Parameters `adversarial_loss_type` (*ashpy.losses.AdversarialLossType*) –
 Type of loss (*ashpy.losses.AdversarialLossType.GAN* or *ashpy.losses.AdversarialLossType.LSGAN*).

Return type *Type[Executor]*

Returns The correct (*ashpy.losses.executor.Executor*) (to be instantiated).

Modules

<i>classifier</i>	The classification losses.
<i>executor</i>	The Executor.
<i>gan</i>	GAN losses.

4.5.22 classifier

The classification losses.

Classes

<i>ClassifierLoss</i>	Classifier Loss Executor using the classifier model, instantiated with a fn.
-----------------------	--

ClassifierLoss

Inheritance Diagram



class `ashpy.losses.classifier.ClassifierLoss` (*fn*)

Bases: *ashpy.losses.executor.Executor*

Classifier Loss Executor using the classifier model, instantiated with a fn.

Methods

<code>__init__(fn)</code>	Initialize <i>ClassifierLoss</i> .
<code>call(*args, **kwargs)</code>	

Attributes

<code>fn</code>	Return the Keras loss function to execute.
<code>global_batch_size</code>	Global batch size comprises the batch size for each cpu.
<code>weight</code>	Return the loss weight.

`__init__(fn)`
Initialize *ClassifierLoss*.

Parameters `fn` (`tf.keras.losses.Loss`) – Classification Loss function, should take as input labels and prediction.

Return type `None`

Returns `None`

class `ashpy.losses.classifier.ClassifierLoss` (*fn*)

Bases: *ashpy.losses.executor.Executor*

Classifier Loss Executor using the classifier model, instantiated with a fn.

`__init__(fn)`
Initialize *ClassifierLoss*.

Parameters `fn` (`tf.keras.losses.Loss`) – Classification Loss function, should take as input labels and prediction.

Return type `None`

Returns `None`

4.5.23 executor

The Executor.

An object that, given an `ashpy.contexts.BaseContext`, carries a function and the way of executing it.

Classes

<i>Executor</i>	Carry a function and the way of executing it.
<i>SumExecutor</i>	The sum executor.

Executor

Inheritance Diagram



class `ashpy.losses.executor.Executor` (*fn=None*)

Bases: `object`

Carry a function and the way of executing it. Given a context.

Methods

<code>__init__</code> (<i>fn</i>)	Initialize the Executor.
<code>call</code> (<i>context</i> , <i>**kwargs</i>)	Execute the function, using the information provided by the context.
<code>reduce_loss</code> (<i>call_fn</i>)	Create a Decorator to reduce Losses.

Attributes

<i>fn</i>	Return the Keras loss function to execute.
<i>global_batch_size</i>	Global batch size comprises the batch size for each cpu.
<i>weight</i>	Return the loss weight.

`__init__` (*fn=None*)

Initialize the Executor.

Parameters *fn* (`tf.keras.losses.Loss`) – A Keras Loss to execute.

Return type `None`

Returns `None`

abstract `call` (*context*, ***kwargs*)

Execute the function, using the information provided by the context.

Parameters *context* (`ashpy.contexts.BaseContext`) – The function execution Context.

Return type `Tensor`

Returns `tf.Tensor` – Output Tensor.

property *fn*

Return the Keras loss function to execute.

Return type `Loss`

Returns `tf.keras.losses.Loss` – Keras Loss.

property global_batch_size

Global batch size comprises the batch size for each cpu.

Calculated as `batch_size_for_replica*replica_numbers`.

Return type `int`

Returns `int` – Global Batch size value.

static reduce_loss (*call_fn*)

Create a Decorator to reduce Losses. Used to simplify things.

Apply a `reduce sum` operation to the loss and divide the result by the batch size.

Parameters `call_fn` (`typing.Callable`) – The executor call method.

Return type `Callable`

Returns `typing.Callable` – The decorated function.

property weight

Return the loss weight.

This weight is multiplied by the loss value. This is useful when working with multiples losses.

Return type `Callable[...float]`

Returns `typing.Callable` – Callable returning the weight (`float`).

SumExecutor

Inheritance Diagram



class `ashpy.losses.executor.SumExecutor` (*executors*)

Bases: `ashpy.losses.executor.Executor`

The sum executor. Executes the call of each fn and weights the losses.

Each Executor gets called (thus reducing its carried function), the results are then summed together.

Methods

<code>__init__</code> (executors)	Initialize the SumExecutor.
<code>call</code> (*args, **kwargs)	Evaluate and sum together the Executors.

Attributes

<code>executors</code>	Return the List of Executors.
<code>fn</code>	Return the Keras loss function to execute.
<code>global_batch_size</code>	Global batch size comprises the batch size for each cpu.
<code>weight</code>	Return the loss weight.

`__init__` (*executors*)

Initialize the SumExecutor.

Parameters `executors` (`list` of [`ashpy.executors.Executor`]) – Array of `ashpy.executors.Executor` to sum evaluate and sum together.

Return type `None`

Returns `None`

`call` (**args, **kwargs*)

Evaluate and sum together the Executors.

Return type `Tensor`

Returns `:py:classes:'tf.Tensor'` – Output Tensor.

property `executors`

Return the List of Executors.

Return type `List[Executor]`

property `global_batch_size`

Global batch size comprises the batch size for each cpu.

Calculated as `batch_size_for_replica*replica_numbers`.

Return type `int`

Returns `int` – Global Batch size value.

class `ashpy.losses.executor.Executor` (*fn=None*)

Bases: `object`

Carry a function and the way of executing it. Given a context.

`__init__` (*fn=None*)

Initialize the Executor.

Parameters `fn` (`tf.keras.losses.Loss`) – A Keras Loss to execute.

Return type `None`

Returns `None`

abstract `call` (*context, **kwargs*)

Execute the function, using the information provided by the context.

Parameters `context` (`ashpy.contexts.BaseContext`) – The function execution Context.

Return type `Tensor`

Returns `tf.Tensor` – Output Tensor.

property `fn`

Return the Keras loss function to execute.

Return type `Loss`

Returns `tf.keras.losses.Loss` – Keras Loss.

property global_batch_size

Global batch size comprises the batch size for each cpu.

Calculated as `batch_size_for_replica*replica_numbers`.

Return type `int`

Returns `int` – Global Batch size value.

static reduce_loss (*call_fn*)

Create a Decorator to reduce Losses. Used to simplify things.

Apply a `reduce` `sum` operation to the loss and divide the result by the batch size.

Parameters `call_fn` (`typing.Callable`) – The executor call method.

Return type `Callable`

Returns `typing.Callable` – The decorated function.

property weight

Return the loss weight.

This weight is multiplied by the loss value. This is useful when working with multiples losses.

Return type `Callable[...float]`

Returns `typing.Callable` – Callable returning the weight (`float`).

class `ashpy.losses.executor.SumExecutor` (*executors*)

Bases: `ashpy.losses.executor.Executor`

The sum executor. Executes the call of each fn and weights the losses.

Each Executor gets called (thus reducing its carried function), the results are then summed together.

__init__ (*executors*)

Initialize the SumExecutor.

Parameters `executors` (`list` of [`ashpy.executors.Executor`]) – Array of `ashpy.executors.Executor` to sum evaluate and sum together.

Return type `None`

Returns `None`

call (**args, **kwargs*)

Evaluate and sum together the Executors.

Return type `Tensor`

Returns **:py:classes:tf.Tensor** – Output Tensor.

property executors

Return the List of Executors.

Return type `List[Executor]`

property global_batch_size

Global batch size comprises the batch size for each cpu.

Calculated as `batch_size_for_replica*replica_numbers`.

Return type `int`

Returns `int` – Global Batch size value.

4.5.24 gan

GAN losses.

Functions

<code>get_adversarial_loss_discriminator</code>	Return the correct loss for the Discriminator.
<code>get_adversarial_loss_generator</code>	Return the correct loss for the Generator.

`ashpy.losses.gan.get_adversarial_loss_discriminator`

`ashpy.losses.gan.get_adversarial_loss_discriminator` (*adversarial_loss_type*=<*AdversarialLossType.GAN*:
I>)

Return the correct loss for the Discriminator.

Parameters *adversarial_loss_type* (*ashpy.losses.gan.AdversarialLossType*) – Type of loss (*ashpy.losses.gan.AdversarialLossType.GAN* or *ashpy.losses.gan.AdversarialLossType.LSGAN*)

Return type *Type[Executor]*

Returns The correct (*ashpy.losses.executor.Executor*) (to be instantiated).

`ashpy.losses.gan.get_adversarial_loss_generator`

`ashpy.losses.gan.get_adversarial_loss_generator` (*adversarial_loss_type*=<*AdversarialLossType.GAN*:
I>)

Return the correct loss for the Generator.

Parameters *adversarial_loss_type* (*ashpy.losses.AdversarialLossType*) – Type of loss (*ashpy.losses.AdversarialLossType.GAN* or *ashpy.losses.AdversarialLossType.LSGAN*).

Return type *Type[Executor]*

Returns The correct (*ashpy.losses.executor.Executor*) (to be instantiated).

Classes

<code>AdversarialLossType</code>	Enumeration for Adversarial Losses.
<code>CategoricalCrossEntropy</code>	Categorical Cross Entropy between generator output and target.
<code>DiscriminatorAdversarialLoss</code>	Base class for the adversarial loss of the discriminator.
<code>DiscriminatorHingeLoss</code>	Hinge loss for the Discriminator.
<code>DiscriminatorLSGAN</code>	Least square Loss for discriminator.
<code>DiscriminatorMinMax</code>	The min-max game played by the discriminator.
<code>EncoderBCE</code>	The Binary Cross Entropy computed among the encoder and the 0 label.
<code>FeatureMatchingLoss</code>	Conditional GAN Feature matching loss.
<code>GANExecutor</code>	Executor for GANs.
<code>GeneratorAdversarialLoss</code>	Base class for the adversarial loss of the generator.

Continued on next page

Table 90 – continued from previous page

<i>GeneratorBCE</i>	The Binary CrossEntropy computed among the generator and the 1 label.
<i>GeneratorHingeLoss</i>	Hinge loss for the Generator.
<i>GeneratorL1</i>	L1 loss between the generator output and the target.
<i>GeneratorLSGAN</i>	Least Square GAN Loss for generator.
<i>Pix2PixLoss</i>	Pix2Pix Loss.
<i>Pix2PixLossSemantic</i>	Semantic Pix2Pix Loss.

AdversarialLossType

Inheritance Diagram



class `ashpy.losses.gan.AdversarialLossType`

Bases: `enum.Enum`

Enumeration for Adversarial Losses. Implemented: GAN and LSGAN.

Attributes

GAN

HINGE_LOSS

LSGAN

CategoricalCrossEntropy

Inheritance Diagram



class `ashpy.losses.gan.CategoricalCrossEntropy`

Bases: `ashpy.losses.executor.Executor`

Categorical Cross Entropy between generator output and target.

Useful when the output of the generator is a distribution over classes.

..note:: The target must be represented in one hot notation.

Methods

<code>__init__()</code>	Initialize the Categorical Cross Entropy Executor.
<code>call(*args, **kwargs)</code>	

Attributes

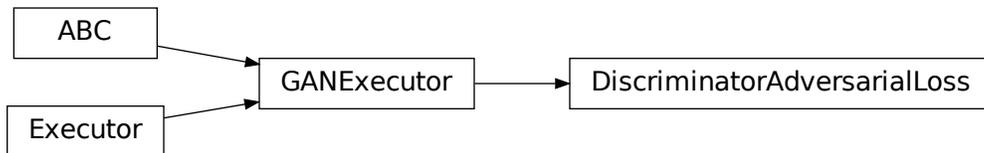
<code>fn</code>	Return the Keras loss function to execute.
<code>global_batch_size</code>	Global batch size comprises the batch size for each cpu.
<code>weight</code>	Return the loss weight.

`__init__()`
Initialize the Categorical Cross Entropy Executor.

Return type None

DiscriminatorAdversarialLoss

Inheritance Diagram



class `ashpy.losses.gan.DiscriminatorAdversarialLoss` (*loss_fn=None*)

Bases: `ashpy.losses.gan.GANExecutor`

Base class for the adversarial loss of the discriminator.

Methods

<code>__init__([loss_fn])</code>	Initialize the Executor.
<code>call(*args, **kwargs)</code>	

Attributes

<code>fn</code>	Return the Keras loss function to execute.
<code>global_batch_size</code>	Global batch size comprises the batch size for each cpu.
<code>weight</code>	Return the loss weight.

`__init__(loss_fn=None)`
 Initialize the Executor.

Parameters

- `loss_fn` (`tf.keras.losses.Loss`) – Loss function call passing
- `d_fake` (`(d_real,)`) –

Return type None

DiscriminatorHingeLoss

Inheritance Diagram



class `ashpy.losses.gan.DiscriminatorHingeLoss`

Bases: `ashpy.losses.gan.DiscriminatorAdversarialLoss`

Hinge loss for the Discriminator.

See Geometric GAN¹ for more details.

Methods

<code>__init__()</code>	Initialize the Least Square Loss for the Generator.
-------------------------	---

Attributes

<code>fn</code>	Return the Keras loss function to execute.
<code>global_batch_size</code>	Global batch size comprises the batch size for each cpu.

Continued on next page

¹ Geometric GAN <https://arxiv.org/abs/1705.02894>

Table 97 – continued from previous page

weight	Return the loss weight.
--------	-------------------------

`__init__()`
 Initialize the Least Square Loss for the Generator.
Return type None

DiscriminatorLSGAN

Inheritance Diagram



class `ashpy.losses.gan.DiscriminatorLSGAN`
 Bases: `ashpy.losses.gan.DiscriminatorAdversarialLoss`

Least square Loss for discriminator.

Reference: Least Squares Generative Adversarial Networks¹ .

Basically the Mean Squared Error between the discriminator output when evaluated in fake samples and 0 and the discriminator output when evaluated in real samples and 1: For the unconditioned case this is:

$$L_D = \frac{1}{2}E[(D(x) - 1)^2 + (0 - D(G(z)))^2]$$

where x are real samples and z is the latent vector.

For the conditioned case this is:

$$L_D = \frac{1}{2}E[(D(x, c) - 1)^2 + (0 - D(G(c), c))^2]$$

where c is the condition and x are real samples.

Methods

<code>__init__()</code>	Initialize loss.
-------------------------	------------------

Attributes

¹ Least Squares Generative Adversarial Networks <https://arxiv.org/abs/1611.04076>

<code>fn</code>	Return the Keras loss function to execute.
<code>global_batch_size</code>	Global batch size comprises the batch size for each cpu.
<code>weight</code>	Return the loss weight.

`__init__()`
Initialize loss.

Return type None

DiscriminatorMinMax

Inheritance Diagram



class `ashpy.losses.gan.DiscriminatorMinMax` (*from_logits=True, label_smoothing=0.0*)
 Bases: `ashpy.losses.gan.DiscriminatorAdversarialLoss`

The min-max game played by the discriminator.

$$L_D = -\frac{1}{2}E[\log(D(x)) + \log(1 - D(G(z)))]$$

Methods

<code>__init__</code> ([from_logits, label_smoothing])	Initialize Loss.
--	------------------

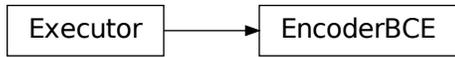
Attributes

<code>fn</code>	Return the Keras loss function to execute.
<code>global_batch_size</code>	Global batch size comprises the batch size for each cpu.
<code>weight</code>	Return the loss weight.

`__init__` (*from_logits=True, label_smoothing=0.0*)
Initialize Loss.

EncoderBCE

Inheritance Diagram



class `ashpy.losses.gan.EncoderBCE` (*from_logits=True*)

Bases: `ashpy.losses.executor.Executor`

The Binary Cross Entropy computed among the encoder and the 0 label.

Methods

<code>__init__</code> ([from_logits])	Initialize the Executor.
<code>call</code> (*args, **kwargs)	

Attributes

<code>fn</code>	Return the Keras loss function to execute.
<code>global_batch_size</code>	Global batch size comprises the batch size for each cpu.
<code>weight</code>	Return the loss weight.

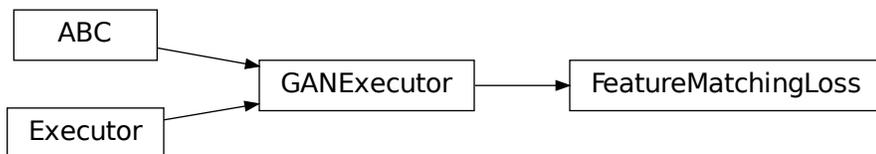
`__init__` (*from_logits=True*)

Initialize the Executor.

Return type None

FeatureMatchingLoss

Inheritance Diagram



class `ashpy.losses.gan.FeatureMatchingLoss`

Bases: `ashpy.losses.gan.GANExecutor`

Conditional GAN Feature matching loss.

The loss is computed for each example and it's the L1 (MAE) of the feature difference. Implementation of pix2pix HD: <https://github.com/NVIDIA/pix2pixHD>

$$FM = \sum_{i=0}^N \frac{1}{M_i} \|D_i(x, c) - D_i(G(c), c)\|_1$$

Where:

- `D_i` is the i-th layer of the discriminator
- `N` is the total number of layer of the discriminator
- `M_i` is the number of components for the i-th layer
- `x` is the target image
- `c` is the condition
- `G(c)` is the generated image from the condition `c`
- `\| \|_1` stands for norm 1.

This is for a single example: basically for each layer of the discriminator we compute the absolute error between the layer evaluated in real examples and in fake examples. Then we average along the batch. In the case where `D_i` is a multidimensional tensor we simply calculate the mean over the axis 1,2,3.

Methods

<code>__init__()</code>	Initialize the Executor.
<code>call(*args, **kwargs)</code>	

Attributes

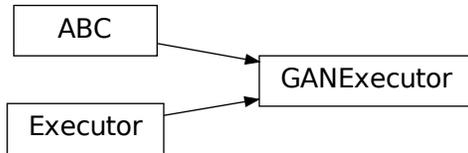
<code>fn</code>	Return the Keras loss function to execute.
<code>global_batch_size</code>	Global batch size comprises the batch size for each cpu.
<code>weight</code>	Return the loss weight.

`__init__()`
Initialize the Executor.

Return type None

GANExecutor

Inheritance Diagram



class `ashpy.losses.gan.GANExecutor` (*fn=None*)
 Bases: `ashpy.losses.executor.Executor`, `abc.ABC`

Executor for GANs.

Implements the basic functions needed by the GAN losses.

Methods

<code>__init__</code> (<i>fn</i>)	Initialize the GANExecutor.
<code>call</code> (<i>context</i> , <i>**kwargs</i>)	Execute the function, using the information provided by the context.
<code>get_discriminator_inputs</code> (<i>context</i> , ...)	Return the discriminator inputs.

Attributes

<code>fn</code>	Return the Keras loss function to execute.
<code>global_batch_size</code>	Global batch size comprises the batch size for each cpu.
<code>weight</code>	Return the loss weight.

`__init__` (*fn=None*)
 Initialize the GANExecutor.

Parameters `fn` (`tf.keras.losses.Loss`) – A Keras Loss to execute.

Returns `None`

abstract `call` (*context*, ***kwargs*)
 Execute the function, using the information provided by the context.

Parameters `context` (`ashpy.contexts.BaseContext`) – The function execution Context.

Returns `tf.Tensor` – Output Tensor.

static `get_discriminator_inputs` (*context*, *fake_or_real*, *condition*, *training*)
 Return the discriminator inputs. If needed it uses the encoder.

The current implementation uses the number of inputs to determine whether the discriminator is conditioned or not.

Parameters

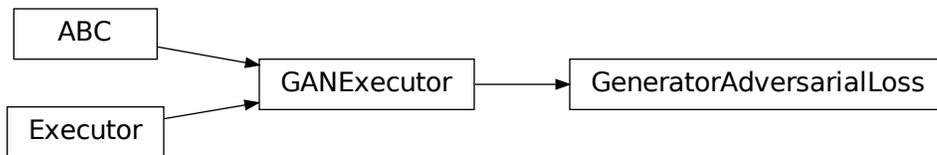
- **context** (*ashpy.contexts.gan.GANContext*) – Context for GAN models.
- **fake_or_real** (*tf.Tensor*) – Discriminator input tensor, it can be fake (generated) or real.
- **condition** (*tf.Tensor*) – Discriminator condition (it can also be generator noise).
- **training** (*bool*) – whether is training phase or not

Return type *Union[Tensor, List[Tensor]]*

Returns The discriminator inputs.

GeneratorAdversarialLoss

Inheritance Diagram



class *ashpy.losses.gan*.**GeneratorAdversarialLoss** (*loss_fn=None*)

Bases: *ashpy.losses.gan.GANExecutor*

Base class for the adversarial loss of the generator.

Methods

<code>__init__</code> (<i>[loss_fn]</i>)	Initialize the Executor.
<code>call</code> (*args, **kwargs)	

Attributes

<code>fn</code>	Return the Keras loss function to execute.
<code>global_batch_size</code>	Global batch size comprises the batch size for each cpu.
<code>weight</code>	Return the loss weight.

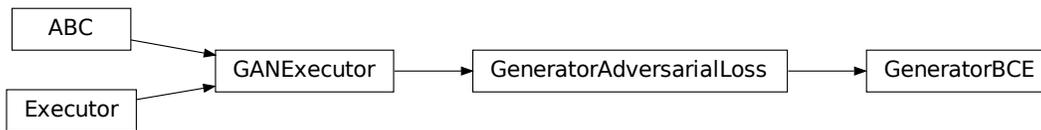
`__init__` (*loss_fn=None*)
Initialize the Executor.

Parameters `loss_fn` (`tf.keras.losses.Loss`) – Keras Loss function to call passing (`tf.ones_like(d_fake_i), d_fake_i`).

Return type None

GeneratorBCE

Inheritance Diagram



class `ashpy.losses.gan.GeneratorBCE` (*from_logits=True*)
Bases: `ashpy.losses.gan.GeneratorAdversarialLoss`

The Binary CrossEntropy computed among the generator and the 1 label.

$$L_G = E[\log(D(G(z)))]$$

Methods

<code>__init__</code> ([from_logits])	Initialize the BCE Loss for the Generator.
---------------------------------------	--

Attributes

<code>fn</code>	Return the Keras loss function to execute.
<code>global_batch_size</code>	Global batch size comprises the batch size for each cpu.
<code>weight</code>	Return the loss weight.

`__init__` (*from_logits=True*)
Initialize the BCE Loss for the Generator.

Return type None

GeneratorHingeLoss

Inheritance Diagram



class `ashpy.losses.gan.GeneratorHingeLoss`

Bases: `ashpy.losses.gan.GeneratorAdversarialLoss`

Hinge loss for the Generator. See Geometric GAN¹ for more details.

Methods

<code>__init__()</code>	Initialize the Least Square Loss for the Generator.
-------------------------	---

Attributes

<code>fn</code>	Return the Keras loss function to execute.
<code>global_batch_size</code>	Global batch size comprises the batch size for each cpu.
<code>weight</code>	Return the loss weight.

`__init__()`
Initialize the Least Square Loss for the Generator.

Return type None

GeneratorL1

¹ Geometric GAN <https://arxiv.org/abs/1705.02894>

Inheritance Diagram



class `ashpy.losses.gan.GeneratorL1`
 Bases: `ashpy.losses.gan.GANExecutor`

L1 loss between the generator output and the target.

$$L_G = E||x - G(z)||_1$$

Where x is the target and G(z) is generated image.

Methods

<code>__init__()</code>	Initialize the Executor.
<code>call(*args, **kwargs)</code>	

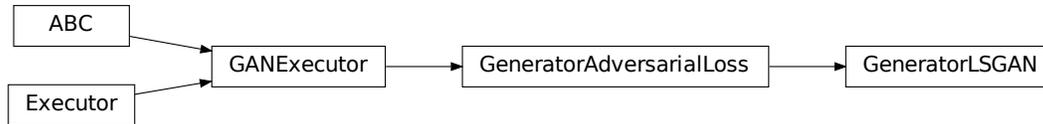
Attributes

<code>fn</code>	Return the Keras loss function to execute.
<code>global_batch_size</code>	Global batch size comprises the batch size for each cpu.
<code>weight</code>	Return the loss weight.

`__init__()`
 Initialize the Executor.
Return type None

GeneratorLSGAN

Inheritance Diagram



class `ashpy.losses.gan.GeneratorLSGAN`

Bases: `ashpy.losses.gan.GeneratorAdversarialLoss`

Least Square GAN Loss for generator.

Reference: <https://arxiv.org/abs/1611.04076>

Note: Basically the Mean Squared Error between the discriminator output when evaluated in fake and 1.

$$L_G = \frac{1}{2} E[(1 - D(G(z)))^2]$$

Methods

<code>__init__()</code>	Initialize the Least Square Loss for the Generator.
-------------------------	---

Attributes

<code>fn</code>	Return the Keras loss function to execute.
<code>global_batch_size</code>	Global batch size comprises the batch size for each cpu.
<code>weight</code>	Return the loss weight.

`__init__()`
Initialize the Least Square Loss for the Generator.

Return type None

Pix2PixLoss

Inheritance Diagram



```

class ashpy.losses.gan.Pix2PixLoss (l1_loss_weight=100.0, adversarial_loss_weight=1.0,
                                   feature_matching_weight=10.0, adversarial_loss_type=<AdversarialLossType.GAN: 1>,
                                   use_feature_matching_loss=False)
  
```

Bases: *ashpy.losses.executor.SumExecutor*

Pix2Pix Loss.

Weighted sum of *ashpy.losses.gan.GeneratorL1*, *ashpy.losses.gan.AdversarialLossG* and *ashpy.losses.gan.FeatureMatchingLoss*.

Used by Pix2Pix [1] and Pix2PixHD [2]

Methods

<code>__init__</code> (l1_loss_weight, ...)	Initialize the loss.
---	----------------------

Attributes

<code>executors</code>	Return the List of Executors.
<code>fn</code>	Return the Keras loss function to execute.
<code>global_batch_size</code>	Global batch size comprises the batch size for each cpu.
<code>weight</code>	Return the loss weight.

```

__init__(l1_loss_weight=100.0, adversarial_loss_weight=1.0, feature_matching_weight=10.0, adversarial_loss_type=<AdversarialLossType.GAN: 1>, use_feature_matching_loss=False)
  Initialize the loss.
  
```

Weighted sum of *ashpy.losses.gan.GeneratorL1*, *ashpy.losses.gan.AdversarialLossG* and *ashpy.losses.gan.FeatureMatchingLoss*.

Parameters

- **l1_loss_weight** (*ashpy.ashtypes.TWeight*) – Weight of L1 loss.
- **adversarial_loss_weight** (*ashpy.ashtypes.TWeight*) – Weight of adversarial loss.
- **feature_matching_weight** (*ashpy.ashtypes.TWeight*) – Weight of the feature matching loss.
- **adversarial_loss_type** (*ashpy.losses.gan.AdversarialLossType*) –

Adversarial loss type (`ashpy.losses.gan.AdversarialLossType.GAN` or `ashpy.losses.gan.AdversarialLossType.LSGAN`).

- `use_feature_matching_loss` (*bool*) – if True use also uses `ashpy.losses.gan.FeatureMatchingLoss`.

Return type `None`

Pix2PixLossSemantic

Inheritance Diagram



```

class ashpy.losses.gan.Pix2PixLossSemantic (cross_entropy_weight=100.0,      ad-
                                             adversarial_loss_weight=1.0,      fea-
                                             ture_matching_weight=10.0,      adversar-
                                             ial_loss_type=<AdversarialLossType.GAN:
                                             1>, use_feature_matching_loss=False)
  
```

Bases: `ashpy.losses.executor.SumExecutor`

Semantic Pix2Pix Loss.

Weighted sum of `ashpy.losses.gan.CategoricalCrossEntropy`, `ashpy.losses.gan.AdversarialLossG` and `ashpy.losses.gan.FeatureMatchingLoss`.

Methods

<code>__init__</code> ([<code>cross_entropy_weight</code> , ...])	Initialize the Executor.
--	--------------------------

Attributes

<code>executors</code>	Return the List of Executors.
<code>fn</code>	Return the Keras loss function to execute.
<code>global_batch_size</code>	Global batch size comprises the batch size for each cpu.
<code>weight</code>	Return the loss weight.

```

__init__(cross_entropy_weight=100.0,      adversarial_loss_weight=1.0,      fea-
         feature_matching_weight=10.0,      adversarial_loss_type=<AdversarialLossType.GAN: 1>,
         use_feature_matching_loss=False)
Initialize the Executor.
  
```

Weighted sum of `ashpy.losses.gan.CategoricalCrossEntropy`, `ashpy.losses.gan.`

AdversarialLossG and *ashpy.losses.gan.FeatureMatchingLoss*

Parameters

- **cross_entropy_weight** (*ashpy.ashtypes.TWeight*) – Weight of the categorical cross entropy loss.
- **adversarial_loss_weight** (*ashpy.ashtypes.TWeight*) – Weight of the adversarial loss.
- **feature_matching_weight** (*ashpy.ashtypes.TWeight*) – Weight of the feature matching loss.
- **adversarial_loss_type** (*ashpy.losses.gan.AdversarialLossType*) – type of adversarial loss, see *ashpy.losses.gan.AdversarialLossType*
- **use_feature_matching_loss** (*bool*) – whether to use feature matching loss or not

class *ashpy.losses.gan.AdversarialLossType*

Bases: *enum.Enum*

Enumeration for Adversarial Losses. Implemented: GAN and LSGAN.

class *ashpy.losses.gan.CategoricalCrossEntropy*

Bases: *ashpy.losses.executor.Executor*

Categorical Cross Entropy between generator output and target.

Useful when the output of the generator is a distribution over classes.

..note:: The target must be represented in one hot notation.

__init__ ()

Initialize the Categorical Cross Entropy Executor.

Return type None

class *ashpy.losses.gan.DiscriminatorAdversarialLoss* (*loss_fn=None*)

Bases: *ashpy.losses.gan.GANExecutor*

Base class for the adversarial loss of the discriminator.

__init__ (*loss_fn=None*)

Initialize the Executor.

Parameters

- **loss_fn** (*tf.keras.losses.Loss*) – Loss function call passing
- **d_fake** (*(d_real,)*) –

Return type None

class *ashpy.losses.gan.DiscriminatorHingeLoss*

Bases: *ashpy.losses.gan.DiscriminatorAdversarialLoss*

Hinge loss for the Discriminator.

See Geometric GAN [1]_ for more details.

__init__ ()

Initialize the Least Square Loss for the Generator.

Return type None

class `ashpy.losses.gan.DiscriminatorLSGAN`

Bases: `ashpy.losses.gan.DiscriminatorAdversarialLoss`

Least square Loss for discriminator.

Reference: Least Squares Generative Adversarial Networks [1].

Basically the Mean Squared Error between the discriminator output when evaluated in fake samples and 0 and the discriminator output when evaluated in real samples and 1: For the unconditioned case this is:

$$L_D = \frac{1}{2}E[(D(x) - 1)^2 + (0 - D(G(z)))^2]$$

where x are real samples and z is the latent vector.

For the conditioned case this is:

$$L_D = \frac{1}{2}E[(D(x, c) - 1)^2 + (0 - D(G(c), c))^2]$$

where c is the condition and x are real samples.

`__init__` ()

Initialize loss.

Return type None

class `ashpy.losses.gan.DiscriminatorMinMax` (*from_logits=True, label_smoothing=0.0*)

Bases: `ashpy.losses.gan.DiscriminatorAdversarialLoss`

The min-max game played by the discriminator.

$$L_D = -\frac{1}{2}E[\log(D(x)) + \log(1 - D(G(z)))]$$

`__init__` (*from_logits=True, label_smoothing=0.0*)

Initialize Loss.

class `ashpy.losses.gan.EncoderBCE` (*from_logits=True*)

Bases: `ashpy.losses.executor.Executor`

The Binary Cross Entropy computed among the encoder and the 0 label.

`__init__` (*from_logits=True*)

Initialize the Executor.

Return type None

class `ashpy.losses.gan.FeatureMatchingLoss`

Bases: `ashpy.losses.gan.GANExecutor`

Conditional GAN Feature matching loss.

The loss is computed for each example and it's the L1 (MAE) of the feature difference. Implementation of pix2pix HD: <https://github.com/NVIDIA/pix2pixHD>

$$FM = \sum_{i=0}^N \frac{1}{M_i} \|D_i(x, c) - D_i(G(c), c)\|_1$$

Where:

- `D_i` is the i-th layer of the discriminator
- `N` is the total number of layer of the discriminator

- M_i is the number of components for the i -th layer
- x is the target image
- c is the condition
- $G(c)$ is the generated image from the condition c
- $\| \cdot \|_1$ stands for norm 1.

This is for a single example: basically for each layer of the discriminator we compute the absolute error between the layer evaluated in real examples and in fake examples. Then we average along the batch. In the case where D_i is a multidimensional tensor we simply calculate the mean over the axis 1,2,3.

`__init__()`

Initialize the Executor.

Return type `None`

class `ashpy.losses.gan.GANExecutor` (*fn=None*)

Bases: `ashpy.losses.executor.Executor`, `abc.ABC`

Executor for GANs.

Implements the basic functions needed by the GAN losses.

`__init__` (*fn=None*)

Initialize the GANExecutor.

Parameters `fn` (`tf.keras.losses.Loss`) – A Keras Loss to execute.

Returns `None`

abstract call (*context, **kwargs*)

Execute the function, using the information provided by the context.

Parameters `context` (`ashpy.contexts.BaseContext`) – The function execution Context.

Returns `tf.Tensor` – Output Tensor.

static get_discriminator_inputs (*context, fake_or_real, condition, training*)

Return the discriminator inputs. If needed it uses the encoder.

The current implementation uses the number of inputs to determine whether the discriminator is conditioned or not.

Parameters

- **context** (`ashpy.contexts.gan.GANContext`) – Context for GAN models.
- **fake_or_real** (`tf.Tensor`) – Discriminator input tensor, it can be fake (generated) or real.
- **condition** (`tf.Tensor`) – Discriminator condition (it can also be generator noise).
- **training** (`bool`) – whether is training phase or not

Return type `Union[Tensor, List[Tensor]]`

Returns The discriminator inputs.

class `ashpy.losses.gan.GeneratorAdversarialLoss` (*loss_fn=None*)

Bases: `ashpy.losses.gan.GANExecutor`

Base class for the adversarial loss of the generator.

`__init__` (*loss_fn=None*)
Initialize the Executor.

Parameters `loss_fn` (`tf.keras.losses.Loss`) – Keras Loss function to call passing (`tf.ones_like(d_fake_i)`, `d_fake_i`).

Return type None

class `ashpy.losses.gan.GeneratorBCE` (*from_logits=True*)
Bases: `ashpy.losses.gan.GeneratorAdversarialLoss`

The Binary CrossEntropy computed among the generator and the 1 label.

$$L_G = E[\log(D(G(z)))]$$

`__init__` (*from_logits=True*)
Initialize the BCE Loss for the Generator.

Return type None

class `ashpy.losses.gan.GeneratorHingeLoss`
Bases: `ashpy.losses.gan.GeneratorAdversarialLoss`

Hinge loss for the Generator. See Geometric GAN [1] for more details.

`__init__` ()
Initialize the Least Square Loss for the Generator.

Return type None

class `ashpy.losses.gan.GeneratorL1`
Bases: `ashpy.losses.gan.GANExecutor`

L1 loss between the generator output and the target.

$$L_G = E\|x - G(z)\|_1$$

Where x is the target and G(z) is generated image.

`__init__` ()
Initialize the Executor.

Return type None

class `ashpy.losses.gan.GeneratorLSGAN`
Bases: `ashpy.losses.gan.GeneratorAdversarialLoss`

Least Square GAN Loss for generator.

Reference: <https://arxiv.org/abs/1611.04076>

Note: Basically the Mean Squared Error between the discriminator output when evaluated in fake and 1.

$$L_G = \frac{1}{2}E[(1 - D(G(z)))^2]$$

`__init__` ()
Initialize the Least Square Loss for the Generator.

Return type None

```
class ashpy.losses.gan.Pix2PixLoss (l1_loss_weight=100.0,      adversarial_loss_weight=1.0,
                                feature_matching_weight=10.0,      adversarial_loss_type=<AdversarialLossType.GAN: 1>,
                                use_feature_matching_loss=False)
```

Bases: *ashpy.losses.executor.SumExecutor*

Pix2Pix Loss.

Weighted sum of *ashpy.losses.gan.GeneratorL1*, *ashpy.losses.gan.AdversarialLossG* and *ashpy.losses.gan.FeatureMatchingLoss*.

Used by Pix2Pix [1] and Pix2PixHD [2]

```
__init__ (l1_loss_weight=100.0, adversarial_loss_weight=1.0, feature_matching_weight=10.0, adversarial_loss_type=<AdversarialLossType.GAN: 1>, use_feature_matching_loss=False)
Initialize the loss.
```

Weighted sum of *ashpy.losses.gan.GeneratorL1*, *ashpy.losses.gan.AdversarialLossG* and *ashpy.losses.gan.FeatureMatchingLoss*.

Parameters

- **l1_loss_weight** (*ashpy.ashtypes.TWeight*) – Weight of L1 loss.
- **adversarial_loss_weight** (*ashpy.ashtypes.TWeight*) – Weight of adversarial loss.
- **feature_matching_weight** (*ashpy.ashtypes.TWeight*) – Weight of the feature matching loss.
- **adversarial_loss_type** (*ashpy.losses.gan.AdversarialLossType*) – Adversarial loss type (*ashpy.losses.gan.AdversarialLossType.GAN* or *ashpy.losses.gan.AdversarialLossType.LSGAN*).
- **use_feature_matching_loss** (*bool*) – if True use also uses *ashpy.losses.gan.FeatureMatchingLoss*.

Return type None

```
class ashpy.losses.gan.Pix2PixLossSemantic (cross_entropy_weight=100.0,      adversarial_loss_weight=1.0,      feature_matching_weight=10.0,      adversarial_loss_type=<AdversarialLossType.GAN: 1>, use_feature_matching_loss=False)
```

Bases: *ashpy.losses.executor.SumExecutor*

Semantic Pix2Pix Loss.

Weighted sum of *ashpy.losses.gan.CategoricalCrossEntropy*, *ashpy.losses.gan.AdversarialLossG* and *ashpy.losses.gan.FeatureMatchingLoss*.

```
__init__ (cross_entropy_weight=100.0,      adversarial_loss_weight=1.0,      feature_matching_weight=10.0,      adversarial_loss_type=<AdversarialLossType.GAN: 1>, use_feature_matching_loss=False)
Initialize the Executor.
```

Weighted sum of *ashpy.losses.gan.CategoricalCrossEntropy*, *ashpy.losses.gan.AdversarialLossG* and *ashpy.losses.gan.FeatureMatchingLoss*

Parameters

- **cross_entropy_weight** (*ashpy.ashtypes.TWeight*) – Weight of the categorical cross entropy loss.

- **adversarial_loss_weight** (`ashpy.ashtypes.TWeight`) – Weight of the adversarial loss.
- **feature_matching_weight** (`ashpy.ashtypes.TWeight`) – Weight of the feature matching loss.
- **adversarial_loss_type** (`ashpy.losses.gan.AdversarialLossType`) – type of adversarial loss, see `ashpy.losses.gan.AdversarialLossType`
- **use_feature_matching_loss** (`bool`) – whether to use feature matching loss or not

`ashpy.losses.gan.get_adversarial_loss_discriminator` (`adversarial_loss_type=<AdversarialLossType.GAN:I>`)

Return the correct loss for the Discriminator.

Parameters **adversarial_loss_type** (`ashpy.losses.gan.AdversarialLossType`) – Type of loss (`ashpy.losses.gan.AdversarialLossType.GAN` or `ashpy.losses.gan.AdversarialLossType.LSGAN`)

Return type `Type[Executor]`

Returns The correct (`ashpy.losses.executor.Executor`) (to be instantiated).

`ashpy.losses.gan.get_adversarial_loss_generator` (`adversarial_loss_type=<AdversarialLossType.GAN:I>`)

Return the correct loss for the Generator.

Parameters **adversarial_loss_type** (`ashpy.losses.AdversarialLossType`) – Type of loss (`ashpy.losses.AdversarialLossType.GAN` or `ashpy.losses.AdversarialLossType.LSGAN`).

Return type `Type[Executor]`

Returns The correct (`ashpy.losses.executor.Executor`) (to be instantiated).

4.6 ashpy.metrics

Collection of Metrics.

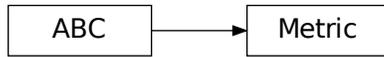
Metric

`metric.Metric`

Metric is the abstract class that every ash Metric must implement.

4.6.1 Metric

Inheritance Diagram



```

class ashpy.metrics.metric.Metric(name, metric, model_selection_operator=None,
                                  logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1.3/

```

Bases: abc.ABC

Metric is the abstract class that every ash Metric must implement.

AshPy Metrics wrap and extend Keras Metrics.

Methods

<code>__init__(name, metric[, ...])</code>	Initialize the Metric object.
<code>json_read(filename)</code>	Read a JSON file.
<code>json_write(filename, what_to_write)</code>	Write inside the specified JSON file the mean and stddev.
<code>log(step)</code>	Log the metric
<code>model_selection(checkpoint, global_step)</code>	Perform model selection.
<code>reset_states()</code>	Reset the state of the metric.
<code>result()</code>	Get the result of the metric.
<code>update_state(context)</code>	Update the internal state of the metric, using the information from the context object.

Attributes

<code>best_folder</code>	Retrieve the folder used to save the best model when doing model selection.
<code>best_model_sel_file</code>	Retrieve the path to JSON file containing the measured performance of the best model.
<code>logdir</code>	Retrieve the log directory.
<code>metric</code>	Retrieve the <code>tf.keras.metrics.Metric</code> object.
<code>model_selection_operator</code>	Retrieve the operator used for model selection.
<code>name</code>	Retrieve the metric name.

```

__init__(name, metric, model_selection_operator=None, logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/
Initialize the Metric object.

```

Parameters

- **name** (*str*) – The name of the metric.
- **metric** (`tf.keras.metrics.Metric`) – The Keras metric to use.

- **model_selection_operator** (`typing.Callable`) – The operation that will be used when `model_selection` is triggered to compare the metrics, used by the `update_state`. Any `typing.Callable` behaving like an `operator` is accepted.

Note: Model selection is done ONLY if an `model_selection_operator` is specified here.

- **logdir** (`str`) – Path to the log dir, defaults to a `log` folder in the current directory.

Return type `None`

property best_folder

Retrieve the folder used to save the best model when doing model selection.

Return type `str`

property best_model_sel_file

Retrieve the path to JSON file containing the measured performance of the best model.

Return type `str`

static json_read (`filename`)

Read a JSON file.

Parameters **filename** (`str`) – The path to the JSON file to read.

Return type `Dict[str, Any]`

Returns `typing.Dict` – Dictionary containing the content of the JSON file.

static json_write (`filename, what_to_write`)

Write inside the specified JSON file the mean and stddev.

Parameters

- **filename** (`str`) – Path to the JSON file to write.
- **what_to_write** (`dict`) – A dictionary containing what to write.

Return type `None`

log (`step`)

Log the metric

Parameters **step** (`int`) – global step of training

Return type `None`

property logdir

Retrieve the log directory.

Return type `str`

property metric

Retrieve the `tf.keras.metrics.Metric` object.

Return type `Metric`

model_selection (`checkpoint, global_step`)

Perform model selection.

Parameters

- **checkpoint** (`tf.train.Checkpoint`) – Checkpoint object that contains the model status.

- `global_step` (`tf.Variable`) – current training step

Return type `None`

property `model_selection_operator`

Retrieve the operator used for model selection.

Return type `Optional[Callable]`

property `name`

Retrieve the metric name.

Return type `str`

reset_states ()

Reset the state of the metric.

Return type `None`

result ()

Get the result of the metric.

Returns `numpy.ndarray` – The current value of the metric.

abstract update_state (`context`)

Update the internal state of the metric, using the information from the context object.

Parameters `context` (`ashpy.contexts.BaseContext`) – An AshPy Context holding all the information the Metric needs.

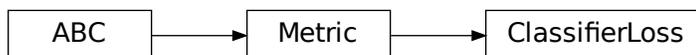
Return type `None`

Classifier

<code>classifier.ClassifierLoss</code>	A handy way to measure the classification loss.
<code>classifier.ClassifierMetric</code>	Wrap a metric using <code>argmax</code> to extract predictions out of a classifier's output.

4.6.2 ClassifierLoss

Inheritance Diagram



class `ashpy.metrics.classifier.ClassifierLoss` (`model_selection_operator=None`,
`logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy'`)
 Bases: `ashpy.metrics.metric.Metric`

A handy way to measure the classification loss.

Methods

<code>__init__</code> ([model_selection_operator, logdir])	Initialize the Metric.
<code>update_state</code> (context)	Update the internal state of the metric, using the information from the context object.

Attributes

<code>best_folder</code>	Retrieve the folder used to save the best model when doing model selection.
<code>best_model_sel_file</code>	Retrieve the path to JSON file containing the measured performance of the best model.
<code>logdir</code>	Retrieve the log directory.
<code>metric</code>	Retrieve the <code>tf.keras.metrics.Metric</code> object.
<code>model_selection_operator</code>	Retrieve the operator used for model selection.
<code>name</code>	Retrieve the metric name.

`__init__` (*model_selection_operator=None, logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1*)
 Initialize the Metric.

Parameters

- **model_selection_operator** (`typing.Callable`) – The operation that will be used when *model_selection* is triggered to compare the metrics, used by the *update_state*. Any `typing.Callable` behaving like an *operator* is accepted.

Note: Model selection is done ONLY if an operator is specified here.

- **logdir** (*str*) – Path to the log dir, defaults to a *log* folder in the current directory.

Return type `None`

update_state (*context*)

Update the internal state of the metric, using the information from the context object.

Parameters **context** (`ashpy.contexts.ClassifierContext`) – An AshPy Context holding all the information the Metric needs.

Return type `None`

4.6.3 ClassifierMetric

Inheritance Diagram



```

class ashpy.metrics.classifier.ClassifierMetric (metric,
                                                model_selection_operator=None,
                                                logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/ashpy/docs',
                                                processing_predictions=None)
    
```

Bases: `ashpy.metrics.metric.Metric`

Wrap a metric using `argmax` to extract predictions out of a classifier's output.

Methods

<code>__init__(metric[, model_selection_operator, ...])</code>	Initialize the Metric.
<code>update_state(context)</code>	Update the internal state of the metric, using the information from the context object.

Attributes

<code>best_folder</code>	Retrieve the folder used to save the best model when doing model selection.
<code>best_model_sel_file</code>	Retrieve the path to JSON file containing the measured performance of the best model.
<code>logdir</code>	Retrieve the log directory.
<code>metric</code>	Retrieve the <code>tf.keras.metrics.Metric</code> object.
<code>model_selection_operator</code>	Retrieve the operator used for model selection.
<code>name</code>	Retrieve the metric name.

```

__init__(metric, model_selection_operator=None, logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/ashpy/docs',
          processing_predictions=None)
    Initialize the Metric.
    
```

Parameters

- **metric** (`tf.keras.metrics.Metric`) – The Keras Metric to use with the classifier (e.g.: `Accuracy()`).
- **model_selection_operator** (`typing.Callable`) – The operation that will be used when `model_selection` is triggered to compare the metrics, used by the `update_state`. Any `typing.Callable` behaving like an `operator` is accepted.

Note: Model selection is done ONLY if an *model_selection_operator* is specified here.

- **logdir** (*str*) – Path to the log dir, defaults to a *log* folder in the current directory.
- **processing_predictions** (*typing.Dict*) – A *dict* in the form of `{“fn”: tf.argmax, “kwargs”: {“axis”: -1}}` with a function “fn” to be used for predictions processing purposes and its “kwargs” as its keyword-arguments. Defaults to `{“fn”: tf.argmax, “kwargs”: {“axis”: -1}}`.

Return type `None`

update_state (*context*)

Update the internal state of the metric, using the information from the context object.

Parameters **context** (`ashpy.contexts.ClassifierContext`) – An AshPy Context holding all the information the Metric needs.

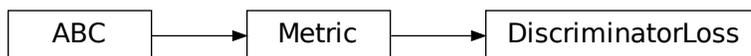
Return type `None`

GAN

<code>gan.DiscriminatorLoss</code>	The Discriminator loss value.
<code>gan.GeneratorLoss</code>	Generator loss value.
<code>gan.EncoderLoss</code>	Encoder Loss value.
<code>gan.InceptionScore</code>	Inception Score Metric.
<code>gan.EncodingAccuracy</code>	Generator and Encoder accuracy performance.
<code>sliced_wasserstein_metric.SlicedWassersteinDistance</code>	Sliced Wasserstein Distance.
<code>ssim_multiscale.SSIM_Multiscale</code>	Multiscale Structural Similarity

4.6.4 DiscriminatorLoss

Inheritance Diagram



class `ashpy.metrics.gan.DiscriminatorLoss` (*model_selection_operator=None*,
logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checko
 Bases: `ashpy.metrics.metric.Metric`
 The Discriminator loss value.

Methods

<code>__init__</code> (<code>model_selection_operator</code> , <code>logdir</code>)	Initialize the Metric.
<code>update_state</code> (<code>context</code>)	Update the internal state of the metric, using the information from the context object.

Attributes

<code>best_folder</code>	Retrieve the folder used to save the best model when doing model selection.
<code>best_model_sel_file</code>	Retrieve the path to JSON file containing the measured performance of the best model.
<code>logdir</code>	Retrieve the log directory.
<code>metric</code>	Retrieve the <code>tf.keras.metrics.Metric</code> object.
<code>model_selection_operator</code>	Retrieve the operator used for model selection.
<code>name</code>	Retrieve the metric name.

`__init__`(`model_selection_operator=None`, `logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1`
Initialize the Metric.

Parameters

- **`model_selection_operator`** (`typing.Callable`) – The operation that will be used when `model_selection` is triggered to compare the metrics, used by the `update_state`. Any `typing.Callable` behaving like an `operator` is accepted.

Note: Model selection is done ONLY if an operator is specified here.

- **`logdir`** (`str`) – Path to the log dir, defaults to a `log` folder in the current directory.

Return type `None`

`update_state` (`context`)

Update the internal state of the metric, using the information from the context object.

Parameters **`context`** (`ashpy.contexts.GANContext`) – An AshPy Context Object that carries all the information the Metric needs.

Return type `None`

4.6.5 GeneratorLoss

Inheritance Diagram



class `ashpy.metrics.gan.GeneratorLoss` (*model_selection_operator=None*,
logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1')

Bases: `ashpy.metrics.metric.Metric`

Generator loss value.

Methods

<code>__init__</code> ([<i>model_selection_operator</i> , <i>logdir</i>])	Initialize the Metric.
<code>update_state</code> (<i>context</i>)	Update the internal state of the metric, using the information from the context object.

Attributes

<code>best_folder</code>	Retrieve the folder used to save the best model when doing model selection.
<code>best_model_sel_file</code>	Retrieve the path to JSON file containing the measured performance of the best model.
<code>logdir</code>	Retrieve the log directory.
<code>metric</code>	Retrieve the <code>tf.keras.metrics.Metric</code> object.
<code>model_selection_operator</code>	Retrieve the operator used for model selection.
<code>name</code>	Retrieve the metric name.

`__init__` (*model_selection_operator=None*, *logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1'*)
 Initialize the Metric.

Parameters

- **model_selection_operator** (`typing.Callable`) – The operation that will be used when *model_selection* is triggered to compare the metrics, used by the *update_state*. Any `typing.Callable` behaving like an *operator* is accepted.

Note: Model selection is done ONLY if an operator is specified here.

- **logdir** (*str*) – Path to the log dir, defaults to a *log* folder in the current directory.

update_state (*context*)

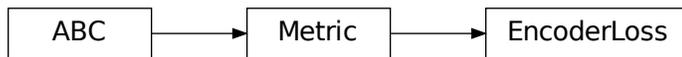
Update the internal state of the metric, using the information from the context object.

Parameters `context` (`ashpy.contexts.GANContext`) – An AshPy Context Object that carries all the information the Metric needs.

Return type `None`

4.6.6 EncoderLoss

Inheritance Diagram



class `ashpy.metrics.gan.EncoderLoss` (`model_selection_operator=None`,
`logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1`
 Bases: `ashpy.metrics.metric.Metric`
 Encoder Loss value.

Methods

<code>__init__</code> (<code>[model_selection_operator, logdir]</code>)	Initialize the Metric.
<code>update_state</code> (<code>context</code>)	Update the internal state of the metric, using the information from the context object.

Attributes

<code>best_folder</code>	Retrieve the folder used to save the best model when doing model selection.
<code>best_model_sel_file</code>	Retrieve the path to JSON file containing the measured performance of the best model.
<code>logdir</code>	Retrieve the log directory.
<code>metric</code>	Retrieve the <code>tf.keras.metrics.Metric</code> object.
<code>model_selection_operator</code>	Retrieve the operator used for model selection.
<code>name</code>	Retrieve the metric name.

`__init__` (`model_selection_operator=None`, `logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1`
 Initialize the Metric.

Parameters

- **`model_selection_operator`** (`typing.Callable`) – The operation that will be used when `model_selection` is triggered to compare the metrics, used by the `update_state`. Any `typing.Callable` behaving like an `operator` is accepted.

Note: Model selection is done ONLY if an operator is specified here.

- **logdir** (*str*) – Path to the log dir, defaults to a *log* folder in the current directory.

Return type `None`

update_state (*context*)

Update the internal state of the metric, using the information from the context object.

Parameters **context** (`ashpy.contexts.GANEncoderContext`) – An AshPy Context Object that carries all the information the Metric needs.

Return type `None`

4.6.7 InceptionScore

Inheritance Diagram



```

class ashpy.metrics.gan.InceptionScore (inception,      model_selection_operator=<built-in
                                     function gt>, logdir='/home/docs/checkouts/readthedocs.org/user_builds/ash
  
```

Bases: `ashpy.metrics.metric.Metric`

Inception Score Metric.

This class is an implementation of the Inception Score technique for evaluating a GAN.

See Improved Techniques for Training GANs¹.

Methods

<code>__init__(inception[, ...])</code>	Initialize the Metric.
<code>get_or_train_inception(dataset, name, ...[, ...])</code>	Restore or train (and save) the Inception model.
<code>inception_score(images)</code>	Compute the Inception Score.
<code>update_state(context)</code>	Update the internal state of the metric, using the information from the context object.

Attributes

¹ Improved Techniques for Training GANs <https://arxiv.org/abs/1606.03498>

<code>best_folder</code>	Retrieve the folder used to save the best model when doing model selection.
<code>best_model_sel_file</code>	Retrieve the path to JSON file containing the measured performance of the best model.
<code>logdir</code>	Retrieve the log directory.
<code>metric</code>	Retrieve the <code>tf.keras.metrics.Metric</code> object.
<code>model_selection_operator</code>	Retrieve the operator used for model selection.
<code>name</code>	Retrieve the metric name.

```
__init__(inception, model_selection_operator=<built-in function gt>,
         logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1.3/docs/source/log')
Initialize the Metric.
```

Parameters

- **inception** (`tf.keras.Model`) – Keras Inception model.
- **model_selection_operator** (`typing.Callable`) – The operation that will be used when `model_selection` is triggered to compare the metrics, used by the `update_state`. Any `typing.Callable` behaving like an `operator` is accepted.

Note: Model selection is done ONLY if an operator is specified here.

- **logdir** (`str`) – Path to the log dir, defaults to a `log` folder in the current directory.

```
static get_or_train_inception(dataset, name, num_classes, epochs, fine_tuning=False,
                             loss_fn=<tensorflow.python.keras.losses.SparseCategoricalCrossentropy
                             object>, optimizer=<tensorflow.python.keras.optimizer_v2.adam.Adam
                             object>, logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkout
```

Restore or train (and save) the Inception model.

Parameters

- **dataset** (`tf.data.Dataset`) – Dataset to re-train Inception Model on.
- **name** (`str`) – Name of this new Inception Model, used for saving it.
- **num_classes** (`int`) – Number of classes to use for classification.
- **epochs** (`int`) – Epochs to train the Inception model for.
- **fine_tuning** (`bool`) – Controls whether the model will be fine-tuned or used as is.
- **loss_fn** (`tf.keras.losses.Loss`) – Keras Loss for the model.
- **optimizer** (`tf.keras.optimizers.Optimizer`) – Keras optimizer for the model.
- **logdir** (`str`) – Path to the log dir, defaults to a `log` folder in the current directory.

Return type `Model`

Returns `tf.keras.Model` – The Inception Model.

```
inception_score(images)
Compute the Inception Score.
```

Parameters **images** (`list` of [`numpy.ndarray`]) – A list of ndarray of generated images of 299x299 of size.

Return type Tensor

Returns tuple of (numpy.ndarray, numpy.ndarray) – Mean and STD.

update_state (context)

Update the internal state of the metric, using the information from the context object.

Parameters context (ashpy.contexts.ClassifierContext) – An AshPy Context holding all the information the Metric needs.

Return type None

4.6.8 EncodingAccuracy

Inheritance Diagram



class ashpy.metrics.gan.**EncodingAccuracy** (classifier, model_selection_operator=None, logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/0.1.3/docs')

Bases: ashpy.metrics.classifier.ClassifierMetric

Generator and Encoder accuracy performance.

Measure the Generator and Encoder performance together, by classifying: $G(E(x))$, y using a pre-trained classifier (on the dataset of x).

Methods

<code>__init__(classifier[, ...])</code>	Measure the Generator and Encoder performance together.
<code>update_state(context)</code>	Update the internal state of the metric, using the information from the context object.

Attributes

<code>best_folder</code>	Retrieve the folder used to save the best model when doing model selection.
<code>best_model_sel_file</code>	Retrieve the path to JSON file containing the measured performance of the best model.
<code>logdir</code>	Retrieve the log directory.
<code>metric</code>	Retrieve the <code>tf.keras.metrics.Metric</code> object.
<code>model_selection_operator</code>	Retrieve the operator used for model selection.

Continued on next page

Table 140 – continued from previous page

name	Retrieve the metric name.
------	---------------------------

`__init__` (*classifier*, *model_selection_operator*=None, *logdir*='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/che...')
 Measure the Generator and Encoder performance together.

This is done by classifying: $G(E(x))$, y using a pre-trained classifier (on the dataset of x).

Parameters

- **classifier** (`tf.keras.Model`) – Keras Model to use as a Classifier to measure the accuracy. Generally assumed to be the Inception Model.
- **model_selection_operator** (`typing.Callable`) – The operation that will be used when *model_selection* is triggered to compare the metrics, used by the *update_state*. Any `typing.Callable` behaving like an `operator` is accepted.

Note: Model selection is done ONLY if an operator is specified here.

- **logdir** (*str*) – Path to the log dir, defaults to a *log* folder in the current directory.

Return type None

update_state (*context*)

Update the internal state of the metric, using the information from the context object.

Parameters **context** (`ashpy.contexts.GANEncoderContext`) – An AshPy Context Object that carries all the information the Metric needs.

Return type None

4.6.9 SlicedWassersteinDistance

Inheritance Diagram



```
class ashpy.metrics.sliced_wasserstein_metric.SlicedWassersteinDistance(model_selection_operator=
    in
    func-
    tion
    lt>,
    logdir='/home/docs/checko
    res-
    o-
    lu-
    tion=128,
    res-
    o-
    lu-
    tion_min=16,
    patches_per_image=64,
    patch_size=7,
    ran-
    dom_sampling_count=1,
    ran-
    dom_projection_dim=147,
    use_svd=False)
```

Bases: *ashpy.metrics.metric.Metric*

Sliced Wasserstein Distance. Used as metric in Progressive Growing of GANs¹.

Methods

<code>__init__([model_selection_operator, logdir, ...])</code>	Initialize the Metric.
<code>log(step)</code>	Log the SWD mean and each sub-metric
<code>model_selection(checkpoint, global_step)</code>	Perform model selection for each sub-metric
<code>reset_states()</code>	Reset the state of the metric and the state of each child metric.
<code>update_state(context)</code>	Update the internal state of the metric, using the information from the context object.

Attributes

<code>best_folder</code>	Retrieve the folder used to save the best model when doing model selection.
<code>best_model_sel_file</code>	Retrieve the path to JSON file containing the measured performance of the best model.
<code>logdir</code>	Retrieve the log directory.
<code>metric</code>	Retrieve the <code>tf.keras.metrics.Metric</code> object.
<code>model_selection_operator</code>	Retrieve the operator used for model selection.
<code>name</code>	Retrieve the metric name.

¹ Progressive Growing of GANs <https://arxiv.org/abs/1710.10196>

`__init__` (*model_selection_operator*=<built-in function lt>, *logdir*='/home/docs/checkouts/readthedocs.org/user_builds/ashpy', *resolution*=128, *resolution_min*=16, *patches_per_image*=64, *patch_size*=7, *random_sampling_count*=1, *random_projection_dim*=147, *use_svd*=False)
Initialize the Metric.

Parameters

- **model_selection_operator** (`typing.Callable`) – The operation that will be used when *model_selection* is triggered to compare the metrics, used by the *update_state*. Any `typing.Callable` behaving like an *operator* is accepted.

Note: Model selection is done ONLY if an operator is specified here.

- **logdir** (*str*) – Path to the log dir, defaults to a *log* folder in the current directory.
- **resolution** (*int*) – Image Resolution, defaults to 128
- **resolution_min** (*int*) – Min Resolution achieved by the metric
- **patches_per_image** (*int*) – Number of patches to extract per image per Laplacian level.
- **patch_size** (*int*) – Width of a square patch.
- **random_sampling_count** (*int*) – Number of random projections to average.
- **random_projection_dim** (*int*) – Dimension of the random projection space.
- **use_svd** (*bool*) – experimental method to compute a more accurate distance.

Return type None

log (*step*)

Log the SWD mean and each sub-metric

model_selection (*checkpoint*, *global_step*)

Perform model selection for each sub-metric

Return type None

reset_states ()

Reset the state of the metric and the state of each child metric.

Return type None

update_state (*context*)

Update the internal state of the metric, using the information from the context object.

Parameters context (`ashpy.contexts.GANContext`) – An AshPy Context Object that carries all the information the Metric needs.

Return type None

4.6.10 SSIM_Multiscale

Inheritance Diagram



```

class ashpy.metrics.ssim_multiscale.SSIM_Multiscale(model_selection_operator=<built-in function lt>,
                                                    logdir='/home/docs/checkouts/readthedocs.org/user_buil
                                                    max_val=2.0,
                                                    power_factors=None, filter
                                                    ter_size=11, filter_sigma=1.5,
                                                    k1=0.01, k2=0.03)
    
```

Bases: *ashpy.metrics.metric.Metric*

Multiscale Structural Similarity

See Multiscale structural similarity for image quality assessment¹

Methods

<code>__init__</code> (<i>model_selection_operator</i> , <i>logdir</i> , ...)	Initialize the Metric.
<code>split_batch</code> (<i>batch</i>)	Split a batch along axis 0 into two tensors having the same size
<code>update_state</code> (<i>context</i>)	Update the internal state of the metric, using the information from the context object.

Attributes

<code>best_folder</code>	Retrieve the folder used to save the best model when doing model selection.
<code>best_model_sel_file</code>	Retrieve the path to JSON file containing the measured performance of the best model.
<code>logdir</code>	Retrieve the log directory.
<code>metric</code>	Retrieve the <code>tf.keras.metrics.Metric</code> object.
<code>model_selection_operator</code>	Retrieve the operator used for model selection.
<code>name</code>	Retrieve the metric name.

```

__init__(model_selection_operator=<built-in function lt>, logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy
        max_val=2.0, power_factors=None, filter_size=11, filter_sigma=1.5, k1=0.01, k2=0.03)
    Initialize the Metric.
    
```

Parameters

¹ Multiscale structural similarity for image quality assessment <https://ieeexplore.ieee.org/document/1292216>

- **model_selection_operator** (`typing.Callable`) – The operation that will be used when `model_selection` is triggered to compare the metrics, used by the `update_state`. Any `typing.Callable` behaving like an `operator` is accepted.

Note: Model selection is done ONLY if an operator is specified here.

- **logdir** (`str`) – Path to the log dir, defaults to a `log` folder in the current directory.
- **max_val** (`float`) – The dynamic range of the images (i.e., the difference between the maximum the and minimum) (see www.tensorflow.org/versions/r2.0/api_docs/python/tf/image/ssim_multiscale)
- **power_factors** (`List[float]`) – Iterable of weights for each of the scales. The number of scales used is the length of the list. Index 0 is the unscaled resolution's weight and each increasing scale corresponds to the image being downsampled by 2. Defaults to (0.0448, 0.2856, 0.3001, 0.2363, 0.1333), which are the values obtained in the original paper.
- **filter_size** (`int`) – Default value 11 (size of gaussian filter).
- **filter_sigma** (`int`) – Default value 1.5 (width of gaussian filter).
- **k1** (`int`) – Default value 0.01
- **k2** (`int`) – Default value 0.03 (SSIM is less sensitivity to K2 for lower values, so it would be better if we taken the values in range of $0 < K2 < 0.4$).

Return type `None`

static split_batch (`batch`)

Split a batch along axis 0 into two tensors having the same size

Parameters `batch` (`tf.Tensor`) – A batch of images

Return type `Tuple[Tensor, Tensor]`

Returns (`Tuple[tf.Tensor, tf.Tensor]`) The batch split in two tensors

Raises `ValueError` – if the batch has size 1

update_state (`context`)

Update the internal state of the metric, using the information from the context object.

Parameters `context` (`ashpy.contexts.GANContext`) – An AshPy Context Object that carries all the information the Metric needs.

Return type `None`

Modules

<code>classifier</code>	The classification metrics.
<code>gan</code>	GAN metrics.
<code>metric</code>	Metric is the abstract class that every ash metric must implement.
<code>sliced_wasserstein_metric</code>	Sliced Wasserstein Distance metric.
<code>ssim_multiscale</code>	Multiscale Structural Similarity metric.

4.6.11 classifier

The classification metrics.

Classes

<code>ClassifierLoss</code>	A handy way to measure the classification loss.
<code>ClassifierMetric</code>	Wrap a metric using <i>argmax</i> to extract predictions out of a classifier's output.

ClassifierLoss

Inheritance Diagram



class `ashpy.metrics.classifier.ClassifierLoss` (*model_selection_operator=None*,
logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy')

Bases: `ashpy.metrics.metric.Metric`

A handy way to measure the classification loss.

Methods

<code>__init__</code> ([<i>model_selection_operator</i> , <i>logdir</i>])	Initialize the Metric.
<code>update_state</code> (context)	Update the internal state of the metric, using the information from the context object.

Attributes

<code>best_folder</code>	Retrieve the folder used to save the best model when doing model selection.
<code>best_model_sel_file</code>	Retrieve the path to JSON file containing the measured performance of the best model.
<code>logdir</code>	Retrieve the log directory.
<code>metric</code>	Retrieve the <code>tf.keras.metrics.Metric</code> object.
<code>model_selection_operator</code>	Retrieve the operator used for model selection.
<code>name</code>	Retrieve the metric name.

`__init__` (*model_selection_operator=None, logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1*
Initialize the Metric.

Parameters

- **model_selection_operator** (*typing.Callable*) – The operation that will be used when *model_selection* is triggered to compare the metrics, used by the *update_state*. Any *typing.Callable* behaving like an *operator* is accepted.

Note: Model selection is done ONLY if an operator is specified here.

- **logdir** (*str*) – Path to the log dir, defaults to a *log* folder in the current directory.

Return type *None*

update_state (*context*)

Update the internal state of the metric, using the information from the context object.

Parameters **context** (*ashpy.contexts.ClassifierContext*) – An AshPy Context holding all the information the Metric needs.

Return type *None*

ClassifierMetric

Inheritance Diagram



class `ashpy.metrics.classifier.ClassifierMetric` (*metric,*
model_selection_operator=None,
logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1
processing_predictions=None)

Bases: *ashpy.metrics.metric.Metric*

Wrap a metric using *argmax* to extract predictions out of a classifier's output.

Methods

<code>__init__</code> (<i>metric[, model_selection_operator, ...]</i>)	Initialize the Metric.
<code>update_state</code> (<i>context</i>)	Update the internal state of the metric, using the information from the context object.

Attributes

<code>best_folder</code>	Retrieve the folder used to save the best model when doing model selection.
<code>best_model_sel_file</code>	Retrieve the path to JSON file containing the measured performance of the best model.
<code>logdir</code>	Retrieve the log directory.
<code>metric</code>	Retrieve the <code>tf.keras.metrics.Metric</code> object.
<code>model_selection_operator</code>	Retrieve the operator used for model selection.
<code>name</code>	Retrieve the metric name.

`__init__` (*metric*, *model_selection_operator*=None, *logdir*='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1.3', *processing_predictions*=None)
 Initialize the Metric.

Parameters

- **metric** (`tf.keras.metrics.Metric`) – The Keras Metric to use with the classifier (e.g.: Accuracy()).
- **model_selection_operator** (`typing.Callable`) – The operation that will be used when *model_selection* is triggered to compare the metrics, used by the *update_state*. Any `typing.Callable` behaving like an *operator* is accepted.

Note: Model selection is done ONLY if an *model_selection_operator* is specified here.

- **logdir** (*str*) – Path to the log dir, defaults to a *log* folder in the current directory.
- **processing_predictions** (`typing.Dict`) – A *dict* in the form of `{"fn": tf.argmax, "kwargs": {"axis": -1}}` with a function “fn” to be used for predictions processing purposes and its “kwargs” as its keyword-arguments. Defaults to `{"fn": tf.argmax, "kwargs": {"axis": -1}}`.

Return type None

update_state (*context*)

Update the internal state of the metric, using the information from the context object.

Parameters **context** (`ashpy.contexts.ClassifierContext`) – An AshPy Context holding all the information the Metric needs.

Return type None

class `ashpy.metrics.classifier.ClassifierLoss` (*model_selection_operator*=None, *logdir*='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1.3')

Bases: `ashpy.metrics.metric.Metric`

A handy way to measure the classification loss.

`__init__` (*model_selection_operator*=None, *logdir*='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1.3', *processing_predictions*=None)
 Initialize the Metric.

Parameters

- **model_selection_operator** (`typing.Callable`) – The operation that will be used when *model_selection* is triggered to compare the metrics, used by the *update_state*. Any `typing.Callable` behaving like an *operator* is accepted.

Note: Model selection is done ONLY if an operator is specified here.

- **logdir** (*str*) – Path to the log dir, defaults to a *log* folder in the current directory.

Return type `None`

update_state (*context*)

Update the internal state of the metric, using the information from the context object.

Parameters **context** (`ashpy.contexts.ClassifierContext`) – An AshPy Context holding all the information the Metric needs.

Return type `None`

```
class ashpy.metrics.classifier.ClassifierMetric (metric,
                                                model_selection_operator=None,
                                                logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/ashpy',
                                                processing_predictions=None)
```

Bases: `ashpy.metrics.metric.Metric`

Wrap a metric using *argmax* to extract predictions out of a classifier's output.

```
__init__ (metric, model_selection_operator=None, logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/ashpy',
          processing_predictions=None)
Initialize the Metric.
```

Parameters

- **metric** (`tf.keras.metrics.Metric`) – The Keras Metric to use with the classifier (e.g.: `Accuracy()`).
- **model_selection_operator** (`typing.Callable`) – The operation that will be used when *model_selection* is triggered to compare the metrics, used by the *update_state*. Any `typing.Callable` behaving like an *operator* is accepted.

Note: Model selection is done ONLY if an *model_selection_operator* is specified here.

- **logdir** (*str*) – Path to the log dir, defaults to a *log* folder in the current directory.
- **processing_predictions** (`typing.Dict`) – A *dict* in the form of `{"fn": tf.argmax, "kwargs": {"axis": -1}}` with a function “*fn*” to be used for predictions processing purposes and its “*kwargs*” as its keyword-arguments. Defaults to `{"fn": tf.argmax, "kwargs": {"axis": -1}}`.

Return type `None`

update_state (*context*)

Update the internal state of the metric, using the information from the context object.

Parameters **context** (`ashpy.contexts.ClassifierContext`) – An AshPy Context holding all the information the Metric needs.

Return type `None`

4.6.12 gan

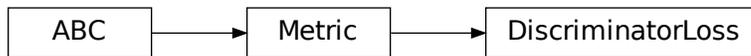
GAN metrics.

Classes

<i>DiscriminatorLoss</i>	The Discriminator loss value.
<i>EncoderLoss</i>	Encoder Loss value.
<i>EncodingAccuracy</i>	Generator and Encoder accuracy performance.
<i>GeneratorLoss</i>	Generator loss value.
<i>InceptionScore</i>	Inception Score Metric.

DiscriminatorLoss

Inheritance Diagram



```

class ashpy.metrics.gan.DiscriminatorLoss (model_selection_operator=None,
                                           logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1.3')
  
```

Bases: *ashpy.metrics.metric.Metric*

The Discriminator loss value.

Methods

<code>__init__</code> ([model_selection_operator, logdir])	Initialize the Metric.
<code>update_state</code> (context)	Update the internal state of the metric, using the information from the context object.

Attributes

<code>best_folder</code>	Retrieve the folder used to save the best model when doing model selection.
<code>best_model_sel_file</code>	Retrieve the path to JSON file containing the measured performance of the best model.
<code>logdir</code>	Retrieve the log directory.
<code>metric</code>	Retrieve the <code>tf.keras.metrics.Metric</code> object.
<code>model_selection_operator</code>	Retrieve the operator used for model selection.
<code>name</code>	Retrieve the metric name.

```

__init__(model_selection_operator=None, logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1.3')
  Initialize the Metric.
  
```

Parameters

- **model_selection_operator** (`typing.Callable`) – The operation that will be used when `model_selection` is triggered to compare the metrics, used by the `update_state`. Any `typing.Callable` behaving like an `operator` is accepted.

Note: Model selection is done ONLY if an operator is specified here.

- **logdir** (`str`) – Path to the log dir, defaults to a `log` folder in the current directory.

Return type `None`

update_state (`context`)

Update the internal state of the metric, using the information from the context object.

Parameters **context** (`ashpy.contexts.GANContext`) – An AshPy Context Object that carries all the information the Metric needs.

Return type `None`

EncoderLoss

Inheritance Diagram



class `ashpy.metrics.gan.EncoderLoss` (`model_selection_operator=None`,
`logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1`
 Bases: `ashpy.metrics.metric.Metric`
 Encoder Loss value.

Methods

<code>__init__</code> ([<code>model_selection_operator</code> , <code>logdir</code>])	Initialize the Metric.
<code>update_state</code> (<code>context</code>)	Update the internal state of the metric, using the information from the context object.

Attributes

<code>best_folder</code>	Retrieve the folder used to save the best model when doing model selection.
<code>best_model_sel_file</code>	Retrieve the path to JSON file containing the measured performance of the best model.
<code>logdir</code>	Retrieve the log directory.

Continued on next page

Table 155 – continued from previous page

<code>metric</code>	Retrieve the <code>tf.keras.metrics.Metric</code> object.
<code>model_selection_operator</code>	Retrieve the operator used for model selection.
<code>name</code>	Retrieve the metric name.

`__init__` (*model_selection_operator=None, logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1*)
 Initialize the Metric.

Parameters

- **model_selection_operator** (`typing.Callable`) – The operation that will be used when *model_selection* is triggered to compare the metrics, used by the *update_state*. Any `typing.Callable` behaving like an *operator* is accepted.

Note: Model selection is done ONLY if an operator is specified here.

- **logdir** (*str*) – Path to the log dir, defaults to a *log* folder in the current directory.

Return type `None`

update_state (*context*)

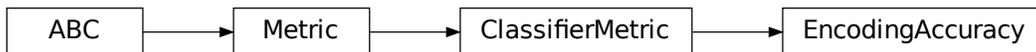
Update the internal state of the metric, using the information from the context object.

Parameters **context** (`ashpy.contexts.GANEncoderContext`) – An AshPy Context Object that carries all the information the Metric needs.

Return type `None`

EncodingAccuracy

Inheritance Diagram



class `ashpy.metrics.gan.EncodingAccuracy` (*classifier, model_selection_operator=None, logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1*)
 Bases: `ashpy.metrics.classifier.ClassifierMetric`

Generator and Encoder accuracy performance.

Measure the Generator and Encoder performance together, by classifying: $G(E(x))$, y using a pre-trained classifier (on the dataset of x).

Methods

<code>__init__(classifier[, ...])</code>	Measure the Generator and Encoder performance together.
<code>update_state(context)</code>	Update the internal state of the metric, using the information from the context object.

Attributes

<code>best_folder</code>	Retrieve the folder used to save the best model when doing model selection.
<code>best_model_sel_file</code>	Retrieve the path to JSON file containing the measured performance of the best model.
<code>logdir</code>	Retrieve the log directory.
<code>metric</code>	Retrieve the <code>tf.keras.metrics.Metric</code> object.
<code>model_selection_operator</code>	Retrieve the operator used for model selection.
<code>name</code>	Retrieve the metric name.

`__init__(classifier, model_selection_operator=None, logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/ashpy')`
 Measure the Generator and Encoder performance together.

This is done by classifying: $G(E(x))$, y using a pre-trained classifier (on the dataset of x).

Parameters

- **classifier** (`tf.keras.Model`) – Keras Model to use as a Classifier to measure the accuracy. Generally assumed to be the Inception Model.
- **model_selection_operator** (`typing.Callable`) – The operation that will be used when `model_selection` is triggered to compare the metrics, used by the `update_state`. Any `typing.Callable` behaving like an `operator` is accepted.

Note: Model selection is done ONLY if an operator is specified here.

- **logdir** (`str`) – Path to the log dir, defaults to a `log` folder in the current directory.

Return type `None`

update_state (`context`)

Update the internal state of the metric, using the information from the context object.

Parameters **context** (`ashpy.contexts.GANEncoderContext`) – An AshPy Context Object that carries all the information the Metric needs.

Return type `None`

GeneratorLoss

Inheritance Diagram



class `ashpy.metrics.gan.GeneratorLoss` (*model_selection_operator=None*,
logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1')

Bases: `ashpy.metrics.metric.Metric`

Generator loss value.

Methods

<code>__init__</code> ([<i>model_selection_operator</i> , <i>logdir</i>])	Initialize the Metric.
<code>update_state</code> (<i>context</i>)	Update the internal state of the metric, using the information from the context object.

Attributes

<code>best_folder</code>	Retrieve the folder used to save the best model when doing model selection.
<code>best_model_sel_file</code>	Retrieve the path to JSON file containing the measured performance of the best model.
<code>logdir</code>	Retrieve the log directory.
<code>metric</code>	Retrieve the <code>tf.keras.metrics.Metric</code> object.
<code>model_selection_operator</code>	Retrieve the operator used for model selection.
<code>name</code>	Retrieve the metric name.

`__init__` (*model_selection_operator=None*, *logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1'*)
 Initialize the Metric.

Parameters

- **model_selection_operator** (`typing.Callable`) – The operation that will be used when *model_selection* is triggered to compare the metrics, used by the *update_state*. Any `typing.Callable` behaving like an *operator* is accepted.

Note: Model selection is done ONLY if an operator is specified here.

- **logdir** (*str*) – Path to the log dir, defaults to a *log* folder in the current directory.

update_state (*context*)

Update the internal state of the metric, using the information from the context object.

Parameters `context` (`ashpy.contexts.GANContext`) – An AshPy Context Object that carries all the information the Metric needs.

Return type `None`

InceptionScore

Inheritance Diagram



class `ashpy.metrics.gan.InceptionScore` (`inception`, `model_selection_operator=<built-in function gt>`, `logdir='<code>/home/docs/checkouts/readthedocs.org/user_builds/ashpy</code>'`)

Bases: `ashpy.metrics.metric.Metric`

Inception Score Metric.

This class is an implementation of the Inception Score technique for evaluating a GAN.

See Improved Techniques for Training GANs¹.

Methods

<code>__init__(inception[, ...])</code>	Initialize the Metric.
<code>get_or_train_inception(dataset, name, ...[, ...])</code>	Restore or train (and save) the Inception model.
<code>inception_score(images)</code>	Compute the Inception Score.
<code>update_state(context)</code>	Update the internal state of the metric, using the information from the context object.

Attributes

<code>best_folder</code>	Retrieve the folder used to save the best model when doing model selection.
<code>best_model_sel_file</code>	Retrieve the path to JSON file containing the measured performance of the best model.
<code>logdir</code>	Retrieve the log directory.
<code>metric</code>	Retrieve the <code>tf.keras.metrics.Metric</code> object.
<code>model_selection_operator</code>	Retrieve the operator used for model selection.
<code>name</code>	Retrieve the metric name.

¹ Improved Techniques for Training GANs <https://arxiv.org/abs/1606.03498>

`__init__(inception, model_selection_operator=<built-in function gt>, logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1.3/docs/source/log')`
Initialize the Metric.

Parameters

- **inception** (`tf.keras.Model`) – Keras Inception model.
- **model_selection_operator** (`typing.Callable`) – The operation that will be used when *model_selection* is triggered to compare the metrics, used by the *update_state*. Any `typing.Callable` behaving like an *operator* is accepted.

Note: Model selection is done ONLY if an operator is specified here.

- **logdir** (`str`) – Path to the log dir, defaults to a *log* folder in the current directory.

`static get_or_train_inception(dataset, name, num_classes, epochs, fine_tuning=False, loss_fn=<tensorflow.python.keras.losses.SparseCategoricalCrossentropy object>, optimizer=<tensorflow.python.keras.optimizer_v2.adam.Adam object>, logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts')`
Restore or train (and save) the Inception model.

Parameters

- **dataset** (`tf.data.Dataset`) – Dataset to re-train Inception Model on.
- **name** (`str`) – Name of this new Inception Model, used for saving it.
- **num_classes** (`int`) – Number of classes to use for classification.
- **epochs** (`int`) – Epochs to train the Inception model for.
- **fine_tuning** (`bool`) – Controls whether the model will be fine-tuned or used as is.
- **loss_fn** (`tf.keras.losses.Loss`) – Keras Loss for the model.
- **optimizer** (`tf.keras.optimizers.Optimizer`) – Keras optimizer for the model.
- **logdir** (`str`) – Path to the log dir, defaults to a *log* folder in the current directory.

Return type `Model`

Returns `tf.keras.Model` – The Inception Model.

`inception_score(images)`

Compute the Inception Score.

Parameters **images** (`list` of [`numpy.ndarray`]) – A list of `numpy.ndarray` of generated images of 299x299 of size.

Return type `Tensor`

Returns `tuple` of (`numpy.ndarray`, `numpy.ndarray`) – Mean and STD.

`update_state(context)`

Update the internal state of the metric, using the information from the context object.

Parameters **context** (`ashpy.contexts.ClassifierContext`) – An AshPy Context holding all the information the Metric needs.

Return type `None`

class `ashpy.metrics.gan.DiscriminatorLoss` (*model_selection_operator=None*,
logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1')

Bases: `ashpy.metrics.metric.Metric`

The Discriminator loss value.

__init__ (*model_selection_operator=None*, *logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1'*)
 Initialize the Metric.

Parameters

- **model_selection_operator** (`typing.Callable`) – The operation that will be used when *model_selection* is triggered to compare the metrics, used by the *update_state*. Any `typing.Callable` behaving like an *operator* is accepted.

Note: Model selection is done ONLY if an operator is specified here.

- **logdir** (*str*) – Path to the log dir, defaults to a *log* folder in the current directory.

Return type `None`

update_state (*context*)

Update the internal state of the metric, using the information from the context object.

Parameters **context** (`ashpy.contexts.GANContext`) – An AshPy Context Object that carries all the information the Metric needs.

Return type `None`

class `ashpy.metrics.gan.EncoderLoss` (*model_selection_operator=None*,
logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1')

Bases: `ashpy.metrics.metric.Metric`

Encoder Loss value.

__init__ (*model_selection_operator=None*, *logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1'*)
 Initialize the Metric.

Parameters

- **model_selection_operator** (`typing.Callable`) – The operation that will be used when *model_selection* is triggered to compare the metrics, used by the *update_state*. Any `typing.Callable` behaving like an *operator* is accepted.

Note: Model selection is done ONLY if an operator is specified here.

- **logdir** (*str*) – Path to the log dir, defaults to a *log* folder in the current directory.

Return type `None`

update_state (*context*)

Update the internal state of the metric, using the information from the context object.

Parameters **context** (`ashpy.contexts.GANEncoderContext`) – An AshPy Context Object that carries all the information the Metric needs.

Return type `None`

class `ashpy.metrics.gan.EncodingAccuracy` (*classifier*, *model_selection_operator=None*,
logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1')

Bases: `ashpy.metrics.classifier.ClassifierMetric`

Generator and Encoder accuracy performance.

Measure the Generator and Encoder performance together, by classifying: $G(E(x))$, y using a pre-trained classifier (on the dataset of x).

`__init__` (*classifier*, *model_selection_operator*=None, *logdir*='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1.3')
 Measure the Generator and Encoder performance together.

This is done by classifying: $G(E(x))$, y using a pre-trained classifier (on the dataset of x).

Parameters

- **classifier** (`tf.keras.Model`) – Keras Model to use as a Classifier to measure the accuracy. Generally assumed to be the Inception Model.
- **model_selection_operator** (`typing.Callable`) – The operation that will be used when *model_selection* is triggered to compare the metrics, used by the *update_state*. Any `typing.Callable` behaving like an `operator` is accepted.

Note: Model selection is done ONLY if an operator is specified here.

- **logdir** (*str*) – Path to the log dir, defaults to a *log* folder in the current directory.

Return type None

`update_state` (*context*)

Update the internal state of the metric, using the information from the context object.

Parameters **context** (`ashpy.contexts.GANEncoderContext`) – An AshPy Context Object that carries all the information the Metric needs.

Return type None

class `ashpy.metrics.gan.GeneratorLoss` (*model_selection_operator*=None, *logdir*='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1.3')

Bases: `ashpy.metrics.metric.Metric`

Generator loss value.

`__init__` (*model_selection_operator*=None, *logdir*='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1.3')
 Initialize the Metric.

Parameters

- **model_selection_operator** (`typing.Callable`) – The operation that will be used when *model_selection* is triggered to compare the metrics, used by the *update_state*. Any `typing.Callable` behaving like an `operator` is accepted.

Note: Model selection is done ONLY if an operator is specified here.

- **logdir** (*str*) – Path to the log dir, defaults to a *log* folder in the current directory.

`update_state` (*context*)

Update the internal state of the metric, using the information from the context object.

Parameters **context** (`ashpy.contexts.GANContext`) – An AshPy Context Object that carries all the information the Metric needs.

Return type None

```
class ashpy.metrics.gan.InceptionScore (inception,      model_selection_operator=<built-in
                                     function gt>, logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy')
```

Bases: `ashpy.metrics.metric.Metric`

Inception Score Metric.

This class is an implementation of the Inception Score technique for evaluating a GAN.

See Improved Techniques for Training GANs¹.

```
__init__ (inception,      model_selection_operator=<built-in      function      gt>,
          logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1.3/docs/source/log')
Initialize the Metric.
```

Parameters

- **inception** (`tf.keras.Model`) – Keras Inception model.
- **model_selection_operator** (`typing.Callable`) – The operation that will be used when `model_selection` is triggered to compare the metrics, used by the `update_state`. Any `typing.Callable` behaving like an `operator` is accepted.

Note: Model selection is done ONLY if an operator is specified here.

- **logdir** (`str`) – Path to the log dir, defaults to a `log` folder in the current directory.

```
static get_or_train_inception (dataset, name, num_classes, epochs, fine_tuning=False,
                                loss_fn=<tensorflow.python.keras.losses.SparseCategoricalCrossentropy
                                object>, optimizer=<tensorflow.python.keras.optimizer_v2.adam.Adam
                                object>, logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1.3/docs/source/log')
Restore or train (and save) the Inception model.
```

Parameters

- **dataset** (`tf.data.Dataset`) – Dataset to re-train Inception Model on.
- **name** (`str`) – Name of this new Inception Model, used for saving it.
- **num_classes** (`int`) – Number of classes to use for classification.
- **epochs** (`int`) – Epochs to train the Inception model for.
- **fine_tuning** (`bool`) – Controls whether the model will be fine-tuned or used as is.
- **loss_fn** (`tf.keras.losses.Loss`) – Keras Loss for the model.
- **optimizer** (`tf.keras.optimizers.Optimizer`) – Keras optimizer for the model.
- **logdir** (`str`) – Path to the log dir, defaults to a `log` folder in the current directory.

Return type `Model`

Returns `tf.keras.Model` – The Inception Model.

```
inception_score (images)
Compute the Inception Score.
```

Parameters **images** (`list` of [`numpy.ndarray`]) – A list of `ndarray` of generated images of 299x299 of size.

Return type `Tensor`

¹ Improved Techniques for Training GANs <https://arxiv.org/abs/1606.03498>

Returns `tuple` of `(numpy.ndarray, numpy.ndarray)` – Mean and STD.

update_state (*context*)

Update the internal state of the metric, using the information from the context object.

Parameters **context** (`ashpy.contexts.ClassifierContext`) – An AshPy Context holding all the information the Metric needs.

Return type `None`

4.6.13 metric

Metric is the abstract class that every ash metric must implement.

Classes

<i>Metric</i>	Metric is the abstract class that every ash Metric must implement.
---------------	--

Metric

Inheritance Diagram



```

class ashpy.metrics.metric.Metric(name, metric, model_selection_operator=None,
                                  logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1.3/

```

Bases: `abc.ABC`

Metric is the abstract class that every ash Metric must implement.

AshPy Metrics wrap and extend Keras Metrics.

Methods

<code>__init__(name, metric[, ...])</code>	Initialize the Metric object.
<code>json_read(filename)</code>	Read a JSON file.
<code>json_write(filename, what_to_write)</code>	Write inside the specified JSON file the mean and stddev.
<code>log(step)</code>	Log the metric
<code>model_selection(checkpoint, global_step)</code>	Perform model selection.
<code>reset_states()</code>	Reset the state of the metric.
<code>result()</code>	Get the result of the metric.

Continued on next page

Table 163 – continued from previous page

<code>update_state(context)</code>	Update the internal state of the metric, using the information from the context object.
------------------------------------	---

Attributes

<code>best_folder</code>	Retrieve the folder used to save the best model when doing model selection.
<code>best_model_sel_file</code>	Retrieve the path to JSON file containing the measured performance of the best model.
<code>logdir</code>	Retrieve the log directory.
<code>metric</code>	Retrieve the <code>tf.keras.metrics.Metric</code> object.
<code>model_selection_operator</code>	Retrieve the operator used for model selection.
<code>name</code>	Retrieve the metric name.

`__init__(name, metric, model_selection_operator=None, logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/')`
Initialize the Metric object.

Parameters

- **name** (`str`) – The name of the metric.
- **metric** (`tf.keras.metrics.Metric`) – The Keras metric to use.
- **model_selection_operator** (`typing.Callable`) – The operation that will be used when `model_selection` is triggered to compare the metrics, used by the `update_state`. Any `typing.Callable` behaving like an `operator` is accepted.

Note: Model selection is done ONLY if an `model_selection_operator` is specified here.

- **logdir** (`str`) – Path to the log dir, defaults to a `log` folder in the current directory.

Return type `None`

property best_folder

Retrieve the folder used to save the best model when doing model selection.

Return type `str`

property best_model_sel_file

Retrieve the path to JSON file containing the measured performance of the best model.

Return type `str`

static json_read(filename)

Read a JSON file.

Parameters **filename** (`str`) – The path to the JSON file to read.

Return type `Dict[str, Any]`

Returns `typing.Dict` – Dictionary containing the content of the JSON file.

static json_write(filename, what_to_write)

Write inside the specified JSON file the mean and stddev.

Parameters

- **filename** (*str*) – Path to the JSON file to write.
- **what_to_write** (*dict*) – A dictionary containing what to write.

Return type None

log (*step*)
Log the metric

Parameters **step** (*int*) – global step of training

Return type None

property logdir
Retrieve the log directory.

Return type *str*

property metric
Retrieve the `tf.keras.metrics.Metric` object.

Return type *Metric*

model_selection (*checkpoint, global_step*)
Perform model selection.

Parameters

- **checkpoint** (`tf.train.Checkpoint`) – Checkpoint object that contains the model status.
- **global_step** (`tf.Variable`) – current training step

Return type None

property model_selection_operator
Retrieve the operator used for model selection.

Return type `Optional[Callable]`

property name
Retrieve the metric name.

Return type *str*

reset_states ()
Reset the state of the metric.

Return type None

result ()
Get the result of the metric.

Returns `numpy.ndarray` – The current value of the metric.

abstract update_state (*context*)
Update the internal state of the metric, using the information from the context object.

Parameters **context** (`ashpy.contexts.BaseContext`) – An AshPy Context holding all the information the Metric needs.

Return type *None*

class `ashpy.metrics.metric.Metric` (*name, metric, model_selection_operator=None, logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1.3/*

Bases: `abc.ABC`

Metric is the abstract class that every ash Metric must implement.

AshPy Metrics wrap and extend Keras Metrics.

__init__ (*name*, *metric*, *model_selection_operator*=None, *logdir*='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/...')
Initialize the Metric object.

Parameters

- **name** (*str*) – The name of the metric.
- **metric** (`tf.keras.metrics.Metric`) – The Keras metric to use.
- **model_selection_operator** (`typing.Callable`) – The operation that will be used when *model_selection* is triggered to compare the metrics, used by the *update_state*. Any `typing.Callable` behaving like an *operator* is accepted.

Note: Model selection is done ONLY if an *model_selection_operator* is specified here.

- **logdir** (*str*) – Path to the log dir, defaults to a *log* folder in the current directory.

Return type None

property best_folder

Retrieve the folder used to save the best model when doing model selection.

Return type `str`

property best_model_sel_file

Retrieve the path to JSON file containing the measured performance of the best model.

Return type `str`

static json_read (*filename*)

Read a JSON file.

Parameters **filename** (*str*) – The path to the JSON file to read.

Return type `Dict[str, Any]`

Returns `typing.Dict` – Dictionary containing the content of the JSON file.

static json_write (*filename*, *what_to_write*)

Write inside the specified JSON file the mean and stddev.

Parameters

- **filename** (*str*) – Path to the JSON file to write.
- **what_to_write** (*dict*) – A dictionary containing what to write.

Return type None

log (*step*)

Log the metric

Parameters **step** (*int*) – global step of training

Return type None

property logdir

Retrieve the log directory.

Return type `str`

property metric

Retrieve the `tf.keras.metrics.Metric` object.

Return type `Metric`

model_selection (*checkpoint*, *global_step*)

Perform model selection.

Parameters

- **checkpoint** (`tf.train.Checkpoint`) – Checkpoint object that contains the model status.
- **global_step** (`tf.Variable`) – current training step

Return type `None`

property model_selection_operator

Retrieve the operator used for model selection.

Return type `Optional[Callable]`

property name

Retrieve the metric name.

Return type `str`

reset_states ()

Reset the state of the metric.

Return type `None`

result ()

Get the result of the metric.

Returns `numpy.ndarray` – The current value of the metric.

abstract update_state (*context*)

Update the internal state of the metric, using the information from the context object.

Parameters context (`ashpy.contexts.BaseContext`) – An AshPy Context holding all the information the Metric needs.

Return type `None`

4.6.14 sliced_wasserstein_metric

Sliced Wasserstein Distance metric.

Classes

<i>SingleSWD</i>	SlicedWassersteinDistance for a certain level of the pyramid.
<i>SlicedWassersteinDistance</i>	Sliced Wasserstein Distance.

SingleSWD

Inheritance Diagram



```

class ashpy.metrics.sliced_wasserstein_metric.SingleSWD (model_selection_operator=<built-
in function lt>,
logdir='/home/docs/checkouts/readthedocs.org/user_
level_of_pyramid=0,
real_or_fake='fake')
  
```

Bases: `ashpy.metrics.metric.Metric`

SlicedWassersteinDistance for a certain level of the pyramid.

Methods

<code>__init__</code> ([model_selection_operator, logdir, ...])	Initialize the Metric.
<code>update_state</code> (context, score)	Update the internal state of the metric, using the information from the context object.

Attributes

<code>best_folder</code>	Retrieve the folder used to save the best model when doing model selection.
<code>best_model_sel_file</code>	Retrieve the path to JSON file containing the measured performance of the best model.
<code>logdir</code>	Retrieve the log directory.
<code>metric</code>	Retrieve the <code>tf.keras.metrics.Metric</code> object.
<code>model_selection_operator</code>	Retrieve the operator used for model selection.
<code>name</code>	Retrieve the metric name.

```

__init__ (model_selection_operator=<built-in function lt>, logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy
level_of_pyramid=0, real_or_fake='fake')
Initialize the Metric.
  
```

Parameters

- **model_selection_operator** (`typing.Callable`) – The operation that will be used when `model_selection` is triggered to compare the metrics, used by the `update_state`. Any `typing.Callable` behaving like an `operator` is accepted.

Note: Model selection is done ONLY if an operator is specified here.

- **logdir** (*str*) – Path to the log dir, defaults to a *log* folder in the current directory.
- **level_of_pyramid** (*int*) – Level of the pyramid related to this metric
- **real_or_fake** (*str*) – string identifying this metric (real or fake distance)

Return type `None`

update_state (*context, score*)

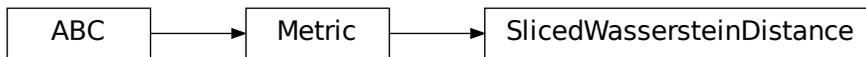
Update the internal state of the metric, using the information from the context object.

Parameters **context** (`ashpy.contexts.GANContext`) – An AshPy Context Object that carries all the information the Metric needs.

Return type `None`

SlicedWassersteinDistance

Inheritance Diagram



```

class ashpy.metrics.sliced_wasserstein_metric.SlicedWassersteinDistance(model_selection_operator=
    in
    func-
    tion
    lt>,
    logdir='/home/docs/checko
    res-
    o-
    lu-
    tion=128,
    res-
    o-
    lu-
    tion_min=16,
    patches_per_image=64,
    patch_size=7,
    ran-
    dom_sampling_count=1,
    ran-
    dom_projection_dim=147,
    use_svd=False)
  
```

Bases: `ashpy.metrics.metric.Metric`

Sliced Wasserstein Distance. Used as metric in Progressive Growing of GANs¹.

¹ Progressive Growing of GANs <https://arxiv.org/abs/1710.10196>

Methods

<code>__init__</code> (<code>model_selection_operator</code> , <code>logdir</code> , ...)	Initialize the Metric.
<code>log</code> (<code>step</code>)	Log the SWD mean and each sub-metric
<code>model_selection</code> (<code>checkpoint</code> , <code>global_step</code>)	Perform model selection for each sub-metric
<code>reset_states</code> ()	Reset the state of the metric and the state of each child metric.
<code>update_state</code> (<code>context</code>)	Update the internal state of the metric, using the information from the context object.

Attributes

<code>best_folder</code>	Retrieve the folder used to save the best model when doing model selection.
<code>best_model_sel_file</code>	Retrieve the path to JSON file containing the measured performance of the best model.
<code>logdir</code>	Retrieve the log directory.
<code>metric</code>	Retrieve the <code>tf.keras.metrics.Metric</code> object.
<code>model_selection_operator</code>	Retrieve the operator used for model selection.
<code>name</code>	Retrieve the metric name.

`__init__` (`model_selection_operator`=<built-in function lt>, `logdir`='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/resolution=128, resolution_min=16, patches_per_image=64, patch_size=7, random_sampling_count=1, random_projection_dim=147, use_svd=False)
Initialize the Metric.

Parameters

- **`model_selection_operator`** (`typing.Callable`) – The operation that will be used when `model_selection` is triggered to compare the metrics, used by the `update_state`. Any `typing.Callable` behaving like an `operator` is accepted.

Note: Model selection is done ONLY if an operator is specified here.

- **`logdir`** (`str`) – Path to the log dir, defaults to a `log` folder in the current directory.
- **`resolution`** (`int`) – Image Resolution, defaults to 128
- **`resolution_min`** (`int`) – Min Resolution achieved by the metric
- **`patches_per_image`** (`int`) – Number of patches to extract per image per Laplacian level.
- **`patch_size`** (`int`) – Width of a square patch.
- **`random_sampling_count`** (`int`) – Number of random projections to average.
- **`random_projection_dim`** (`int`) – Dimension of the random projection space.
- **`use_svd`** (`bool`) – experimental method to compute a more accurate distance.

Return type None

`log` (`step`)
Log the SWD mean and each sub-metric

model_selection (*checkpoint, global_step*)
 Perform model selection for each sub-metric

Return type None

reset_states ()
 Reset the state of the metric and the state of each child metric.

Return type None

update_state (*context*)
 Update the internal state of the metric, using the information from the context object.

Parameters context (`ashpy.contexts.GANContext`) – An AshPy Context Object that carries all the information the Metric needs.

Return type None

class `ashpy.metrics.sliced_wasserstein_metric.SinglesWD` (*model_selection_operator=<built-in function lt>, logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/level_of_pyramid=0, real_or_fake='fake'*)

Bases: `ashpy.metrics.metric.Metric`

SlicedWassersteinDistance for a certain level of the pyramid.

__init__ (*model_selection_operator=<built-in function lt>, logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/level_of_pyramid=0, real_or_fake='fake'*)
 Initialize the Metric.

Parameters

- **model_selection_operator** (`typing.Callable`) – The operation that will be used when *model_selection* is triggered to compare the metrics, used by the *update_state*. Any `typing.Callable` behaving like an `operator` is accepted.

Note: Model selection is done ONLY if an operator is specified here.

- **logdir** (*str*) – Path to the log dir, defaults to a *log* folder in the current directory.
- **level_of_pyramid** (*int*) – Level of the pyramid related to this metric
- **real_or_fake** (*str*) – string identifying this metric (real or fake distance)

Return type None

update_state (*context, score*)
 Update the internal state of the metric, using the information from the context object.

Parameters context (`ashpy.contexts.GANContext`) – An AshPy Context Object that carries all the information the Metric needs.

Return type None

```

class ashpy.metrics.sliced_wasserstein_metric.SlicedWassersteinDistance (model_selection_operator=
    in
    func-
    tion
    lt>,
    logdir='/home/docs/checko
    res-
    o-
    lu-
    tion=128,
    res-
    o-
    lu-
    tion_min=16,
    patches_per_image=64,
    patch_size=7,
    ran-
    dom_sampling_count=1,
    ran-
    dom_projection_dim=147,
    use_svd=False)

```

Bases: `ashpy.metrics.metric.Metric`

Sliced Wasserstein Distance. Used as metric in Progressive Growing of GANs¹.

```

__init__(model_selection_operator=<built-in function lt>, logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy
    resolution=128, resolution_min=16, patches_per_image=64, patch_size=7, ran-
    dom_sampling_count=1, random_projection_dim=147, use_svd=False)
Initialize the Metric.

```

Parameters

- **model_selection_operator** (`typing.Callable`) – The operation that will be used when `model_selection` is triggered to compare the metrics, used by the `update_state`. Any `typing.Callable` behaving like an `operator` is accepted.

Note: Model selection is done ONLY if an operator is specified here.

- **logdir** (`str`) – Path to the log dir, defaults to a `log` folder in the current directory.
- **resolution** (`int`) – Image Resolution, defaults to 128
- **resolution_min** (`int`) – Min Resolution achieved by the metric
- **patches_per_image** (`int`) – Number of patches to extract per image per Laplacian level.
- **patch_size** (`int`) – Width of a square patch.
- **random_sampling_count** (`int`) – Number of random projections to average.
- **random_projection_dim** (`int`) – Dimension of the random projection space.
- **use_svd** (`bool`) – experimental method to compute a more accurate distance.

Return type None

¹ Progressive Growing of GANs <https://arxiv.org/abs/1710.10196>

log (*step*)

Log the SWD mean and each sub-metric

model_selection (*checkpoint, global_step*)

Perform model selection for each sub-metric

Return type None

reset_states ()

Reset the state of the metric and the state of each child metric.

Return type None

update_state (*context*)

Update the internal state of the metric, using the information from the context object.

Parameters **context** (`ashpy.contexts.GANContext`) – An AshPy Context Object that carries all the information the Metric needs.

Return type None

4.6.15 ssim_multiscale

Multiscale Structural Similarity metric.

Classes

SSIM_Multiscale

Multiscale Structural Similarity

SSIM_Multiscale

Inheritance Diagram



```

class ashpy.metrics.ssim_multiscale.SSIM_Multiscale(model_selection_operator=<built-in function lt>,
                                                    logdir='/home/docs/checkouts/readthedocs.org/user_buil,
                                                    max_val=2.0,
                                                    power_factors=None, filter_size=11, filter_sigma=1.5,
                                                    k1=0.01, k2=0.03)
  
```

Bases: `ashpy.metrics.metric.Metric`

Multiscale Structural Similarity

See Multiscale structural similarity for image quality assessment¹

Methods

<code>__init__</code> ([<i>model_selection_operator</i> , <i>logdir</i> , ...])	Initialize the Metric.
<code>split_batch</code> (<i>batch</i>)	Split a batch along axis 0 into two tensors having the same size
<code>update_state</code> (<i>context</i>)	Update the internal state of the metric, using the information from the context object.

Attributes

<code>best_folder</code>	Retrieve the folder used to save the best model when doing model selection.
<code>best_model_sel_file</code>	Retrieve the path to JSON file containing the measured performance of the best model.
<code>logdir</code>	Retrieve the log directory.
<code>metric</code>	Retrieve the <code>tf.keras.metrics.Metric</code> object.
<code>model_selection_operator</code>	Retrieve the operator used for model selection.
<code>name</code>	Retrieve the metric name.

`__init__` (*model_selection_operator*=<built-in function lt>, *logdir*='/home/docs/checkouts/readthedocs.org/user_builds/ashpy', *max_val*=2.0, *power_factors*=None, *filter_size*=11, *filter_sigma*=1.5, *k1*=0.01, *k2*=0.03)
Initialize the Metric.

Parameters

- **model_selection_operator** (`typing.Callable`) – The operation that will be used when *model_selection* is triggered to compare the metrics, used by the *update_state*. Any `typing.Callable` behaving like an *operator* is accepted.

Note: Model selection is done ONLY if an operator is specified here.

- **logdir** (`str`) – Path to the log dir, defaults to a *log* folder in the current directory.
- **max_val** (`float`) – The dynamic range of the images (i.e., the difference between the maximum the and minimum) (see www.tensorflow.org/versions/r2.0/api_docs/python/tf/image/ssim_multiscale)
- **power_factors** (`List[float]`) – Iterable of weights for each of the scales. The number of scales used is the length of the list. Index 0 is the unscaled resolution's weight and each increasing scale corresponds to the image being downsampled by 2. Defaults to (0.0448, 0.2856, 0.3001, 0.2363, 0.1333), which are the values obtained in the original paper.
- **filter_size** (`int`) – Default value 11 (size of gaussian filter).
- **filter_sigma** (`int`) – Default value 1.5 (width of gaussian filter).
- **k1** (`int`) – Default value 0.01

¹ Multiscale structural similarity for image quality assessment <https://ieeexplore.ieee.org/document/1292216>

- **k2** (*int*) – Default value 0.03 (SSIM is less sensitivity to K2 for lower values, so it would be better if we taken the values in range of $0 < K2 < 0.4$).

Return type `None`

static split_batch (*batch*)

Split a batch along axis 0 into two tensors having the same size

Parameters **batch** (*tf.Tensor*) – A batch of images

Return type `Tuple[Tensor, Tensor]`

Returns (`Tuple[tf.Tensor, tf.Tensor]`) The batch split in two tensors

Raises **ValueError** – if the batch has size 1

update_state (*context*)

Update the internal state of the metric, using the information from the context object.

Parameters **context** (`ashpy.contexts.GANContext`) – An AshPy Context Object that carries all the information the Metric needs.

Return type `None`

class `ashpy.metrics.ssim_multiscale.SSIM_Multiscale` (*model_selection_operator=<built-in function lt>*, *logdir='/home/docs/checkouts/readthedocs.org/user_buil*, *max_val=2.0*, *power_factors=None*, *filter_size=11*, *filter_sigma=1.5*, *k1=0.01*, *k2=0.03*)

Bases: `ashpy.metrics.metric.Metric`

Multiscale Structural Similarity

See Multiscale structural similarity for image quality assessment¹

__init__ (*model_selection_operator=<built-in function lt>*, *logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy*, *max_val=2.0*, *power_factors=None*, *filter_size=11*, *filter_sigma=1.5*, *k1=0.01*, *k2=0.03*)
Initialize the Metric.

Parameters

- **model_selection_operator** (`typing.Callable`) – The operation that will be used when *model_selection* is triggered to compare the metrics, used by the *update_state*. Any `typing.Callable` behaving like an *operator* is accepted.

Note: Model selection is done ONLY if an operator is specified here.

- **logdir** (*str*) – Path to the log dir, defaults to a *log* folder in the current directory.
- **max_val** (*float*) – The dynamic range of the images (i.e., the difference between the maximum the and minimum) (see www.tensorflow.org/versions/r2.0/api_docs/python/tf/image/ssim_multiscale)
- **power_factors** (`List[float]`) – Iterable of weights for each of the scales. The number of scales used is the length of the list. Index 0 is the unscaled resolution's weight and each increasing scale corresponds to the image being downsampled by 2. Defaults to (0.0448, 0.2856, 0.3001, 0.2363, 0.1333), which are the values obtained in the original paper.

¹ Multiscale structural similarity for image quality assessment <https://ieeexplore.ieee.org/document/1292216>

- **filter_size** (*int*) – Default value 11 (size of gaussian filter).
- **filter_sigma** (*int*) – Default value 1.5 (width of gaussian filter).
- **k1** (*int*) – Default value 0.01
- **k2** (*int*) – Default value 0.03 (SSIM is less sensitivity to K2 for lower values, so it would be better if we taken the values in range of $0 < K2 < 0.4$).

Return type None

static split_batch (*batch*)

Split a batch along axis 0 into two tensors having the same size

Parameters **batch** (*tf.Tensor*) – A batch of images

Return type Tuple[*Tensor*, *Tensor*]

Returns (Tuple[*tf.Tensor*, *tf.Tensor*]) The batch split in two tensors

Raises **ValueError** – if the batch has size 1

update_state (*context*)

Update the internal state of the metric, using the information from the context object.

Parameters **context** (*ashpy.contexts.GANContext*) – An AshPy Context Object that carries all the information the Metric needs.

Return type None

4.7 ashpy.models

Collection of Models.

Architectures

<i>convolutional</i>	Collection of Convolutional Models constructors.
<i>fc</i>	Collection of Fully Connected Models constructors.

4.7.1 convolutional

Collection of Convolutional Models constructors.

Interfaces

<i>interfaces.Conv2DInterface</i>	Primitive Interface to be used by all <i>ashpy.models</i> .
-----------------------------------	---

Conv2DInterface

Inheritance Diagram



class `ashpy.models.convolutional.interfaces.Conv2DInterface`

Bases: `tensorflow.python.keras.engine.training.Model`

Primitive Interface to be used by all `ashpy.models`.

Methods

<code>__init__()</code>	Primitive Interface to be used by all <code>ashpy.models</code> .
<code>call(inputs[, training, return_features])</code>	Execute the model on input data.

Attributes

<code>activity_regularizer</code>	Optional regularizer function for the output of this layer.
<code>dtype</code>	
<code>dynamic</code>	
<code>inbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>input</code>	Retrieves the input tensor(s) of a layer.
<code>input_mask</code>	Retrieves the input mask tensor(s) of a layer.
<code>input_shape</code>	Retrieves the input shape(s) of a layer.
<code>input_spec</code>	Gets the network's input specs.
<code>layers</code>	
<code>losses</code>	Losses which are associated with this <i>Layer</i> .
<code>metrics</code>	Returns the model's metrics added using <code>compile</code> , <code>add_metric</code> APIs.
<code>metrics_names</code>	Returns the model's display labels for all outputs.
<code>name</code>	Returns the name of this module as passed or determined in the ctor.
<code>name_scope</code>	Returns a <code>tf.name_scope</code> instance for this class.
<code>non_trainable_variables</code>	
<code>non_trainable_weights</code>	
<code>outbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>output</code>	Retrieves the output tensor(s) of a layer.
<code>output_mask</code>	Retrieves the output mask tensor(s) of a layer.
<code>output_shape</code>	Retrieves the output shape(s) of a layer.

Continued on next page

Table 176 – continued from previous page

<code>run_eagerly</code>	Settable attribute indicating whether the model should run eagerly.
<code>sample_weights</code>	
<code>state_updates</code>	Returns the <i>updates</i> from all layers that are stateful.
<code>stateful</code>	
<code>submodules</code>	Sequence of all sub-modules.
<code>trainable</code>	
<code>trainable_variables</code>	Sequence of variables owned by this module and it's submodules.
<code>trainable_weights</code>	
<code>updates</code>	
<code>variables</code>	Returns the list of all layer variables/weights.
<code>weights</code>	Returns the list of all layer variables/weights.

`__init__()`

Primitive Interface to be used by all *ashpy.models*. Declares the *self.model_layers* list.

static `_get_layer_spec` (*initial_filters*, *filters_cap*, *input_res*, *target_res*)

Compose the *layer_spec*, the building block of a convolutional model.

The *layer_spec* is an iterator. Every element returned is the number of filters to learn for the current layer. The generated sequence of filters starts from *initial_filters* and halve/double the number of filters depending on the *input_res* and *target_res*. If *input_res* > *target_res* the number of filters increases, else it decreases. The progression is always a power of 2.

Parameters

- **initial_filters** (*int*) – Depth of the first convolutional layer.
- **filters_cap** (*int*) – Maximum number of filters per layer.
- **input_res** (*tuple* of (*int*, *int*)) – Input resolution.
- **target_res** (*tuple* of (*int*, *int*)) – Output resolution.

Yields *int* – Number of filters to use for the conv layer.

Examples

```
# Encoder
class T(Conv2DInterface):
    pass
spec = T._get_layer_spec(
    initial_filters=16,
    filters_cap=128,
    input_res=(512, 256),
    target_res=(32, 16)
)
print([s for s in spec])

spec = T._get_layer_spec(
    initial_filters=16,
    filters_cap=128,
    input_res=(28, 28),
    target_res=(7, 7)
)
```

(continues on next page)

(continued from previous page)

```
print([s for s in spec])

# Decoder
spec = T._get_layer_spec(
    initial_filters=128,
    filters_cap=16,
    input_res=(32, 16),
    target_res=(512, 256)
)
print([s for s in spec])

spec = T._get_layer_spec(
    initial_filters=128,
    filters_cap=16,
    input_res=(7, 7),
    target_res=(28, 28)
)
print([s for s in spec])
```

```
[32, 64, 128, 128]
[32, 64]
[64, 32, 16, 16]
[64, 32]
```

Notes

This is useful since it enables us to dynamically redefine models sharing an underlying architecture but with different resolutions.

call (*inputs*, *training=True*, *return_features=False*)
Execute the model on input data.

Parameters

- **inputs** (*tf.Tensor*) – Input tensor(s).
- **training** (*bool*) – Training flag.
- **return_features** (*bool*) – If True returns the features.

Returns *tf.Tensor* – The model output.

Decoders

decoders.BaseDecoder

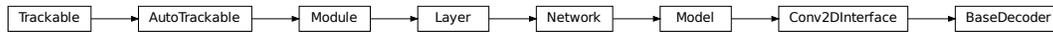
Primitive Model for all decoder (i.e., transpose convolution) based architecture.

decoders.FCNNBaseDecoder

Fully Convolutional Decoder.

BaseDecoder

Inheritance Diagram



class `ashpy.models.convolutional.decoders.BaseDecoder` (*layer_spec_input_res*,
layer_spec_target_res,
kernel_size, *initial_filters*,
filters_cap, *channels*,
use_dropout=True,
dropout_prob=0.3,
non_linearity=<class 'tensorflow.python.keras.layers.advanced_activations.LeakyReLU'>)

Bases: `ashpy.models.convolutional.interfaces.Conv2DInterface`

Primitive Model for all decoder (i.e., transpose convolution) based architecture.

Notes

Default to DCGAN Generator architecture.

Examples

- Direct Usage:

```

dummy_generator = BaseDecoder(
    layer_spec_input_res=(8, 8),
    layer_spec_target_res=(64, 64),
    kernel_size=(5, 5),
    initial_filters=1024,
    filters_cap=16,
    channels=3,
)
  
```

- Subclassing

```

class DummyGenerator(BaseDecoder):
    def call(self, input, training=True):
        print("Dummy Generator!")
        return input

dummy_generator = DummyGenerator(
    layer_spec_input_res=(8, 8),
    layer_spec_target_res=(32, 32),
    kernel_size=(5, 5),
    initial_filters=1024,
    filters_cap=16,
    channels=3,
)
  
```

(continues on next page)

(continued from previous page)

```
)
dummy_generator(tf.random.normal((1, 100)))
```

```
Dummy Generator!
```

Methods

<code>__init__(layer_spec_input_res, ..., [...])</code>	Instantiate the <i>BaseDecoder</i> .
---	--------------------------------------

Attributes

<code>activity_regularizer</code>	Optional regularizer function for the output of this layer.
<code>dtype</code>	
<code>dynamic</code>	
<code>inbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>input</code>	Retrieves the input tensor(s) of a layer.
<code>input_mask</code>	Retrieves the input mask tensor(s) of a layer.
<code>input_shape</code>	Retrieves the input shape(s) of a layer.
<code>input_spec</code>	Gets the network's input specs.
<code>layers</code>	
<code>losses</code>	Losses which are associated with this <i>Layer</i> .
<code>metrics</code>	Returns the model's metrics added using <i>compile</i> , <i>add_metric</i> APIs.
<code>metrics_names</code>	Returns the model's display labels for all outputs.
<code>name</code>	Returns the name of this module as passed or determined in the ctor.
<code>name_scope</code>	Returns a <i>tf.name_scope</i> instance for this class.
<code>non_trainable_variables</code>	
<code>non_trainable_weights</code>	
<code>outbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>output</code>	Retrieves the output tensor(s) of a layer.
<code>output_mask</code>	Retrieves the output mask tensor(s) of a layer.
<code>output_shape</code>	Retrieves the output shape(s) of a layer.
<code>run_eagerly</code>	Settable attribute indicating whether the model should run eagerly.
<code>sample_weights</code>	
<code>state_updates</code>	Returns the <i>updates</i> from all layers that are stateful.
<code>stateful</code>	
<code>submodules</code>	Sequence of all sub-modules.
<code>trainable</code>	
<code>trainable_variables</code>	Sequence of variables owned by this module and it's submodules.
<code>trainable_weights</code>	
<code>updates</code>	
<code>variables</code>	Returns the list of all layer variables/weights.

Continued on next page

Table 179 – continued from previous page

weights	Returns the list of all layer variables/weights.
---------	--

`__init__` (*layer_spec_input_res*, *layer_spec_target_res*, *kernel_size*, *initial_filters*, *filters_cap*, *channels*, *use_dropout=True*, *dropout_prob=0.3*, *non_linearity=<class 'tensorflow.python.keras.layers.advanced_activations.LeakyReLU'>*)
 Instantiate the *BaseDecoder*.

Model Assembly: 1. `_add_initial_block()`: Ingest the `tf.keras.Model` inputs and prepare them for `_add_building_block()`.

2. `_add_building_block()`: Core of the model, the layers specified here get added to the `tf.keras.Model` multiple times consuming the hyperparameters generated in the `_get_layer_spec()`.

3. `_add_final_block()`: Final block of our `tf.keras.Model`, take the model after `_add_building_block()` and prepare them for the for the final output.

Parameters

- **layer_spec_input_res** (tuple of (int, int)) – Shape of the `_get_layer_spec()` input tensors.
- **layer_spec_target_res** – (tuple of (int, int)): Shape of tensor desired as output of `_get_layer_spec()`.
- **kernel_size** (tuple of (int, int)) – Kernel used by the convolution layers.
- **initial_filters** (int) – Numbers of filters at the end of the first block.
- **filters_cap** (int) – Cap filters to a set amount, in the case of Decoder is a floor value AKA the minimum amount of filters.
- **channels** (int) – Channels of the output images (1 for Grayscale, 3 for RGB).

Returns None

Raises **ValueError** – If *filters_cap* > *initial_filters*

`_add_building_block` (*filters*)

Construct the core of the `tf.keras.Model`.

The layers specified here get added to the `tf.keras.Model` multiple times consuming the hyperparameters generated in the `_get_layer_spec()`.

Parameters **filters** (int) – Number of filters to use for this iteration of the Building Block.

`_add_final_block` (*channels*)

Take the results of `_add_building_block()` and prepare them for the for the final output.

Parameters

- **channels** (int) – Channels of the output images (1 for Grayscale, 3 for RGB).
- **kernel_size** (tuple of (int, int)) – Kernel used by the convolution layers.

`_add_initial_block` (*initial_filters*, *input_res*)

Ingest the `tf.keras.Model` inputs and prepare them for `_add_building_block()`.

Parameters

- **initial_filters** (int) – Numbers of filters to used as a base value.

- **input_res** (tuple of (int, int)) – Shape of the `_get_layer_spec()` input tensors.

FCNNBaseDecoder

Inheritance Diagram



class `ashpy.models.convolutional.decoders.FCNNBaseDecoder` (*layer_spec_input_res*, *layer_spec_target_res*, *kernel_size*, *initial_filters*, *filters_cap*, *channels*, *use_dropout=True*, *dropout_prob=0.3*, *non_linearity=<class 'tensorflow.python.keras.layers.advanced_activations.L*

Bases: `ashpy.models.convolutional.decoders.BaseDecoder`

Fully Convolutional Decoder. Expected input is a feature map.

Examples

- Direct Usage:

```

dummy_generator = FCNNBaseDecoder(
    layer_spec_input_res=(8, 8),
    layer_spec_target_res=(64, 64),
    kernel_size=(5, 5),
    initial_filters=1024,
    filters_cap=16,
    channels=3,
)

print(dummy_generator(tf.zeros((1, 1, 1, 100))).shape)
  
```

```
(1, 64, 64, 3)
```

Methods

<code>__init__(layer_spec_input_res, ...[, ...])</code>	Instantiate the <i>BaseDecoder</i> .
---	--------------------------------------

Attributes

<code>activity_regularizer</code>	Optional regularizer function for the output of this layer.
<code>dtype</code>	
<code>dynamic</code>	
<code>inbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>input</code>	Retrieves the input tensor(s) of a layer.
<code>input_mask</code>	Retrieves the input mask tensor(s) of a layer.
<code>input_shape</code>	Retrieves the input shape(s) of a layer.
<code>input_spec</code>	Gets the network's input specs.
<code>layers</code>	
<code>losses</code>	Losses which are associated with this <i>Layer</i> .
<code>metrics</code>	Returns the model's metrics added using <i>compile</i> , <i>add_metric</i> APIs.
<code>metrics_names</code>	Returns the model's display labels for all outputs.
<code>name</code>	Returns the name of this module as passed or determined in the ctor.
<code>name_scope</code>	Returns a <i>tf.name_scope</i> instance for this class.
<code>non_trainable_variables</code>	
<code>non_trainable_weights</code>	
<code>outbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>output</code>	Retrieves the output tensor(s) of a layer.
<code>output_mask</code>	Retrieves the output mask tensor(s) of a layer.
<code>output_shape</code>	Retrieves the output shape(s) of a layer.
<code>run_eagerly</code>	Settable attribute indicating whether the model should run eagerly.
<code>sample_weights</code>	
<code>state_updates</code>	Returns the <i>updates</i> from all layers that are stateful.
<code>stateful</code>	
<code>submodules</code>	Sequence of all sub-modules.
<code>trainable</code>	
<code>trainable_variables</code>	Sequence of variables owned by this module and it's submodules.
<code>trainable_weights</code>	
<code>updates</code>	
<code>variables</code>	Returns the list of all layer variables/weights.
<code>weights</code>	Returns the list of all layer variables/weights.

`_add_initial_block` (*initial_filters*, *input_res*)

Ingest the `tf.keras.Model` inputs and prepare them for `_add_building_block()`.

Parameters

- **`initial_filters`** (*int*) – Numbers of filters to used as a base value.
- **`input_res`** (*tuple of (int, int)*) – Shape of the `_get_layer_spec()` input tensors.

Encoders

<code>encoders.BaseEncoder</code>	Primitive Model for all encoder (i.e., convolution) based architecture.
<code>encoders.FCNNBaseEncoder</code>	Fully Convolutional Encoder.

BaseEncoder

Inheritance Diagram



```

class ashpy.models.convolutional.encoders.BaseEncoder (layer_spec_input_res,
                                                         layer_spec_target_res,
                                                         kernel_size,    initial_filters,
                                                         filters_cap,    output_shape,
                                                         use_dropout=True,
                                                         dropout_prob=0.3,
                                                         non_linearity=<class 'tensorflow.python.keras.layers.advanced_activations.LeakyReLU'>)
  
```

Bases: `ashpy.models.convolutional.interfaces.Conv2DInterface`

Primitive Model for all encoder (i.e., convolution) based architecture.

Notes

Default to DCGAN Discriminator architecture.

Examples

- Direct Usage:

```

dummy_generator = BaseEncoder(
    layer_spec_input_res=(64, 64),
    layer_spec_target_res=(8, 8),
    kernel_size=5,
    initial_filters=4,
    filters_cap=128,
    output_shape=1,
)
  
```

- Subclassing

```

class DummyDiscriminator(BaseEncoder):
    def call(self, inputs, training=True):
        print("Dummy Discriminator!")
        # build the model using
        # self._layers and inputs
        return inputs

dummy_discriminator = DummyDiscriminator(
    layer_spec_input_res=(64, 64),
    layer_spec_target_res=(8, 8),
    kernel_size=5,
    initial_filters=16,
    filters_cap=128,
    output_shape=1,
)
dummy_discriminator(tf.zeros((1, 28, 28, 3)))
    
```

```
Dummy Discriminator!
```

Methods

<code>__init__(layer_spec_input_res, ...[, ...])</code>	Instantiate the BaseDecoder.
---	------------------------------

Attributes

<code>activity_regularizer</code>	Optional regularizer function for the output of this layer.
<code>dtype</code>	
<code>dynamic</code>	
<code>inbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>input</code>	Retrieves the input tensor(s) of a layer.
<code>input_mask</code>	Retrieves the input mask tensor(s) of a layer.
<code>input_shape</code>	Retrieves the input shape(s) of a layer.
<code>input_spec</code>	Gets the network's input specs.
<code>layers</code>	
<code>losses</code>	Losses which are associated with this <i>Layer</i> .
<code>metrics</code>	Returns the model's metrics added using <i>compile</i> , <i>add_metric</i> APIs.
<code>metrics_names</code>	Returns the model's display labels for all outputs.
<code>name</code>	Returns the name of this module as passed or determined in the ctor.
<code>name_scope</code>	Returns a <i>tf.name_scope</i> instance for this class.
<code>non_trainable_variables</code>	
<code>non_trainable_weights</code>	
<code>outbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>output</code>	Retrieves the output tensor(s) of a layer.
<code>output_mask</code>	Retrieves the output mask tensor(s) of a layer.
<code>output_shape</code>	Retrieves the output shape(s) of a layer.

Continued on next page

Table 184 – continued from previous page

<code>run_eagerly</code>	Settable attribute indicating whether the model should run eagerly.
<code>sample_weights</code>	
<code>state_updates</code>	Returns the <i>updates</i> from all layers that are stateful.
<code>stateful</code>	
<code>submodules</code>	Sequence of all sub-modules.
<code>trainable</code>	
<code>trainable_variables</code>	Sequence of variables owned by this module and it's submodules.
<code>trainable_weights</code>	
<code>updates</code>	
<code>variables</code>	Returns the list of all layer variables/weights.
<code>weights</code>	Returns the list of all layer variables/weights.

`__init__`(*layer_spec_input_res*, *layer_spec_target_res*, *kernel_size*, *initial_filters*, *filters_cap*, *output_shape*, *use_dropout=True*, *dropout_prob=0.3*, *non_linearity=<class 'tensorflow.python.keras.layers.advanced_activations.LeakyReLU'>*)
 Instantiate the BaseDecoder.

Parameters

- **layer_spec_input_res** (tuple of (int, int)) – Shape of the input tensors.
- **layer_spec_target_res** (Union[int, Tuple[int, int]]) – (tuple of (int, int)): Shape of tensor desired as output of `_get_layer_spec()`.
- **kernel_size** (int) – Kernel used by the convolution layers.
- **initial_filters** (int) – Numbers of filters to used as a base value.
- **filters_cap** (int) – Cap filters to a set amount, in the case of an Encoder is a ceil value AKA the max amount of filters.
- **output_shape** (int) – Amount of units of the last `tf.keras.layers.Dense`.

Returns None

Raises **ValueError** – If `filters_cap < initial_filters`

`_add_building_block`(*filters*)

Construct the core of the `tf.keras.Model`.

The layers specified here get added to the `tf.keras.Model` multiple times consuming the hyper-parameters generated in the `_get_layer_spec()`.

Parameters **filters** (int) – Number of filters to use for this iteration of the Building Block.

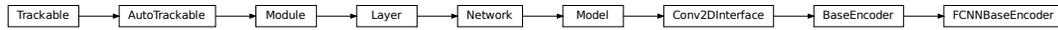
`_add_final_block`(*output_shape*)

Take the results of `_add_building_block()` and prepare them for the for the final output.

Parameters **output_shape** (int) – Amount of units of the last `tf.keras.layers.Dense`

FCNNBaseEncoder

Inheritance Diagram



class `ashpy.models.convolutional.encoders.FCNNBaseEncoder` (*layer_spec_input_res*,
layer_spec_target_res,
kernel_size, *initial_filters*, *filters_cap*,
encoding_dimension)

Bases: `ashpy.models.convolutional.encoders.BaseEncoder`

Fully Convolutional Encoder. Output a 1x1*encoding_size* vector. The output neurons are linear.

Examples

- Direct Usage:

```

dummy_generator = FCNNBaseEncoder(
    layer_spec_input_res=(64, 64),
    layer_spec_target_res=(8, 8),
    kernel_size=5,
    initial_filters=4,
    filters_cap=128,
    encoding_dimension=100,
)
print(dummy_generator(tf.zeros((1, 64, 64, 3))).shape)
  
```

```
(1, 1, 1, 100)
```

Methods

<code>__init__(layer_spec_input_res, ...)</code>	Instantiate the FCNNBaseDecoder.
--	----------------------------------

Attributes

<code>activity_regularizer</code>	Optional regularizer function for the output of this layer.
<code>dtype</code>	
<code>dynamic</code>	
<code>inbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>input</code>	Retrieves the input tensor(s) of a layer.
<code>input_mask</code>	Retrieves the input mask tensor(s) of a layer.

Continued on next page

Table 186 – continued from previous page

<code>input_shape</code>	Retrieves the input shape(s) of a layer.
<code>input_spec</code>	Gets the network's input specs.
<code>layers</code>	
<code>losses</code>	Losses which are associated with this <i>Layer</i> .
<code>metrics</code>	Returns the model's metrics added using <i>compile</i> , <i>add_metric</i> APIs.
<code>metrics_names</code>	Returns the model's display labels for all outputs.
<code>name</code>	Returns the name of this module as passed or determined in the ctor.
<code>name_scope</code>	Returns a <i>tf.name_scope</i> instance for this class.
<code>non_trainable_variables</code>	
<code>non_trainable_weights</code>	
<code>outbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>output</code>	Retrieves the output tensor(s) of a layer.
<code>output_mask</code>	Retrieves the output mask tensor(s) of a layer.
<code>output_shape</code>	Retrieves the output shape(s) of a layer.
<code>run_eagerly</code>	Settable attribute indicating whether the model should run eagerly.
<code>sample_weights</code>	
<code>state_updates</code>	Returns the <i>updates</i> from all layers that are stateful.
<code>stateful</code>	
<code>submodules</code>	Sequence of all sub-modules.
<code>trainable</code>	
<code>trainable_variables</code>	Sequence of variables owned by this module and it's submodules.
<code>trainable_weights</code>	
<code>updates</code>	
<code>variables</code>	Returns the list of all layer variables/weights.
<code>weights</code>	Returns the list of all layer variables/weights.

`__init__` (*layer_spec_input_res*, *layer_spec_target_res*, *kernel_size*, *initial_filters*, *filters_cap*, *encoding_dimension*)

Instantiate the FCNNBaseDecoder.

Parameters

- **layer_spec_input_res** (tuple of (int, int)) – Shape of the input tensors.
- **layer_spec_target_res** – (tuple of (int, int)): Shape of tensor desired as output of `_get_layer_spec()`.
- **kernel_size** (*int*) – Kernel used by the convolution layers.
- **initial_filters** (*int*) – Numbers of filters to used as a base value.
- **filters_cap** (*int*) – Cap filters to a set amount, in the case of an Encoder is a ceil value AKA the max amount of filters.
- **encoding_dimension** (*int*) – encoding dimension.

Returns None

Raises **ValueError** – If *filters_cap* < *initial_filters*

`_add_final_block` (*output_shape*)

Take the results of `_add_building_block()` and prepare them for the for the final output.

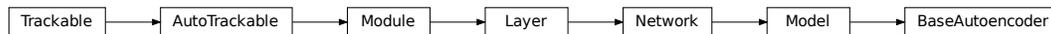
Parameters `output_shape` (*int*) – Amount of units of the last `tf.keras.layers.Dense`

Autoencoders

<code>autoencoders.BaseAutoencoder</code>	Primitive Model for all convolutional autoencoders.
<code>autoencoders.FCNNBaseAutoencoder</code>	Primitive Model for all fully convolutional autoencoders.

BaseAutoencoder

Inheritance Diagram



class `ashpy.models.convolutional.autoencoders.BaseAutoencoder` (*layer_spec_input_res, layer_spec_target_res, kernel_size, initial_filters, filters_cap, encoding_dimension, channels*)

Bases: `tensorflow.python.keras.engine.training.Model`

Primitive Model for all convolutional autoencoders.

Examples

- Direct Usage:

```

autoencoder = BaseAutoencoder(
    layer_spec_input_res=(64, 64),
    layer_spec_target_res=(8, 8),
    kernel_size=5,
    initial_filters=32,
    filters_cap=128,
    encoding_dimension=100,
    channels=3,
)

encoding, reconstruction = autoencoder(tf.zeros((1, 64, 64, 3)))
print(encoding.shape)
print(reconstruction.shape)
  
```



Methods

<code>__init__(layer_spec_input_res, ...)</code>	Instantiate the <code>BaseAutoEncoder</code> .
<code>call(inputs[, training])</code>	Execute the model on input data.

Attributes

<code>activity_regularizer</code>	Optional regularizer function for the output of this layer.
<code>dtype</code>	
<code>dynamic</code>	
<code>inbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>input</code>	Retrieves the input tensor(s) of a layer.
<code>input_mask</code>	Retrieves the input mask tensor(s) of a layer.
<code>input_shape</code>	Retrieves the input shape(s) of a layer.
<code>input_spec</code>	Gets the network's input specs.
<code>layers</code>	
<code>losses</code>	Losses which are associated with this <i>Layer</i> .
<code>metrics</code>	Returns the model's metrics added using <i>compile</i> , <i>add_metric</i> APIs.
<code>metrics_names</code>	Returns the model's display labels for all outputs.
<code>name</code>	Returns the name of this module as passed or determined in the ctor.
<code>name_scope</code>	Returns a <i>tf.name_scope</i> instance for this class.
<code>non_trainable_variables</code>	
<code>non_trainable_weights</code>	
<code>outbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>output</code>	Retrieves the output tensor(s) of a layer.
<code>output_mask</code>	Retrieves the output mask tensor(s) of a layer.
<code>output_shape</code>	Retrieves the output shape(s) of a layer.
<code>run_eagerly</code>	Settable attribute indicating whether the model should run eagerly.
<code>sample_weights</code>	
<code>state_updates</code>	Returns the <i>updates</i> from all layers that are stateful.
<code>stateful</code>	
<code>submodules</code>	Sequence of all sub-modules.
<code>trainable</code>	
<code>trainable_variables</code>	Sequence of variables owned by this module and it's submodules.
<code>trainable_weights</code>	
<code>updates</code>	
<code>variables</code>	Returns the list of all layer variables/weights.
<code>weights</code>	Returns the list of all layer variables/weights.

`__init__` (*layer_spec_input_res*, *layer_spec_target_res*, *kernel_size*, *initial_filters*, *filters_cap*, *encoding_dimension*, *channels*)

Instantiate the BaseAutoEncoder.

Parameters

- **layer_spec_input_res** (tuple of (int, int)) – Shape of the input tensors.
- **layer_spec_target_res** – (tuple of (int, int)): Shape of tensor desired as output of `_get_layer_spec()`.
- **kernel_size** (*int*) – Kernel used by the convolution layers.
- **initial_filters** (*int*) – Numbers of filters to used as a base value.
- **filters_cap** (*int*) – Cap filters to a set amount, in the case of an Encoder is a ceil value AKA the max amount of filters.
- **encoding_dimension** (*int*) – encoding dimension.
- **channels** (*int*) – Number of channels for the reconstructed image.

Returns None

call (*inputs*, *training=True*)

Execute the model on input data.

Parameters

- **inputs** (`tf.Tensor`) – Input tensors.
- **training** (`bool`) – Training flag.

Returns (*encoding*, *reconstruction*) – Pair of tensors.

FCNNBaseAutoencoder

Inheritance Diagram



class `ashpy.models.convolutional.autoencoders.FCNNBaseAutoencoder` (*layer_spec_input_res*, *layer_spec_target_res*, *kernel_size*, *initial_filters*, *filters_cap*, *encoding_dimension*, *channels*)

Bases: `tensorflow.python.keras.engine.training.Model`

Primitive Model for all fully convolutional autoencoders.

Examples

- Direct Usage:

```

autoencoder = FCNNBaseAutoencoder(
    layer_spec_input_res=(64, 64),
    layer_spec_target_res=(8, 8),
    kernel_size=5,
    initial_filters=32,
    filters_cap=128,
    encoding_dimension=100,
    channels=3,
)

encoding, reconstruction = autoencoder(tf.zeros((1, 64, 64, 3)))
print(encoding.shape)
print(reconstruction.shape)

```

Methods

<code>__init__(layer_spec_input_res, ...)</code>	Instantiate the FCNNBaseAutoEncoder.
<code>call(inputs[, training])</code>	Execute the model on input data.

Attributes

<code>activity_regularizer</code>	Optional regularizer function for the output of this layer.
<code>dtype</code>	
<code>dynamic</code>	
<code>inbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>input</code>	Retrieves the input tensor(s) of a layer.
<code>input_mask</code>	Retrieves the input mask tensor(s) of a layer.
<code>input_shape</code>	Retrieves the input shape(s) of a layer.
<code>input_spec</code>	Gets the network's input specs.
<code>layers</code>	
<code>losses</code>	Losses which are associated with this <i>Layer</i> .
<code>metrics</code>	Returns the model's metrics added using <i>compile</i> , <i>add_metric</i> APIs.
<code>metrics_names</code>	Returns the model's display labels for all outputs.
<code>name</code>	Returns the name of this module as passed or determined in the ctor.
<code>name_scope</code>	Returns a <i>tf.name_scope</i> instance for this class.
<code>non_trainable_variables</code>	
<code>non_trainable_weights</code>	
<code>outbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>output</code>	Retrieves the output tensor(s) of a layer.

Continued on next page

Table 191 – continued from previous page

<code>output_mask</code>	Retrieves the output mask tensor(s) of a layer.
<code>output_shape</code>	Retrieves the output shape(s) of a layer.
<code>run_eagerly</code>	Settable attribute indicating whether the model should run eagerly.
<code>sample_weights</code>	
<code>state_updates</code>	Returns the <i>updates</i> from all layers that are stateful.
<code>stateful</code>	
<code>submodules</code>	Sequence of all sub-modules.
<code>trainable</code>	
<code>trainable_variables</code>	Sequence of variables owned by this module and it's submodules.
<code>trainable_weights</code>	
<code>updates</code>	
<code>variables</code>	Returns the list of all layer variables/weights.
<code>weights</code>	Returns the list of all layer variables/weights.

`__init__` (*layer_spec_input_res, layer_spec_target_res, kernel_size, initial_filters, filters_cap, encoding_dimension, channels*)

Instantiate the FCNNBaseAutoEncoder.

Parameters

- **layer_spec_input_res** (tuple of (*int, int*)) – Shape of the input tensors.
- **layer_spec_target_res** – (tuple of (*int, int*)): Shape of tensor desired as output of `_get_layer_spec()`.
- **kernel_size** (*int*) – Kernel used by the convolution layers.
- **initial_filters** (*int*) – Numbers of filters to used as a base value.
- **filters_cap** (*int*) – Cap filters to a set amount, in the case of an Encoder is a ceil value AKA the max amount of filters.
- **encoding_dimension** (*int*) – encoding dimension.
- **channels** (*int*) – Number of channels for the reconstructed image.

Returns `None`

call (*inputs, training=True*)

Execute the model on input data.

Parameters

- **inputs** (`tf.Tensor`) – Input tensors.
- **training** (`bool`) – Training flag.

Returns (*encoding, reconstruction*) – Pair of tensors.

UNet

<code>unet.UNet</code>	UNet Architecture
<code>unet.SUNet</code>	Semantic UNet
<code>unet.FUNet</code>	Functional UNET Implementation

UNet

Inheritance Diagram



```

class ashpy.models.convolutional.unet.UNet(input_res,      min_res,      kernel_size,
      initial_filters,    filters_cap,    chan-
      nels,                use_dropout_encoder=True,
      use_dropout_decoder=True,
      dropout_prob=0.3,          en-
      coder_non_linearity=<class      'tensor-
      flow.python.keras.layers.advanced_activations.LeakyReLU'>,
      decoder_non_linearity=<class      'tensor-
      flow.python.keras.layers.advanced_activations.ReLU'>,
      normalization_layer=<class
      'ashpy.layers.instance_normalization.InstanceNormalization'>,
      last_activation=<function      tanh>,
      use_attention=False)
  
```

Bases: *ashpy.models.convolutional.interfaces.Conv2DInterface*

UNet Architecture

Used in Image-to-Image Translation with Conditional Adversarial Nets¹

Examples

- Direct Usage:

```

x = tf.ones((1, 512, 512, 3))
u_net = UNet(input_res = 512,
             min_res=4,
             kernel_size=4,
             initial_filters=64,
             filters_cap=512,
             channels=3)

y = u_net(x)
print(y.shape)
print(len(u_net.trainable_variables)>0)
  
```

```

(1, 512, 512, 3)
True
  
```

Methods

¹ Image-to-Image Translation with Conditional Adversarial Nets <https://arxiv.org/abs/1611.04076>

<code>__init__(input_res, min_res, kernel_size, ...)</code>	type <code>input_res</code> <code>int</code>
<code>call(inputs[, training])</code>	Execute the model on input data.
<code>get_decoder_block(filters[, use_bn, ...])</code>	Returns a block to be used in the decoder part of the UNET :param filters: number of filters :param use_bn: whether to use batch normalization :param use_dropout: whether to use dropout :param use_attention: whether to use attention
<code>get_encoder_block(filters[, use_bn, ...])</code>	Returns a block to be used in the encoder part of the UNET :param filters: number of filters :param use_bn: whether to use batch normalization :param use_attention: whether to use attention

Attributes

<code>activity_regularizer</code>	Optional regularizer function for the output of this layer.
<code>dtype</code>	
<code>dynamic</code>	
<code>inbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>input</code>	Retrieves the input tensor(s) of a layer.
<code>input_mask</code>	Retrieves the input mask tensor(s) of a layer.
<code>input_shape</code>	Retrieves the input shape(s) of a layer.
<code>input_spec</code>	Gets the network's input specs.
<code>layers</code>	
<code>losses</code>	Losses which are associated with this <i>Layer</i> .
<code>metrics</code>	Returns the model's metrics added using <i>compile</i> , <i>add_metric</i> APIs.
<code>metrics_names</code>	Returns the model's display labels for all outputs.
<code>name</code>	Returns the name of this module as passed or determined in the ctor.
<code>name_scope</code>	Returns a <i>tf.name_scope</i> instance for this class.
<code>non_trainable_variables</code>	
<code>non_trainable_weights</code>	
<code>outbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>output</code>	Retrieves the output tensor(s) of a layer.
<code>output_mask</code>	Retrieves the output mask tensor(s) of a layer.
<code>output_shape</code>	Retrieves the output shape(s) of a layer.
<code>run_eagerly</code>	Settable attribute indicating whether the model should run eagerly.
<code>sample_weights</code>	
<code>state_updates</code>	Returns the <i>updates</i> from all layers that are stateful.
<code>stateful</code>	
<code>submodules</code>	Sequence of all sub-modules.
<code>trainable</code>	
<code>trainable_variables</code>	Sequence of variables owned by this module and it's submodules.

Continued on next page

Table 194 – continued from previous page

<code>trainable_weights</code>	
<code>updates</code>	
<code>variables</code>	Returns the list of all layer variables/weights.
<code>weights</code>	Returns the list of all layer variables/weights.

```
__init__(input_res, min_res, kernel_size, initial_filters, filters_cap, chan-
nels, use_dropout_encoder=True, use_dropout_decoder=True,
dropout_prob=0.3, encoder_non_linearity=<class 'tensor-
flow.python.keras.layers.advanced_activations.LeakyReLU'>, de-
coder_non_linearity=<class 'tensorflow.python.keras.layers.advanced_activations.ReLU'>,
normalization_layer=<class 'ashpy.layers.instance_normalization.InstanceNormalization'>,
last_activation=<function tanh>, use_attention=False)
```

Parameters

- **input_res** (`int`) – input resolution
- **min_res** (`int`) – minimum resolution reached after decode
- **kernel_size** (`int`) – kernel size used in the network
- **initial_filters** (`int`) – number of filter of the initial convolution
- **filters_cap** (`int`) – maximum number of filters
- **channels** (`int`) – number of output channels
- **use_dropout_encoder** (`bool`) – whether to use dropout in the encoder module
- **use_dropout_decoder** (`bool`) – whether to use dropout in the decoder module
- **dropout_prob** (`float`) – probability of dropout
- **encoder_non_linearity** (`Type[Layer]`) – non linearity of encoder
- **decoder_non_linearity** (`Type[Layer]`) – non linearity of decoder
- **last_activation** (`<module 'tensorflow.python.keras.api._v2.keras.activations' from '/home/docs/checkouts/readthedocs.org/user_builds/ashpy/envs/v0.1.3/lib/python3.7/site-packages/tensorflow/python/keras/api/_v2/keras/activations/__init__.py'>`) – last activation function, tanh or softmax (for semantic images)
- **use_attention** (`bool`) – whether to use attention

call (`inputs, training=False`)

Execute the model on input data.

Parameters

- **inputs** (`tf.Tensor`) – Input tensor(s).
- **training** (`bool`) – Training flag.
- **return_features** (`bool`) – If True returns the features.

Returns `tf.Tensor` – The model output.

get_decoder_block (`filters, use_bn=True, use_dropout=False, use_attention=False`)

Returns a block to be used in the decoder part of the UNET :param filters: number of filters :param use_bn: whether to use batch normalization :param use_dropout: whether to use dropout :param use_attention: whether to use attention

Returns A block to be used in the decoder part

get_encoder_block (*filters, use_bn=True, use_attention=False*)

Returns a block to be used in the encoder part of the UNET :param filters: number of filters :param use_bn: whether to use batch normalization :param use_attention: whether to use attention

Returns A block to be used in the encoder part

SUNet

Inheritance Diagram



```

class ashpy.models.convolutional.unet.SUNet (input_res, min_res, kernel_size,
initial_filters, filters_cap, channels,
use_dropout_encoder=True, use_dropout_decoder=True,
dropout_prob=0.3, encoder_non_linearity=<class 'tensorflow.python.keras.layers.advanced_activations.LeakyReLU'>,
decoder_non_linearity=<class 'tensorflow.python.keras.layers.advanced_activations.ReLU'>,
use_attention=False)

```

Bases: *ashpy.models.convolutional.unet.UNet*

Semantic UNet

Methods

<code>__init__(input_res, min_res, kernel_size, ...)</code>	<p>Args: input_res: input resolution min_res: minimum resolution reached after decode kernel_size: kernel size used in the network initial_filters: number of filter of the initial convolution filters_cap: maximum number of filters channels: number of output channels use_dropout_encoder: whether to use dropout in the encoder module use_dropout_decoder: whether to use dropout in the decoder module dropout_prob: probability of dropout encoder_non_linearity: non linearity of encoder decoder_non_linearity: non linearity of decoder last_activation: last activation function, tanh or softmax (for semantic images) use_attention: whether to use attention</p>
---	--

Attributes

<code>activity_regularizer</code>	Optional regularizer function for the output of this layer.
<code>dtype</code>	
<code>dynamic</code>	
<code>inbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>input</code>	Retrieves the input tensor(s) of a layer.
<code>input_mask</code>	Retrieves the input mask tensor(s) of a layer.
<code>input_shape</code>	Retrieves the input shape(s) of a layer.
<code>input_spec</code>	Gets the network's input specs.
<code>layers</code>	
<code>losses</code>	Losses which are associated with this <i>Layer</i> .
<code>metrics</code>	Returns the model's metrics added using <i>compile</i> , <i>add_metric</i> APIs.
<code>metrics_names</code>	Returns the model's display labels for all outputs.
<code>name</code>	Returns the name of this module as passed or determined in the ctor.
<code>name_scope</code>	Returns a <i>tf.name_scope</i> instance for this class.
<code>non_trainable_variables</code>	
<code>non_trainable_weights</code>	
<code>outbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>output</code>	Retrieves the output tensor(s) of a layer.
<code>output_mask</code>	Retrieves the output mask tensor(s) of a layer.
<code>output_shape</code>	Retrieves the output shape(s) of a layer.
<code>run_eagerly</code>	Settable attribute indicating whether the model should run eagerly.
<code>sample_weights</code>	
<code>state_updates</code>	Returns the <i>updates</i> from all layers that are stateful.
<code>stateful</code>	
<code>submodules</code>	Sequence of all sub-modules.
<code>trainable</code>	
<code>trainable_variables</code>	Sequence of variables owned by this module and it's submodules.
<code>trainable_weights</code>	
<code>updates</code>	
<code>variables</code>	Returns the list of all layer variables/weights.
<code>weights</code>	Returns the list of all layer variables/weights.

ashpy.models.convolutional.unet.FUNet

```
ashpy.models.convolutional.unet.FUNet(input_res, min_res, kernel_size, initial_filters, filters_cap, channels, input_channels=3, use_dropout_encoder=True, use_dropout_decoder=True, dropout_prob=0.3, encoder_non_linearity=<class 'tensorflow.python.keras.layers.advanced_activations.LeakyReLU'>, decoder_non_linearity=<class 'tensorflow.python.keras.layers.advanced_activations.ReLU'>, last_activation=<function tanh>, use_attention=False)
```

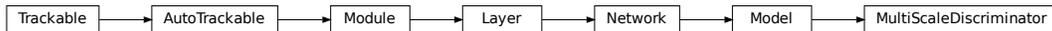
Functional UNET Implementation

Discriminators

<code>discriminators.MultiScaleDiscriminator</code>	Multi-Scale discriminator.
<code>discriminators.PatchDiscriminator</code>	Pix2Pix discriminator: The last layer is an image in which each pixels is the probability of being fake or real.

MultiScaleDiscriminator

Inheritance Diagram



```

class ashpy.models.convolutional.discriminators.MultiScaleDiscriminator(input_res,
                                                                    min_res,
                                                                    ker-
                                                                    nel_size,
                                                                    ini-
                                                                    tial_filters,
                                                                    fil-
                                                                    ters_cap,
                                                                    use_dropout=True,
                                                                    dropout_prob=0.3,
                                                                    non_linearity=<class
                                                                    'ten-
                                                                    sor-
                                                                    flow.python.keras.layers.ad-
                                                                    nor-
                                                                    mal-
                                                                    iza-
                                                                    tion_layer=<class
                                                                    'ashpy.layers.instance_norm
                                                                    use_attention=False,
                                                                    n_discriminators=1)

```

Bases: tensorflow.python.keras.engine.training.Model

Multi-Scale discriminator. This discriminator architecture is composed by multiple discriminators working at different scales. Each discriminator is a `ashpy.models.convolutional.discriminators.PatchDiscriminator`.

Examples

```

x = tf.ones((1, 256, 256, 3))

# instantiate the PatchDiscriminator
multiScaleDiscriminator = MultiScaleDiscriminator(input_res=256,
                                                    min_res=16,
                                                    kernel_size=5,
                                                    initial_filters=64,
                                                    filters_cap=512,
                                                    n_discriminators=3
                                                    )

# evaluate passing x
outputs = multiScaleDiscriminator(x)

# the output shape is
# the same as the input shape
print(len(outputs))
for output in outputs:
    print(output.shape)

```

```

3
(1, 12, 12, 1)
(1, 12, 12, 1)
(1, 12, 12, 1)

```

Methods

<code>__init__(input_res, min_res, kernel_size, ...)</code>	Multi Scale Discriminator.
<code>build_discriminator(input_res)</code>	Build a single discriminator using parameters defined in this object :param input_res: input resolution of the discriminator
<code>call(inputs[, training, return_features])</code>	type inputs Union[List[~T], Tensor]

Attributes

<code>activity_regularizer</code>	Optional regularizer function for the output of this layer.
<code>dtype</code>	
<code>dynamic</code>	
<code>inbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>input</code>	Retrieves the input tensor(s) of a layer.
<code>input_mask</code>	Retrieves the input mask tensor(s) of a layer.
<code>input_shape</code>	Retrieves the input shape(s) of a layer.
<code>input_spec</code>	Gets the network's input specs.
<code>layers</code>	
<code>losses</code>	Losses which are associated with this <i>Layer</i> .
<code>metrics</code>	Returns the model's metrics added using <i>compile</i> , <i>add_metric</i> APIs.
<code>metrics_names</code>	Returns the model's display labels for all outputs.
<code>name</code>	Returns the name of this module as passed or determined in the ctor.
<code>name_scope</code>	Returns a <i>tf.name_scope</i> instance for this class.
<code>non_trainable_variables</code>	
<code>non_trainable_weights</code>	
<code>outbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>output</code>	Retrieves the output tensor(s) of a layer.
<code>output_mask</code>	Retrieves the output mask tensor(s) of a layer.
<code>output_shape</code>	Retrieves the output shape(s) of a layer.
<code>run_eagerly</code>	Settable attribute indicating whether the model should run eagerly.
<code>sample_weights</code>	
<code>state_updates</code>	Returns the <i>updates</i> from all layers that are stateful.
<code>stateful</code>	
<code>submodules</code>	Sequence of all sub-modules.
<code>trainable</code>	
<code>trainable_variables</code>	Sequence of variables owned by this module and it's submodules.
<code>trainable_weights</code>	
<code>updates</code>	
<code>variables</code>	Returns the list of all layer variables/weights.

Continued on next page

Table 199 – continued from previous page

weights	Returns the list of all layer variables/weights.
---------	--

```
__init__(input_res, min_res, kernel_size, initial_filters, filters_cap,
         use_dropout=True, dropout_prob=0.3, non_linearity=<class 'tensorflow.python.keras.layers.advanced_activations.LeakyReLU'>, normalization_layer=<class 'ashpy.layers.instance_normalization.InstanceNormalization'>, use_attention=False,
         n_discriminators=1)
```

Multi Scale Discriminator. Different generator for different scales of the input image.

Used by Pix2PixHD¹.

Parameters

- **input_res** (*int*) – input resolution
- **min_res** (*int*) – minimum resolution reached by the discriminators
- **kernel_size** (*int*) – kernel size of discriminators
- **initial_filters** (*int*) – number of initial filters in the first layer of the discriminators
- **filters_cap** (*int*) – maximum number of filters in the discriminators
- **use_dropout** (*bool*) – whether to use dropout
- **dropout_prob** (*float*) – probability of dropout
- **non_linearity** (*tf.keras.layers.Layer*) – non linearity used in discriminators
- **normalization_layer** (*tf.keras.layers.Layer*) – normalization used by the discriminators
- **use_attention** (*bool*) – whether to use attention
- **n_discriminators** (*int*) – Number of discriminators

build_discriminator (*input_res*)

Build a single discriminator using parameters defined in this object :param input_res: input resolution of the discriminator

Return type *BaseEncoder*

Returns A Discriminator (PatchDiscriminator).

call (*inputs, training=True, return_features=False*)

Parameters

- **inputs** (*tf.Tensor*) – input tensor
- **training** (*bool*) – whether is training or not
- **return_features** (*bool*) – whether to return features or not

Return type *Union[List[Tensor], Tuple[List[Tensor], List[Tensor]]]*

Returns

(*[tf.Tensor]*) –

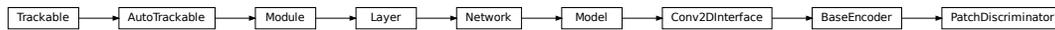
A List of Tensors containing the value of D_i for each input

¹ High-Resolution Image Synthesis and Semantic Manipulation with Conditional GANs <https://arxiv.org/abs/1711.11585>

(`tf.Tensor`): A List of features for each discriminator if `return_features`

PatchDiscriminator

Inheritance Diagram



```

class ashpy.models.convolutional.discriminators.PatchDiscriminator(input_res,
                                                                    min_res,
                                                                    ker-
                                                                    nel_size,
                                                                    ini-
                                                                    tial_filters,
                                                                    fil-
                                                                    ters_cap,
                                                                    use_dropout=True,
                                                                    dropout_prob=0.3,
                                                                    non_linearity=<class
                                                                    'tensor-
                                                                    flow.python.keras.layers.advanced_
                                                                    nor-
                                                                    maliza-
                                                                    tion_layer=<class
                                                                    'ashpy.layers.instance_normalizati
                                                                    use_attention=False)
  
```

Bases: `ashpy.models.convolutional.encoders.BaseEncoder`

Pix2Pix discriminator: The last layer is an image in which each pixels is the probability of being fake or real.

Examples

```

x = tf.ones((1, 64, 64, 3))

# instantiate the PatchDiscriminator
patchDiscriminator = PatchDiscriminator(input_res=64,
                                       min_res=16,
                                       kernel_size=5,
                                       initial_filters=64,
                                       filters_cap=512,
                                       )

# evaluate passing x
output = patchDiscriminator(x)

# the output shape is the same as the input shape
print(output.shape)
  
```

(1, 12, 12, 1)

Methods

<code>__init__(input_res, min_res, kernel_size, ...)</code>	PatchDiscriminator used by pix2pix, when min_res=1 this is the same as a standard fully convolutional discriminator
<code>call(inputs[, training, return_features])</code>	Execute the model on input data.

Attributes

<code>activity_regularizer</code>	Optional regularizer function for the output of this layer.
<code>dtype</code>	
<code>dynamic</code>	
<code>inbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>input</code>	Retrieves the input tensor(s) of a layer.
<code>input_mask</code>	Retrieves the input mask tensor(s) of a layer.
<code>input_shape</code>	Retrieves the input shape(s) of a layer.
<code>input_spec</code>	Gets the network's input specs.
<code>layers</code>	
<code>losses</code>	Losses which are associated with this <i>Layer</i> .
<code>metrics</code>	Returns the model's metrics added using <i>compile</i> , <i>add_metric</i> APIs.
<code>metrics_names</code>	Returns the model's display labels for all outputs.
<code>name</code>	Returns the name of this module as passed or determined in the ctor.
<code>name_scope</code>	Returns a <i>tf.name_scope</i> instance for this class.
<code>non_trainable_variables</code>	
<code>non_trainable_weights</code>	
<code>outbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>output</code>	Retrieves the output tensor(s) of a layer.
<code>output_mask</code>	Retrieves the output mask tensor(s) of a layer.
<code>output_shape</code>	Retrieves the output shape(s) of a layer.
<code>run_eagerly</code>	Settable attribute indicating whether the model should run eagerly.
<code>sample_weights</code>	
<code>state_updates</code>	Returns the <i>updates</i> from all layers that are stateful.
<code>stateful</code>	
<code>submodules</code>	Sequence of all sub-modules.
<code>trainable</code>	
<code>trainable_variables</code>	Sequence of variables owned by this module and it's submodules.
<code>trainable_weights</code>	
<code>updates</code>	
<code>variables</code>	Returns the list of all layer variables/weights.

Continued on next page

Table 201 – continued from previous page

weights	Returns the list of all layer variables/weights.
<p>__init__ (<i>input_res</i>, <i>min_res</i>, <i>kernel_size</i>, <i>initial_filters</i>, <i>filters_cap</i>, <i>use_dropout=True</i>, <i>dropout_prob=0.3</i>, <i>non_linearity=<class 'tensorflow.python.keras.layers.advanced_activations.LeakyReLU'></i>, <i>normalization_layer=<class 'ashpy.layers.instance_normalization.InstanceNormalization'></i>, <i>use_attention=False</i>)</p> <p>PatchDiscriminator used by pix2pix, when <i>min_res=1</i> this is the same as a standard fully convolutional discriminator</p> <p>Parameters</p> <ul style="list-style-type: none"> • input_res (<i>int</i>) – Input Resolution • min_res (<i>int</i>) – Minimum Resolution reached by the discriminator • kernel_size (<i>int</i>) – Kernel Size used in Conv Layer • initial_filters (<i>int</i>) – number of filters in the first convolutional layer • filters_cap (<i>int</i>) – Maximum number of filters • use_dropout (<i>bool</i>) – whether to use dropout • dropout_prob (<i>float</i>) – probability of dropout • non_linearity (<i>tf.keras.layers.Layer</i>) – non linearity used in the model • normalization_layer (<i>tf.keras.layers.Layer</i>) – normalization layer used in the model • use_attention (<i>bool</i>) – whether to use attention <p>_add_building_block (<i>filters</i>, <i>use_bn=False</i>)</p> <p>Construct the core of the <i>tf.keras.Model</i>.</p> <p>The layers specified here get added to the <i>tf.keras.Model</i> multiple times consuming the hyper-parameters generated in the <i>_get_layer_spec()</i>.</p> <p>Parameters filters (<i>int</i>) – Number of filters to use for this iteration of the Building Block.</p> <p>_add_final_block (<i>output_shape</i>)</p> <p>Take the results of <i>_add_building_block()</i> and prepare them for the for the final output.</p> <p>Parameters output_shape (<i>int</i>) – Amount of units of the last <i>tf.keras.layers.Dense</i></p> <p>call (<i>inputs</i>, <i>training=False</i>, <i>return_features=False</i>)</p> <p>Execute the model on input data.</p> <p>Parameters</p> <ul style="list-style-type: none"> • inputs (<i>tf.Tensor</i>) – Input tensor(s). • training (<i>bool</i>) – Training flag. • return_features (<i>bool</i>) – If True returns the features. <p>Returns <i>tf.Tensor</i> – The model output.</p>	

Pix2PixHD

<code>pix2pixhd.LocalEnhancer</code>	Local Enhancer module of the Pix2PixHD architecture.
<code>pix2pixhd.GlobalGenerator</code>	Global Generator from pix2pixHD paper:

LocalEnhancer

Inheritance Diagram



```

class ashpy.models.convolutional.pix2pixhd.LocalEnhancer(input_res=512,
                                                         min_res=64,      ini-
                                                         tial_filters=64,   fil-
                                                         ters_cap=512,     chan-
                                                         nels=3,          normaliza-
                                                         tion_layer=<class
                                                         'ashpy.layers.instance_normalization.InstanceNormaliza-
                                                         tion'>,
                                                         non_linearity=<class
                                                         'tensorflow.python.keras.layers.advanced_activations.ReLU'>,
                                                         num_resnet_blocks_global=9,
                                                         num_resnet_blocks_local=3,
                                                         kernel_size_resnet=3,
                                                         kernel_size_front_back=7,
                                                         num_internal_resnet_blocks=2)
  
```

Bases: tensorflow.python.keras.engine.training.Model

Local Enhancer module of the Pix2PixHD architecture.

Example

```

# instantiate the model
model = LocalEnhancer()

# call the model passing inputs
inputs = tf.ones((1, 512, 512, 3))
output = model(inputs)

# the output shape is
# the same as the input shape
print(output.shape)
  
```

```
(1, 512, 512, 3)
```

Methods

<code>__init__</code> ([input_res, min_res, ...])	Build the LocalEnhancer module of the Pix2PixHD architecture.
<code>call</code> (inputs[, training])	LocalEnhancer call.

Attributes

<code>activity_regularizer</code>	Optional regularizer function for the output of this layer.
<code>dtype</code>	
<code>dynamic</code>	
<code>inbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>input</code>	Retrieves the input tensor(s) of a layer.
<code>input_mask</code>	Retrieves the input mask tensor(s) of a layer.
<code>input_shape</code>	Retrieves the input shape(s) of a layer.
<code>input_spec</code>	Gets the network's input specs.
<code>layers</code>	
<code>losses</code>	Losses which are associated with this <i>Layer</i> .
<code>metrics</code>	Returns the model's metrics added using <i>compile</i> , <i>add_metric</i> APIs.
<code>metrics_names</code>	Returns the model's display labels for all outputs.
<code>name</code>	Returns the name of this module as passed or determined in the ctor.
<code>name_scope</code>	Returns a <i>tf.name_scope</i> instance for this class.
<code>non_trainable_variables</code>	
<code>non_trainable_weights</code>	
<code>outbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>output</code>	Retrieves the output tensor(s) of a layer.
<code>output_mask</code>	Retrieves the output mask tensor(s) of a layer.
<code>output_shape</code>	Retrieves the output shape(s) of a layer.
<code>run_eagerly</code>	Settable attribute indicating whether the model should run eagerly.
<code>sample_weights</code>	
<code>state_updates</code>	Returns the <i>updates</i> from all layers that are stateful.
<code>stateful</code>	
<code>submodules</code>	Sequence of all sub-modules.
<code>trainable</code>	
<code>trainable_variables</code>	Sequence of variables owned by this module and it's submodules.
<code>trainable_weights</code>	
<code>updates</code>	
<code>variables</code>	Returns the list of all layer variables/weights.
<code>weights</code>	Returns the list of all layer variables/weights.

```
__init__(input_res=512, min_res=64, initial_filters=64, filters_cap=512, channels=3, normalization_layer=<class 'ashpy.layers.instance_normalization.InstanceNormalization'>, non_linearity=<class 'tensorflow.python.keras.layers.advanced_activations.ReLU'>, num_resnet_blocks_global=9, num_resnet_blocks_local=3, kernel_size_resnet=3, kernel_size_front_back=7, num_internal_resnet_blocks=2)
```

Build the LocalEnhancer module of the Pix2PixHD architecture.

See High-Resolution Image Synthesis and Semantic Manipulation with Conditional GANs² for more details.

Parameters

- **input_res** (*int*) – input resolution
- **min_res** (*int*) – minimum resolution reached by the global generator
- **initial_filters** (*int*) – number of initial filters
- **filters_cap** (*int*) – maximum number of filters
- **channels** (*int*) – number of channels
- **normalization_layer** (`tf.keras.layers.Layer`) – layer of normalization
- **Instance Normalization or BatchNormalization or LayerNormalization** ((*e.g.*)) –
- **non_linearity** (`tf.keras.layers.Layer`) – non linearity used in Pix2Pix HD
- **num_resnet_blocks_global** (*int*) – number of residual blocks used in the global generator
- **num_resnet_blocks_local** (*int*) – number of residual blocks used in the local generator
- **kernel_size_resnet** (*int*) – kernel size used in resnets
- **kernel_size_front_back** (*int*) – kernel size used for the front and back convolution
- **num_internal_resnet_blocks** (*int*) – number of internal blocks of the resnet

call (*inputs*, *training=False*)

LocalEnhancer call. :param inputs: Input Tensors :type inputs: `tf.Tensor` :param training: Whether it is training phase or not :type training: bool

Returns

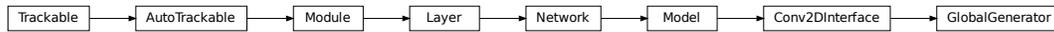
(`tf.Tensor`) –

Image of size (**input_res**, **input_res**, **channels**) as specified in the init call

GlobalGenerator

² High-Resolution Image Synthesis and Semantic Manipulation with Conditional GANs <https://arxiv.org/abs/1711.11585>

Inheritance Diagram



```

class ashpy.models.convolutional.pix2pixhd.GlobalGenerator (input_res=512,
                                                            min_res=64,      ini-
                                                            tial_filters=64,
                                                            filters_cap=512,
                                                            channels=3,  normal-
                                                            ization_layer=<class
                                                            'ashpy.layers.instance_normalization.Instance
                                                            non_linearity=<class
                                                            'tensor-
                                                            flow.python.keras.layers.advanced_activations.
                                                            num_resnet_blocks=9,
                                                            kernel_size_resnet=3,
                                                            ker-
                                                            nel_size_front_back=7,
                                                            num_internal_resnet_blocks=2)
  
```

Bases: *ashpy.models.convolutional.interfaces.Conv2DInterface*

Global Generator from pix2pixHD paper:

- G1^F: Convolutional frontend (downsampling)
- G1^R: ResNet Block
- G1^B: Convolutional backend (upsampling)

Methods

<code>__init__</code> ([input_res, min_res, ...])	Global Generator from Pix2PixHD
<code>call</code> (inputs[, training])	Call of the Pix2Pix HD model :param inputs: input tensor(s) :param training: If True training phase

Attributes

<code>activity_regularizer</code>	Optional regularizer function for the output of this layer.
<code>dtype</code>	
<code>dynamic</code>	
<code>inbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>input</code>	Retrieves the input tensor(s) of a layer.
<code>input_mask</code>	Retrieves the input mask tensor(s) of a layer.

Continued on next page

Table 206 – continued from previous page

<code>input_shape</code>	Retrieves the input shape(s) of a layer.
<code>input_spec</code>	Gets the network's input specs.
<code>layers</code>	
<code>losses</code>	Losses which are associated with this <i>Layer</i> .
<code>metrics</code>	Returns the model's metrics added using <i>compile</i> , <i>add_metric</i> APIs.
<code>metrics_names</code>	Returns the model's display labels for all outputs.
<code>name</code>	Returns the name of this module as passed or determined in the ctor.
<code>name_scope</code>	Returns a <i>tf.name_scope</i> instance for this class.
<code>non_trainable_variables</code>	
<code>non_trainable_weights</code>	
<code>outbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>output</code>	Retrieves the output tensor(s) of a layer.
<code>output_mask</code>	Retrieves the output mask tensor(s) of a layer.
<code>output_shape</code>	Retrieves the output shape(s) of a layer.
<code>run_eagerly</code>	Settable attribute indicating whether the model should run eagerly.
<code>sample_weights</code>	
<code>state_updates</code>	Returns the <i>updates</i> from all layers that are stateful.
<code>stateful</code>	
<code>submodules</code>	Sequence of all sub-modules.
<code>trainable</code>	
<code>trainable_variables</code>	Sequence of variables owned by this module and its submodules.
<code>trainable_weights</code>	
<code>updates</code>	
<code>variables</code>	Returns the list of all layer variables/weights.
<code>weights</code>	Returns the list of all layer variables/weights.

```

__init__(input_res=512, min_res=64, initial_filters=64, filters_cap=512, channels=3, normal-
        ization_layer=<class 'ashpy.layers.instance_normalization.InstanceNormalization'>,
        non_linearity=<class 'tensorflow.python.keras.layers.advanced_activations.ReLU'>,
        num_resnet_blocks=9, kernel_size_resnet=3, kernel_size_front_back=7,
        num_internal_resnet_blocks=2)

```

Global Generator from Pix2PixHD

Parameters

- **input_res** (*int*) – Input Resolution
- **min_res** (*int*) – Minimum resolution reached by the downsampling
- **initial_filters** (*int*) – number of initial filters
- **filters_cap** (*int*) – maximum number of filters
- **channels** (*int*) – output channels
- **normalization_layer** (*tf.keras.layers.Layer*) – normalization layer used by the global generator, can be Instance Norm, Layer Norm, Batch Norm
- **non_linearity** (*tf.keras.layers.Layer*) – non linearity used in the global generator
- **num_resnet_blocks** (*int*) – number of resnet blocks

- **kernel_size_resnet** (*int*) – kernel size used in resnets conv layers
- **kernel_size_front_back** (*int*) – kernel size used by the convolutional frontend and backend
- **num_internal_resnet_blocks** (*int*) – number of blocks used by internal resnet

call (*inputs, training=True*)

Call of the Pix2Pix HD model :param inputs: input tensor(s) :param training: If True training phase

Returns Tuple – Generated images.

Modules

<i>autoencoders</i>	Collection of Fully Convolutional Autoencoders
<i>discriminators</i>	Discriminators
<i>decoders</i>	Collection of Decoders (i.e., GANs’ Generators) models.
<i>encoders</i>	Collection of Encoders (i.e., GANs’ Discriminators) models.
<i>interfaces</i>	Primitive Convolutional interfaces.
<i>UNET</i>	UNET implementations.
<i>pix2pixhd</i>	Pix2Pix HD Implementation See: “High-Resolution Image Synthesis and Semantic Manipulation with Conditional GANs” [1] .

autoencoders

Collection of Fully Convolutional Autoencoders

Classes

<i>BaseAutoencoder</i>	Primitive Model for all convolutional autoencoders.
<i>FCNNBaseAutoencoder</i>	Primitive Model for all fully convolutional autoencoders.

BaseAutoencoder

Inheritance Diagram



class `ashpy.models.convolutional.autoencoders.BaseAutoencoder` (*layer_spec_input_res, layer_spec_target_res, kernel_size, initial_filters, filters_cap, encoding_dimension, channels*)

Bases: `tensorflow.python.keras.engine.training.Model`

Primitive Model for all convolutional autoencoders.

Examples

- Direct Usage:

```

autoencoder = BaseAutoencoder(
    layer_spec_input_res=(64, 64),
    layer_spec_target_res=(8, 8),
    kernel_size=5,
    initial_filters=32,
    filters_cap=128,
    encoding_dimension=100,
    channels=3,
)

encoding, reconstruction = autoencoder(tf.zeros((1, 64, 64, 3)))
print(encoding.shape)
print(reconstruction.shape)

```

Methods

<code>__init__(layer_spec_input_res, ...)</code>	Instantiate the BaseAutoEncoder.
<code>call(inputs[, training])</code>	Execute the model on input data.

Attributes

<code>activity_regularizer</code>	Optional regularizer function for the output of this layer.
<code>dtype</code>	
<code>dynamic</code>	
<code>inbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>input</code>	Retrieves the input tensor(s) of a layer.
<code>input_mask</code>	Retrieves the input mask tensor(s) of a layer.
<code>input_shape</code>	Retrieves the input shape(s) of a layer.
<code>input_spec</code>	Gets the network's input specs.
<code>layers</code>	
<code>losses</code>	Losses which are associated with this <i>Layer</i> .

Continued on next page

Table 210 – continued from previous page

<code>metrics</code>	Returns the model's metrics added using <code>compile</code> , <code>add_metric</code> APIs.
<code>metrics_names</code>	Returns the model's display labels for all outputs.
<code>name</code>	Returns the name of this module as passed or determined in the ctor.
<code>name_scope</code>	Returns a <code>tf.name_scope</code> instance for this class.
<code>non_trainable_variables</code>	
<code>non_trainable_weights</code>	
<code>outbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>output</code>	Retrieves the output tensor(s) of a layer.
<code>output_mask</code>	Retrieves the output mask tensor(s) of a layer.
<code>output_shape</code>	Retrieves the output shape(s) of a layer.
<code>run_eagerly</code>	Settable attribute indicating whether the model should run eagerly.
<code>sample_weights</code>	
<code>state_updates</code>	Returns the <code>updates</code> from all layers that are stateful.
<code>stateful</code>	
<code>submodules</code>	Sequence of all sub-modules.
<code>trainable</code>	
<code>trainable_variables</code>	Sequence of variables owned by this module and it's submodules.
<code>trainable_weights</code>	
<code>updates</code>	
<code>variables</code>	Returns the list of all layer variables/weights.
<code>weights</code>	Returns the list of all layer variables/weights.

`__init__` (*layer_spec_input_res, layer_spec_target_res, kernel_size, initial_filters, filters_cap, encoding_dimension, channels*)

Instantiate the `BaseAutoEncoder`.

Parameters

- `layer_spec_input_res` (*tuple of (int, int)*) – Shape of the input tensors.
- `layer_spec_target_res` – (*tuple of (int, int)*): Shape of tensor desired as output of `_get_layer_spec()`.
- `kernel_size` (*int*) – Kernel used by the convolution layers.
- `initial_filters` (*int*) – Numbers of filters to used as a base value.
- `filters_cap` (*int*) – Cap filters to a set amount, in the case of an Encoder is a ceil value AKA the max amount of filters.
- `encoding_dimension` (*int*) – encoding dimension.
- `channels` (*int*) – Number of channels for the reconstructed image.

Returns `None`

`call` (*inputs, training=True*)

Execute the model on input data.

Parameters

- `inputs` (`tf.Tensor`) – Input tensors.
- `training` (`bool`) – Training flag.

Returns (*encoding, reconstruction*) – Pair of tensors.

FCNNBaseAutoencoder

Inheritance Diagram



class `ashpy.models.convolutional.autoencoders.FCNNBaseAutoencoder` (*layer_spec_input_res, layer_spec_target_res, kernel_size, initial_filters, filters_cap, encoding_dimension, channels*)

Bases: `tensorflow.python.keras.engine.training.Model`

Primitive Model for all fully convolutional autoencoders.

Examples

- Direct Usage:

```

autoencoder = FCNNBaseAutoencoder(
    layer_spec_input_res=(64, 64),
    layer_spec_target_res=(8, 8),
    kernel_size=5,
    initial_filters=32,
    filters_cap=128,
    encoding_dimension=100,
    channels=3,
)

encoding, reconstruction = autoencoder(tf.zeros((1, 64, 64, 3)))
print(encoding.shape)
print(reconstruction.shape)
  
```

Methods

<code>__init__(layer_spec_input_res, ...)</code>	Instantiate the <code>FCNNBaseAutoEncoder</code> .
<code>call(inputs[, training])</code>	Execute the model on input data.

Attributes

<code>activity_regularizer</code>	Optional regularizer function for the output of this layer.
<code>dtype</code>	
<code>dynamic</code>	
<code>inbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>input</code>	Retrieves the input tensor(s) of a layer.
<code>input_mask</code>	Retrieves the input mask tensor(s) of a layer.
<code>input_shape</code>	Retrieves the input shape(s) of a layer.
<code>input_spec</code>	Gets the network's input specs.
<code>layers</code>	
<code>losses</code>	Losses which are associated with this <i>Layer</i> .
<code>metrics</code>	Returns the model's metrics added using <i>compile</i> , <i>add_metric</i> APIs.
<code>metrics_names</code>	Returns the model's display labels for all outputs.
<code>name</code>	Returns the name of this module as passed or determined in the ctor.
<code>name_scope</code>	Returns a <i>tf.name_scope</i> instance for this class.
<code>non_trainable_variables</code>	
<code>non_trainable_weights</code>	
<code>outbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>output</code>	Retrieves the output tensor(s) of a layer.
<code>output_mask</code>	Retrieves the output mask tensor(s) of a layer.
<code>output_shape</code>	Retrieves the output shape(s) of a layer.
<code>run_eagerly</code>	Settable attribute indicating whether the model should run eagerly.
<code>sample_weights</code>	
<code>state_updates</code>	Returns the <i>updates</i> from all layers that are stateful.
<code>stateful</code>	
<code>submodules</code>	Sequence of all sub-modules.
<code>trainable</code>	
<code>trainable_variables</code>	Sequence of variables owned by this module and it's submodules.
<code>trainable_weights</code>	
<code>updates</code>	
<code>variables</code>	Returns the list of all layer variables/weights.
<code>weights</code>	Returns the list of all layer variables/weights.

`__init__(layer_spec_input_res, layer_spec_target_res, kernel_size, initial_filters, filters_cap, encoding_dimension, channels)`
 Instantiate the `FCNNBaseAutoEncoder`.

Parameters

- `layer_spec_input_res` (tuple of (int, int)) – Shape of the input tensors.

- **layer_spec_target_res** – (tuple of (int, int)): Shape of tensor desired as output of `_get_layer_spec()`.
- **kernel_size** (int) – Kernel used by the convolution layers.
- **initial_filters** (int) – Numbers of filters to used as a base value.
- **filters_cap** (int) – Cap filters to a set amount, in the case of an Encoder is a ceil value AKA the max amount of filters.
- **encoding_dimension** (int) – encoding dimension.
- **channels** (int) – Number of channels for the reconstructed image.

Returns None

call (inputs, training=True)

Execute the model on input data.

Parameters

- **inputs** (tf.Tensor) – Input tensors.
- **training** (bool) – Training flag.

Returns (encoding, reconstruction) – Pair of tensors.

```
class ashpy.models.convolutional.autoencoders.BaseAutoencoder (layer_spec_input_res,
                                                             layer_spec_target_res,
                                                             kernel_size,  initial_filters,  fil-
                                                             ters_cap,  encod-
                                                             ing_dimension,
                                                             channels)
```

Bases: tensorflow.python.keras.engine.training.Model

Primitive Model for all convolutional autoencoders.

Examples

- Direct Usage:

```
autoencoder = BaseAutoencoder(
    layer_spec_input_res=(64, 64),
    layer_spec_target_res=(8, 8),
    kernel_size=5,
    initial_filters=32,
    filters_cap=128,
    encoding_dimension=100,
    channels=3,
)

encoding, reconstruction = autoencoder(tf.zeros((1, 64, 64, 3)))
print(encoding.shape)
print(reconstruction.shape)
```

```
__init__ (layer_spec_input_res, layer_spec_target_res, kernel_size, initial_filters, filters_cap, encod-
          ing_dimension, channels)
```

Instantiate the BaseAutoEncoder.

Parameters

- **layer_spec_input_res** (*tuple of (int, int)*) – Shape of the input tensors.
- **layer_spec_target_res** – (*tuple of (int, int)*): Shape of tensor desired as output of `_get_layer_spec()`.
- **kernel_size** (*int*) – Kernel used by the convolution layers.
- **initial_filters** (*int*) – Numbers of filters to used as a base value.
- **filters_cap** (*int*) – Cap filters to a set amount, in the case of an Encoder is a ceil value AKA the max amount of filters.
- **encoding_dimension** (*int*) – encoding dimension.
- **channels** (*int*) – Number of channels for the reconstructed image.

Returns `None`**call** (*inputs, training=True*)

Execute the model on input data.

Parameters

- **inputs** (*tf.Tensor*) – Input tensors.
- **training** (*bool*) – Training flag.

Returns (*encoding, reconstruction*) – Pair of tensors.

```
class ashpy.models.convolutional.autoencoders.FCNNBaseAutoencoder(layer_spec_input_res,
                                                                    layer_spec_target_res,
                                                                    kernel_size,
                                                                    ini-
                                                                    tial_filters,
                                                                    filters_cap,
                                                                    encod-
                                                                    ing_dimension,
                                                                    channels)
```

Bases: `tensorflow.python.keras.engine.training.Model`

Primitive Model for all fully convolutional autoencoders.

Examples

- Direct Usage:

```
autoencoder = FCNNBaseAutoencoder(
    layer_spec_input_res=(64, 64),
    layer_spec_target_res=(8, 8),
    kernel_size=5,
    initial_filters=32,
    filters_cap=128,
    encoding_dimension=100,
    channels=3,
)

encoding, reconstruction = autoencoder(tf.zeros((1, 64, 64, 3)))
print(encoding.shape)
print(reconstruction.shape)
```



`__init__` (*layer_spec_input_res, layer_spec_target_res, kernel_size, initial_filters, filters_cap, encoding_dimension, channels*)

Instantiate the FCNNBaseAutoEncoder.

Parameters

- **layer_spec_input_res** (tuple of (int, int)) – Shape of the input tensors.
- **layer_spec_target_res** – (tuple of (int, int)): Shape of tensor desired as output of `_get_layer_spec()`.
- **kernel_size** (int) – Kernel used by the convolution layers.
- **initial_filters** (int) – Numbers of filters to used as a base value.
- **filters_cap** (int) – Cap filters to a set amount, in the case of an Encoder is a ceil value AKA the max amount of filters.
- **encoding_dimension** (int) – encoding dimension.
- **channels** (int) – Number of channels for the reconstructed image.

Returns None

`call` (*inputs, training=True*)

Execute the model on input data.

Parameters

- **inputs** (tf.Tensor) – Input tensors.
- **training** (bool) – Training flag.

Returns (*encoding, reconstruction*) – Pair of tensors.

discriminators

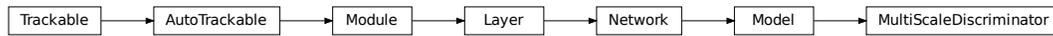
Discriminators

Classes

<i>MultiScaleDiscriminator</i>	Multi-Scale discriminator.
<i>PatchDiscriminator</i>	Pix2Pix discriminator: The last layer is an image in which each pixels is the probability of being fake or real.

MultiScaleDiscriminator

Inheritance Diagram



```

class ashpy.models.convolutional.discriminators.MultiScaleDiscriminator (input_res,
                                                                    min_res,
                                                                    ker-
                                                                    nel_size,
                                                                    ini-
                                                                    tial_filters,
                                                                    fil-
                                                                    ters_cap,
                                                                    use_dropout=True,
                                                                    dropout_prob=0.3,
                                                                    non_linearity=<class
                                                                    'ten-
                                                                    sor-
                                                                    flow.python.keras.layers.ad-
                                                                    nor-
                                                                    mal-
                                                                    iza-
                                                                    tion_layer=<class
                                                                    'ashpy.layers.instance_norm
                                                                    use_attention=False,
                                                                    n_discriminators=1)
  
```

Bases: tensorflow.python.keras.engine.training.Model

Multi-Scale discriminator. This discriminator architecture is composed by multiple discriminators working at different scales. Each discriminator is a *ashpy.models.convolutional.discriminators.PatchDiscriminator*.

Examples

```

x = tf.ones((1, 256, 256, 3))

# instantiate the PatchDiscriminator
multiScaleDiscriminator = MultiScaleDiscriminator(input_res=256,
                                                  min_res=16,
                                                  kernel_size=5,
                                                  initial_filters=64,
                                                  filters_cap=512,
                                                  n_discriminators=3
                                                  )

# evaluate passing x
outputs = multiScaleDiscriminator(x)
  
```

(continues on next page)

(continued from previous page)

```
# the output shape is
# the same as the input shape
print(len(outputs))
for output in outputs:
    print(output.shape)
```

```
3
(1, 12, 12, 1)
(1, 12, 12, 1)
(1, 12, 12, 1)
```

Methods

<code>__init__(input_res, min_res, kernel_size, ...)</code>	Multi Scale Discriminator.
<code>build_discriminator(input_res)</code>	Build a single discriminator using parameters defined in this object :param input_res: input resolution of the discriminator
<code>call(inputs[, training, return_features])</code>	type inputs Union[List[~T], Tensor]

Attributes

<code>activity_regularizer</code>	Optional regularizer function for the output of this layer.
<code>dtype</code>	
<code>dynamic</code>	
<code>inbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>input</code>	Retrieves the input tensor(s) of a layer.
<code>input_mask</code>	Retrieves the input mask tensor(s) of a layer.
<code>input_shape</code>	Retrieves the input shape(s) of a layer.
<code>input_spec</code>	Gets the network's input specs.
<code>layers</code>	
<code>losses</code>	Losses which are associated with this <i>Layer</i> .
<code>metrics</code>	Returns the model's metrics added using <i>compile</i> , <i>add_metric</i> APIs.
<code>metrics_names</code>	Returns the model's display labels for all outputs.
<code>name</code>	Returns the name of this module as passed or determined in the ctor.
<code>name_scope</code>	Returns a <i>tf.name_scope</i> instance for this class.
<code>non_trainable_variables</code>	
<code>non_trainable_weights</code>	
<code>outbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>output</code>	Retrieves the output tensor(s) of a layer.
<code>output_mask</code>	Retrieves the output mask tensor(s) of a layer.

Continued on next page

Table 215 – continued from previous page

<code>output_shape</code>	Retrieves the output shape(s) of a layer.
<code>run_eagerly</code>	Settable attribute indicating whether the model should run eagerly.
<code>sample_weights</code>	
<code>state_updates</code>	Returns the <i>updates</i> from all layers that are stateful.
<code>stateful</code>	
<code>submodules</code>	Sequence of all sub-modules.
<code>trainable</code>	
<code>trainable_variables</code>	Sequence of variables owned by this module and it's submodules.
<code>trainable_weights</code>	
<code>updates</code>	
<code>variables</code>	Returns the list of all layer variables/weights.
<code>weights</code>	Returns the list of all layer variables/weights.

```
__init__(input_res, min_res, kernel_size, initial_filters, filters_cap,
         use_dropout=True, dropout_prob=0.3, non_linearity=<class 'tensorflow.python.keras.layers.advanced_activations.LeakyReLU'>,
         normalization_layer=<class 'ashpy.layers.instance_normalization.InstanceNormalization'>, use_attention=False,
         n_discriminators=1)
```

Multi Scale Discriminator. Different generator for different scales of the input image.

Used by Pix2PixHD¹.

Parameters

- **input_res** (*int*) – input resolution
- **min_res** (*int*) – minimum resolution reached by the discriminators
- **kernel_size** (*int*) – kernel size of discriminators
- **initial_filters** (*int*) – number of initial filters in the first layer of the discriminators
- **filters_cap** (*int*) – maximum number of filters in the discriminators
- **use_dropout** (*bool*) – whether to use dropout
- **dropout_prob** (*float*) – probability of dropout
- **non_linearity** (`tf.keras.layers.Layer`) – non linearity used in discriminators
- **normalization_layer** (`tf.keras.layers.Layer`) – normalization used by the discriminators
- **use_attention** (*bool*) – whether to use attention
- **n_discriminators** (*int*) – Number of discriminators

build_discriminator (*input_res*)

Build a single discriminator using parameters defined in this object :param input_res: input resolution of the discriminator

Return type *BaseEncoder*

Returns A Discriminator (PatchDiscriminator).

call (*inputs*, *training=True*, *return_features=False*)

¹ High-Resolution Image Synthesis and Semantic Manipulation with Conditional GANs <https://arxiv.org/abs/1711.11585>

Parameters

- **inputs** (`tf.Tensor`) – input tensor
- **training** (`bool`) – whether is training or not
- **return_features** (`bool`) – whether to return features or not

Return type `Union[List[Tensor], Tuple[List[Tensor], List[Tensor]]]`

Returns

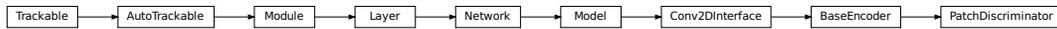
`([tf.Tensor])` –

A List of Tensors containing the value of D_i for each input

`([tf.Tensor])`: A List of features for each discriminator if `return_features`

PatchDiscriminator

Inheritance Diagram



```

class ashpy.models.convolutional.discriminators.PatchDiscriminator(input_res,
                                                                    min_res,
                                                                    ker-
                                                                    nel_size,
                                                                    ini-
                                                                    tial_filters,
                                                                    fil-
                                                                    ters_cap,
                                                                    use_dropout=True,
                                                                    dropout_prob=0.3,
                                                                    non_linearity=<class
                                                                    'tensor-
                                                                    flow.python.keras.layers.advanced_
                                                                    nor-
                                                                    maliza-
                                                                    tion_layer=<class
                                                                    'ashpy.layers.instance_normalizati
                                                                    use_attention=False)
  
```

Bases: `ashpy.models.convolutional.encoders.BaseEncoder`

Pix2Pix discriminator: The last layer is an image in which each pixels is the probability of being fake or real.

Examples

```
x = tf.ones((1, 64, 64, 3))

# instantiate the PatchDiscriminator
patchDiscriminator = PatchDiscriminator(input_res=64,
                                       min_res=16,
                                       kernel_size=5,
                                       initial_filters=64,
                                       filters_cap=512,
                                       )

# evaluate passing x
output = patchDiscriminator(x)

# the output shape is the same as the input shape
print(output.shape)
```

```
(1, 12, 12, 1)
```

Methods

<code>__init__(input_res, min_res, kernel_size, ...)</code>	PatchDiscriminator used by pix2pix, when min_res=1 this is the same as a standard fully convolutional discriminator
<code>call(inputs[, training, return_features])</code>	Execute the model on input data.

Attributes

<code>activity_regularizer</code>	Optional regularizer function for the output of this layer.
<code>dtype</code>	
<code>dynamic</code>	
<code>inbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>input</code>	Retrieves the input tensor(s) of a layer.
<code>input_mask</code>	Retrieves the input mask tensor(s) of a layer.
<code>input_shape</code>	Retrieves the input shape(s) of a layer.
<code>input_spec</code>	Gets the network's input specs.
<code>layers</code>	
<code>losses</code>	Losses which are associated with this <i>Layer</i> .
<code>metrics</code>	Returns the model's metrics added using <i>compile</i> , <i>add_metric</i> APIs.
<code>metrics_names</code>	Returns the model's display labels for all outputs.
<code>name</code>	Returns the name of this module as passed or determined in the ctor.
<code>name_scope</code>	Returns a <i>tf.name_scope</i> instance for this class.
<code>non_trainable_variables</code>	
<code>non_trainable_weights</code>	
<code>outbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>output</code>	Retrieves the output tensor(s) of a layer.

Continued on next page

Table 217 – continued from previous page

<code>output_mask</code>	Retrieves the output mask tensor(s) of a layer.
<code>output_shape</code>	Retrieves the output shape(s) of a layer.
<code>run_eagerly</code>	Settable attribute indicating whether the model should run eagerly.
<code>sample_weights</code>	
<code>state_updates</code>	Returns the <i>updates</i> from all layers that are stateful.
<code>stateful</code>	
<code>submodules</code>	Sequence of all sub-modules.
<code>trainable</code>	
<code>trainable_variables</code>	Sequence of variables owned by this module and it's submodules.
<code>trainable_weights</code>	
<code>updates</code>	
<code>variables</code>	Returns the list of all layer variables/weights.
<code>weights</code>	Returns the list of all layer variables/weights.

```
__init__(input_res, min_res, kernel_size, initial_filters, filters_cap,
         use_dropout=True, dropout_prob=0.3, non_linearity=<class 'tensorflow.python.keras.layers.advanced_activations.LeakyReLU'>,
         normalization_layer=<class 'ashpy.layers.instance_normalization.InstanceNormalization'>, use_attention=False)
```

PatchDiscriminator used by pix2pix, when `min_res=1` this is the same as a standard fully convolutional discriminator

Parameters

- `input_res` (*int*) – Input Resolution
- `min_res` (*int*) – Minimum Resolution reached by the discriminator
- `kernel_size` (*int*) – Kernel Size used in Conv Layer
- `initial_filters` (*int*) – number of filters in the first convolutional layer
- `filters_cap` (*int*) – Maximum number of filters
- `use_dropout` (*bool*) – whether to use dropout
- `dropout_prob` (*float*) – probability of dropout
- `non_linearity` (`tf.keras.layers.Layer`) – non linearity used in the model
- `normalization_layer` (`tf.keras.layers.Layer`) – normalization layer used in the model
- `use_attention` (*bool*) – whether to use attention

```
_add_building_block(filters, use_bn=False)
```

Construct the core of the `tf.keras.Model`.

The layers specified here get added to the `tf.keras.Model` multiple times consuming the hyper-parameters generated in the `_get_layer_spec()`.

Parameters `filters` (*int*) – Number of filters to use for this iteration of the Building Block.

```
_add_final_block(output_shape)
```

Take the results of `_add_building_block()` and prepare them for the for the final output.

Parameters `output_shape` (*int*) – Amount of units of the last `tf.keras.layers.Dense`

call (*inputs*, *training=False*, *return_features=False*)
Execute the model on input data.

Parameters

- **inputs** (`tf.Tensor`) – Input tensor(s).
- **training** (`bool`) – Training flag.
- **return_features** (`bool`) – If True returns the features.

Returns `tf.Tensor` – The model output.

```
class ashpy.models.convolutional.discriminators.MultiScaleDiscriminator (input_res,
                                                                    min_res,
                                                                    ker-
                                                                    nel_size,
                                                                    ini-
                                                                    tial_filters,
                                                                    fil-
                                                                    ters_cap,
                                                                    use_dropout=True,
                                                                    dropout_prob=0.3,
                                                                    non_linearity=<class
                                                                    'ten-
                                                                    sor-
                                                                    flow.python.keras.layers.ad
                                                                    nor-
                                                                    mal-
                                                                    iza-
                                                                    tion_layer=<class
                                                                    'ashpy.layers.instance_norm
                                                                    use_attention=False,
                                                                    n_discriminators=1)
```

Bases: `tensorflow.python.keras.engine.training.Model`

Multi-Scale discriminator. This discriminator architecture is composed by multiple discriminators working at different scales. Each discriminator is a `ashpy.models.convolutional.discriminators.PatchDiscriminator`.

Examples

```
x = tf.ones((1, 256, 256, 3))

# instantiate the PatchDiscriminator
multiScaleDiscriminator = MultiScaleDiscriminator(input_res=256,
                                                  min_res=16,
                                                  kernel_size=5,
                                                  initial_filters=64,
                                                  filters_cap=512,
                                                  n_discriminators=3
                                                  )

# evaluate passing x
outputs = multiScaleDiscriminator(x)

# the output shape is
```

(continues on next page)

(continued from previous page)

```
# the same as the input shape
print(len(outputs))
for output in outputs:
    print(output.shape)
```

```
3
(1, 12, 12, 1)
(1, 12, 12, 1)
(1, 12, 12, 1)
```

```
__init__(input_res, min_res, kernel_size, initial_filters, filters_cap,
         use_dropout=True, dropout_prob=0.3, non_linearity=<class 'tensorflow.python.keras.layers.advanced_activations.LeakyReLU'>,
         normalization_layer=<class 'ashpy.layers.instance_normalization.InstanceNormalization'>, use_attention=False,
         n_discriminators=1)
```

Multi Scale Discriminator. Different generator for different scales of the input image.

Used by Pix2PixHD¹.

Parameters

- **input_res** (*int*) – input resolution
- **min_res** (*int*) – minimum resolution reached by the discriminators
- **kernel_size** (*int*) – kernel size of discriminators
- **initial_filters** (*int*) – number of initial filters in the first layer of the discriminators
- **filters_cap** (*int*) – maximum number of filters in the discriminators
- **use_dropout** (*bool*) – whether to use dropout
- **dropout_prob** (*float*) – probability of dropout
- **non_linearity** (*tf.keras.layers.Layer*) – non linearity used in discriminators
- **normalization_layer** (*tf.keras.layers.Layer*) – normalization used by the discriminators
- **use_attention** (*bool*) – whether to use attention
- **n_discriminators** (*int*) – Number of discriminators

build_discriminator (*input_res*)

Build a single discriminator using parameters defined in this object :param input_res: input resolution of the discriminator

Return type *BaseEncoder*

Returns A Discriminator (PatchDiscriminator).

call (*inputs, training=True, return_features=False*)

Parameters

- **inputs** (*tf.Tensor*) – input tensor
- **training** (*bool*) – whether is training or not

¹ High-Resolution Image Synthesis and Semantic Manipulation with Conditional GANs <https://arxiv.org/abs/1711.11585>

- **return_features** (*bool*) – whether to return features or not

Return type `Union[List[Tensor], Tuple[List[Tensor], List[Tensor]]]`

Returns

`([tf.Tensor])` –

A List of Tensors containing the value of `Di` for each input

`([tf.Tensor])`: **A List of features for each discriminator if** `return_features`

```
class ashpy.models.convolutional.discriminators.PatchDiscriminator(input_res,
                                                                    min_res,
                                                                    ker-
                                                                    nel_size,
                                                                    ini-
                                                                    tial_filters,
                                                                    fil-
                                                                    ters_cap,
                                                                    use_dropout=True,
                                                                    dropout_prob=0.3,
                                                                    non_linearity=<class
                                                                    'tensor-
                                                                    flow.python.keras.layers.advanced_
                                                                    nor-
                                                                    maliza-
                                                                    tion_layer=<class
                                                                    'ashpy.layers.instance_normalizati
                                                                    use_attention=False)
```

Bases: `ashpy.models.convolutional.encoders.BaseEncoder`

Pix2Pix discriminator: The last layer is an image in which each pixels is the probability of being fake or real.

Examples

```
x = tf.ones((1, 64, 64, 3))

# instantiate the PatchDiscriminator
patchDiscriminator = PatchDiscriminator(input_res=64,
                                        min_res=16,
                                        kernel_size=5,
                                        initial_filters=64,
                                        filters_cap=512,
                                        )

# evaluate passing x
output = patchDiscriminator(x)

# the output shape is the same as the input shape
print(output.shape)
```

```
(1, 12, 12, 1)
```

```
__init__(input_res, min_res, kernel_size, initial_filters, filters_cap,
          use_dropout=True, dropout_prob=0.3, non_linearity=<class 'tensorflow.python.keras.layers.advanced_activations.LeakyReLU'>, normalization_layer=<class 'ashpy.layers.instance_normalization.InstanceNormalization'>, use_attention=False)
```

PatchDiscriminator used by pix2pix, when min_res=1 this is the same as a standard fully convolutional discriminator

Parameters

- **input_res** (*int*) – Input Resolution
- **min_res** (*int*) – Minimum Resolution reached by the discriminator
- **kernel_size** (*int*) – Kernel Size used in Conv Layer
- **initial_filters** (*int*) – number of filters in the first convolutional layer
- **filters_cap** (*int*) – Maximum number of filters
- **use_dropout** (*bool*) – whether to use dropout
- **dropout_prob** (*float*) – probability of dropout
- **non_linearity** (*tf.keras.layers.Layer*) – non linearity used in the model
- **normalization_layer** (*tf.keras.layers.Layer*) – normalization layer used in the model
- **use_attention** (*bool*) – whether to use attention

_add_building_block (*filters, use_bn=False*)

Construct the core of the *tf.keras.Model*.

The layers specified here get added to the *tf.keras.Model* multiple times consuming the hyper-parameters generated in the *_get_layer_spec()*.

Parameters filters (*int*) – Number of filters to use for this iteration of the Building Block.

_add_final_block (*output_shape*)

Take the results of *_add_building_block()* and prepare them for the for the final output.

Parameters output_shape (*int*) – Amount of units of the last *tf.keras.layers.Dense*

call (*inputs, training=False, return_features=False*)

Execute the model on input data.

Parameters

- **inputs** (*tf.Tensor*) – Input tensor(s).
- **training** (*bool*) – Training flag.
- **return_features** (*bool*) – If True returns the features.

Returns *tf.Tensor* – The model output.

decoders

Collection of Decoders (i.e., GANs’ Generators) models.

Classes

<i>BaseDecoder</i>	Primitive Model for all decoder (i.e., transpose convolution) based architecture.
<i>FCNNBaseDecoder</i>	Fully Convolutional Decoder.

BaseDecoder

Inheritance Diagram



```

class ashpy.models.convolutional.decoders.BaseDecoder(layer_spec_input_res,
                                                    layer_spec_target_res,
                                                    kernel_size,    initial_filters,
                                                    filters_cap,      channels,
                                                    use_dropout=True,
                                                    dropout_prob=0.3,
                                                    non_linearity=<class 'tensorflow.python.keras.layers.advanced_activations.LeakyReLU'>)

```

Bases: *ashpy.models.convolutional.interfaces.Conv2DInterface*

Primitive Model for all decoder (i.e., transpose convolution) based architecture.

Notes

Default to DCGAN Generator architecture.

Examples

- Direct Usage:

```

dummy_generator = BaseDecoder(
    layer_spec_input_res=(8, 8),
    layer_spec_target_res=(64, 64),
    kernel_size=(5, 5),
    initial_filters=1024,
    filters_cap=16,
    channels=3,
)

```

- Subclassing

```

class DummyGenerator(BaseDecoder):
    def call(self, input, training=True):
        print("Dummy Generator!")
        return input

dummy_generator = DummyGenerator(
    layer_spec_input_res=(8, 8),
    layer_spec_target_res=(32, 32),
    kernel_size=(5, 5),
    initial_filters=1024,
)

```

(continues on next page)

(continued from previous page)

```

filters_cap=16,
channels=3,
)
dummy_generator(tf.random.normal((1, 100)))

```

Dummy Generator!

Methods

<code>__init__(layer_spec_input_res, ...[, ...])</code>	Instantiate the <i>BaseDecoder</i> .
---	--------------------------------------

Attributes

<code>activity_regularizer</code>	Optional regularizer function for the output of this layer.
<code>dtype</code>	
<code>dynamic</code>	
<code>inbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>input</code>	Retrieves the input tensor(s) of a layer.
<code>input_mask</code>	Retrieves the input mask tensor(s) of a layer.
<code>input_shape</code>	Retrieves the input shape(s) of a layer.
<code>input_spec</code>	Gets the network's input specs.
<code>layers</code>	
<code>losses</code>	Losses which are associated with this <i>Layer</i> .
<code>metrics</code>	Returns the model's metrics added using <i>compile</i> , <i>add_metric</i> APIs.
<code>metrics_names</code>	Returns the model's display labels for all outputs.
<code>name</code>	Returns the name of this module as passed or determined in the ctor.
<code>name_scope</code>	Returns a <i>tf.name_scope</i> instance for this class.
<code>non_trainable_variables</code>	
<code>non_trainable_weights</code>	
<code>outbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>output</code>	Retrieves the output tensor(s) of a layer.
<code>output_mask</code>	Retrieves the output mask tensor(s) of a layer.
<code>output_shape</code>	Retrieves the output shape(s) of a layer.
<code>run_eagerly</code>	Settable attribute indicating whether the model should run eagerly.
<code>sample_weights</code>	
<code>state_updates</code>	Returns the <i>updates</i> from all layers that are stateful.
<code>stateful</code>	
<code>submodules</code>	Sequence of all sub-modules.
<code>trainable</code>	
<code>trainable_variables</code>	Sequence of variables owned by this module and its submodules.
<code>trainable_weights</code>	

Continued on next page

Table 220 – continued from previous page

updates	
variables	Returns the list of all layer variables/weights.
weights	Returns the list of all layer variables/weights.

```
__init__(layer_spec_input_res, layer_spec_target_res, kernel_size, initial_filters, filters_cap,
         channels, use_dropout=True, dropout_prob=0.3, non_linearity=<class 'tensorflow.python.keras.layers.advanced_activations.LeakyReLU'>)
```

Instantiate the *BaseDecoder*.

Model Assembly: 1. *_add_initial_block()*: Ingest the `tf.keras.Model` inputs and prepare them for *_add_building_block()*.

2. *_add_building_block()*: Core of the model, the layers specified here get added to the `tf.keras.Model` multiple times consuming the hyperparameters generated in the *_get_layer_spec()*.

3. *_add_final_block()*: Final block of our `tf.keras.Model`, take the model after *_add_building_block()* and prepare them for the for the final output.

Parameters

- **layer_spec_input_res** (tuple of (int, int)) – Shape of the *_get_layer_spec()* input tensors.
- **layer_spec_target_res** – (tuple of (int, int)): Shape of tensor desired as output of *_get_layer_spec()*.
- **kernel_size** (tuple of (int, int)) – Kernel used by the convolution layers.
- **initial_filters** (int) – Numbers of filters at the end of the first block.
- **filters_cap** (int) – Cap filters to a set amount, in the case of Decoder is a floor value AKA the minimum amount of filters.
- **channels** (int) – Channels of the output images (1 for Grayscale, 3 for RGB).

Returns None

Raises **ValueError** – If *filters_cap > initial_filters*

```
_add_building_block(filters)
```

Construct the core of the `tf.keras.Model`.

The layers specified here get added to the `tf.keras.Model` multiple times consuming the hyperparameters generated in the *_get_layer_spec()*.

Parameters **filters** (int) – Number of filters to use for this iteration of the Building Block.

```
_add_final_block(channels)
```

Take the results of *_add_building_block()* and prepare them for the for the final output.

Parameters

- **channels** (int) – Channels of the output images (1 for Grayscale, 3 for RGB).
- **kernel_size** (tuple of (int, int)) – Kernel used by the convolution layers.

```
_add_initial_block(initial_filters, input_res)
```

Ingest the `tf.keras.Model` inputs and prepare them for *_add_building_block()*.

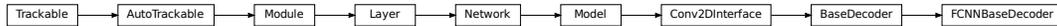
Parameters

- **initial_filters** (int) – Numbers of filters to used as a base value.

- **input_res** (tuple of (int, int)) – Shape of the `_get_layer_spec()` input tensors.

FCNNBaseDecoder

Inheritance Diagram



class `ashpy.models.convolutional.decoders.FCNNBaseDecoder` (*layer_spec_input_res*, *layer_spec_target_res*, *kernel_size*, *initial_filters*, *filters_cap*, *channels*, *use_dropout=True*, *dropout_prob=0.3*, *non_linearity=<class 'tensorflow.python.keras.layers.advanced_activations.L*

Bases: `ashpy.models.convolutional.decoders.BaseDecoder`

Fully Convolutional Decoder. Expected input is a feature map.

Examples

- Direct Usage:

```

dummy_generator = FCNNBaseDecoder(
    layer_spec_input_res=(8, 8),
    layer_spec_target_res=(64, 64),
    kernel_size=(5, 5),
    initial_filters=1024,
    filters_cap=16,
    channels=3,
)

print(dummy_generator(tf.zeros((1, 1, 1, 100))).shape)
  
```

```

(1, 64, 64, 3)
  
```

Methods

<code>__init__(layer_spec_input_res, ...[, ...])</code>	Instantiate the <i>BaseDecoder</i> .
---	--------------------------------------

Attributes

<code>activity_regularizer</code>	Optional regularizer function for the output of this layer.
<code>dtype</code>	
<code>dynamic</code>	
<code>inbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>input</code>	Retrieves the input tensor(s) of a layer.
<code>input_mask</code>	Retrieves the input mask tensor(s) of a layer.
<code>input_shape</code>	Retrieves the input shape(s) of a layer.
<code>input_spec</code>	Gets the network's input specs.
<code>layers</code>	
<code>losses</code>	Losses which are associated with this <i>Layer</i> .
<code>metrics</code>	Returns the model's metrics added using <i>compile</i> , <i>add_metric</i> APIs.
<code>metrics_names</code>	Returns the model's display labels for all outputs.
<code>name</code>	Returns the name of this module as passed or determined in the ctor.
<code>name_scope</code>	Returns a <i>tf.name_scope</i> instance for this class.
<code>non_trainable_variables</code>	
<code>non_trainable_weights</code>	
<code>outbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>output</code>	Retrieves the output tensor(s) of a layer.
<code>output_mask</code>	Retrieves the output mask tensor(s) of a layer.
<code>output_shape</code>	Retrieves the output shape(s) of a layer.
<code>run_eagerly</code>	Settable attribute indicating whether the model should run eagerly.
<code>sample_weights</code>	
<code>state_updates</code>	Returns the <i>updates</i> from all layers that are stateful.
<code>stateful</code>	
<code>submodules</code>	Sequence of all sub-modules.
<code>trainable</code>	
<code>trainable_variables</code>	Sequence of variables owned by this module and it's submodules.
<code>trainable_weights</code>	
<code>updates</code>	
<code>variables</code>	Returns the list of all layer variables/weights.
<code>weights</code>	Returns the list of all layer variables/weights.

`_add_initial_block` (*initial_filters*, *input_res*)

Ingest the `tf.keras.Model` inputs and prepare them for `_add_building_block()`.

Parameters

- **`initial_filters`** (*int*) – Numbers of filters to used as a base value.
- **`input_res`** (*tuple of (int, int)*) – Shape of the `_get_layer_spec()` input tensors.

```
class ashpy.models.convolutional.decoders.BaseDecoder (layer_spec_input_res,
                                                    layer_spec_target_res,
                                                    kernel_size, initial_filters,
                                                    filters_cap, channels,
                                                    use_dropout=True,
                                                    dropout_prob=0.3,
                                                    non_linearity=<class 'tensorflow.python.keras.layers.advanced_activations.LeakyReLU'
```

Bases: *ashpy.models.convolutional.interfaces.Conv2DInterface*

Primitive Model for all decoder (i.e., transpose convolution) based architecture.

Notes

Default to DCGAN Generator architecture.

Examples

- Direct Usage:

```
dummy_generator = BaseDecoder(
    layer_spec_input_res=(8, 8),
    layer_spec_target_res=(64, 64),
    kernel_size=(5, 5),
    initial_filters=1024,
    filters_cap=16,
    channels=3,
)
```

- Subclassing

```
class DummyGenerator(BaseDecoder):
    def call(self, input, training=True):
        print("Dummy Generator!")
        return input

dummy_generator = DummyGenerator(
    layer_spec_input_res=(8, 8),
    layer_spec_target_res=(32, 32),
    kernel_size=(5, 5),
    initial_filters=1024,
    filters_cap=16,
    channels=3,
)
dummy_generator(tf.random.normal((1, 100)))
```

```
Dummy Generator!
```

```
__init__(layer_spec_input_res, layer_spec_target_res, kernel_size, initial_filters, filters_cap,
          channels, use_dropout=True, dropout_prob=0.3, non_linearity=<class 'tensorflow.python.keras.layers.advanced_activations.LeakyReLU'>)
```

Instantiate the *BaseDecoder*.

Model Assembly: 1. *_add_initial_block()*: Ingest the `tf.keras.Model` inputs and prepare them for *_add_building_block()*.

2. `_add_building_block()`: Core of the model, the layers specified here get added to the `tf.keras.Model` multiple times consuming the hyperparameters generated in the `_get_layer_spec()`.
3. `_add_final_block()`: Final block of our `tf.keras.Model`, take the model after `_add_building_block()` and prepare them for the for the final output.

Parameters

- **layer_spec_input_res** (tuple of (int, int)) – Shape of the `_get_layer_spec()` input tensors.
- **layer_spec_target_res** – (tuple of (int, int)): Shape of tensor desired as output of `_get_layer_spec()`.
- **kernel_size** (tuple of (int, int)) – Kernel used by the convolution layers.
- **initial_filters** (int) – Numbers of filters at the end of the first block.
- **filters_cap** (int) – Cap filters to a set amount, in the case of Decoder is a floor value AKA the minimum amount of filters.
- **channels** (int) – Channels of the output images (1 for Grayscale, 3 for RGB).

Returns None

Raises `ValueError` – If `filters_cap > initial_filters`

`_add_building_block(filters)`

Construct the core of the `tf.keras.Model`.

The layers specified here get added to the `tf.keras.Model` multiple times consuming the hyperparameters generated in the `_get_layer_spec()`.

Parameters **filters** (int) – Number of filters to use for this iteration of the Building Block.

`_add_final_block(channels)`

Take the results of `_add_building_block()` and prepare them for the for the final output.

Parameters

- **channels** (int) – Channels of the output images (1 for Grayscale, 3 for RGB).
- **kernel_size** (tuple of (int, int)) – Kernel used by the convolution layers.

`_add_initial_block(initial_filters, input_res)`

Ingest the `tf.keras.Model` inputs and prepare them for `_add_building_block()`.

Parameters

- **initial_filters** (int) – Numbers of filters to used as a base value.
- **input_res** (tuple of (int, int)) – Shape of the `_get_layer_spec()` input tensors.

class `ashpy.models.convolutional.decoders.FCNNBaseDecoder` (*layer_spec_input_res*,
layer_spec_target_res,
kernel_size, *initial_filters*, *filters_cap*, *channels*,
use_dropout=True,
dropout_prob=0.3,
non_linearity=<class
'tensorflow.python.keras.layers.advanced_activations.L

Bases: *ashpy.models.convolutional.decoders.BaseDecoder*

Fully Convolutional Decoder. Expected input is a feature map.

Examples

- Direct Usage:

```
dummy_generator = FCNNBaseDecoder(
    layer_spec_input_res=(8, 8),
    layer_spec_target_res=(64, 64),
    kernel_size=(5, 5),
    initial_filters=1024,
    filters_cap=16,
    channels=3,
)

print(dummy_generator(tf.zeros((1, 1, 1, 100))).shape)

(1, 64, 64, 3)
```

`_add_initial_block` (*initial_filters*, *input_res*)

Ingest the `tf.keras.Model` inputs and prepare them for `_add_building_block()`.

Parameters

- **`initial_filters`** (*int*) – Numbers of filters to used as a base value.
- **`input_res`** (*tuple of (int, int)*) – Shape of the `_get_layer_spec()` input tensors.

encoders

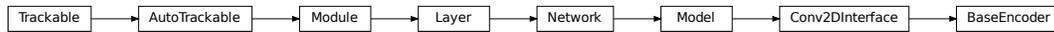
Collection of Encoders (i.e., GANs' Discriminators) models.

Classes

<i>BaseEncoder</i>	Primitive Model for all encoder (i.e., convolution) based architecture.
<i>FCNNBaseEncoder</i>	Fully Convolutional Encoder.

BaseEncoder

Inheritance Diagram



class `ashpy.models.convolutional.encoders.BaseEncoder` (*layer_spec_input_res*,
layer_spec_target_res,
kernel_size, *initial_filters*,
filters_cap, *output_shape*,
use_dropout=True,
dropout_prob=0.3,
non_linearity=<class 'tensorflow.python.keras.layers.advanced_activations.LeakyReLU'>)

Bases: `ashpy.models.convolutional.interfaces.Conv2DInterface`

Primitive Model for all encoder (i.e., convolution) based architecture.

Notes

Default to DCGAN Discriminator architecture.

Examples

- Direct Usage:

```

dummy_generator = BaseEncoder(
    layer_spec_input_res=(64, 64),
    layer_spec_target_res=(8, 8),
    kernel_size=5,
    initial_filters=4,
    filters_cap=128,
    output_shape=1,
)
  
```

- Subclassing

```

class DummyDiscriminator(BaseEncoder):
    def call(self, inputs, training=True):
        print("Dummy Discriminator!")
        # build the model using
        # self._layers and inputs
        return inputs

dummy_discriminator = DummyDiscriminator(
    layer_spec_input_res=(64, 64),
    layer_spec_target_res=(8, 8),
    kernel_size=5,
    initial_filters=16,
)
  
```

(continues on next page)

(continued from previous page)

```

        filters_cap=128,
        output_shape=1,
    )
    dummy_discriminator(tf.zeros((1, 28, 28, 3)))
    
```

Dummy Discriminator!

Methods

<code>__init__(layer_spec_input_res, ...[, ...])</code>	Instantiate the BaseDecoder.
---	------------------------------

Attributes

<code>activity_regularizer</code>	Optional regularizer function for the output of this layer.
<code>dtype</code>	
<code>dynamic</code>	
<code>inbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>input</code>	Retrieves the input tensor(s) of a layer.
<code>input_mask</code>	Retrieves the input mask tensor(s) of a layer.
<code>input_shape</code>	Retrieves the input shape(s) of a layer.
<code>input_spec</code>	Gets the network's input specs.
<code>layers</code>	
<code>losses</code>	Losses which are associated with this <i>Layer</i> .
<code>metrics</code>	Returns the model's metrics added using <i>compile</i> , <i>add_metric</i> APIs.
<code>metrics_names</code>	Returns the model's display labels for all outputs.
<code>name</code>	Returns the name of this module as passed or determined in the ctor.
<code>name_scope</code>	Returns a <i>tf.name_scope</i> instance for this class.
<code>non_trainable_variables</code>	
<code>non_trainable_weights</code>	
<code>outbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>output</code>	Retrieves the output tensor(s) of a layer.
<code>output_mask</code>	Retrieves the output mask tensor(s) of a layer.
<code>output_shape</code>	Retrieves the output shape(s) of a layer.
<code>run_eagerly</code>	Settable attribute indicating whether the model should run eagerly.
<code>sample_weights</code>	
<code>state_updates</code>	Returns the <i>updates</i> from all layers that are stateful.
<code>stateful</code>	
<code>submodules</code>	Sequence of all sub-modules.
<code>trainable</code>	
<code>trainable_variables</code>	Sequence of variables owned by this module and it's submodules.
<code>trainable_weights</code>	

Continued on next page

Table 225 – continued from previous page

updates	
variables	Returns the list of all layer variables/weights.
weights	Returns the list of all layer variables/weights.

`__init__` (*layer_spec_input_res*, *layer_spec_target_res*, *kernel_size*, *initial_filters*, *filters_cap*, *output_shape*, *use_dropout=True*, *dropout_prob=0.3*, *non_linearity=<class 'tensorflow.python.keras.layers.advanced_activations.LeakyReLU'>*)
 Instantiate the BaseDecoder.

Parameters

- **layer_spec_input_res** (tuple of (int, int)) – Shape of the input tensors.
- **layer_spec_target_res** (Union[int, Tuple[int, int]]) – (tuple of (int, int)): Shape of tensor desired as output of `_get_layer_spec()`.
- **kernel_size** (int) – Kernel used by the convolution layers.
- **initial_filters** (int) – Numbers of filters to used as a base value.
- **filters_cap** (int) – Cap filters to a set amount, in the case of an Encoder is a ceil value AKA the max amount of filters.
- **output_shape** (int) – Amount of units of the last `tf.keras.layers.Dense`.

Returns None

Raises **ValueError** – If `filters_cap < initial_filters`

`_add_building_block` (*filters*)

Construct the core of the `tf.keras.Model`.

The layers specified here get added to the `tf.keras.Model` multiple times consuming the hyper-parameters generated in the `_get_layer_spec()`.

Parameters **filters** (int) – Number of filters to use for this iteration of the Building Block.

`_add_final_block` (*output_shape*)

Take the results of `_add_building_block()` and prepare them for the for the final output.

Parameters **output_shape** (int) – Amount of units of the last `tf.keras.layers.Dense`

FCNNBaseEncoder

Inheritance Diagram



class `ashpy.models.convolutional.encoders.FCNNBaseEncoder` (*layer_spec_input_res, layer_spec_target_res, kernel_size, initial_filters, filters_cap, encoding_dimension*)

Bases: `ashpy.models.convolutional.encoders.BaseEncoder`

Fully Convolutional Encoder. Output a 1x1encoding_size vector. The output neurons are linear.

Examples

- Direct Usage:

```
dummy_generator = FCNNBaseEncoder(
    layer_spec_input_res=(64, 64),
    layer_spec_target_res=(8, 8),
    kernel_size=5,
    initial_filters=4,
    filters_cap=128,
    encoding_dimension=100,
)
print(dummy_generator(tf.zeros((1, 64, 64, 3))).shape)
```

```
(1, 1, 1, 100)
```

Methods

<code>__init__(layer_spec_input_res, ...)</code>	Instantiate the FCNNBaseDecoder.
--	----------------------------------

Attributes

<code>activity_regularizer</code>	Optional regularizer function for the output of this layer.
<code>dtype</code>	
<code>dynamic</code>	
<code>inbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>input</code>	Retrieves the input tensor(s) of a layer.
<code>input_mask</code>	Retrieves the input mask tensor(s) of a layer.
<code>input_shape</code>	Retrieves the input shape(s) of a layer.
<code>input_spec</code>	Gets the network's input specs.
<code>layers</code>	
<code>losses</code>	Losses which are associated with this <i>Layer</i> .
<code>metrics</code>	Returns the model's metrics added using <i>compile</i> , <i>add_metric</i> APIs.
<code>metrics_names</code>	Returns the model's display labels for all outputs.
<code>name</code>	Returns the name of this module as passed or determined in the ctor.
<code>name_scope</code>	Returns a <i>tf.name_scope</i> instance for this class.
<code>non_trainable_variables</code>	

Continued on next page

Table 227 – continued from previous page

<code>non_trainable_weights</code>	
<code>outbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>output</code>	Retrieves the output tensor(s) of a layer.
<code>output_mask</code>	Retrieves the output mask tensor(s) of a layer.
<code>output_shape</code>	Retrieves the output shape(s) of a layer.
<code>run_eagerly</code>	Settable attribute indicating whether the model should run eagerly.
<code>sample_weights</code>	
<code>state_updates</code>	Returns the <i>updates</i> from all layers that are stateful.
<code>stateful</code>	
<code>submodules</code>	Sequence of all sub-modules.
<code>trainable</code>	
<code>trainable_variables</code>	Sequence of variables owned by this module and it's submodules.
<code>trainable_weights</code>	
<code>updates</code>	
<code>variables</code>	Returns the list of all layer variables/weights.
<code>weights</code>	Returns the list of all layer variables/weights.

`__init__` (*layer_spec_input_res*, *layer_spec_target_res*, *kernel_size*, *initial_filters*, *filters_cap*, *encoding_dimension*)
 Instantiate the FCNNBaseDecoder.

Parameters

- **layer_spec_input_res** (tuple of (*int*, *int*)) – Shape of the input tensors.
- **layer_spec_target_res** – (tuple of (*int*, *int*)): Shape of tensor desired as output of `_get_layer_spec()`.
- **kernel_size** (*int*) – Kernel used by the convolution layers.
- **initial_filters** (*int*) – Numbers of filters to used as a base value.
- **filters_cap** (*int*) – Cap filters to a set amount, in the case of an Encoder is a ceil value AKA the max amount of filters.
- **encoding_dimension** (*int*) – encoding dimension.

Returns `None`

Raises `ValueError` – If *filters_cap* < *initial_filters*

`_add_final_block` (*output_shape*)

Take the results of `_add_building_block()` and prepare them for the for the final output.

Parameters **output_shape** (*int*) – Amount of units of the last `tf.keras.layers.Dense`

```
class ashpy.models.convolutional.encoders.BaseEncoder(layer_spec_input_res,
                                                    layer_spec_target_res,
                                                    kernel_size, initial_filters,
                                                    filters_cap, output_shape,
                                                    use_dropout=True,
                                                    dropout_prob=0.3,
                                                    non_linearity=<class 'tensorflow.python.keras.layers.advanced_activations.LeakyReLU'>)
```

Bases: `ashpy.models.convolutional.interfaces.Conv2DInterface`

Primitive Model for all encoder (i.e., convolution) based architecture.

Notes

Default to DCGAN Discriminator architecture.

Examples

- Direct Usage:

```
dummy_generator = BaseEncoder(
    layer_spec_input_res=(64, 64),
    layer_spec_target_res=(8, 8),
    kernel_size=5,
    initial_filters=4,
    filters_cap=128,
    output_shape=1,
)
```

- Subclassing

```
class DummyDiscriminator(BaseEncoder):
    def call(self, inputs, training=True):
        print("Dummy Discriminator!")
        # build the model using
        # self._layers and inputs
        return inputs

dummy_discriminator = DummyDiscriminator(
    layer_spec_input_res=(64, 64),
    layer_spec_target_res=(8, 8),
    kernel_size=5,
    initial_filters=16,
    filters_cap=128,
    output_shape=1,
)
dummy_discriminator(tf.zeros((1, 28, 28, 3)))
```

```
Dummy Discriminator!
```

```
__init__(layer_spec_input_res, layer_spec_target_res, kernel_size, initial_filters, filters_cap,
         output_shape, use_dropout=True, dropout_prob=0.3, non_linearity=<class 'tensorflow.python.keras.layers.advanced_activations.LeakyReLU'>)
```

Instantiate the BaseDecoder.

Parameters

- **layer_spec_input_res** (tuple of (int, int)) – Shape of the input tensors.
- **layer_spec_target_res** (Union[int, Tuple[int, int]]) – (tuple of (int, int)): Shape of tensor desired as output of `_get_layer_spec()`.
- **kernel_size** (int) – Kernel used by the convolution layers.
- **initial_filters** (int) – Numbers of filters to used as a base value.
- **filters_cap** (int) – Cap filters to a set amount, in the case of an Encoder is a ceil value AKA the max amount of filters.

- **output_shape** (*int*) – Amount of units of the last `tf.keras.layers.Dense`.

Returns `None`

Raises `ValueError` – If `filters_cap < initial_filters`

`_add_building_block` (*filters*)

Construct the core of the `tf.keras.Model`.

The layers specified here get added to the `tf.keras.Model` multiple times consuming the hyper-parameters generated in the `_get_layer_spec()`.

Parameters **filters** (*int*) – Number of filters to use for this iteration of the Building Block.

`_add_final_block` (*output_shape*)

Take the results of `_add_building_block()` and prepare them for the for the final output.

Parameters **output_shape** (*int*) – Amount of units of the last `tf.keras.layers.Dense`

```
class ashpy.models.convolutional.encoders.FCNNBaseEncoder (layer_spec_input_res,
                                                         layer_spec_target_res,
                                                         kernel_size,           ini-
                                                         tial_filters, filters_cap,
                                                         encoding_dimension)
```

Bases: `ashpy.models.convolutional.encoders.BaseEncoder`

Fully Convolutional Encoder. Output a `1x1xencoding_size` vector. The output neurons are linear.

Examples

- Direct Usage:

```
dummy_generator = FCNNBaseEncoder(
    layer_spec_input_res=(64, 64),
    layer_spec_target_res=(8, 8),
    kernel_size=5,
    initial_filters=4,
    filters_cap=128,
    encoding_dimension=100,
)
print(dummy_generator(tf.zeros((1, 64, 64, 3))).shape)
```

```
(1, 1, 1, 100)
```

`__init__` (*layer_spec_input_res, layer_spec_target_res, kernel_size, initial_filters, filters_cap, encoding_dimension*)

Instantiate the `FCNNBaseDecoder`.

Parameters

- **layer_spec_input_res** (*tuple of (int, int)*) – Shape of the input tensors.
- **layer_spec_target_res** – (*tuple of (int, int)*): Shape of tensor desired as output of `_get_layer_spec()`.
- **kernel_size** (*int*) – Kernel used by the convolution layers.
- **initial_filters** (*int*) – Numbers of filters to used as a base value.

- **filters_cap** (*int*) – Cap filters to a set amount, in the case of an Encoder is a ceil value AKA the max amount of filters.
- **encoding_dimension** (*int*) – encoding dimension.

Returns `None`

Raises `ValueError` – If `filters_cap < initial_filters`

_add_final_block (*output_shape*)

Take the results of `_add_building_block()` and prepare them for the for the final output.

Parameters **output_shape** (*int*) – Amount of units of the last `tf.keras.layers.Dense`

interfaces

Primitive Convolutional interfaces.

Classes

Conv2DInterface

Primitive Interface to be used by all *ashpy.models*.

Conv2DInterface

Inheritance Diagram



class `ashpy.models.convolutional.interfaces.Conv2DInterface`

Bases: `tensorflow.python.keras.engine.training.Model`

Primitive Interface to be used by all *ashpy.models*.

Methods

`__init__()`

Primitive Interface to be used by all *ashpy.models*.

`call(inputs[, training, return_features])`

Execute the model on input data.

Attributes

<code>activity_regularizer</code>	Optional regularizer function for the output of this layer.
<code>dtype</code>	
<code>dynamic</code>	
<code>inbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>input</code>	Retrieves the input tensor(s) of a layer.
<code>input_mask</code>	Retrieves the input mask tensor(s) of a layer.
<code>input_shape</code>	Retrieves the input shape(s) of a layer.
<code>input_spec</code>	Gets the network's input specs.
<code>layers</code>	
<code>losses</code>	Losses which are associated with this <i>Layer</i> .
<code>metrics</code>	Returns the model's metrics added using <i>compile</i> , <i>add_metric</i> APIs.
<code>metrics_names</code>	Returns the model's display labels for all outputs.
<code>name</code>	Returns the name of this module as passed or determined in the ctor.
<code>name_scope</code>	Returns a <i>tf.name_scope</i> instance for this class.
<code>non_trainable_variables</code>	
<code>non_trainable_weights</code>	
<code>outbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>output</code>	Retrieves the output tensor(s) of a layer.
<code>output_mask</code>	Retrieves the output mask tensor(s) of a layer.
<code>output_shape</code>	Retrieves the output shape(s) of a layer.
<code>run_eagerly</code>	Settable attribute indicating whether the model should run eagerly.
<code>sample_weights</code>	
<code>state_updates</code>	Returns the <i>updates</i> from all layers that are stateful.
<code>stateful</code>	
<code>submodules</code>	Sequence of all sub-modules.
<code>trainable</code>	
<code>trainable_variables</code>	Sequence of variables owned by this module and it's submodules.
<code>trainable_weights</code>	
<code>updates</code>	
<code>variables</code>	Returns the list of all layer variables/weights.
<code>weights</code>	Returns the list of all layer variables/weights.

`__init__()`

Primitive Interface to be used by all *ashpy.models*. Declares the *self.model_layers* list.

`static _get_layer_spec(initial_filters, filters_cap, input_res, target_res)`

Compose the *layer_spec*, the building block of a convolutional model.

The *layer_spec* is an iterator. Every element returned is the number of filters to learn for the current layer. The generated sequence of filters starts from *initial_filters* and halve/double the number of filters depending on the *input_res* and *target_res*. If *input_res* > *target_res* the number of filters increases, else it decreases. The progression is always a power of 2.

Parameters

- **initial_filters** (*int*) – Depth of the first convolutional layer.

- **filters_cap** (*int*) – Maximum number of filters per layer.
- **input_res** (*tuple* of (*int*, *int*)) – Input resolution.
- **target_res** (*tuple* of (*int*, *int*)) – Output resolution.

Yields *int* – Number of filters to use for the conv layer.

Examples

```
# Encoder
class T(Conv2DInterface):
    pass
spec = T._get_layer_spec(
    initial_filters=16,
    filters_cap=128,
    input_res=(512, 256),
    target_res=(32, 16)
)
print([s for s in spec])

spec = T._get_layer_spec(
    initial_filters=16,
    filters_cap=128,
    input_res=(28, 28),
    target_res=(7, 7)
)
print([s for s in spec])

# Decoder
spec = T._get_layer_spec(
    initial_filters=128,
    filters_cap=16,
    input_res=(32, 16),
    target_res=(512, 256)
)
print([s for s in spec])

spec = T._get_layer_spec(
    initial_filters=128,
    filters_cap=16,
    input_res=(7, 7),
    target_res=(28, 28)
)
print([s for s in spec])
```

```
[32, 64, 128, 128]
[32, 64]
[64, 32, 16, 16]
[64, 32]
```

Notes

This is useful since it enables us to dynamically redefine models sharing an underlying architecture but with different resolutions.

call (*inputs*, *training=True*, *return_features=False*)
Execute the model on input data.

Parameters

- **inputs** (`tf.Tensor`) – Input tensor(s).
- **training** (`bool`) – Training flag.
- **return_features** (`bool`) – If True returns the features.

Returns `tf.Tensor` – The model output.

class `ashpy.models.convolutional.interfaces.Conv2DInterface`

Bases: `tensorflow.python.keras.engine.training.Model`

Primitive Interface to be used by all `ashpy.models`.

__init__ ()

Primitive Interface to be used by all `ashpy.models`. Declares the `self.model_layers` list.

static **_get_layer_spec** (*initial_filters*, *filters_cap*, *input_res*, *target_res*)

Compose the `layer_spec`, the building block of a convolutional model.

The `layer_spec` is an iterator. Every element returned is the number of filters to learn for the current layer. The generated sequence of filters starts from `initial_filters` and halve/double the number of filters depending on the `input_res` and `target_res`. If `input_res > target_res` the number of filters increases, else it decreases. The progression is always a power of 2.

Parameters

- **initial_filters** (`int`) – Depth of the first convolutional layer.
- **filters_cap** (`int`) – Maximum number of filters per layer.
- **input_res** (`tuple of (int, int)`) – Input resolution.
- **target_res** (`tuple of (int, int)`) – Output resolution.

Yields `int` – Number of filters to use for the conv layer.

Examples

```
# Encoder
class T(Conv2DInterface):
    pass
spec = T._get_layer_spec(
    initial_filters=16,
    filters_cap=128,
    input_res=(512, 256),
    target_res=(32, 16)
)
print([s for s in spec])

spec = T._get_layer_spec(
    initial_filters=16,
    filters_cap=128,
    input_res=(28, 28),
    target_res=(7, 7)
)
print([s for s in spec])
```

(continues on next page)

(continued from previous page)

```
# Decoder
spec = T._get_layer_spec(
    initial_filters=128,
    filters_cap=16,
    input_res=(32, 16),
    target_res=(512, 256)
)
print([s for s in spec])

spec = T._get_layer_spec(
    initial_filters=128,
    filters_cap=16,
    input_res=(7, 7),
    target_res=(28, 28)
)
print([s for s in spec])
```

```
[32, 64, 128, 128]
[32, 64]
[64, 32, 16, 16]
[64, 32]
```

Notes

This is useful since it enables us to dynamically redefine models sharing an underlying architecture but with different resolutions.

call (*inputs*, *training=True*, *return_features=False*)

Execute the model on input data.

Parameters

- **inputs** (*tf.Tensor*) – Input tensor(s).
- **training** (*bool*) – Training flag.
- **return_features** (*bool*) – If True returns the features.

Returns *tf.Tensor* – The model output.

unet

UNET implementations.

Functions

*FUNet*Functional UNET Implementation

ashpy.models.convolutional.unet.FUNet

```
ashpy.models.convolutional.unet.FUNet(input_res, min_res, kernel_size, initial_filters, filters_cap, channels, input_channels=3, use_dropout_encoder=True, use_dropout_decoder=True, dropout_prob=0.3, encoder_non_linearity=<class 'tensorflow.python.keras.layers.advanced_activations.LeakyReLU'>, decoder_non_linearity=<class 'tensorflow.python.keras.layers.advanced_activations.ReLU'>, last_activation=<function tanh>, use_attention=False)
```

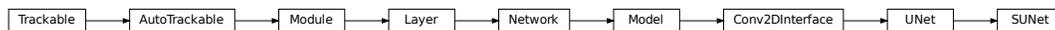
Functional UNET Implementation

Classes

<i>SUNet</i>	Semantic UNet
<i>UNet</i>	UNet Architecture

SUNet

Inheritance Diagram



```
class ashpy.models.convolutional.unet.SUNet(input_res, min_res, kernel_size, initial_filters, filters_cap, channels, use_dropout_encoder=True, use_dropout_decoder=True, dropout_prob=0.3, encoder_non_linearity=<class 'tensorflow.python.keras.layers.advanced_activations.LeakyReLU'>, decoder_non_linearity=<class 'tensorflow.python.keras.layers.advanced_activations.ReLU'>, use_attention=False)
```

Bases: *ashpy.models.convolutional.unet.UNet*

Semantic UNet

Methods

<code>__init__(input_res, min_res, kernel_size, ...)</code>	Args: <code>input_res</code> : input resolution <code>min_res</code> : minimum resolution reached after decode <code>kernel_size</code> : kernel size used in the network <code>initial_filters</code> : number of filter of the initial convolution <code>filters_cap</code> : maximum number of filters <code>channels</code> : number of output channels <code>use_dropout_encoder</code> : whether to use dropout in the encoder module <code>use_dropout_decoder</code> : whether to use dropout in the decoder module <code>dropout_prob</code> : probability of dropout <code>encoder_non_linearity</code> : non linearity of encoder <code>decoder_non_linearity</code> : non linearity of decoder <code>last_activation</code> : last activation function, tanh or softmax (for semantic images) <code>use_attention</code> : whether to use attention
---	---

Attributes

<code>activity_regularizer</code>	Optional regularizer function for the output of this layer.
<code>dtype</code>	
<code>dynamic</code>	
<code>inbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>input</code>	Retrieves the input tensor(s) of a layer.
<code>input_mask</code>	Retrieves the input mask tensor(s) of a layer.
<code>input_shape</code>	Retrieves the input shape(s) of a layer.
<code>input_spec</code>	Gets the network's input specs.
<code>layers</code>	
<code>losses</code>	Losses which are associated with this <i>Layer</i> .
<code>metrics</code>	Returns the model's metrics added using <i>compile</i> , <i>add_metric</i> APIs.
<code>metrics_names</code>	Returns the model's display labels for all outputs.
<code>name</code>	Returns the name of this module as passed or determined in the ctor.
<code>name_scope</code>	Returns a <i>tf.name_scope</i> instance for this class.
<code>non_trainable_variables</code>	
<code>non_trainable_weights</code>	
<code>outbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>output</code>	Retrieves the output tensor(s) of a layer.
<code>output_mask</code>	Retrieves the output mask tensor(s) of a layer.
<code>output_shape</code>	Retrieves the output shape(s) of a layer.
<code>run_eagerly</code>	Settable attribute indicating whether the model should run eagerly.
<code>sample_weights</code>	
<code>state_updates</code>	Returns the <i>updates</i> from all layers that are stateful.
<code>stateful</code>	
<code>submodules</code>	Sequence of all sub-modules.
<code>trainable</code>	
<code>trainable_variables</code>	Sequence of variables owned by this module and it's submodules.

Continued on next page

Table 234 – continued from previous page

<code>trainable_weights</code>	
<code>updates</code>	
<code>variables</code>	Returns the list of all layer variables/weights.
<code>weights</code>	Returns the list of all layer variables/weights.

UNet

Inheritance Diagram



```

class ashpy.models.convolutional.unet.UNet (input_res,      min_res,      kernel_size,
                                           initial_filters,  filters_cap,  chan-
                                           nels,          use_dropout_encoder=True,
                                           use_dropout_decoder=True,
                                           dropout_prob=0.3,          en-
                                           coder_non_linearity=<class      'tensor-
                                           flow.python.keras.layers.advanced_activations.LeakyReLU'>,
                                           decoder_non_linearity=<class      'tensor-
                                           flow.python.keras.layers.advanced_activations.ReLU'>,
                                           normalization_layer=<class
                                           'ashpy.layers.instance_normalization.InstanceNormalization'>,
                                           last_activation=<function      tanh>,
                                           use_attention=False)
  
```

Bases: `ashpy.models.convolutional.interfaces.Conv2DInterface`

UNet Architecture

Used in Image-to-Image Translation with Conditional Adversarial Nets¹

Examples

- Direct Usage:

```

x = tf.ones((1, 512, 512, 3))
u_net = UNet(input_res = 512,
             min_res=4,
             kernel_size=4,
             initial_filters=64,
             filters_cap=512,
             channels=3)

y = u_net(x)
print(y.shape)
print(len(u_net.trainable_variables)>0)
  
```

¹ Image-to-Image Translation with Conditional Adversarial Nets <https://arxiv.org/abs/1611.04076>

```
(1, 512, 512, 3)
True
```

Methods

<code>__init__(input_res, min_res, kernel_size, ...)</code>	<code>type input_res int</code>
<code>call(inputs[, training])</code>	Execute the model on input data.
<code>get_decoder_block(filters[, use_bn, ...])</code>	Returns a block to be used in the decoder part of the UNET :param filters: number of filters :param use_bn: whether to use batch normalization :param use_dropout: whether to use dropout :param use_attention: whether to use attention
<code>get_encoder_block(filters[, use_bn, ...])</code>	Returns a block to be used in the encoder part of the UNET :param filters: number of filters :param use_bn: whether to use batch normalization :param use_attention: whether to use attention

Attributes

<code>activity_regularizer</code>	Optional regularizer function for the output of this layer.
<code>dtype</code>	
<code>dynamic</code>	
<code>inbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>input</code>	Retrieves the input tensor(s) of a layer.
<code>input_mask</code>	Retrieves the input mask tensor(s) of a layer.
<code>input_shape</code>	Retrieves the input shape(s) of a layer.
<code>input_spec</code>	Gets the network's input specs.
<code>layers</code>	
<code>losses</code>	Losses which are associated with this <i>Layer</i> .
<code>metrics</code>	Returns the model's metrics added using <i>compile</i> , <i>add_metric</i> APIs.
<code>metrics_names</code>	Returns the model's display labels for all outputs.
<code>name</code>	Returns the name of this module as passed or determined in the ctor.
<code>name_scope</code>	Returns a <i>tf.name_scope</i> instance for this class.
<code>non_trainable_variables</code>	
<code>non_trainable_weights</code>	
<code>outbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>output</code>	Retrieves the output tensor(s) of a layer.
<code>output_mask</code>	Retrieves the output mask tensor(s) of a layer.
<code>output_shape</code>	Retrieves the output shape(s) of a layer.
<code>run_eagerly</code>	Settable attribute indicating whether the model should run eagerly.
<code>sample_weights</code>	

Continued on next page

Table 236 – continued from previous page

<code>state_updates</code>	Returns the <i>updates</i> from all layers that are stateful.
<code>stateful</code>	
<code>submodules</code>	Sequence of all sub-modules.
<code>trainable</code>	
<code>trainable_variables</code>	Sequence of variables owned by this module and it's submodules.
<code>trainable_weights</code>	
<code>updates</code>	
<code>variables</code>	Returns the list of all layer variables/weights.
<code>weights</code>	Returns the list of all layer variables/weights.

```
__init__(input_res, min_res, kernel_size, initial_filters, filters_cap, chan-
nels, use_dropout_encoder=True, use_dropout_decoder=True,
dropout_prob=0.3, encoder_non_linearity=<class 'tensor-
flow.python.keras.layers.advanced_activations.LeakyReLU'>, de-
coder_non_linearity=<class 'tensorflow.python.keras.layers.advanced_activations.ReLU'>,
normalization_layer=<class 'ashpy.layers.instance_normalization.InstanceNormalization'>,
last_activation=<function tanh>, use_attention=False)
```

Parameters

- **input_res** (*int*) – input resolution
- **min_res** (*int*) – minimum resolution reached after decode
- **kernel_size** (*int*) – kernel size used in the network
- **initial_filters** (*int*) – number of filter of the initial convolution
- **filters_cap** (*int*) – maximum number of filters
- **channels** (*int*) – number of output channels
- **use_dropout_encoder** (*bool*) – whether to use dropout in the encoder module
- **use_dropout_decoder** (*bool*) – whether to use dropout in the decoder module
- **dropout_prob** (*float*) – probability of dropout
- **encoder_non_linearity** (*Type[Layer]*) – non linearity of encoder
- **decoder_non_linearity** (*Type[Layer]*) – non linearity of decoder
- **last_activation** (*<module 'tensorflow.python.keras.api._v2.keras.activations' from '/home/docs/checkouts/readthedocs.org/user_builds/ashpy/envs/v0.1.3/lib/python3.7/site-packages/tensorflow/python/keras/api/_v2/keras/activations/__init__.py'>*) – last activation function, tanh or softmax (for semantic images)
- **use_attention** (*bool*) – whether to use attention

call (*inputs, training=False*)
Execute the model on input data.

Parameters

- **inputs** (*tf.Tensor*) – Input tensor(s).
- **training** (*bool*) – Training flag.
- **return_features** (*bool*) – If True returns the features.

Returns `tf.Tensor` – The model output.

get_decoder_block (*filters, use_bn=True, use_dropout=False, use_attention=False*)

Returns a block to be used in the decoder part of the UNET :param filters: number of filters :param use_bn: whether to use batch normalization :param use_dropout: whether to use dropout :param use_attention: whether to use attention

Returns A block to be used in the decoder part

get_encoder_block (*filters, use_bn=True, use_attention=False*)

Returns a block to be used in the encoder part of the UNET :param filters: number of filters :param use_bn: whether to use batch normalization :param use_attention: whether to use attention

Returns A block to be used in the encoder part

```
ashpy.models.convolutional.unet.FUNet (input_res,      min_res,      kernel_size,      ini-
                                     tial_filters,   filters_cap,   channels,   in-
                                     put_channels=3,   use_dropout_encoder=True,
                                     use_dropout_decoder=True,   dropout_prob=0.3,
                                     encoder_non_linearity=<class      'tensor-
                                     flow.python.keras.layers.advanced_activations.LeakyReLU'>,
                                     decoder_non_linearity=<class      'tensor-
                                     flow.python.keras.layers.advanced_activations.ReLU'>,
                                     last_activation=<function      tanh>,
                                     use_attention=False)
```

Functional UNET Implementation

```
class ashpy.models.convolutional.unet.SUNet (input_res,      min_res,      kernel_size,
                                             initial_filters,   filters_cap,   chan-
                                             nels,           use_dropout_encoder=True,
                                             use_dropout_decoder=True,
                                             dropout_prob=0.3,           en-
                                             coder_non_linearity=<class      'tensor-
                                             flow.python.keras.layers.advanced_activations.LeakyReLU'>,
                                             decoder_non_linearity=<class      'tensor-
                                             flow.python.keras.layers.advanced_activations.ReLU'>,
                                             use_attention=False)
```

Bases: `ashpy.models.convolutional.unet.UNet`

Semantic UNet

```
class ashpy.models.convolutional.unet.UNet (input_res,      min_res,      kernel_size,
                                             initial_filters,   filters_cap,   chan-
                                             nels,           use_dropout_encoder=True,
                                             use_dropout_decoder=True,
                                             dropout_prob=0.3,           en-
                                             coder_non_linearity=<class      'tensor-
                                             flow.python.keras.layers.advanced_activations.LeakyReLU'>,
                                             decoder_non_linearity=<class      'tensor-
                                             flow.python.keras.layers.advanced_activations.ReLU'>,
                                             normalization_layer=<class
                                             'ashpy.layers.instance_normalization.InstanceNormalization'>,
                                             last_activation=<function      tanh>,
                                             use_attention=False)
```

Bases: `ashpy.models.convolutional.interfaces.Conv2DInterface`

UNet Architecture

Used in Image-to-Image Translation with Conditional Adversarial Nets¹

Examples

- Direct Usage:

```
x = tf.ones((1, 512, 512, 3))
u_net = UNet(input_res = 512,
            min_res=4,
            kernel_size=4,
            initial_filters=64,
            filters_cap=512,
            channels=3)
y = u_net(x)
print(y.shape)
print(len(u_net.trainable_variables)>0)
```

```
(1, 512, 512, 3)
True
```

```
__init__(input_res, min_res, kernel_size, initial_filters, filters_cap, chan-
nels, use_dropout_encoder=True, use_dropout_decoder=True,
dropout_prob=0.3, encoder_non_linearity=<class 'tensor-
flow.python.keras.layers.advanced_activations.LeakyReLU'>, de-
coder_non_linearity=<class 'tensorflow.python.keras.layers.advanced_activations.ReLU'>,
normalization_layer=<class 'ashpy.layers.instance_normalization.InstanceNormalization'>,
last_activation=<function tanh>, use_attention=False)
```

Parameters

- **input_res** (*int*) – input resolution
- **min_res** (*int*) – minimum resolution reached after decode
- **kernel_size** (*int*) – kernel size used in the network
- **initial_filters** (*int*) – number of filter of the initial convolution
- **filters_cap** (*int*) – maximum number of filters
- **channels** (*int*) – number of output channels
- **use_dropout_encoder** (*bool*) – whether to use dropout in the encoder module
- **use_dropout_decoder** (*bool*) – whether to use dropout in the decoder module
- **dropout_prob** (*float*) – probability of dropout
- **encoder_non_linearity** (*Type[Layer]*) – non linearity of encoder
- **decoder_non_linearity** (*Type[Layer]*) – non linearity of decoder
- **last_activation** (*<module 'tensorflow.python.keras.api._v2.keras.activations' from '/home/docs/checkouts/readthedocs.org/user_builds/ashpy/envs/v0.1.3/lib/python3.7/site-packages/tensorflow/python/keras/api/_v2/keras/activations/__init__.py'>*) – last activation function, tanh or softmax (for semantic images)

¹ Image-to-Image Translation with Conditional Adversarial Nets <https://arxiv.org/abs/1611.04076>

- **use_attention** (`bool`) – whether to use attention

call (*inputs, training=False*)

Execute the model on input data.

Parameters

- **inputs** (`tf.Tensor`) – Input tensor(s).
- **training** (`bool`) – Training flag.
- **return_features** (`bool`) – If True returns the features.

Returns `tf.Tensor` – The model output.

get_decoder_block (*filters, use_bn=True, use_dropout=False, use_attention=False*)

Returns a block to be used in the decoder part of the UNET :param filters: number of filters :param use_bn: whether to use batch normalization :param use_dropout: whether to use dropout :param use_attention: whether to use attention

Returns A block to be used in the decoder part

get_encoder_block (*filters, use_bn=True, use_attention=False*)

Returns a block to be used in the encoder part of the UNET :param filters: number of filters :param use_bn: whether to use batch normalization :param use_attention: whether to use attention

Returns A block to be used in the encoder part

pix2pixhd

Pix2Pix HD Implementation See: “High-Resolution Image Synthesis and Semantic Manipulation with Conditional GANs”¹

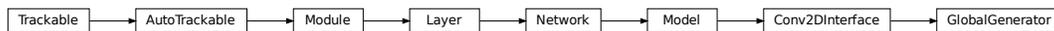
Global Generator + Local Enhancer

Classes

<i>GlobalGenerator</i>	Global Generator from pix2pixHD paper:
<i>LocalEnhancer</i>	Local Enhancer module of the Pix2PixHD architecture.
<i>ResNetBlock</i>	ResNet Blocks: the input filters is the same as the output filters.

GlobalGenerator

Inheritance Diagram



¹ High-Resolution Image Synthesis and Semantic Manipulation with Conditional GANs: <https://arxiv.org/abs/1711.11585>

```
class ashpy.models.convolutional.pix2pixhd.GlobalGenerator (input_res=512,
                                                           min_res=64,      ini-
                                                           tial_filters=64,
                                                           filters_cap=512,
                                                           channels=3, normal-
                                                           ization_layer=<class
                                                           'ashpy.layers.instance_normalization.Instance
                                                           non_linearity=<class
                                                           'tensor-
                                                           flow.python.keras.layers.advanced_activations.
                                                           num_resnet_blocks=9,
                                                           kernel_size_resnet=3,
                                                           ker-
                                                           nel_size_front_back=7,
                                                           num_internal_resnet_blocks=2)
```

Bases: *ashpy.models.convolutional.interfaces.Conv2DInterface*

Global Generator from pix2pixHD paper:

- G1^F: Convolutional frontend (downsampling)
- G1^R: ResNet Block
- G1^B: Convolutional backend (upsampling)

Methods

<code>__init__</code> ([input_res, min_res, ...])	Global Generator from Pix2PixHD
<code>call</code> (inputs[, training])	Call of the Pix2Pix HD model :param inputs: input tensor(s) :param training: If True training phase

Attributes

<code>activity_regularizer</code>	Optional regularizer function for the output of this layer.
<code>dtype</code>	
<code>dynamic</code>	
<code>inbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>input</code>	Retrieves the input tensor(s) of a layer.
<code>input_mask</code>	Retrieves the input mask tensor(s) of a layer.
<code>input_shape</code>	Retrieves the input shape(s) of a layer.
<code>input_spec</code>	Gets the network's input specs.
<code>layers</code>	
<code>losses</code>	Losses which are associated with this <i>Layer</i> .
<code>metrics</code>	Returns the model's metrics added using <i>compile</i> , <i>add_metric</i> APIs.
<code>metrics_names</code>	Returns the model's display labels for all outputs.
<code>name</code>	Returns the name of this module as passed or determined in the ctor.
<code>name_scope</code>	Returns a <i>tf.name_scope</i> instance for this class.
<code>non_trainable_variables</code>	

Continued on next page

Table 239 – continued from previous page

<code>non_trainable_weights</code>	
<code>outbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>output</code>	Retrieves the output tensor(s) of a layer.
<code>output_mask</code>	Retrieves the output mask tensor(s) of a layer.
<code>output_shape</code>	Retrieves the output shape(s) of a layer.
<code>run_eagerly</code>	Settable attribute indicating whether the model should run eagerly.
<code>sample_weights</code>	
<code>state_updates</code>	Returns the <i>updates</i> from all layers that are stateful.
<code>stateful</code>	
<code>submodules</code>	Sequence of all sub-modules.
<code>trainable</code>	
<code>trainable_variables</code>	Sequence of variables owned by this module and it's submodules.
<code>trainable_weights</code>	
<code>updates</code>	
<code>variables</code>	Returns the list of all layer variables/weights.
<code>weights</code>	Returns the list of all layer variables/weights.

```
__init__(input_res=512, min_res=64, initial_filters=64, filters_cap=512, channels=3, normalization_layer=<class 'ashpy.layers.instance_normalization.InstanceNormalization'>, non_linearity=<class 'tensorflow.python.keras.layers.advanced_activations.ReLU'>, num_resnet_blocks=9, kernel_size_resnet=3, kernel_size_front_back=7, num_internal_resnet_blocks=2)
Global Generator from Pix2PixHD
```

Parameters

- **input_res** (*int*) – Input Resolution
- **min_res** (*int*) – Minimum resolution reached by the downsampling
- **initial_filters** (*int*) – number of initial filters
- **filters_cap** (*int*) – maximum number of filters
- **channels** (*int*) – output channels
- **normalization_layer** (`tf.keras.layers.Layer`) – normalization layer used by the global generator, can be Instance Norm, Layer Norm, Batch Norm
- **non_linearity** (`tf.keras.layers.Layer`) – non linearity used in the global generator
- **num_resnet_blocks** (*int*) – number of resnet blocks
- **kernel_size_resnet** (*int*) – kernel size used in resnets conv layers
- **kernel_size_front_back** (*int*) – kernel size used by the convolutional frontend and backend
- **num_internal_resnet_blocks** (*int*) – number of blocks used by internal resnet

call (*inputs*, *training=True*)

Call of the Pix2Pix HD model :param inputs: input tensor(s) :param training: If True training phase

Returns Tuple – Generated images.

LocalEnhancer

Inheritance Diagram



```

class ashpy.models.convolutional.pix2pixhd.LocalEnhancer (input_res=512,
                                                           min_res=64,      ini-
                                                           tial_filters=64,   fil-
                                                           ters_cap=512,     chan-
                                                           nels=3,          normaliza-
                                                           tion_layer=<class
                                                           'ashpy.layers.instance_normalization.InstanceNor
                                                           non_linearity=<class
                                                           'tensor-
                                                           flow.python.keras.layers.advanced_activations.Re
                                                           num_resnet_blocks_global=9,
                                                           num_resnet_blocks_local=3,
                                                           kernel_size_resnet=3,
                                                           ker-
                                                           nel_size_front_back=7,
                                                           num_internal_resnet_blocks=2)

```

Bases: tensorflow.python.keras.engine.training.Model

Local Enhancer module of the Pix2PixHD architecture.

Example

```

# instantiate the model
model = LocalEnhancer()

# call the model passing inputs
inputs = tf.ones((1, 512, 512, 3))
output = model(inputs)

# the output shape is
# the same as the input shape
print(output.shape)

```

```
(1, 512, 512, 3)
```

Methods

<code>__init__</code> ([input_res, min_res, ...])	Build the LocalEnhancer module of the Pix2PixHD architecture.
<code>call</code> (inputs[, training])	LocalEnhancer call.

Attributes

<code>activity_regularizer</code>	Optional regularizer function for the output of this layer.
<code>dtype</code>	
<code>dynamic</code>	
<code>inbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>input</code>	Retrieves the input tensor(s) of a layer.
<code>input_mask</code>	Retrieves the input mask tensor(s) of a layer.
<code>input_shape</code>	Retrieves the input shape(s) of a layer.
<code>input_spec</code>	Gets the network's input specs.
<code>layers</code>	
<code>losses</code>	Losses which are associated with this <i>Layer</i> .
<code>metrics</code>	Returns the model's metrics added using <i>compile</i> , <i>add_metric</i> APIs.
<code>metrics_names</code>	Returns the model's display labels for all outputs.
<code>name</code>	Returns the name of this module as passed or determined in the ctor.
<code>name_scope</code>	Returns a <i>tf.name_scope</i> instance for this class.
<code>non_trainable_variables</code>	
<code>non_trainable_weights</code>	
<code>outbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>output</code>	Retrieves the output tensor(s) of a layer.
<code>output_mask</code>	Retrieves the output mask tensor(s) of a layer.
<code>output_shape</code>	Retrieves the output shape(s) of a layer.
<code>run_eagerly</code>	Settable attribute indicating whether the model should run eagerly.
<code>sample_weights</code>	
<code>state_updates</code>	Returns the <i>updates</i> from all layers that are stateful.
<code>stateful</code>	
<code>submodules</code>	Sequence of all sub-modules.
<code>trainable</code>	
<code>trainable_variables</code>	Sequence of variables owned by this module and it's submodules.
<code>trainable_weights</code>	
<code>updates</code>	
<code>variables</code>	Returns the list of all layer variables/weights.
<code>weights</code>	Returns the list of all layer variables/weights.

`__init__`(input_res=512, min_res=64, initial_filters=64, filters_cap=512, channels=3, normalization_layer=<class 'ashpy.layers.instance_normalization.InstanceNormalization'>, non_linearity=<class 'tensorflow.python.keras.layers.advanced_activations.ReLU'>, num_resnet_blocks_global=9, num_resnet_blocks_local=3, kernel_size_resnet=3, kernel_size_front_back=7, num_internal_resnet_blocks=2)
Build the LocalEnhancer module of the Pix2PixHD architecture.

See High-Resolution Image Synthesis and Semantic Manipulation with Conditional GANs² for more details.

Parameters

- **input_res** (*int*) – input resolution
- **min_res** (*int*) – minimum resolution reached by the global generator
- **initial_filters** (*int*) – number of initial filters
- **filters_cap** (*int*) – maximum number of filters
- **channels** (*int*) – number of channels
- **normalization_layer** (`tf.keras.layers.Layer`) – layer of normalization
- **Instance Normalization or BatchNormalization or LayerNormalization** ((*e.g.*) –
- **non_linearity** (`tf.keras.layers.Layer`) – non linearity used in Pix2Pix HD
- **num_resnet_blocks_global** (*int*) – number of residual blocks used in the global generator
- **num_resnet_blocks_local** (*int*) – number of residual blocks used in the local generator
- **kernel_size_resnet** (*int*) – kernel size used in resnets
- **kernel_size_front_back** (*int*) – kernel size used for the front and back convolution
- **num_internal_resnet_blocks** (*int*) – number of internal blocks of the resnet

call (*inputs*, *training=False*)

LocalEnhancer call. :param inputs: Input Tensors :type inputs: `tf.Tensor` :param training: Whether it is training phase or not :type training: bool

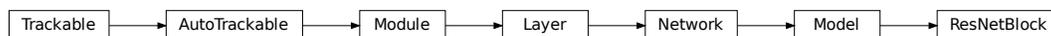
Returns

(`tf.Tensor`) –

Image of size (input_res, input_res, channels) as specified in the init call

ResNetBlock

Inheritance Diagram



² High-Resolution Image Synthesis and Semantic Manipulation with Conditional GANs <https://arxiv.org/abs/1711.11585>

```
class ashpy.models.convolutional.pix2pixhd.ResNetBlock (filters, normalization_layer=<class
'ashpy.layers.instance_normalization.InstanceNormaliza-
tion_layer=<class
'ashpy.layers.instance_normalization.InstanceNormaliza-
non_linearity=<class
'tensor-
flow.python.keras.layers.advanced_activations.ReLU
kernel_size=3,
num_blocks=2)
```

Bases: tensorflow.python.keras.engine.training.Model

ResNet Blocks: the input filters is the same as the output filters.

Methods

<code>__init__(filters[, normalization_layer, ...])</code>	ResNet block composed by num_blocks.
<code>call(inputs[, training])</code>	Forward pass :param inputs: input tensor :param training: whether is training or not

Attributes

<code>activity_regularizer</code>	Optional regularizer function for the output of this layer.
<code>dtype</code>	
<code>dynamic</code>	
<code>inbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>input</code>	Retrieves the input tensor(s) of a layer.
<code>input_mask</code>	Retrieves the input mask tensor(s) of a layer.
<code>input_shape</code>	Retrieves the input shape(s) of a layer.
<code>input_spec</code>	Gets the network's input specs.
<code>layers</code>	
<code>losses</code>	Losses which are associated with this <i>Layer</i> .
<code>metrics</code>	Returns the model's metrics added using <i>compile</i> , <i>add_metric</i> APIs.
<code>metrics_names</code>	Returns the model's display labels for all outputs.
<code>name</code>	Returns the name of this module as passed or determined in the ctor.
<code>name_scope</code>	Returns a <i>tf.name_scope</i> instance for this class.
<code>non_trainable_variables</code>	
<code>non_trainable_weights</code>	
<code>outbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>output</code>	Retrieves the output tensor(s) of a layer.
<code>output_mask</code>	Retrieves the output mask tensor(s) of a layer.
<code>output_shape</code>	Retrieves the output shape(s) of a layer.
<code>run_eagerly</code>	Settable attribute indicating whether the model should run eagerly.
<code>sample_weights</code>	
<code>state_updates</code>	Returns the <i>updates</i> from all layers that are stateful.
<code>stateful</code>	

Continued on next page

Table 243 – continued from previous page

<code>submodules</code>	Sequence of all sub-modules.
<code>trainable</code>	
<code>trainable_variables</code>	Sequence of variables owned by this module and it's submodules.
<code>trainable_weights</code>	
<code>updates</code>	
<code>variables</code>	Returns the list of all layer variables/weights.
<code>weights</code>	Returns the list of all layer variables/weights.

`__init__` (*filters*, *normalization_layer*=<class 'ashpy.layers.instance_normalization.InstanceNormalization'>, *non_linearity*=<class 'tensorflow.python.keras.layers.advanced_activations.ReLU'>, *kernel_size*=3, *num_blocks*=2)

ResNet block composed by `num_blocks`. Each block is composed by

- Conv2D with strides 1 and padding “same”
- Normalization Layer
- Non Linearity

The final result is the output of the ResNet + input

Parameters

- **filters** (*int*) – initial filters (same as the output filters)
- **normalization_layer** (`tf.keras.layers.Layer`) – layer of normalization used by the residual block
- **non_linearity** (`tf.keras.layers.Layer`) – non linearity used in the resnet block
- **kernel_size** (*int*) – kernel size used in the resnet block
- **num_blocks** (*int*) – number of blocks, each block is composed by conv, normalization and non linearity

`call` (*inputs*, *training*=False)

Forward pass :param inputs: input tensor :param training: whether is training or not

Returns A Tensor of the same shape as the inputs. The input passed through `num_blocks` blocks.

```
class ashpy.models.convolutional.pix2pixhd.GlobalGenerator (input_res=512,
                                                            min_res=64,      ini-
                                                            tial_filters=64,
                                                            filters_cap=512,
                                                            channels=3, normal-
                                                            ization_layer=<class
                                                            'ashpy.layers.instance_normalization.InstanceNormali-
                                                            zation'>,
                                                            non_linearity=<class
                                                            'tensorflow.python.keras.layers.advanced_activations.ReLU'>,
                                                            num_resnet_blocks=9,
                                                            kernel_size_resnet=3,
                                                            kernel_size_front_back=7,
                                                            num_internal_resnet_blocks=2)
```

Bases: `ashpy.models.convolutional.interfaces.Conv2DInterface`

Global Generator from pix2pixHD paper:

- G1^F: Convolutional frontend (downsampling)
- G1^R: ResNet Block
- G1^B: Convolutional backend (upsampling)

```
__init__(input_res=512, min_res=64, initial_filters=64, filters_cap=512, channels=3, normalization_layer=<class 'ashpy.layers.instance_normalization.InstanceNormalization'>, non_linearity=<class 'tensorflow.python.keras.layers.advanced_activations.ReLU'>, num_resnet_blocks=9, kernel_size_resnet=3, kernel_size_front_back=7, num_internal_resnet_blocks=2)
```

Global Generator from Pix2PixHD

Parameters

- **input_res** (*int*) – Input Resolution
- **min_res** (*int*) – Minimum resolution reached by the downsampling
- **initial_filters** (*int*) – number of initial filters
- **filters_cap** (*int*) – maximum number of filters
- **channels** (*int*) – output channels
- **normalization_layer** (*tf.keras.layers.Layer*) – normalization layer used by the global generator, can be Instance Norm, Layer Norm, Batch Norm
- **non_linearity** (*tf.keras.layers.Layer*) – non linearity used in the global generator
- **num_resnet_blocks** (*int*) – number of resnet blocks
- **kernel_size_resnet** (*int*) – kernel size used in resnets conv layers
- **kernel_size_front_back** (*int*) – kernel size used by the convolutional frontend and backend
- **num_internal_resnet_blocks** (*int*) – number of blocks used by internal resnet

call (*inputs, training=True*)

Call of the Pix2Pix HD model :param inputs: input tensor(s) :param training: If True training phase

Returns *tuple* – Generated images.

```
class ashpy.models.convolutional.pix2pixhd.LocalEnhancer(input_res=512, min_res=64, initial_filters=64, filters_cap=512, channels=3, normalization_layer=<class 'ashpy.layers.instance_normalization.InstanceNormalization'>, non_linearity=<class 'tensorflow.python.keras.layers.advanced_activations.ReLU'>, num_resnet_blocks_global=9, num_resnet_blocks_local=3, kernel_size_resnet=3, kernel_size_front_back=7, num_internal_resnet_blocks=2)
```

Bases: *tensorflow.python.keras.engine.training.Model*

Local Enhancer module of the Pix2PixHD architecture.

Example

```
# instantiate the model
model = LocalEnhancer()

# call the model passing inputs
inputs = tf.ones((1, 512, 512, 3))
output = model(inputs)

# the output shape is
# the same as the input shape
print(output.shape)
```

```
(1, 512, 512, 3)
```

```
__init__(input_res=512, min_res=64, initial_filters=64, filters_cap=512, channels=3, normalization_layer=<class 'ashpy.layers.instance_normalization.InstanceNormalization'>, non_linearity=<class 'tensorflow.python.keras.layers.advanced_activations.ReLU'>, num_resnet_blocks_global=9, num_resnet_blocks_local=3, kernel_size_resnet=3, kernel_size_front_back=7, num_internal_resnet_blocks=2)
```

Build the LocalEnhancer module of the Pix2PixHD architecture.

See High-Resolution Image Synthesis and Semantic Manipulation with Conditional GANs² for more details.

Parameters

- **input_res** (*int*) – input resolution
- **min_res** (*int*) – minimum resolution reached by the global generator
- **initial_filters** (*int*) – number of initial filters
- **filters_cap** (*int*) – maximum number of filters
- **channels** (*int*) – number of channels
- **normalization_layer** (*tf.keras.layers.Layer*) – layer of normalization
- **Instance Normalization or BatchNormalization or LayerNormalization** ((*e.g.*)) –
- **non_linearity** (*tf.keras.layers.Layer*) – non linearity used in Pix2Pix HD
- **num_resnet_blocks_global** (*int*) – number of residual blocks used in the global generator
- **num_resnet_blocks_local** (*int*) – number of residual blocks used in the local generator
- **kernel_size_resnet** (*int*) – kernel size used in resnets
- **kernel_size_front_back** (*int*) – kernel size used for the front and back convolution
- **num_internal_resnet_blocks** (*int*) – number of internal blocks of the resnet

call (*inputs, training=False*)

LocalEnhancer call. :param inputs: Input Tensors :type inputs: *tf.Tensor* :param training: Whether it is training phase or not :type training: bool

² High-Resolution Image Synthesis and Semantic Manipulation with Conditional GANs <https://arxiv.org/abs/1711.11585>

Returns

(`tf.Tensor`) –

Image of size (input_res, input_res, channels) as specified in the init call

```
class ashpy.models.convolutional.pix2pixhd.ResNetBlock (filters, normalization_layer=<class 'ashpy.layers.instance_normalization.InstanceNormalization'>, non_linearity=<class 'tensorflow.python.keras.layers.advanced_activations.ReLU'>, kernel_size=3, num_blocks=2)
```

Bases: `tensorflow.python.keras.engine.training.Model`

ResNet Blocks: the input filters is the same as the output filters.

```
__init__ (filters, normalization_layer=<class 'ashpy.layers.instance_normalization.InstanceNormalization'>, non_linearity=<class 'tensorflow.python.keras.layers.advanced_activations.ReLU'>, kernel_size=3, num_blocks=2)
```

ResNet block composed by `num_blocks`. Each block is composed by

- Conv2D with strides 1 and padding “same”
- Normalization Layer
- Non Linearity

The final result is the output of the ResNet + input

Parameters

- **filters** (*int*) – initial filters (same as the output filters)
- **normalization_layer** (`tf.keras.layers.Layer`) – layer of normalization used by the residual block
- **non_linearity** (`tf.keras.layers.Layer`) – non linearity used in the resnet block
- **kernel_size** (*int*) – kernel size used in the resnet block
- **num_blocks** (*int*) – number of blocks, each block is composed by conv, normalization and non linearity

```
call (inputs, training=False)
```

Forward pass :param inputs: input tensor :param training: whether is training or not

Returns A Tensor of the same shape as the inputs. The input passed through `num_blocks` blocks.

4.7.2 fc

Collection of Fully Connected Models constructors.

Interfaces

`interfaces.FCInterface`

Primitive Interface to be used by all `ashpy.models`.

FCInterface

Inheritance Diagram



class ashpy.models.fc.interfaces.FCInterface

Bases: tensorflow.python.keras.engine.training.Model

Primitive Interface to be used by all *ashpy.models*.

Methods

<code>__init__()</code>	Primitive Interface to be used by all <i>ashpy.models</i> .
<code>call(inputs[, training])</code>	Execute the model on input data.

Attributes

<code>activity_regularizer</code>	Optional regularizer function for the output of this layer.
<code>dtype</code>	
<code>dynamic</code>	
<code>inbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>input</code>	Retrieves the input tensor(s) of a layer.
<code>input_mask</code>	Retrieves the input mask tensor(s) of a layer.
<code>input_shape</code>	Retrieves the input shape(s) of a layer.
<code>input_spec</code>	Gets the network's input specs.
<code>layers</code>	
<code>losses</code>	Losses which are associated with this <i>Layer</i> .
<code>metrics</code>	Returns the model's metrics added using <i>compile</i> , <i>add_metric</i> APIs.
<code>metrics_names</code>	Returns the model's display labels for all outputs.
<code>name</code>	Returns the name of this module as passed or determined in the ctor.
<code>name_scope</code>	Returns a <i>tf.name_scope</i> instance for this class.
<code>non_trainable_variables</code>	
<code>non_trainable_weights</code>	
<code>outbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>output</code>	Retrieves the output tensor(s) of a layer.
<code>output_mask</code>	Retrieves the output mask tensor(s) of a layer.
<code>output_shape</code>	Retrieves the output shape(s) of a layer.

Continued on next page

Table 246 – continued from previous page

<code>run_eagerly</code>	Settable attribute indicating whether the model should run eagerly.
<code>sample_weights</code>	
<code>state_updates</code>	Returns the <i>updates</i> from all layers that are stateful.
<code>stateful</code>	
<code>submodules</code>	Sequence of all sub-modules.
<code>trainable</code>	
<code>trainable_variables</code>	Sequence of variables owned by this module and it's submodules.
<code>trainable_weights</code>	
<code>updates</code>	
<code>variables</code>	Returns the list of all layer variables/weights.
<code>weights</code>	Returns the list of all layer variables/weights.

`__init__()`

Primitive Interface to be used by all *ashpy.models*.

Declares the `self._layers` list.

Returns `None`

`call(inputs, training=True)`

Execute the model on input data.

Parameters

- **inputs** (`tf.Tensor`) – Input tensors.
- **training** (`bool`) – Training flag.

Returns `:py_class:'tf.Tensor'`

Decoders

`decoders.BaseDecoder`

Primitive Model for all fully connected decoder based architecture.

BaseDecoder

Inheritance Diagram



class `ashpy.models.fc.decoders.BaseDecoder` (*hidden_units, output_shape*)

Bases: `ashpy.models.fc.interfaces.FCInterface`

Primitive Model for all fully connected decoder based architecture.

Examples

```
decoder = BaseDecoder(
    hidden_units=[64,128,256],
    output_shape=55)
print(decoder(tf.zeros((1,10))).shape)
```

```
(1, 55)
```

Methods

<code>__init__(hidden_units, output_shape)</code>	Instantiate the <i>BaseDecoder</i> .
---	--------------------------------------

Attributes

<code>activity_regularizer</code>	Optional regularizer function for the output of this layer.
<code>dtype</code>	
<code>dynamic</code>	
<code>inbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>input</code>	Retrieves the input tensor(s) of a layer.
<code>input_mask</code>	Retrieves the input mask tensor(s) of a layer.
<code>input_shape</code>	Retrieves the input shape(s) of a layer.
<code>input_spec</code>	Gets the network's input specs.
<code>layers</code>	
<code>losses</code>	Losses which are associated with this <i>Layer</i> .
<code>metrics</code>	Returns the model's metrics added using <i>compile</i> , <i>add_metric</i> APIs.
<code>metrics_names</code>	Returns the model's display labels for all outputs.
<code>name</code>	Returns the name of this module as passed or determined in the ctor.
<code>name_scope</code>	Returns a <i>tf.name_scope</i> instance for this class.
<code>non_trainable_variables</code>	
<code>non_trainable_weights</code>	
<code>outbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>output</code>	Retrieves the output tensor(s) of a layer.
<code>output_mask</code>	Retrieves the output mask tensor(s) of a layer.
<code>output_shape</code>	Retrieves the output shape(s) of a layer.
<code>run_eagerly</code>	Settable attribute indicating whether the model should run eagerly.
<code>sample_weights</code>	
<code>state_updates</code>	Returns the <i>updates</i> from all layers that are stateful.
<code>stateful</code>	
<code>submodules</code>	Sequence of all sub-modules.

Continued on next page

Table 249 – continued from previous page

<code>trainable</code>	
<code>trainable_variables</code>	Sequence of variables owned by this module and it's submodules.
<code>trainable_weights</code>	
<code>updates</code>	
<code>variables</code>	Returns the list of all layer variables/weights.
<code>weights</code>	Returns the list of all layer variables/weights.

`__init__` (*hidden_units*, *output_shape*)
 Instantiate the *BaseDecoder*.

Parameters

- **hidden_units** (tuple of int) – Number of units per hidden layer.
- **output_shape** (int) – Amount of units of the last `tf.keras.layers.Dense`.

Returns None

Encoders

<code>encoders.BaseEncoder</code>	Primitive Model for all fully connected encoder based architecture.
-----------------------------------	---

BaseEncoder

Inheritance Diagram



class `ashpy.models.fc.encoders.BaseEncoder` (*hidden_units*, *output_shape*)
 Bases: `ashpy.models.fc.interfaces.FCInterface`
 Primitive Model for all fully connected encoder based architecture.

Examples

```

encoder = BaseEncoder(
    hidden_units=[256, 128, 64],
    output_shape=10)
print(encoder(tf.zeros((1, 55))).shape)
  
```

(1, 10)

Methods

<code>__init__(hidden_units, output_shape)</code>	Instantiate the BaseDecoder.
---	------------------------------

Attributes

<code>activity_regularizer</code>	Optional regularizer function for the output of this layer.
<code>dtype</code>	
<code>dynamic</code>	
<code>inbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>input</code>	Retrieves the input tensor(s) of a layer.
<code>input_mask</code>	Retrieves the input mask tensor(s) of a layer.
<code>input_shape</code>	Retrieves the input shape(s) of a layer.
<code>input_spec</code>	Gets the network's input specs.
<code>layers</code>	
<code>losses</code>	Losses which are associated with this <i>Layer</i> .
<code>metrics</code>	Returns the model's metrics added using <i>compile</i> , <i>add_metric</i> APIs.
<code>metrics_names</code>	Returns the model's display labels for all outputs.
<code>name</code>	Returns the name of this module as passed or determined in the ctor.
<code>name_scope</code>	Returns a <i>tf.name_scope</i> instance for this class.
<code>non_trainable_variables</code>	
<code>non_trainable_weights</code>	
<code>outbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>output</code>	Retrieves the output tensor(s) of a layer.
<code>output_mask</code>	Retrieves the output mask tensor(s) of a layer.
<code>output_shape</code>	Retrieves the output shape(s) of a layer.
<code>run_eagerly</code>	Settable attribute indicating whether the model should run eagerly.
<code>sample_weights</code>	
<code>state_updates</code>	Returns the <i>updates</i> from all layers that are stateful.
<code>stateful</code>	
<code>submodules</code>	Sequence of all sub-modules.
<code>trainable</code>	
<code>trainable_variables</code>	Sequence of variables owned by this module and it's submodules.
<code>trainable_weights</code>	
<code>updates</code>	
<code>variables</code>	Returns the list of all layer variables/weights.
<code>weights</code>	Returns the list of all layer variables/weights.

`__init__(hidden_units, output_shape)`
 Instantiate the BaseDecoder.

Parameters

- **hidden_units** (*tuple of int*) – Number of units per hidden layer.
- **output_shape** (*int*) – Amount of units of the last `tf.keras.layers.Dense`.

Returns `None`

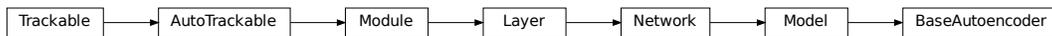
Autoencoders

`autoencoders.BaseAutoencoder`

Primitive Model for all fully connected autoencoders.

BaseAutoencoder

Inheritance Diagram



class `ashpy.models.fc.autoencoders.BaseAutoencoder` (*hidden_units*, *encoding_dimension*, *output_shape*)

Bases: `tensorflow.python.keras.engine.training.Model`

Primitive Model for all fully connected autoencoders.

Examples

- Direct Usage:

```

autoencoder = BaseAutoencoder(
    hidden_units=[256, 128, 64],
    encoding_dimension=100,
    output_shape=55
)

encoding, reconstruction = autoencoder(tf.zeros((1, 55)))
print(encoding.shape)
print(reconstruction.shape)
  
```

Methods

<code>__init__(hidden_units, encoding_dimension, ...)</code>	Instantiate the BaseDecoder.
<code>call(inputs[, training])</code>	Execute the model on input data.

Attributes

<code>activity_regularizer</code>	Optional regularizer function for the output of this layer.
<code>dtype</code>	
<code>dynamic</code>	
<code>inbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>input</code>	Retrieves the input tensor(s) of a layer.
<code>input_mask</code>	Retrieves the input mask tensor(s) of a layer.
<code>input_shape</code>	Retrieves the input shape(s) of a layer.
<code>input_spec</code>	Gets the network's input specs.
<code>layers</code>	
<code>losses</code>	Losses which are associated with this <i>Layer</i> .
<code>metrics</code>	Returns the model's metrics added using <i>compile</i> , <i>add_metric</i> APIs.
<code>metrics_names</code>	Returns the model's display labels for all outputs.
<code>name</code>	Returns the name of this module as passed or determined in the ctor.
<code>name_scope</code>	Returns a <i>tf.name_scope</i> instance for this class.
<code>non_trainable_variables</code>	
<code>non_trainable_weights</code>	
<code>outbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>output</code>	Retrieves the output tensor(s) of a layer.
<code>output_mask</code>	Retrieves the output mask tensor(s) of a layer.
<code>output_shape</code>	Retrieves the output shape(s) of a layer.
<code>run_eagerly</code>	Settable attribute indicating whether the model should run eagerly.
<code>sample_weights</code>	
<code>state_updates</code>	Returns the <i>updates</i> from all layers that are stateful.
<code>stateful</code>	
<code>submodules</code>	Sequence of all sub-modules.
<code>trainable</code>	
<code>trainable_variables</code>	Sequence of variables owned by this module and it's submodules.
<code>trainable_weights</code>	
<code>updates</code>	
<code>variables</code>	Returns the list of all layer variables/weights.
<code>weights</code>	Returns the list of all layer variables/weights.

`__init__(hidden_units, encoding_dimension, output_shape)`
 Instantiate the BaseDecoder.

Parameters

- `hidden_units` (tuple of int) – Number of units per hidden layer.

- **encoding_dimension** (*int*) – encoding dimension.
- **output_shape** (*int*) – output shape, usual equal to the input shape.

Returns *None*

call (*inputs, training=True*)

Execute the model on input data.

Parameters

- **inputs** (*tf.Tensor*) – Input tensors.
- **training** (*bool*) – Training flag.

Returns (*encoding, reconstruction*) – Pair of tensors.

Modules

<i>autoencoders</i>	Collection of Fully Connected Autoencoders
<i>decoders</i>	Collection of Decoders (i.e., GANs' Discriminators) models.
<i>encoders</i>	Collection of Encoders (i.e., GANs' Discriminators) models.
<i>interfaces</i>	Primitive Fully Connected interfaces.

autoencoders

Collection of Fully Connected Autoencoders

Classes

<i>BaseAutoencoder</i>	Primitive Model for all fully connected autoencoders.
------------------------	---

BaseAutoencoder

Inheritance Diagram



class `ashpy.models.fc.autoencoders.BaseAutoencoder` (*hidden_units*, *encoding_dimension*, *output_shape*)

Bases: `tensorflow.python.keras.engine.training.Model`

Primitive Model for all fully connected autoencoders.

Examples

- Direct Usage:

```

autoencoder = BaseAutoencoder(
    hidden_units=[256,128,64],
    encoding_dimension=100,
    output_shape=55
)

encoding, reconstruction = autoencoder(tf.zeros((1, 55)))
print(encoding.shape)
print(reconstruction.shape)

```

Methods

<code>__init__(hidden_units, encoding_dimension, ...)</code>	Instantiate the BaseDecoder.
<code>call(inputs[, training])</code>	Execute the model on input data.

Attributes

<code>activity_regularizer</code>	Optional regularizer function for the output of this layer.
<code>dtype</code>	
<code>dynamic</code>	
<code>inbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>input</code>	Retrieves the input tensor(s) of a layer.
<code>input_mask</code>	Retrieves the input mask tensor(s) of a layer.
<code>input_shape</code>	Retrieves the input shape(s) of a layer.
<code>input_spec</code>	Gets the network's input specs.
<code>layers</code>	
<code>losses</code>	Losses which are associated with this <i>Layer</i> .
<code>metrics</code>	Returns the model's metrics added using <i>compile</i> , <i>add_metric</i> APIs.
<code>metrics_names</code>	Returns the model's display labels for all outputs.
<code>name</code>	Returns the name of this module as passed or determined in the ctor.
<code>name_scope</code>	Returns a <i>tf.name_scope</i> instance for this class.
<code>non_trainable_variables</code>	
<code>non_trainable_weights</code>	
<code>outbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>output</code>	Retrieves the output tensor(s) of a layer.
<code>output_mask</code>	Retrieves the output mask tensor(s) of a layer.
<code>output_shape</code>	Retrieves the output shape(s) of a layer.

Continued on next page

Table 259 – continued from previous page

<code>run_eagerly</code>	Settable attribute indicating whether the model should run eagerly.
<code>sample_weights</code>	
<code>state_updates</code>	Returns the <i>updates</i> from all layers that are stateful.
<code>stateful</code>	
<code>submodules</code>	Sequence of all sub-modules.
<code>trainable</code>	
<code>trainable_variables</code>	Sequence of variables owned by this module and it's submodules.
<code>trainable_weights</code>	
<code>updates</code>	
<code>variables</code>	Returns the list of all layer variables/weights.
<code>weights</code>	Returns the list of all layer variables/weights.

`__init__` (*hidden_units, encoding_dimension, output_shape*)

Instantiate the BaseDecoder.

Parameters

- **hidden_units** (*tuple of int*) – Number of units per hidden layer.
- **encoding_dimension** (*int*) – encoding dimension.
- **output_shape** (*int*) – output shape, usual equal to the input shape.

Returns *None*

`call` (*inputs, training=True*)

Execute the model on input data.

Parameters

- **inputs** (*tf.Tensor*) – Input tensors.
- **training** (*bool*) – Training flag.

Returns (*encoding, reconstruction*) – Pair of tensors.

class `ashpy.models.fc.autoencoders.BaseAutoencoder` (*hidden_units, encoding_dimension, output_shape*)

Bases: `tensorflow.python.keras.engine.training.Model`

Primitive Model for all fully connected autoencoders.

Examples

- Direct Usage:

```

autoencoder = BaseAutoencoder(
    hidden_units=[256, 128, 64],
    encoding_dimension=100,
    output_shape=55
)

encoding, reconstruction = autoencoder(tf.zeros((1, 55)))
print(encoding.shape)
print(reconstruction.shape)

```



`__init__` (*hidden_units, encoding_dimension, output_shape*)

Instantiate the `BaseDecoder`.

Parameters

- **hidden_units** (*tuple of int*) – Number of units per hidden layer.
- **encoding_dimension** (*int*) – encoding dimension.
- **output_shape** (*int*) – output shape, usual equal to the input shape.

Returns `None`

`call` (*inputs, training=True*)

Execute the model on input data.

Parameters

- **inputs** (`tf.Tensor`) – Input tensors.
- **training** (`bool`) – Training flag.

Returns (*encoding, reconstruction*) – Pair of tensors.

decoders

Collection of Decoders (i.e., GANs' Discriminators) models.

Classes

BaseDecoder

Primitive Model for all fully connected decoder based architecture.

BaseDecoder

Inheritance Diagram



class `ashpy.models.fc.decoders.BaseDecoder` (*hidden_units, output_shape*)

Bases: `ashpy.models.fc.interfaces.FCInterface`

Primitive Model for all fully connected decoder based architecture.

Examples

```
decoder = BaseDecoder(
    hidden_units=[64,128,256],
    output_shape=55)
print(decoder(tf.zeros((1,10))).shape)
```

```
(1, 55)
```

Methods

<code>__init__(hidden_units, output_shape)</code>	Instantiate the <i>BaseDecoder</i> .
---	--------------------------------------

Attributes

<code>activity_regularizer</code>	Optional regularizer function for the output of this layer.
<code>dtype</code>	
<code>dynamic</code>	
<code>inbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>input</code>	Retrieves the input tensor(s) of a layer.
<code>input_mask</code>	Retrieves the input mask tensor(s) of a layer.
<code>input_shape</code>	Retrieves the input shape(s) of a layer.
<code>input_spec</code>	Gets the network's input specs.
<code>layers</code>	
<code>losses</code>	Losses which are associated with this <i>Layer</i> .
<code>metrics</code>	Returns the model's metrics added using <i>compile</i> , <i>add_metric</i> APIs.
<code>metrics_names</code>	Returns the model's display labels for all outputs.
<code>name</code>	Returns the name of this module as passed or determined in the ctor.
<code>name_scope</code>	Returns a <i>tf.name_scope</i> instance for this class.
<code>non_trainable_variables</code>	
<code>non_trainable_weights</code>	
<code>outbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>output</code>	Retrieves the output tensor(s) of a layer.
<code>output_mask</code>	Retrieves the output mask tensor(s) of a layer.
<code>output_shape</code>	Retrieves the output shape(s) of a layer.
<code>run_eagerly</code>	Settable attribute indicating whether the model should run eagerly.
<code>sample_weights</code>	
<code>state_updates</code>	Returns the <i>updates</i> from all layers that are stateful.
<code>stateful</code>	
<code>submodules</code>	Sequence of all sub-modules.
<code>trainable</code>	
<code>trainable_variables</code>	Sequence of variables owned by this module and it's submodules.

Continued on next page

Table 262 – continued from previous page

<code>trainable_weights</code>	
<code>updates</code>	
<code>variables</code>	Returns the list of all layer variables/weights.
<code>weights</code>	Returns the list of all layer variables/weights.

`__init__` (*hidden_units*, *output_shape*)
 Instantiate the *BaseDecoder*.

Parameters

- **hidden_units** (tuple of int) – Number of units per hidden layer.
- **output_shape** (int) – Amount of units of the last `tf.keras.layers.Dense`.

Returns None

class `ashpy.models.fc.decoders.BaseDecoder` (*hidden_units*, *output_shape*)
 Bases: *ashpy.models.fc.interfaces.FCInterface*

Primitive Model for all fully connected decoder based architecture.

Examples

```
decoder = BaseDecoder(
    hidden_units=[64,128,256],
    output_shape=55)
print(decoder(tf.zeros((1,10))).shape)
```

```
(1, 55)
```

`__init__` (*hidden_units*, *output_shape*)
 Instantiate the *BaseDecoder*.

Parameters

- **hidden_units** (tuple of int) – Number of units per hidden layer.
- **output_shape** (int) – Amount of units of the last `tf.keras.layers.Dense`.

Returns None

encoders

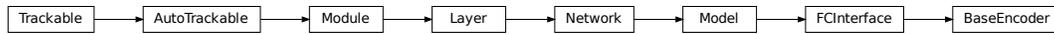
Collection of Encoders (i.e., GANs’ Discriminators) models.

Classes

<i>BaseEncoder</i>	Primitive Model for all fully connected encoder based architecture.
--------------------	---

BaseEncoder

Inheritance Diagram



class `ashpy.models.fc.encoders.BaseEncoder` (*hidden_units, output_shape*)

Bases: `ashpy.models.fc.interfaces.FCInterface`

Primitive Model for all fully connected encoder based architecture.

Examples

```

encoder = BaseEncoder(
    hidden_units=[256, 128, 64],
    output_shape=10)
print(encoder(tf.zeros((1, 55))).shape)
    
```

```
(1, 10)
```

Methods

<code>__init__(hidden_units, output_shape)</code>	Instantiate the BaseDecoder.
---	------------------------------

Attributes

<code>activity_regularizer</code>	Optional regularizer function for the output of this layer.
<code>dtype</code>	
<code>dynamic</code>	
<code>inbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>input</code>	Retrieves the input tensor(s) of a layer.
<code>input_mask</code>	Retrieves the input mask tensor(s) of a layer.
<code>input_shape</code>	Retrieves the input shape(s) of a layer.
<code>input_spec</code>	Gets the network's input specs.
<code>layers</code>	
<code>losses</code>	Losses which are associated with this <i>Layer</i> .
<code>metrics</code>	Returns the model's metrics added using <i>compile</i> , <i>add_metric</i> APIs.
<code>metrics_names</code>	Returns the model's display labels for all outputs.
<code>name</code>	Returns the name of this module as passed or determined in the ctor.
<code>name_scope</code>	Returns a <i>tf.name_scope</i> instance for this class.

Continued on next page

Table 265 – continued from previous page

non_trainable_variables	
non_trainable_weights	
outbound_nodes	Deprecated, do NOT use! Only for compatibility with external Keras.
output	Retrieves the output tensor(s) of a layer.
output_mask	Retrieves the output mask tensor(s) of a layer.
output_shape	Retrieves the output shape(s) of a layer.
run_eagerly	Settable attribute indicating whether the model should run eagerly.
sample_weights	
state_updates	Returns the <i>updates</i> from all layers that are stateful.
stateful	
submodules	Sequence of all sub-modules.
trainable	
trainable_variables	Sequence of variables owned by this module and it's submodules.
trainable_weights	
updates	
variables	Returns the list of all layer variables/weights.
weights	Returns the list of all layer variables/weights.

`__init__(hidden_units, output_shape)`

Instantiate the BaseDecoder.

Parameters

- **hidden_units** (tuple of int) – Number of units per hidden layer.
- **output_shape** (int) – Amount of units of the last `tf.keras.layers.Dense`.

Returns None

class `ashpy.models.fc.encoders.BaseEncoder` (*hidden_units, output_shape*)

Bases: `ashpy.models.fc.interfaces.FCInterface`

Primitive Model for all fully connected encoder based architecture.

Examples

```
encoder = BaseEncoder(
    hidden_units=[256, 128, 64],
    output_shape=10)
print(encoder(tf.zeros((1, 55))).shape)
```

```
(1, 10)
```

`__init__(hidden_units, output_shape)`

Instantiate the BaseDecoder.

Parameters

- **hidden_units** (tuple of int) – Number of units per hidden layer.
- **output_shape** (int) – Amount of units of the last `tf.keras.layers.Dense`.

Returns None

interfaces

Primitive Fully Connected interfaces.

Classes

FCInterface

Primitive Interface to be used by all *ashpy.models*.

FCInterface

Inheritance Diagram



class `ashpy.models.fc.interfaces.FCInterface`

Bases: `tensorflow.python.keras.engine.training.Model`

Primitive Interface to be used by all *ashpy.models*.

Methods

<code>__init__()</code>	Primitive Interface to be used by all <i>ashpy.models</i> .
<code>call(inputs[, training])</code>	Execute the model on input data.

Attributes

<code>activity_regularizer</code>	Optional regularizer function for the output of this layer.
<code>dtype</code>	
<code>dynamic</code>	
<code>inbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>input</code>	Retrieves the input tensor(s) of a layer.
<code>input_mask</code>	Retrieves the input mask tensor(s) of a layer.
<code>input_shape</code>	Retrieves the input shape(s) of a layer.
<code>input_spec</code>	Gets the network's input specs.
<code>layers</code>	
<code>losses</code>	Losses which are associated with this <i>Layer</i> .
<code>metrics</code>	Returns the model's metrics added using <i>compile</i> , <i>add_metric</i> APIs.

Continued on next page

Table 268 – continued from previous page

<code>metrics_names</code>	Returns the model's display labels for all outputs.
<code>name</code>	Returns the name of this module as passed or determined in the ctor.
<code>name_scope</code>	Returns a <i>tf.name_scope</i> instance for this class.
<code>non_trainable_variables</code>	
<code>non_trainable_weights</code>	
<code>outbound_nodes</code>	Deprecated, do NOT use! Only for compatibility with external Keras.
<code>output</code>	Retrieves the output tensor(s) of a layer.
<code>output_mask</code>	Retrieves the output mask tensor(s) of a layer.
<code>output_shape</code>	Retrieves the output shape(s) of a layer.
<code>run_eagerly</code>	Settable attribute indicating whether the model should run eagerly.
<code>sample_weights</code>	
<code>state_updates</code>	Returns the <i>updates</i> from all layers that are stateful.
<code>stateful</code>	
<code>submodules</code>	Sequence of all sub-modules.
<code>trainable</code>	
<code>trainable_variables</code>	Sequence of variables owned by this module and it's submodules.
<code>trainable_weights</code>	
<code>updates</code>	
<code>variables</code>	Returns the list of all layer variables/weights.
<code>weights</code>	Returns the list of all layer variables/weights.

`__init__()`
 Primitive Interface to be used by all *ashpy.models*.
 Declares the `self._layers` list.
Returns `None`

call (*inputs*, *training=True*)
 Execute the model on input data.

- Parameters**
- **inputs** (`tf.Tensor`) – Input tensors.
 - **training** (`bool`) – Training flag.

Returns `:py_class:'tf.Tensor'`

class `ashpy.models.fc.interfaces.FCInterface`
 Bases: `tensorflow.python.keras.engine.training.Model`

Primitive Interface to be used by all *ashpy.models*.

`__init__()`
 Primitive Interface to be used by all *ashpy.models*.
 Declares the `self._layers` list.
Returns `None`

call (*inputs*, *training=True*)
 Execute the model on input data.

Parameters

- **inputs** (`tf.Tensor`) – Input tensors.
- **training** (`bool`) – Training flag.

Returns :py_class:`tf.Tensor`

gans

GANs Models.

4.7.3 gans

GANs Models.

Generators

ConvGenerator

alias of `ashpy.models.convolutional.decoders.BaseDecoder`

DenseGenerator

alias of `ashpy.models.fc.decoders.BaseDecoder`

ConvGenerator

Inheritance Diagram



`ashpy.models.gans.ConvGenerator`

alias of `ashpy.models.convolutional.decoders.BaseDecoder`

DenseGenerator

Inheritance Diagram



`ashpy.models.gans.DenseGenerator`

alias of `ashpy.models.fc.decoders.BaseDecoder`

Discriminators

<i>ConvDiscriminator</i>	alias of <i>ashpy.models.convolutional.encoders.BaseEncoder</i>
<i>DenseDiscriminator</i>	alias of <i>ashpy.models.fc.encoders.BaseEncoder</i>

ConvDiscriminator

Inheritance Diagram



`ashpy.models.gans.ConvDiscriminator`
 alias of *ashpy.models.convolutional.encoders.BaseEncoder*

DenseDiscriminator

Inheritance Diagram



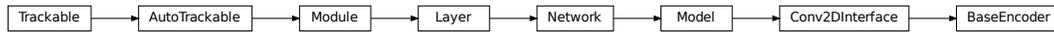
`ashpy.models.gans.DenseDiscriminator`
 alias of *ashpy.models.fc.encoders.BaseEncoder*

Encoders

<i>ConvEncoder</i>	alias of <i>ashpy.models.convolutional.encoders.BaseEncoder</i>
<i>DenseEncoder</i>	alias of <i>ashpy.models.fc.encoders.BaseEncoder</i>

ConvEncoder

Inheritance Diagram



`ashpy.models.gans.ConvEncoder`
alias of `ashpy.models.convolutional.encoders.BaseEncoder`

DenseEncoder

Inheritance Diagram



`ashpy.models.gans.DenseEncoder`
alias of `ashpy.models.fc.encoders.BaseEncoder`

4.8 ashpy.modes

Various modalities used to configure certain ash behaviours.

Classes

<code>LogEvalMode</code>	Mode to use when logging end evaluating a model.
--------------------------	--

4.8.1 LogEvalMode

Inheritance Diagram



class `ashpy.modes.LogEvalMode`

Bases: `enum.Enum`

Mode to use when logging end evaluating a model. Models often have the same behaviour (or very similar) when evaluated in trainable and non trainable setting.

There are some other models, see `pix2pix`, that instead require to be tested and trained in TRAIN mode (Model with `trainable=True`).

Attributes

TEST

TRAIN

class `ashpy.modes.LogEvalMode`

Bases: `enum.Enum`

Mode to use when logging end evaluating a model. Models often have the same behaviour (or very similar) when evaluated in trainable and non trainable setting.

There are some other models, see `pix2pix`, that instead require to be tested and trained in TRAIN mode (Model with `trainable=True`).

4.9 ashpy.trainers

Trainers help reducing boilerplate code by bootstrapping models training.

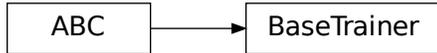
The module contains a primitive Interface and specific trainers that inherits from it.

Classes

<i>BaseTrainer</i>	<i>BaseTrainer</i> provide an interface for all trainers to inherit from.
<i>AdversarialTrainer</i>	Primitive Trainer for GANs subclassed from <i>ashpy.trainers.BaseTrainer</i> .
<i>EncoderTrainer</i>	Primitive Trainer for GANs using an Encoder sub-network.

4.9.1 BaseTrainer

Inheritance Diagram



```

class ashpy.trainers.BaseTrainer (epochs, logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1.3/docs/source/log',
                                  log_eval_mode=<LogEvalMode.TEST: 0>,
                                  global_step=<tf.Variable 'global_step:0' shape=() dtype=int64, numpy=0>, post_process_callback=None)
    
```

Bases: `abc.ABC`

`BaseTrainer` provide an interface for all trainers to inherit from.

Methods

<code>__init__</code> (epochs[, logdir, log_eval_mode, ...])	Primitive trainer interface.
<code>call</code> (dataset)	Execute the training process.

```

__init__(epochs, logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1.3/docs/source/log',
          log_eval_mode=<LogEvalMode.TEST: 0>, global_step=<tf.Variable 'global_step:0'
          shape=() dtype=int64, numpy=0>, post_process_callback=None)
    Primitive trainer interface. Handles model saving and restore.
    
```

Parameters

- **epochs** (*int*) – Number of training epochs.
- **logdir** (*str*) – Checkpoint and log directory.
- **log_eval_mode** – models' mode to use when evaluating and logging.
- **global_step** – `tf.Variable` that keeps track of the training steps.
- **post_process_callback** – the function to postprocess the model output, if needed.

```

__current_epoch()
    Get the current epoch using the (restored) variables.
    
```

Returns `current_epoch` (*int*)

```

__epoch_completed(epoch)
    Handle the end of the training epoch.
    
```

Parameters `epoch` (*int*) – the just completed training epoch.

```

__log(name, out)
    Log the out tensor using name as its name in tensorboard.
    
```

Parameters

- **name** – summary name.
- **out** – the tensor to log.

_reduce (*per_replica_tensor, reduce_op*)

Given the input tensor, reduces it in a distributed fashion, using the specified op.

_restore_or_init ()

Restores or initializes the persistence layer (checkpoint).

_save ()

Save the current checkpointable object status.

_update_global_batch_size (*dataset, executors=None*)

Given a dataset and the current distribution strategy sets the self._global_batch_size variable where needed.
:param dataset: a dataset from which the batch size will be extracted. :param executors: a list of executor with the property “global_batch_size” settable.

abstract call (*dataset*)

Execute the training process.

Iterate over the elements of a `tf.data.Dataset`. The dataset must contain everything needed to train the model.

Parameters dataset (`DatasetV2`) – A `tf.data.Dataset` to loop on to train the model.

4.9.2 AdversarialTrainer

Inheritance Diagram



```

class ashpy.trainers.AdversarialTrainer(generator, discriminator, generator_optimizer,
                                         discriminator_optimizer, generator_loss,
                                         discriminator_loss, epochs, metrics=None,
                                         logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkout',
                                         post_process_callback=None,
                                         log_eval_mode=<LogEvalMode.TEST: 0>,
                                         global_step=<tf.Variable 'global_step:0' shape=()
                                         dtype=int64, numpy=0>)
  
```

Bases: `ashpy.trainers.base_trainer.BaseTrainer`

Primitive Trainer for GANs subclassed from `ashpy.trainers.BaseTrainer`.

Examples

```
import shutil
import operator
from ashpy.models.gans import ConvGenerator, ConvDiscriminator
from ashpy.metrics import InceptionScore
from ashpy.losses.gan import DiscriminatorMinMax, GeneratorBCE

generator = ConvGenerator(
    layer_spec_input_res=(7, 7),
    layer_spec_target_res=(28, 28),
    kernel_size=(5, 5),
    initial_filters=32,
    filters_cap=16,
    channels=1,
)

discriminator = ConvDiscriminator(
    layer_spec_input_res=(28, 28),
    layer_spec_target_res=(7, 7),
    kernel_size=(5, 5),
    initial_filters=16,
    filters_cap=32,
    output_shape=1,
)

# Losses
generator_bce = GeneratorBCE()
minmax = DiscriminatorMinMax()

# Real data
batch_size = 2
mnist_x, mnist_y = tf.zeros((100,28,28)), tf.zeros((100,))

# Trainer
epochs = 2
logdir = "testlog/adversarial"
metrics = [
    InceptionScore(
        # Fake inception model
        ConvDiscriminator(
            layer_spec_input_res=(299, 299),
            layer_spec_target_res=(7, 7),
            kernel_size=(5, 5),
            initial_filters=16,
            filters_cap=32,
            output_shape=10,
        ),
        model_selection_operator=operator.gt,
        logdir=logdir,
    )
]
trainer = AdversarialTrainer(
    generator,
    discriminator,
    tf.optimizers.Adam(1e-4),
    tf.optimizers.Adam(1e-4),
    generator_bce,
    minmax,
```

(continues on next page)

(continued from previous page)

```

epochs,
metrics,
logdir,
)

# take only 2 samples to speed up tests
real_data = (
    tf.data.Dataset.from_tensor_slices(
        (tf.expand_dims(mnist_x, -1), tf.expand_dims(mnist_y, -1))).take(batch_size)
        .batch(batch_size)
        .prefetch(1)
)

# Add noise in the same dataset, just by mapping.
# The return type of the dataset must be: tuple(tuple(a,b), noise)
dataset = real_data.map(
    lambda x, y: ((x, y), tf.random.normal(shape=(batch_size, 100)))
)

trainer(dataset)
shutil.rmtree(logdir)

```

```

Initializing checkpoint.
[1] Saved checkpoint: testlog/adversarial/ckpts/ckpt-1
Epoch 1 completed.
[2] Saved checkpoint: testlog/adversarial/ckpts/ckpt-2
Epoch 2 completed.

```

Methods

<code>__init__(generator, discriminator, ...[, ...])</code>	Instantiate a <i>AdversarialTrainer</i> .
<code>call(dataset)</code>	Perform the adversarial training.
<code>train_step(real_xy, g_inputs)</code>	Train step for the AdversarialTrainer.

```

__init__(generator, discriminator, generator_optimizer, discrimina-
tor_optimizer, generator_loss, discriminator_loss, epochs, metrics=None,
logdir='home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1.3/docs/source/log',
post_process_callback=None, log_eval_mode=<LogEvalMode.TEST: 0>,
global_step=<tf.Variable 'global_step:0' shape=() dtype=int64, numpy=0>)
Instantiate a AdversarialTrainer.

```

Parameters

- **generator** (`tf.keras.Model`) – A `tf.keras.Model` describing the Generator part of a GAN.
- **discriminator** (`tf.keras.Model`) – A `tf.keras.Model` describing the Discriminator part of a GAN.
- **generator_optimizer** (`tf.optimizers.Optimizer`) – A `tf.optimizers` to use for the Generator.
- **discriminator_optimizer** (`tf.optimizers.Optimizer`) – A `tf.optimizers` to use for the Discriminator.

- **generator_loss** (*ashpy.losses.executor.Executor*) – A ash Executor to compute the loss of the Generator.
- **discriminator_loss** (*ashpy.losses.executor.Executor*) – A ash Executor to compute the loss of the Discriminator.
- **epochs** (*int*) – number of training epochs.
- **metrics** – (List): list of `tf.metrics` to measure on training and validation data
- **logdir** – checkpoint and log directory.
- **post_process_callback** (*callable*) – the function to postprocess the model output, if needed
- **log_eval_mode** – models' mode to use when evaluating and logging.
- **global_step** – `tf.Variable` that keeps track of the training steps.

Returns `None`

`_measure_performance` (*dataset*)
Measure performance on dataset.

Parameters `dataset` (`tf.data.Dataset`) –

`_train_step`
Training step with the distribution strategy.

`call` (*dataset*)
Perform the adversarial training.

Parameters `dataset` (`tf.data.Dataset`) – The adversarial training dataset.

`train_step` (*real_xy, g_inputs*)
Train step for the AdversarialTrainer.

Parameters

- **real_xy** – input batch as extracted from the input dataset. (features, label) pair.
- **g_inputs** – batch of generator_input as generated from the input dataset.

Returns

d_loss, g_loss, fake –

discriminator, generator loss values. fake is the generator output.

4.9.3 EncoderTrainer

Inheritance Diagram



```

class ashpy.trainers.EncoderTrainer(generator, discriminator, encoder, genera-
tor_optimizer, discriminator_optimizer, en-
coder_optimizer, generator_loss, discrimi-
nator_loss, encoder_loss, epochs, metrics,
logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1
post_process_callback=None,
log_eval_mode=<LogEvalMode.TEST: 0>,
global_step=<tf.Variable 'global_step:0' shape=()
dtype=int64, numpy=0>)

```

Bases: *ashpy.trainers.gan.AdversarialTrainer*

Primitive Trainer for GANs using an Encoder sub-network. The implementation is thought to be used with the BCE losses. To use another loss function consider subclassing the model and overriding the `train_step` method.

Examples

```

import shutil
import operator
from ashpy.metrics import EncodingAccuracy
from ashpy.losses.gan import DiscriminatorMinMax, GeneratorBCE, EncoderBCE

def real_gen():
    label = 0
    for _ in tf.range(100):
        yield ((10.0,), (label,))

latent_dim = 100

generator = tf.keras.Sequential([tf.keras.layers.Dense(1)])

left_input = tf.keras.layers.Input(shape=(1,))
left = tf.keras.layers.Dense(10, activation=tf.nn.elu)(left_input)

right_input = tf.keras.layers.Input(shape=(latent_dim,))
right = tf.keras.layers.Dense(10, activation=tf.nn.elu)(right_input)

net = tf.keras.layers.Concatenate()([left, right])
out = tf.keras.layers.Dense(1)(net)

discriminator = tf.keras.Model(inputs=[left_input, right_input], outputs=[out])

encoder = tf.keras.Sequential([tf.keras.layers.Dense(latent_dim)])

# Losses
generator_bce = GeneratorBCE()
encoder_bce = EncoderBCE()
minmax = DiscriminatorMinMax()

epochs = 2

# Fake pre-trained classifier
num_classes = 1
classifier = tf.keras.Sequential(
    [tf.keras.layers.Dense(10), tf.keras.layers.Dense(num_classes)]
)

```

(continues on next page)

(continued from previous page)

```

logdir = "testlog/adversarial/encoder"

metrics = [
    EncodingAccuracy(
        classifier,
        # model_selection_operator=operator.gt,
        logdir=logdir
    )
]

trainer = EncoderTrainer(
    generator=generator,
    discriminator=discriminator,
    encoder=encoder,
    discriminator_optimizer=tf.optimizers.Adam(1e-4),
    generator_optimizer=tf.optimizers.Adam(1e-5),
    encoder_optimizer=tf.optimizers.Adam(1e-6),
    generator_loss=generator_bce,
    discriminator_loss=minmax,
    encoder_loss=encoder_bce,
    epochs=epochs,
    metrics=metrics,
    logdir=logdir,
)

batch_size = 10
discriminator_input = tf.data.Dataset.from_generator(
    real_gen, (tf.float32, tf.int64), ((1), (1))
).batch(batch_size)

dataset = discriminator_input.map(
    lambda x, y: ((x, y), tf.random.normal(shape=(batch_size, latent_dim)))
)

trainer(dataset)

shutil.rmtree(logdir)

```

```

Initializing checkpoint.
[10] Saved checkpoint: testlog/adversarial/encoder/ckpts/ckpt-1
Epoch 1 completed.
[20] Saved checkpoint: testlog/adversarial/encoder/ckpts/ckpt-2
Epoch 2 completed.

```

Methods

<code>__init__(generator, discriminator, encoder, ...)</code>	Instantiate a <i>EncoderTrainer</i> .
<code>call(dataset)</code>	Perform the adversarial training.
<code>train_step(real_xy, g_inputs)</code>	Adversarial training step.

```
__init__(generator, discriminator, encoder, generator_optimizer, discriminator_optimizer, encoder_optimizer, generator_loss, discriminator_loss, encoder_loss, epochs, metrics, logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1.3/docs/source/log', post_process_callback=None, log_eval_mode=<LogEvalMode.TEST: 0>, global_step=<tf.Variable 'global_step:0' shape=() dtype=int64, numpy=0>)
```

Instantiate a `EncoderTrainer`.

Parameters

- **generator** (`tf.keras.Model`) – A `tf.keras.Model` describing the Generator part of a GAN.
- **discriminator** (`tf.keras.Model`) – A `tf.keras.Model` describing the Discriminator part of a GAN.
- **encoder** (`tf.keras.Model`) – A `tf.keras.Model` describing the Encoder part of a GAN.
- **generator_optimizer** (`tf.optimizers.Optimizer`) – A `tf.optimizers` to use for the Generator.
- **discriminator_optimizer** (`tf.optimizers.Optimizer`) – A `tf.optimizers` to use for the Discriminator.
- **encoder_optimizer** (`tf.optimizers.Optimizer`) – A `tf.optimizers` to use for the Encoder.
- **generator_loss** (`ashpy.losses.executor.Executor`) – A `ash` `Executor` to compute the loss of the Generator.
- **discriminator_loss** (`ashpy.losses.executor.Executor`) – A `ash` `Executor` to compute the loss of the Discriminator.
- **encoder_loss** (`ashpy.losses.executor.Executor`) – A `ash` `Executor` to compute the loss of the Discriminator.
- **epochs** (`int`) – number of training epochs.
- **metrics** – (List): list of `tf.metrics` to measure on training and validation data.
- **logdir** – checkpoint and log directory.
- **post_process_callback** (`callable`) – a function to post-process the output.
- **log_eval_mode** – models' mode to use when evaluating and logging.
- **global_step** – `tf.Variable` that keeps track of the training steps.

_measure_performance (`dataset`)
Measure performance on dataset.

Parameters `dataset` (`tf.data.Dataset`) –

_train_step
The training step that uses the distribution strategy.

call (`dataset`)
Perform the adversarial training.

Parameters `dataset` (`tf.data.Dataset`) – The adversarial training dataset.

train_step (`real_xy, g_inputs`)
Adversarial training step.

Parameters

- **real_xy** – input batch as extracted from the discriminator input dataset. (features, label) pair
 - **g_inputs** – batch of noise as generated by the generator input dataset.
- Returns** *d_loss, g_loss, e_loss* – discriminator, generator, encoder loss values.
-

Modules

<i>base_trainer</i>	Primitive Trainer Interface.
<i>classifier</i>	Primitive Trainer Interface.
<i>gan</i>	Collection of GANs trainers.

4.9.4 base_trainer

Primitive Trainer Interface.

Classes

<i>BaseTrainer</i>	<i>BaseTrainer</i> provide an interface for all trainers to inherit from.
--------------------	---

BaseTrainer

Inheritance Diagram



```

class ashpy.trainers.base_trainer.BaseTrainer (epochs, logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/0.1.3/docs',
log_eval_mode=<LogEvalMode.TEST:
0>, global_step=<tf.Variable
'global_step:0' shape=()
dtype=int64, numpy=0>,
post_process_callback=None)
  
```

Bases: `abc.ABC`

BaseTrainer provide an interface for all trainers to inherit from.

Methods

<code>__init__(epochs[, logdir, log_eval_mode, ...])</code>	Primitive trainer interface.
<code>call(dataset)</code>	Execute the training process.

`__init__` (*epochs*, *logdir*='~/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1.3/docs/source/log', *log_eval_mode*=<LogEvalMode.TEST: 0>, *global_step*=<tf.Variable 'global_step:0' shape=() dtype=int64, numpy=0>, *post_process_callback*=None)
 Primitive trainer interface. Handles model saving and restore.

Parameters

- **epochs** (*int*) – Number of training epochs.
- **logdir** (*str*) – Checkpoint and log directory.
- **log_eval_mode** – models' mode to use when evaluating and logging.
- **global_step** – tf.Variable that keeps track of the training steps.
- **post_process_callback** – the function to postprocess the model output, if needed.

`__current_epoch` ()
 Get the current epoch using the (restored) variables.

Returns `current_epoch` (*int*)

`__epoch_completed` (*epoch*)
 Handle the end of the training epoch.

Parameters `epoch` (*int*) – the just completed training epoch.

`__log` (*name*, *out*)
 Log the out tensor using name as its name in tensorboard.

Parameters

- **name** – summary name.
- **out** – the tensor to log.

`__reduce` (*per_replica_tensor*, *reduce_op*)
 Given the input tensor, reduces it in a distributed fashion, using the specified op.

`__restore_or_init` ()
 Restores or initializes the persistence layer (checkpoint).

`__save` ()
 Save the current checkpointable object status.

`__update_global_batch_size` (*dataset*, *executors*=None)
 Given a dataset and the current distribution strategy sets the self_ `global_batch_size` variable where needed.
 :param dataset: a dataset from wich the batch size will be extracted. :param executors: a list of executor with the property “`global_batch_size`” settable.

abstract call (*dataset*)
 Execute the training process.

Iterate over the elements of a `tf.data.Dataset`. The dataset must contain everything needed to train the model.

Parameters `dataset` (`DatasetV2`) – A `tf.data.Dataset` to loop on to train the model.

```
class ashpy.trainers.base_trainer.BaseTrainer (epochs, logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1.3/docs/source/log',
log_eval_mode=<LogEvalMode.TEST: 0>, global_step=<tf.Variable 'global_step:0'
dtype=int64, numpy=0>, post_process_callback=None)
shape=()
```

Bases: `abc.ABC`

`BaseTrainer` provide an interface for all trainers to inherit from.

```
__init__ (epochs, logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1.3/docs/source/log',
log_eval_mode=<LogEvalMode.TEST: 0>, global_step=<tf.Variable 'global_step:0'
dtype=int64, numpy=0>, post_process_callback=None)
Primitive trainer interface. Handles model saving and restore.
```

Parameters

- **epochs** (*int*) – Number of training epochs.
- **logdir** (*str*) – Checkpoint and log directory.
- **log_eval_mode** – models' mode to use when evaluating and logging.
- **global_step** – `tf.Variable` that keeps track of the training steps.
- **post_process_callback** – the function to postprocess the model output, if needed.

```
_current_epoch ()
Get the current epoch using the (restored) variables.
```

Returns `current_epoch` (*int*)

```
_epoch_completed (epoch)
Handle the end of the training epoch.
```

Parameters **epoch** (*int*) – the just completed training epoch.

```
_log (name, out)
Log the out tensor using name as its name in tensorboard.
```

Parameters

- **name** – summary name.
- **out** – the tensor to log.

```
_reduce (per_replica_tensor, reduce_op)
Given the input tensor, reduces it in a distributed fashion, using the specified op.
```

```
_restore_or_init ()
Restores or initializes the persistence layer (checkpoint).
```

```
_save ()
Save the current checkpointable object status.
```

```
_update_global_batch_size (dataset, executors=None)
Given a dataset and the current distribution strategy sets the self._global_batch_size variable where needed.
:param dataset: a dataset from wich the batch size will be extracted.
:param executors: a list of executor with the property “global_batch_size” settable.
```

```
abstract call (dataset)
Execute the training process.
```

Iterate over the elements of a `tf.data.Dataset`. The dataset must contain everything needed to train the model.

Parameters `dataset` (`DatasetV2`) – A `tf.data.Dataset` to loop on to train the model.

4.9.5 classifier

Primitive Trainer Interface.

Classes

<code>ClassifierTrainer</code>	<code>ClassifierTrainer</code> provide the standard training loop for a classifier.
--------------------------------	---

ClassifierTrainer

Inheritance Diagram



```

class ashpy.trainers.classifier.ClassifierTrainer(model, optimizer, loss, epochs, metrics=None, logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/ashpy', global_step=<tf.Variable 'global_step:0' shape=() dtype=int64, numpy=0>, post_process_callback=None)
  
```

Bases: `ashpy.trainers.base_trainer.BaseTrainer`

`ClassifierTrainer` provide the standard training loop for a classifier.

Methods

<code>__init__(model, optimizer, loss, epochs[, ...])</code>	Instantiate the <code>ClassifierTrainer</code> trainer.
<code>call(train_set, validation_set)</code>	Start the training.
<code>train_step(features, labels)</code>	Train step.

```

__init__(model, optimizer, loss, epochs, metrics=None, logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/ashpy', global_step=<tf.Variable 'global_step:0' shape=() dtype=int64, numpy=0>, post_process_callback=None)
  
```

Instantiate the `ClassifierTrainer` trainer.

Parameters

- `model` (`tf.keras.Model`) – A `tf.keras.Model` model.

- **optimizer** (`tf.optimizers.Optimizer`) – A `tf.optimizers.Optimizer`.
- **loss** (`callable`) – A loss function built following `tf.losses`.
- **epochs** (`int`) – Number of training epochs.
- **metrics** – (List): List of python objects (dictionaries or `tf.metrics` objects) to measure on training and validation data.
- **logdir** (`str`) – Checkpoint and log directory.
- **global_step** – `tf.Variable` that keeps track of the training steps.
- **post_process_callback** (`callable`) – the function to postprocess the model output, if needed.

Examples

```
def toy_dataset():
    inputs = tf.expand_dims(tf.range(1, 1000.0), -1)
    labels = tf.expand_dims(
        [1 if tf.equal(tf.math.mod(tf.squeeze(i), 2), 0) else 0 for i in_
↪inputs], -1
    )
    return tf.data.Dataset.from_tensor_slices((inputs, labels)).shuffle(10).
↪batch(2)

model = tf.keras.Sequential(
    [tf.keras.layers.Dense(10, activation=tf.nn.sigmoid), tf.keras.layers.
↪Dense(2)]
)
optimizer = tf.optimizers.Adam(1e-3)

loss = ClassifierLoss(tf.losses.SparseCategoricalCrossentropy(from_
↪logits=True))
logdir = "testlog"
epochs = 2

metrics = [
    ClassifierMetric(tf.metrics.Accuracy()),
    ClassifierMetric(tf.metrics.BinaryAccuracy()),
]

trainer = ClassifierTrainer(model, optimizer, loss, epochs, metrics,
↪logdir=logdir)
train, validation = toy_dataset(), toy_dataset()
trainer.train(validation)
shutil.rmtree(logdir)
```

```
Initializing checkpoint.
[500] Saved checkpoint: testlog/ckpts/ckpt-1
Epoch 1 completed.
[1000] Saved checkpoint: testlog/ckpts/ckpt-2
Epoch 2 completed.
```

`_measure_performance` (`dataset`)
 Measure and log metrics on the dataset.

_train_step

The training step that uses the distribution strategy.

call (*train_set*, *validation_set*)

Start the training.

Parameters

- **train_set** (*tf.data.Dataset*) – Training dataset.
- **validation_set** (*tf.data.Dataset*) – Validation dataset.

train_step (*features*, *labels*)

Train step.

Parameters

- **features** – Input features.
- **labels** – The labels.

Returns Loss value.

```
class ashpy.trainers.classifier.ClassifierTrainer (model, optimizer, loss,
                                                epochs, metrics=None,
                                                logdir='/home/docs/checkouts/readthedocs.org/user_builds/
                                                global_step=<tf.Variable
                                                'global_step:0' shape=()
                                                dtype=int64, numpy=0>,
                                                post_process_callback=None)
```

Bases: *ashpy.trainers.base_trainer.BaseTrainer*

ClassifierTrainer provide the standard training loop for a classifier.

```
__init__ (model, optimizer, loss, epochs, metrics=None, logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/ch
                                                global_step=<tf.Variable 'global_step:0' shape=() dtype=int64, numpy=0>,
                                                post_process_callback=None)
```

Instantiate the *ClassifierTrainer* trainer.

Parameters

- **model** (*tf.keras.Model*) – A *tf.keras.Model* model.
- **optimizer** (*tf.optimizers.Optimizer*) – A *tf.optimizers.Optimizer*.
- **loss** (*callable*) – A loss function built following *tf.losses*.
- **epochs** (*int*) – Number of training epochs.
- **metrics** – (List): List of python objects (dictionaries or *tf.metrics* objects) to measure on training and validation data.
- **logdir** (*str*) – Checkpoint and log directory.
- **global_step** – *tf.Variable* that keeps track of the training steps.
- **post_process_callback** (*callable*) – the function to postprocess the model output, if needed.

Examples

```

def toy_dataset():
    inputs = tf.expand_dims(tf.range(1, 1000.0), -1)
    labels = tf.expand_dims(
        [1 if tf.equal(tf.math.mod(tf.squeeze(i), 2), 0) else 0 for i in_
↪inputs], -1
    )
    return tf.data.Dataset.from_tensor_slices((inputs, labels)).shuffle(10).
↪batch(2)

model = tf.keras.Sequential(
    [tf.keras.layers.Dense(10, activation=tf.nn.sigmoid), tf.keras.layers.
↪Dense(2)]
)
optimizer = tf.optimizers.Adam(1e-3)

loss = ClassifierLoss(tf.losses.SparseCategoricalCrossentropy(from_
↪logits=True))
logdir = "testlog"
epochs = 2

metrics = [
    ClassifierMetric(tf.metrics.Accuracy()),
    ClassifierMetric(tf.metrics.BinaryAccuracy()),
]

trainer = ClassifierTrainer(model, optimizer, loss, epochs, metrics,
↪logdir=logdir)
train, validation = toy_dataset(), toy_dataset()
trainer(train, validation)
shutil.rmtree(logdir)

```

```

Initializing checkpoint.
[500] Saved checkpoint: testlog/ckpts/ckpt-1
Epoch 1 completed.
[1000] Saved checkpoint: testlog/ckpts/ckpt-2
Epoch 2 completed.

```

`_measure_performance` (*dataset*)

Measure and log metrics on the dataset.

`_train_step`

The training step that uses the distribution strategy.

`call` (*train_set, validation_set*)

Start the training.

Parameters

- **`train_set`** (*tf.data.Dataset*) – Training dataset.
- **`validation_set`** (*tf.data.Dataset*) – Validation dataset.

`train_step` (*features, labels*)

Train step.

Parameters

- **`features`** – Input features.
- **`labels`** – The labels.

Returns Loss value.

4.9.6 gan

Collection of GANs trainers.

Classes

<i>AdversarialTrainer</i>	Primitive Trainer for GANs subclassed from <i>ashpy.trainers.BaseTrainer</i> .
<i>EncoderTrainer</i>	Primitive Trainer for GANs using an Encoder sub-network.

AdversarialTrainer

Inheritance Diagram



```

class ashpy.trainers.gan.AdversarialTrainer(generator, discriminator, generator_optimizer, discriminator_optimizer, generator_loss, discriminator_loss, epochs, metrics=None, logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/ch... post_process_callback=None, log_eval_mode=<LogEvalMode.TEST: 0>, global_step=<tf.Variable 'global_step:0' shape=() dtype=int64, numpy=0>)
  
```

Bases: *ashpy.trainers.base_trainer.BaseTrainer*

Primitive Trainer for GANs subclassed from *ashpy.trainers.BaseTrainer*.

Examples

```

import shutil
import operator
from ashpy.models.gans import ConvGenerator, ConvDiscriminator
from ashpy.metrics import InceptionScore
from ashpy.losses.gan import DiscriminatorMinMax, GeneratorBCE

generator = ConvGenerator(
    layer_spec_input_res=(7, 7),
  
```

(continues on next page)

(continued from previous page)

```

        layer_spec_target_res=(28, 28),
        kernel_size=(5, 5),
        initial_filters=32,
        filters_cap=16,
        channels=1,
    )

discriminator = ConvDiscriminator(
    layer_spec_input_res=(28, 28),
    layer_spec_target_res=(7, 7),
    kernel_size=(5, 5),
    initial_filters=16,
    filters_cap=32,
    output_shape=1,
)

# Losses
generator_bce = GeneratorBCE()
minmax = DiscriminatorMinMax()

# Real data
batch_size = 2
mnist_x, mnist_y = tf.zeros((100,28,28)), tf.zeros((100,))

# Trainer
epochs = 2
logdir = "testlog/adversarial"
metrics = [
    InceptionScore(
        # Fake inception model
        ConvDiscriminator(
            layer_spec_input_res=(299, 299),
            layer_spec_target_res=(7, 7),
            kernel_size=(5, 5),
            initial_filters=16,
            filters_cap=32,
            output_shape=10,
        ),
        model_selection_operator=operator.gt,
        logdir=logdir,
    )
]
trainer = AdversarialTrainer(
    generator,
    discriminator,
    tf.optimizers.Adam(1e-4),
    tf.optimizers.Adam(1e-4),
    generator_bce,
    minmax,
    epochs,
    metrics,
    logdir,
)

# take only 2 samples to speed up tests
real_data = (
    tf.data.Dataset.from_tensor_slices(

```

(continues on next page)

(continued from previous page)

```

(tf.expand_dims(mnist_x, -1), tf.expand_dims(mnist_y, -1)).take(batch_size)
.batch(batch_size)
.prefetch(1)
)

# Add noise in the same dataset, just by mapping.
# The return type of the dataset must be: tuple(tuple(a,b), noise)
dataset = real_data.map(
    lambda x, y: ((x, y), tf.random.normal(shape=(batch_size, 100)))
)

trainer(dataset)
shutil.rmtree(logdir)

```

```

Initializing checkpoint.
[1] Saved checkpoint: testlog/adversarial/ckpts/ckpt-1
Epoch 1 completed.
[2] Saved checkpoint: testlog/adversarial/ckpts/ckpt-2
Epoch 2 completed.

```

Methods

<code>__init__(generator, discriminator, ...[, ...])</code>	Instantiate a <i>AdversarialTrainer</i> .
<code>call(dataset)</code>	Perform the adversarial training.
<code>train_step(real_xy, g_inputs)</code>	Train step for the <i>AdversarialTrainer</i> .

```

__init__(generator, discriminator, generator_optimizer, discrimina-
tor_optimizer, generator_loss, discriminator_loss, epochs, metrics=None,
logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1.3/docs/source/log',
post_process_callback=None, log_eval_mode=<LogEvalMode.TEST: 0>,
global_step=<tf.Variable 'global_step:0' shape=() dtype=int64, numpy=0>)
Instantiate a AdversarialTrainer.

```

Parameters

- **generator** (`tf.keras.Model`) – A `tf.keras.Model` describing the Generator part of a GAN.
- **discriminator** (`tf.keras.Model`) – A `tf.keras.Model` describing the Discriminator part of a GAN.
- **generator_optimizer** (`tf.optimizers.Optimizer`) – A `tf.optimizers` to use for the Generator.
- **discriminator_optimizer** (`tf.optimizers.Optimizer`) – A `tf.optimizers` to use for the Discriminator.
- **generator_loss** (`ashpy.losses.executor.Executor`) – A `ash Executor` to compute the loss of the Generator.
- **discriminator_loss** (`ashpy.losses.executor.Executor`) – A `ash Executor` to compute the loss of the Discriminator.
- **epochs** (`int`) – number of training epochs.
- **metrics** – (List): list of `tf.metrics` to measure on training and validation data

- **logdir** – checkpoint and log directory.
- **post_process_callback** (*callable*) – the function to postprocess the model output, if needed
- **log_eval_mode** – models' mode to use when evaluating and logging.
- **global_step** – `tf.Variable` that keeps track of the training steps.

Returns *None*

_measure_performance (*dataset*)
Measure performance on dataset.

Parameters *dataset* (`tf.data.Dataset`) –

_train_step
Training step with the distribution strategy.

call (*dataset*)
Perform the adversarial training.

Parameters *dataset* (`tf.data.Dataset`) – The adversarial training dataset.

train_step (*real_xy, g_inputs*)
Train step for the AdversarialTrainer.

Parameters

- **real_xy** – input batch as extracted from the input dataset. (features, label) pair.
- **g_inputs** – batch of generator_input as generated from the input dataset.

Returns

d_loss, g_loss, fake –

discriminator, generator loss values. fake is the generator output.

EncoderTrainer

Inheritance Diagram



```

class ashpy.trainers.gan.EncoderTrainer(generator, discriminator, encoder, generator_optimizer, discriminator_optimizer, encoder_optimizer, generator_loss, discriminator_loss, encoder_loss, epochs, metrics, logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkout', post_process_callback=None, log_eval_mode=<LogEvalMode.TEST: 0>, global_step=<tf.Variable 'global_step:0' shape=() dtype=int64, numpy=0>)
  
```

Bases: `ashpy.trainers.gan.AdversarialTrainer`

Primitive Trainer for GANs using an Encoder sub-network. The implementation is thought to be used with the BCE losses. To use another loss function consider subclassing the model and overriding the `train_step` method.

Examples

```
import shutil
import operator
from ashpy.metrics import EncodingAccuracy
from ashpy.losses.gan import DiscriminatorMinMax, GeneratorBCE, EncoderBCE

def real_gen():
    label = 0
    for _ in tf.range(100):
        yield ((10.0,), (label,))

latent_dim = 100

generator = tf.keras.Sequential([tf.keras.layers.Dense(1)])

left_input = tf.keras.layers.Input(shape=(1,))
left = tf.keras.layers.Dense(10, activation=tf.nn.elu)(left_input)

right_input = tf.keras.layers.Input(shape=(latent_dim,))
right = tf.keras.layers.Dense(10, activation=tf.nn.elu)(right_input)

net = tf.keras.layers.Concatenate()([left, right])
out = tf.keras.layers.Dense(1)(net)

discriminator = tf.keras.Model(inputs=[left_input, right_input], outputs=[out])

encoder = tf.keras.Sequential([tf.keras.layers.Dense(latent_dim)])

# Losses
generator_bce = GeneratorBCE()
encoder_bce = EncoderBCE()
minmax = DiscriminatorMinMax()

epochs = 2

# Fake pre-trained classifier
num_classes = 1
classifier = tf.keras.Sequential(
    [tf.keras.layers.Dense(10), tf.keras.layers.Dense(num_classes)]
)

logdir = "testlog/adversarial/encoder"

metrics = [
    EncodingAccuracy(
        classifier,
        # model_selection_operator=operator.gt,
        logdir=logdir
    )
]
```

(continues on next page)

(continued from previous page)

```

trainer = EncoderTrainer(
    generator=generator,
    discriminator=discriminator,
    encoder=encoder,
    discriminator_optimizer=tf.optimizers.Adam(1e-4),
    generator_optimizer=tf.optimizers.Adam(1e-5),
    encoder_optimizer=tf.optimizers.Adam(1e-6),
    generator_loss=generator_bce,
    discriminator_loss=minmax,
    encoder_loss=encoder_bce,
    epochs=epochs,
    metrics=metrics,
    logdir=logdir,
)

batch_size = 10
discriminator_input = tf.data.Dataset.from_generator(
    real_gen, (tf.float32, tf.int64), ((1), (1))
).batch(batch_size)

dataset = discriminator_input.map(
    lambda x, y: ((x, y), tf.random.normal(shape=(batch_size, latent_dim)))
)

trainer(dataset)

shutil.rmtree(logdir)

```

```

Initializing checkpoint.
[10] Saved checkpoint: testlog/adversarial/encoder/ckpts/ckpt-1
Epoch 1 completed.
[20] Saved checkpoint: testlog/adversarial/encoder/ckpts/ckpt-2
Epoch 2 completed.

```

Methods

<code>__init__(generator, discriminator, encoder, ...)</code>	Instantiate a <i>EncoderTrainer</i> .
<code>call(dataset)</code>	Perform the adversarial training.
<code>train_step(real_xy, g_inputs)</code>	Adversarial training step.

`__init__(generator, discriminator, encoder, generator_optimizer, discriminator_optimizer, encoder_optimizer, generator_loss, discriminator_loss, encoder_loss, epochs, metrics, logdir='~/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1.3/docs/source/log', post_process_callback=None, log_eval_mode=<LogEvalMode.TEST: 0>, global_step=<tf.Variable 'global_step:0' shape=() dtype=int64, numpy=0>)`
 Instantiate a *EncoderTrainer*.

Parameters

- **generator** (`tf.keras.Model`) – A `tf.keras.Model` describing the Generator part of a GAN.
- **discriminator** (`tf.keras.Model`) – A `tf.keras.Model` describing the Discriminator part of a GAN.

- **encoder** (`tf.keras.Model`) – A `tf.keras.Model` describing the Encoder part of a GAN.
- **generator_optimizer** (`tf.optimizers.Optimizer`) – A `tf.optimizers` to use for the Generator.
- **discriminator_optimizer** (`tf.optimizers.Optimizer`) – A `tf.optimizers` to use for the Discriminator.
- **encoder_optimizer** (`tf.optimizers.Optimizer`) – A `tf.optimizers` to use for the Encoder.
- **generator_loss** (`ashpy.losses.executor.Executor`) – A `ash Executor` to compute the loss of the Generator.
- **discriminator_loss** (`ashpy.losses.executor.Executor`) – A `ash Executor` to compute the loss of the Discriminator.
- **encoder_loss** (`ashpy.losses.executor.Executor`) – A `ash Executor` to compute the loss of the Discriminator.
- **epochs** (`int`) – number of training epochs.
- **metrics** – (List): list of `tf.metrics` to measure on training and validation data.
- **logdir** – checkpoint and log directory.
- **post_process_callback** (`callable`) – a function to post-process the output.
- **log_eval_mode** – models' mode to use when evaluating and logging.
- **global_step** – `tf.Variable` that keeps track of the training steps.

_measure_performance (`dataset`)

Measure performance on dataset.

Parameters `dataset` (`tf.data.Dataset`) –

_train_step

The training step that uses the distribution strategy.

call (`dataset`)

Perform the adversarial training.

Parameters `dataset` (`tf.data.Dataset`) – The adversarial training dataset.

train_step (`real_xy`, `g_inputs`)

Adversarial training step.

Parameters

- **real_xy** – input batch as extracted from the discriminator input dataset. (features, label) pair
- **g_inputs** – batch of noise as generated by the generator input dataset.

Returns `d_loss`, `g_loss`, `e_loss` – discriminator, generator, encoder loss values.

```
class ashpy.trainers.gan.AdversarialTrainer(generator, discriminator, generator_optimizer, discriminator_optimizer, generator_loss, discriminator_loss, epochs, metrics=None, logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/ch...
post_process_callback=None, log_eval_mode=<LogEvalMode.TEST: 0>, global_step=<tf.Variable 'global_step:0' shape=() dtype=int64, numpy=0>)
```

Bases: *ashpy.trainers.base_trainer.BaseTrainer*

Primitive Trainer for GANs subclassed from *ashpy.trainers.BaseTrainer*.

Examples

```
import shutil
import operator
from ashpy.models.gans import ConvGenerator, ConvDiscriminator
from ashpy.metrics import InceptionScore
from ashpy.losses.gan import DiscriminatorMinMax, GeneratorBCE

generator = ConvGenerator(
    layer_spec_input_res=(7, 7),
    layer_spec_target_res=(28, 28),
    kernel_size=(5, 5),
    initial_filters=32,
    filters_cap=16,
    channels=1,
)

discriminator = ConvDiscriminator(
    layer_spec_input_res=(28, 28),
    layer_spec_target_res=(7, 7),
    kernel_size=(5, 5),
    initial_filters=16,
    filters_cap=32,
    output_shape=1,
)

# Losses
generator_bce = GeneratorBCE()
minmax = DiscriminatorMinMax()

# Real data
batch_size = 2
mnist_x, mnist_y = tf.zeros((100,28,28)), tf.zeros((100,))

# Trainer
epochs = 2
logdir = "testlog/adversarial"
metrics = [
    InceptionScore(
        # Fake inception model
        ConvDiscriminator(
            layer_spec_input_res=(299, 299),
            layer_spec_target_res=(7, 7),
```

(continues on next page)

(continued from previous page)

```

        kernel_size=(5, 5),
        initial_filters=16,
        filters_cap=32,
        output_shape=10,
    ),
    model_selection_operator=operator.gt,
    logdir=logdir,
)
]
trainer = AdversarialTrainer(
    generator,
    discriminator,
    tf.optimizers.Adam(1e-4),
    tf.optimizers.Adam(1e-4),
    generator_bce,
    minmax,
    epochs,
    metrics,
    logdir,
)

# take only 2 samples to speed up tests
real_data = (
    tf.data.Dataset.from_tensor_slices(
        (tf.expand_dims(mnist_x, -1), tf.expand_dims(mnist_y, -1))).take(batch_size)
    .batch(batch_size)
    .prefetch(1)
)

# Add noise in the same dataset, just by mapping.
# The return type of the dataset must be: tuple(tuple(a,b), noise)
dataset = real_data.map(
    lambda x, y: ((x, y), tf.random.normal(shape=(batch_size, 100)))
)

trainer(dataset)
shutil.rmtree(logdir)

```

```

Initializing checkpoint.
[1] Saved checkpoint: testlog/adversarial/ckpts/ckpt-1
Epoch 1 completed.
[2] Saved checkpoint: testlog/adversarial/ckpts/ckpt-2
Epoch 2 completed.

```

```

__init__(generator, discriminator, generator_optimizer, discrimina-
tor_optimizer, generator_loss, discriminator_loss, epochs, metrics=None,
logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1.3/docs/source/log',
post_process_callback=None, log_eval_mode=<LogEvalMode.TEST: 0>,
global_step=<tf.Variable 'global_step:0' shape=() dtype=int64, numpy=0>)
Instantiate a AdversarialTrainer.

```

Parameters

- **generator** (`tf.keras.Model`) – A `tf.keras.Model` describing the Generator part of a GAN.
- **discriminator** (`tf.keras.Model`) – A `tf.keras.Model` describing the Discriminator part of a GAN.

- **generator_optimizer** (`tf.optimizers.Optimizer`) – A `tf.optimizers` to use for the Generator.
- **discriminator_optimizer** (`tf.optimizers.Optimizer`) – A `tf.optimizers` to use for the Discriminator.
- **generator_loss** (`ashpy.losses.executor.Executor`) – A `ash` `Executor` to compute the loss of the Generator.
- **discriminator_loss** (`ashpy.losses.executor.Executor`) – A `ash` `Executor` to compute the loss of the Discriminator.
- **epochs** (`int`) – number of training epochs.
- **metrics** – (List): list of `tf.metrics` to measure on training and validation data
- **logdir** – checkpoint and log directory.
- **post_process_callback** (`callable`) – the function to postprocess the model output, if needed
- **log_eval_mode** – models' mode to use when evaluating and logging.
- **global_step** – `tf.Variable` that keeps track of the training steps.

Returns `None`

`_measure_performance` (`dataset`)
Measure performance on dataset.

Parameters `dataset` (`tf.data.Dataset`) –

`_train_step`
Training step with the distribution strategy.

`call` (`dataset`)
Perform the adversarial training.

Parameters `dataset` (`tf.data.Dataset`) – The adversarial training dataset.

`train_step` (`real_xy`, `g_inputs`)
Train step for the `AdversarialTrainer`.

Parameters

- **`real_xy`** – input batch as extracted from the input dataset. (features, label) pair.
- **`g_inputs`** – batch of generator_input as generated from the input dataset.

Returns

`d_loss`, `g_loss`, `fake` –

discriminator, generator loss values. fake is the generator output.

```
class ashpy.trainers.gan.EncoderTrainer(generator, discriminator, encoder, generator_optimizer, discriminator_optimizer, encoder_optimizer, generator_loss, discriminator_loss, encoder_loss, epochs, metrics, logdir='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts', post_process_callback=None, log_eval_mode=<LogEvalMode.TEST: 0>, global_step=<tf.Variable 'global_step:0' shape=() dtype=int64, numpy=0>)
```

Bases: `ashpy.trainers.gan.AdversarialTrainer`

Primitive Trainer for GANs using an Encoder sub-network. The implementation is thought to be used with the BCE losses. To use another loss function consider subclassing the model and overriding the `train_step` method.

Examples

```
import shutil
import operator
from ashpy.metrics import EncodingAccuracy
from ashpy.losses.gan import DiscriminatorMinMax, GeneratorBCE, EncoderBCE

def real_gen():
    label = 0
    for _ in tf.range(100):
        yield ((10.0,), (label,))

latent_dim = 100

generator = tf.keras.Sequential([tf.keras.layers.Dense(1)])

left_input = tf.keras.layers.Input(shape=(1,))
left = tf.keras.layers.Dense(10, activation=tf.nn.elu)(left_input)

right_input = tf.keras.layers.Input(shape=(latent_dim,))
right = tf.keras.layers.Dense(10, activation=tf.nn.elu)(right_input)

net = tf.keras.layers.Concatenate()([left, right])
out = tf.keras.layers.Dense(1)(net)

discriminator = tf.keras.Model(inputs=[left_input, right_input], outputs=[out])

encoder = tf.keras.Sequential([tf.keras.layers.Dense(latent_dim)])

# Losses
generator_bce = GeneratorBCE()
encoder_bce = EncoderBCE()
minmax = DiscriminatorMinMax()

epochs = 2

# Fake pre-trained classifier
num_classes = 1
classifier = tf.keras.Sequential(
    [tf.keras.layers.Dense(10), tf.keras.layers.Dense(num_classes)]
)

logdir = "testlog/adversarial/encoder"

metrics = [
    EncodingAccuracy(
        classifier,
        # model_selection_operator=operator.gt,
        logdir=logdir
    )
]

trainer = EncoderTrainer(
```

(continues on next page)

(continued from previous page)

```

generator=generator,
discriminator=discriminator,
encoder=encoder,
discriminator_optimizer=tf.optimizers.Adam(1e-4),
generator_optimizer=tf.optimizers.Adam(1e-5),
encoder_optimizer=tf.optimizers.Adam(1e-6),
generator_loss=generator_bce,
discriminator_loss=minmax,
encoder_loss=encoder_bce,
epochs=epochs,
metrics=metrics,
logdir=logdir,
)

batch_size = 10
discriminator_input = tf.data.Dataset.from_generator(
    real_gen, (tf.float32, tf.int64), ((1), (1))
).batch(batch_size)

dataset = discriminator_input.map(
    lambda x, y: ((x, y), tf.random.normal(shape=(batch_size, latent_dim)))
)

trainer(dataset)

shutil.rmtree(logdir)

```

```

Initializing checkpoint.
[10] Saved checkpoint: testlog/adversarial/encoder/ckpts/ckpt-1
Epoch 1 completed.
[20] Saved checkpoint: testlog/adversarial/encoder/ckpts/ckpt-2
Epoch 2 completed.

```

`__init__`(*generator*, *discriminator*, *encoder*, *generator_optimizer*, *discriminator_optimizer*, *encoder_optimizer*, *generator_loss*, *discriminator_loss*, *encoder_loss*, *epochs*, *metrics*, *logdir*='/home/docs/checkouts/readthedocs.org/user_builds/ashpy/checkouts/v0.1.3/docs/source/log', *post_process_callback*=None, *log_eval_mode*=<LogEvalMode.TEST: 0>, *global_step*=<tf.Variable 'global_step:0' shape=() dtype=int64, numpy=0>)

Instantiate a `EncoderTrainer`.

Parameters

- **generator** (`tf.keras.Model`) – A `tf.keras.Model` describing the Generator part of a GAN.
- **discriminator** (`tf.keras.Model`) – A `tf.keras.Model` describing the Discriminator part of a GAN.
- **encoder** (`tf.keras.Model`) – A `tf.keras.Model` describing the Encoder part of a GAN.
- **generator_optimizer** (`tf.optimizers.Optimizer`) – A `tf.optimizers` to use for the Generator.
- **discriminator_optimizer** (`tf.optimizers.Optimizer`) – A `tf.optimizers` to use for the Discriminator.
- **encoder_optimizer** (`tf.optimizers.Optimizer`) – A `tf.optimizers` to use for the Encoder.

- **generator_loss** (*ashpy.losses.executor.Executor*) – A ash Executor to compute the loss of the Generator.
- **discriminator_loss** (*ashpy.losses.executor.Executor*) – A ash Executor to compute the loss of the Discriminator.
- **encoder_loss** (*ashpy.losses.executor.Executor*) – A ash Executor to compute the loss of the Discriminator.
- **epochs** (*int*) – number of training epochs.
- **metrics** – (List): list of tf.metrics to measure on training and validation data.
- **logdir** – checkpoint and log directory.
- **post_process_callback** (*callable*) – a function to post-process the output.
- **log_eval_mode** – models' mode to use when evaluating and logging.
- **global_step** – tf.Variable that keeps track of the training steps.

_measure_performance (*dataset*)

Measure performance on dataset.

Parameters **dataset** (*tf.data.Dataset*) –

_train_step

The training step that uses the distribution strategy.

call (*dataset*)

Perform the adversarial training.

Parameters **dataset** (*tf.data.Dataset*) – The adversarial training dataset.

train_step (*real_xy, g_inputs*)

Adversarial training step.

Parameters

- **real_xy** – input batch as extracted from the discriminator input dataset. (features, label) pair
- **g_inputs** – batch of noise as generated by the generator input dataset.

Returns *d_loss, g_loss, e_loss* – discriminator, generator, encoder loss values.

ASHPY INTERNALS

The two main concepts of AshPy internals are *BaseContext* and *Executor*.

5.1 Context

A *BaseContext* is an object that contains all the needed information. Here needed depends on the application. In AshPy the *Context* concept links a generic training loop with the loss function calculation and the model evaluation. A *Context* is a useful class in which all the models, metrics, dataset and mode of your network are set. Passing the context around means that you can any time access to all what you need in order to perform any type of computation.

In AshPy we have (until now) three types of contexts:

- *Classifier Context*
- *GAN Context*
- *GANEncoder Context*

5.1.1 Classifier Context

The *ClassifierContext* is very simple, it contains only:

- classifier_model
- loss
- dataset
- metrics
- log_eval_mode
- global_step
- ckpt

In this way the loss function (*Executor*) can use the context in order to get the model and the needed information in order to correctly feed the model.

5.1.2 GAN Context

The basic *GANContext* is composed by:

- dataset
- generator_model

- discriminator_model
- generator_loss
- discriminator_loss
- metrics
- log_eval_mode
- global_step
- ckpt

As we can see we have all information needed to define our training and evaluation loop.

5.1.3 GANEncoder Context

The `GANEncoderContext` extends the `GANContext`, contains all the information of the base class plus:

- Encoder Model
- Encoder Loss

5.2 Executor

The `Executor` is the main concept behind the loss function implementation in AshPy. An Executor is a class that helps in order to better generalize a training loop. With an Executor you can construct, for example, a custom loss function and put every computation you need inside it. You should define a `call` function inside your class and decorate it with `@Executor.reduce` header, if needed.

Inside the `call` function you can take advantage of a context.

Executors can be summed up, subtracted and multiplied by scalars.

An executor takes also care of the distribution strategy by reducing appropriately the loss (see [Tensorflow Guide](#)).

In this example we will see the implementation of the Generator Binary CrossEntropy loss.

The `__init__` method is straightforward, we need only to instantiate `tf.losses.BinaryCrossentropy` object and then we pass it to our parent:

```
class GeneratorBCE(GANExecutor):

    def __init__(self, from_logits=True):
        self.name = "GeneratorBCE"
        # call super passing the BinaryCrossentropy as function
        super().__init__(tf.losses.BinaryCrossentropy(from_logits=from_logits))
```

Then we need to implement the `call` function respecting the signature:

```
def call(self, context, *, fake, condition, training, **kwargs):

    # we need a function that gives us the correct inputs given the discriminator_
    ↪model
    fake_inputs = self.get_discriminator_inputs(
        context=context, fake_or_real=fake, condition=condition, training=training
    )
```

(continues on next page)

(continued from previous page)

```
# get the discriminator predictions from the discriminator model
d_fake = context.discriminator_model(fake_inputs, training=training)

# get the target prediction for the generator
value = self._fn(tf.ones_like(d_fake), d_fake)

# mean everything
return tf.reduce_mean(value)
```

The function `get_discriminator_inputs()` returns the correct discriminator inputs using the context. The discriminator input can be the output of the generator (unconditioned case) or the output of the generator together with the condition (conditioned case).

The `call()` uses the discriminator model inside the context in order to obtain the output of the discriminator when evaluated in the `fake_inputs`.

After that the `self._fn()` (`BinaryCrossentropy`) is used to get the value of the loss. This loss is then averaged.

In this way the executor computes correctly the loss function.

This is ok if we do not want use our code in a distribution strategy.

If we want to use our executor in a distribution strategy the only modifications are:

```
@Executor.reduce_loss
def call(self, context, *, fake, condition, training, **kwargs):

    # we need a function that gives us the correct inputs given the discriminator_
    ↪model
    fake_inputs = self.get_discriminator_inputs(
        context=context, fake_or_real=fake, condition=condition, training=training
    )

    # get the discriminator predictions from the discriminator model
    d_fake = context.discriminator_model(fake_inputs, training=training)

    # get the target prediction for the generator
    value = self._fn(tf.ones_like(d_fake), d_fake)

    # mean only over the axis 1
    return tf.reduce_mean(value, axis=1)
```

The important things are:

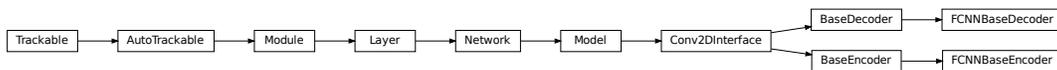
- `Executor.reduce_loss` decoration: uses the `Executor` decorator in order to correctly reduce the loss
- `tf.reduce_mean(value, axis=1)` (last line), we perform only the mean over the axis 1. The output of the `call` function

should be a `tf.Tensor` with shape $(N, 1)$ or $(N,)$. This is because the decorator performs the mean over the axis 0.

DEPENDENCIES GRAPH

6.1 ashpy.models

6.1.1 Convolutional



6.1.2 GANs

6.2 ashpy.trainers

6.2.1 Adversarial

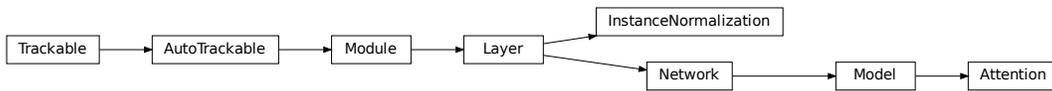


6.2.2 Classifier



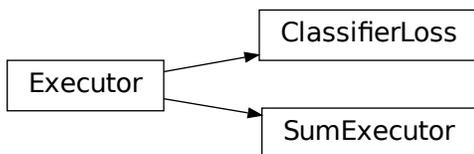
6.3 ashpy.layers

6.3.1 Layers

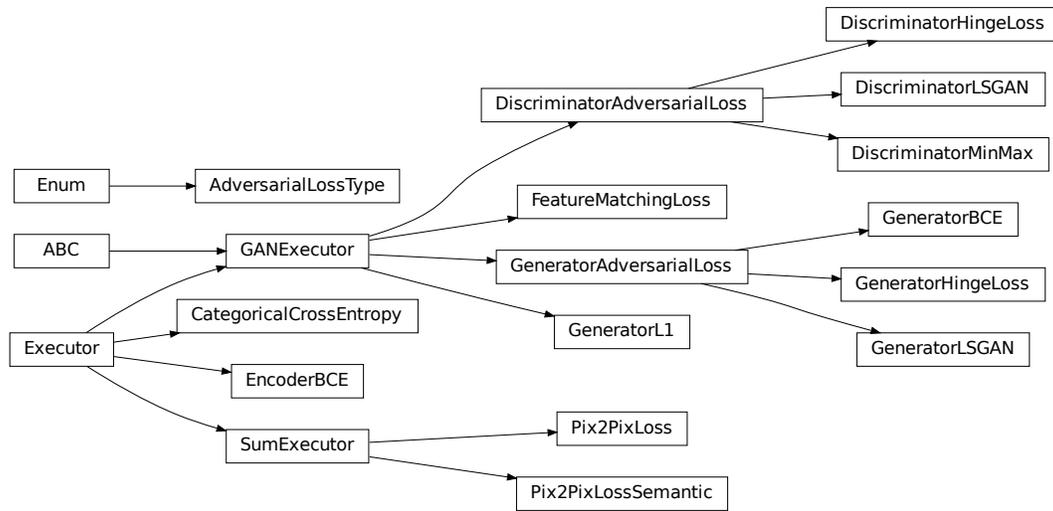


6.4 ashpy.losses

6.4.1 Classifier Losses

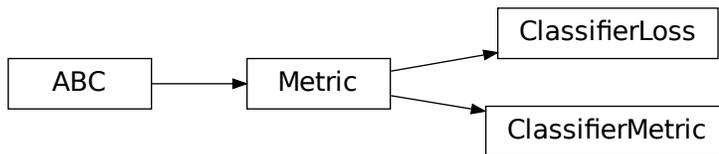


6.4.2 GAN Losses

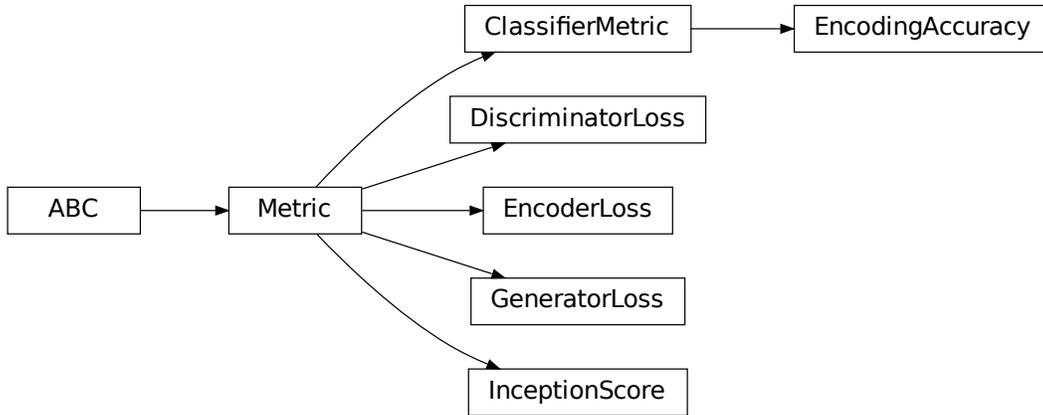


6.5 ashpy.metrics

6.5.1 Classifier Metrics



6.5.2 GAN Metrics



ABOUT

AshPy is an open-source project available on [Github](#) under APACHE licence.

The Framework is created and maintained primarily by the ML & CV Lab @ [Zuru Tech](#).

Please contact ml@zuru.tech for doubt, information or suggestions.

INDICES AND TABLES

- genindex
- modindex
- search

PYTHON MODULE INDEX

a

- ashpy.ashtypes, 31
- ashpy.contexts, 31
- ashpy.contexts.base_context, 38
- ashpy.contexts.classifier, 41
- ashpy.contexts.gan, 43
- ashpy.keras, 49
- ashpy.keras.losses, 49
- ashpy.layers, 57
- ashpy.losses, 62
- ashpy.losses.classifier, 83
- ashpy.losses.executor, 84
- ashpy.losses.gan, 89
- ashpy.metrics, 110
- ashpy.metrics.classifier, 128
- ashpy.metrics.gan, 131
- ashpy.metrics.metric, 142
- ashpy.metrics.sliced_wasserstein_metric, 146
- ashpy.metrics.ssim_multiscale, 152
- ashpy.models, 155
- ashpy.models.convolutional, 155
- ashpy.models.convolutional.autoencoders, 191
- ashpy.models.convolutional.decoders, 208
- ashpy.models.convolutional.discriminators, 198
- ashpy.models.convolutional.encoders, 216
- ashpy.models.convolutional.interfaces, 224
- ashpy.models.convolutional.pix2pixhd, 236
- ashpy.models.convolutional.unet, 228
- ashpy.models.fc, 246
- ashpy.models.fc.autoencoders, 254
- ashpy.models.fc.decoders, 257
- ashpy.models.fc.encoders, 259
- ashpy.models.fc.interfaces, 262
- ashpy.models.gans, 264
- ashpy.modes, 266
- ashpy.trainers, 267
- ashpy.trainers.base_trainer, 276
- ashpy.trainers.classifier, 279
- ashpy.trainers.gan, 283

Symbols

- `__init__()` (*ashpy.contexts.base_context.BaseContext method*), 32, 39, 40
- `__init__()` (*ashpy.contexts.classifier.ClassifierContext method*), 34, 42, 43
- `__init__()` (*ashpy.contexts.gan.GANContext method*), 35, 44, 47
- `__init__()` (*ashpy.contexts.gan.GANEncoderContext method*), 37, 46, 48
- `__init__()` (*ashpy.keras.losses.DHingeLoss method*), 50, 55
- `__init__()` (*ashpy.keras.losses.DLeastSquare method*), 51, 55
- `__init__()` (*ashpy.keras.losses.DMinMax method*), 52, 56
- `__init__()` (*ashpy.keras.losses.GHingeLoss method*), 53, 56
- `__init__()` (*ashpy.keras.losses.L1 method*), 54, 57
- `__init__()` (*ashpy.layers.attention.Attention method*), 62
- `__init__()` (*ashpy.layers.instance_normalization.InstanceNormalization method*), 59
- `__init__()` (*ashpy.losses.classifier.ClassifierLoss method*), 66, 84
- `__init__()` (*ashpy.losses.executor.Executor method*), 63, 85, 87
- `__init__()` (*ashpy.losses.executor.SumExecutor method*), 65, 87, 88
- `__init__()` (*ashpy.losses.gan.CategoricalCrossEntropy method*), 76, 91, 105
- `__init__()` (*ashpy.losses.gan.DiscriminatorAdversarialLoss method*), 70, 92, 105
- `__init__()` (*ashpy.losses.gan.DiscriminatorHingeLoss method*), 82, 93, 105
- `__init__()` (*ashpy.losses.gan.DiscriminatorLSGAN method*), 81, 94, 106
- `__init__()` (*ashpy.losses.gan.DiscriminatorMinMax method*), 80, 94, 106
- `__init__()` (*ashpy.losses.gan.EncoderBCE method*), 80, 95, 106
- `__init__()` (*ashpy.losses.gan.FeatureMatchingLoss method*), 76, 96, 107
- `__init__()` (*ashpy.losses.gan.GANExecutor method*), 68, 97, 107
- `__init__()` (*ashpy.losses.gan.GeneratorAdversarialLoss method*), 69, 99, 107
- `__init__()` (*ashpy.losses.gan.GeneratorBCE method*), 71, 99, 108
- `__init__()` (*ashpy.losses.gan.GeneratorHingeLoss method*), 74, 100, 108
- `__init__()` (*ashpy.losses.gan.GeneratorL1 method*), 73, 101, 108
- `__init__()` (*ashpy.losses.gan.GeneratorLSGAN method*), 72, 102, 108
- `__init__()` (*ashpy.losses.gan.Pix2PixLoss method*), 77, 103, 109
- `__init__()` (*ashpy.losses.gan.Pix2PixLossSemantic method*), 79, 104, 109
- `__init__()` (*ashpy.metrics.classifier.ClassifierLoss method*), 114, 129, 130
- `__init__()` (*ashpy.metrics.classifier.ClassifierMetric method*), 115, 130, 131
- `__init__()` (*ashpy.metrics.gan.DiscriminatorLoss method*), 117, 132, 139
- `__init__()` (*ashpy.metrics.gan.EncoderLoss method*), 119, 134, 139
- `__init__()` (*ashpy.metrics.gan.EncodingAccuracy method*), 123, 135, 140
- `__init__()` (*ashpy.metrics.gan.GeneratorLoss method*), 118, 136, 140
- `__init__()` (*ashpy.metrics.gan.InceptionScore method*), 121, 138, 141
- `__init__()` (*ashpy.metrics.metric.Metric method*), 111, 143, 145
- `__init__()` (*ashpy.metrics.sliced_wasserstein_metric.SingleSWD method*), 147, 150
- `__init__()` (*ashpy.metrics.sliced_wasserstein_metric.SlicedWasserstein method*), 125, 149, 151
- `__init__()` (*ashpy.metrics.ssim_multiscale.SSIM_Multiscale method*), 126, 153, 154
- `__init__()` (*ashpy.models.convolutional.autoencoders.BaseAutoencoder method*), 171, 193, 196
- `__init__()` (*ashpy.models.convolutional.autoencoders.FCNNBaseAutoencoder method*), 173, 195, 198

method), 277, 278
 _train_step (*ashpy.trainers.AdversarialTrainer attribute*), 272
 _train_step (*ashpy.trainers.EncoderTrainer attribute*), 275
 _train_step (*ashpy.trainers.classifier.ClassifierTrainer attribute*), 281, 282
 _train_step (*ashpy.trainers.gan.AdversarialTrainer attribute*), 286, 292
 _train_step (*ashpy.trainers.gan.EncoderTrainer attribute*), 289, 295
 _update_global_batch_size() (*ashpy.trainers.BaseTrainer method*), 269
 _update_global_batch_size() (*ashpy.trainers.base_trainer.BaseTrainer method*), 277, 278
 _validate_metrics() (*ashpy.contexts.base_context.BaseContext method*), 33, 39, 40

A

AdversarialLossType (*class in ashpy.losses.gan*), 68, 90, 105
 AdversarialTrainer (*class in ashpy.trainers*), 269
 AdversarialTrainer (*class in ashpy.trainers.gan*), 283, 289
 ashpy.ashtypes (*module*), 31
 ashpy.contexts (*module*), 31
 ashpy.contexts.base_context (*module*), 38
 ashpy.contexts.classifier (*module*), 41
 ashpy.contexts.gan (*module*), 43
 ashpy.keras (*module*), 49
 ashpy.keras.losses (*module*), 49
 ashpy.layers (*module*), 57
 ashpy.losses (*module*), 62
 ashpy.losses.classifier (*module*), 83
 ashpy.losses.executor (*module*), 84
 ashpy.losses.gan (*module*), 89
 ashpy.metrics (*module*), 110
 ashpy.metrics.classifier (*module*), 128
 ashpy.metrics.gan (*module*), 131
 ashpy.metrics.metric (*module*), 142
 ashpy.metrics.sliced_wasserstein_metric (*module*), 146
 ashpy.metrics.ssim_multiscale (*module*), 152
 ashpy.models (*module*), 155
 ashpy.models.convolutional (*module*), 155
 ashpy.models.convolutional.autoencoders (*module*), 191
 ashpy.models.convolutional.decoders (*module*), 208
 ashpy.models.convolutional.discriminators (*module*), 198

ashpy.models.convolutional.encoders (*module*), 216
 ashpy.models.convolutional.interfaces (*module*), 224
 ashpy.models.convolutional.pix2pixhd (*module*), 236
 ashpy.models.convolutional.unet (*module*), 228
 ashpy.models.fc (*module*), 246
 ashpy.models.fc.autoencoders (*module*), 254
 ashpy.models.fc.decoders (*module*), 257
 ashpy.models.fc.encoders (*module*), 259
 ashpy.models.fc.interfaces (*module*), 262
 ashpy.models.gans (*module*), 264
 ashpy.modes (*module*), 266
 ashpy.trainers (*module*), 267
 ashpy.trainers.base_trainer (*module*), 276
 ashpy.trainers.classifier (*module*), 279
 ashpy.trainers.gan (*module*), 283
 Attention (*class in ashpy.layers.attention*), 60

B

BaseAutoencoder (*class in ashpy.models.convolutional.autoencoders*), 169, 191, 196
 BaseAutoencoder (*class in ashpy.models.fc.autoencoders*), 252, 254, 256
 BaseContext (*class in ashpy.contexts.base_context*), 32, 39, 40
 BaseDecoder (*class in ashpy.models.convolutional.decoders*), 159, 209, 213
 BaseDecoder (*class in ashpy.models.fc.decoders*), 248, 257, 259
 BaseEncoder (*class in ashpy.models.convolutional.encoders*), 164, 217, 221
 BaseEncoder (*class in ashpy.models.fc.encoders*), 250, 260, 261
 BaseTrainer (*class in ashpy.trainers*), 268
 BaseTrainer (*class in ashpy.trainers.base_trainer*), 276, 277
 best_folder() (*ashpy.metrics.metric.Metric property*), 112, 143, 145
 best_model_sel_file() (*ashpy.metrics.metric.Metric property*), 112, 143, 145
 build() (*ashpy.layers.instance_normalization.InstanceNormalization method*), 59
 build_discriminator() (*ashpy.models.convolutional.discriminators.MultiScaleDiscriminator method*), 182, 201, 206

C

call () (*ashpy.keras.losses.DHingeLoss* method), 50, 55
 call () (*ashpy.keras.losses.DLeastSquare* method), 51, 55
 call () (*ashpy.keras.losses.DMinMax* method), 52, 56
 call () (*ashpy.keras.losses.GHingeLoss* method), 53, 57
 call () (*ashpy.keras.losses.L1* method), 54, 57
 call () (*ashpy.layers.attention.Attention* method), 62
 call () (*ashpy.layers.instance_normalization.InstanceNormalization* method), 59
 call () (*ashpy.losses.executor.Executor* method), 63, 85, 87
 call () (*ashpy.losses.executor.SumExecutor* method), 65, 87, 88
 call () (*ashpy.losses.gan.GANExecutor* method), 68, 97, 107
 call () (*ashpy.models.convolutional.autoencoders.BaseAutoencoder* method), 171, 193, 197
 call () (*ashpy.models.convolutional.autoencoders.FCNBaseAutoencoder* method), 173, 196, 198
 call () (*ashpy.models.convolutional.discriminators.MultiScaleDiscriminator* method), 182, 201, 206
 call () (*ashpy.models.convolutional.discriminators.PatchDiscriminator* method), 185, 204, 208
 call () (*ashpy.models.convolutional.interfaces.Conv2DInterface* method), 158, 226, 228
 call () (*ashpy.models.convolutional.pix2pixhd.GlobalGenerator* method), 191, 238, 244
 call () (*ashpy.models.convolutional.pix2pixhd.LocalEnhancer* method), 188, 241, 245
 call () (*ashpy.models.convolutional.pix2pixhd.ResNetBlock* method), 243, 246
 call () (*ashpy.models.convolutional.unet.UNet* method), 176, 233, 236
 call () (*ashpy.models.fc.autoencoders.BaseAutoencoder* method), 254, 256, 257
 call () (*ashpy.models.fc.interfaces.FCInterface* method), 248, 263
 call () (*ashpy.trainers.AdversarialTrainer* method), 272
 call () (*ashpy.trainers.base_trainer.BaseTrainer* method), 277, 278
 call () (*ashpy.trainers.BaseTrainer* method), 269
 call () (*ashpy.trainers.classifier.ClassifierTrainer* method), 281, 282
 call () (*ashpy.trainers.EncoderTrainer* method), 275
 call () (*ashpy.trainers.gan.AdversarialTrainer* method), 286, 292
 call () (*ashpy.trainers.gan.EncoderTrainer* method), 289, 295
 CategoricalCrossEntropy (class in *ashpy.losses.gan*), 76, 90, 105

classifier_model () (*ashpy.contexts.classifier.ClassifierContext* property), 34, 42, 43
 ClassifierContext (class in *ashpy.contexts.classifier*), 33, 41, 43
 ClassifierLoss (class in *ashpy.losses.classifier*), 66, 83, 84
 ClassifierLoss (class in *ashpy.metrics.classifier*), 113, 128, 130
 ClassifierMetric (class in *ashpy.metrics.classifier*), 115, 129, 131
 ClassifierTrainer (class in *ashpy.trainers.classifier*), 279, 281
 Conv2DInterface (class in *ashpy.models.convolutional.interfaces*), 156, 224, 227
 ConvDiscriminator (in module *ashpy.models.gans*), 265
 ConvEncoder (in module *ashpy.models.gans*), 266
 ConvGenerator (in module *ashpy.models.gans*), 264
D
 BaseDiscriminator (*ashpy.contexts.base_context.BaseContext* property), 33, 39, 40
 Discriminator (in module *ashpy.models.gans*), 265
 DenseEncoder (in module *ashpy.models.gans*), 266
 DenseGenerator (in module *ashpy.models.gans*), 264
 DHingeLoss (class in *ashpy.keras.losses*), 49, 54
 discriminator_loss () (*ashpy.contexts.gan.GANContext* property), 36, 45, 47
 discriminator_model () (*ashpy.contexts.gan.GANContext* property), 36, 45, 47
 DiscriminatorAdversarialLoss (class in *ashpy.losses.gan*), 70, 91, 105
 DiscriminatorHingeLoss (class in *ashpy.losses.gan*), 82, 92, 105
 DiscriminatorLoss (class in *ashpy.metrics.gan*), 116, 132, 138
 DiscriminatorLSGAN (class in *ashpy.losses.gan*), 81, 93, 105
 DiscriminatorMinMax (class in *ashpy.losses.gan*), 80, 94, 106
 DLeastSquare (class in *ashpy.keras.losses*), 50, 55
 DMinMax (class in *ashpy.keras.losses*), 52, 56
E
 encoder_loss () (*ashpy.contexts.gan.GANEncoderContext* property), 38, 47, 48
 encoder_model () (*ashpy.contexts.gan.GANEncoderContext* property), 38, 47, 49
 EncoderBCE (class in *ashpy.losses.gan*), 79, 95, 106

- EncoderLoss (class in *ashpy.metrics.gan*), 119, 133, 139
- EncoderTrainer (class in *ashpy.trainers*), 272
- EncoderTrainer (class in *ashpy.trainers.gan*), 286, 292
- EncodingAccuracy (class in *ashpy.metrics.gan*), 122, 134, 139
- Executor (class in *ashpy.losses.executor*), 63, 85, 87
- executors() (*ashpy.losses.executor.SumExecutor* property), 65, 87, 88
- ## F
- FCInterface (class in *ashpy.models.fc.interfaces*), 247, 262, 263
- FCNNBaseAutoencoder (class in *ashpy.models.convolutional.autoencoders*), 171, 194, 197
- FCNNBaseDecoder (class in *ashpy.models.convolutional.decoders*), 162, 212, 215
- FCNNBaseEncoder (class in *ashpy.models.convolutional.encoders*), 167, 219, 223
- FeatureMatchingLoss (class in *ashpy.losses.gan*), 75, 96, 106
- fn() (*ashpy.losses.executor.Executor* property), 63, 85, 87
- FUNet() (in module *ashpy.models.convolutional.unet*), 179, 229, 234
- ## G
- GANContext (class in *ashpy.contexts.gan*), 35, 44, 47
- GANEncoderContext (class in *ashpy.contexts.gan*), 36, 45, 48
- GANExecutor (class in *ashpy.losses.gan*), 67, 97, 107
- generator_loss() (*ashpy.contexts.gan.GANContext* property), 36, 45, 48
- generator_model() (*ashpy.contexts.gan.GANContext* property), 36, 45, 48
- GeneratorAdversarialLoss (class in *ashpy.losses.gan*), 69, 98, 107
- GeneratorBCE (class in *ashpy.losses.gan*), 71, 99, 108
- GeneratorHingeLoss (class in *ashpy.losses.gan*), 74, 100, 108
- GeneratorL1 (class in *ashpy.losses.gan*), 73, 101, 108
- GeneratorLoss (class in *ashpy.metrics.gan*), 118, 136, 140
- GeneratorLSGAN (class in *ashpy.losses.gan*), 72, 102, 108
- get_adversarial_loss_discriminator() (in module *ashpy.losses.gan*), 82, 89, 110
- get_adversarial_loss_generator() (in module *ashpy.losses.gan*), 83, 89, 110
- get_decoder_block() (*ashpy.models.convolutional.unet.UNet* method), 176, 234, 236
- get_discriminator_inputs() (*ashpy.losses.gan.GANExecutor* static method), 68, 97, 107
- get_encoder_block() (*ashpy.models.convolutional.unet.UNet* method), 177, 234, 236
- get_or_train_inception() (*ashpy.metrics.gan.InceptionScore* static method), 121, 138, 141
- GHingeLoss (class in *ashpy.keras.losses*), 53, 56
- global_batch_size() (*ashpy.losses.executor.Executor* property), 64, 85, 88
- global_batch_size() (*ashpy.losses.executor.SumExecutor* property), 65, 87, 88
- global_step() (*ashpy.contexts.base_context.BaseContext* property), 33, 39, 40
- GlobalGenerator (class in *ashpy.models.convolutional.pix2pixhd*), 189, 236, 243
- ## I
- inception_score() (*ashpy.metrics.gan.InceptionScore* method), 121, 138, 141
- InceptionScore (class in *ashpy.metrics.gan*), 120, 137, 140
- InstanceNormalization (class in *ashpy.layers.instance_normalization*), 58
- ## J
- json_read() (*ashpy.metrics.metric.Metric* static method), 112, 143, 145
- json_write() (*ashpy.metrics.metric.Metric* static method), 112, 143, 145
- ## L
- L1 (class in *ashpy.keras.losses*), 54, 57
- LocalEnhancer (class in *ashpy.models.convolutional.pix2pixhd*), 186, 239, 244
- log() (*ashpy.metrics.metric.Metric* method), 112, 144, 145
- log() (*ashpy.metrics.sliced_wasserstein_metric.SlicedWassersteinDistance* method), 125, 149, 151
- log_eval_mode() (*ashpy.contexts.base_context.BaseContext* property), 33, 40, 41
- logdir() (*ashpy.metrics.metric.Metric* property), 112, 144, 145
- LogEvalMode (class in *ashpy.modes*), 267

loss() (*ashpy.contexts.classifier.ClassifierContext property*), 34, 42, 43

M

measure_metrics() (*ashpy.contexts.base_context.BaseContext method*), 33, 40, 41

Metric (*class in ashpy.metrics.metric*), 111, 142, 144

metric() (*ashpy.metrics.metric.Metric property*), 112, 144, 145

metrics() (*ashpy.contexts.base_context.BaseContext property*), 33, 40, 41

model_selection() (*ashpy.contexts.base_context.BaseContext method*), 33, 40, 41

model_selection() (*ashpy.metrics.metric.Metric method*), 112, 144, 146

model_selection() (*ashpy.metrics.sliced_wasserstein_metric.SlicedWassersteinDistance method*), 125, 149, 152

model_selection_operator() (*ashpy.metrics.metric.Metric property*), 113, 144, 146

MultiScaleDiscriminator (*class in ashpy.models.convolutional.discriminators*), 179, 199, 205

N

name() (*ashpy.metrics.metric.Metric property*), 113, 144, 146

P

PatchDiscriminator (*class in ashpy.models.convolutional.discriminators*), 183, 202, 207

Pix2PixLoss (*class in ashpy.losses.gan*), 77, 103, 109

Pix2PixLossSemantic (*class in ashpy.losses.gan*), 78, 104, 109

R

reduce_loss() (*ashpy.losses.executor.Executor static method*), 64, 86, 88

reduction() (*ashpy.keras.losses.DHingeLoss property*), 50, 55

reduction() (*ashpy.keras.losses.DLeastSquare property*), 51, 56

reduction() (*ashpy.keras.losses.DMinMax property*), 52, 56

reduction() (*ashpy.keras.losses.GHingeLoss property*), 54, 57

reduction() (*ashpy.keras.losses.L1 property*), 54, 57

reset_states() (*ashpy.metrics.metric.Metric method*), 113, 144, 146

reset_states() (*ashpy.metrics.sliced_wasserstein_metric.SlicedWassersteinDistance method*), 125, 150, 152

ResNetBlock (*class in ashpy.models.convolutional.pix2pixhd*), 241, 246

result() (*ashpy.metrics.metric.Metric method*), 113, 144, 146

S

SingleSWD (*class in ashpy.metrics.sliced_wasserstein_metric*), 147, 150

SlicedWassersteinDistance (*class in ashpy.metrics.sliced_wasserstein_metric*), 123, 148, 150

split_batch() (*ashpy.metrics.ssim_multiscale.SSIM_Multiscale static method*), 127, 154, 155

SSIM_Multiscale (*class in ashpy.metrics.ssim_multiscale*), 126, 152, 154

SumExecutor (*class in ashpy.losses.executor*), 64, 86, 88

SUNet (*class in ashpy.models.convolutional.unet*), 177, 229, 234

T

train_step() (*ashpy.trainers.AdversarialTrainer method*), 272

train_step() (*ashpy.trainers.classifier.ClassifierTrainer method*), 281, 282

train_step() (*ashpy.trainers.EncoderTrainer method*), 275

train_step() (*ashpy.trainers.gan.AdversarialTrainer method*), 286, 292

train_step() (*ashpy.trainers.gan.EncoderTrainer method*), 289, 295

U

UNet (*class in ashpy.models.convolutional.unet*), 174, 231, 234

update_state() (*ashpy.metrics.classifier.ClassifierLoss method*), 114, 129, 131

update_state() (*ashpy.metrics.classifier.ClassifierMetric method*), 116, 130, 131

update_state() (*ashpy.metrics.gan.DiscriminatorLoss method*), 117, 133, 139

update_state() (*ashpy.metrics.gan.EncoderLoss method*), 120, 134, 139

update_state() (*ashpy.metrics.gan.EncodingAccuracy method*), 123, 135, 140

update_state() (*ashpy.metrics.gan.GeneratorLoss method*), 118, 136, 140

update_state() (*ashpy.metrics.gan.InceptionScore method*), 122, 138, 142

`update_state()` (*ashpy.metrics.metric.Metric*
method), 113, 144, 146

`update_state()` (*ashpy.metrics.sliced_wasserstein_metric.SingleSWD*
method), 148, 150

`update_state()` (*ashpy.metrics.sliced_wasserstein_metric.SlicedWassersteinDistance*
method), 125, 150, 152

`update_state()` (*ashpy.metrics.ssim_multiscale.SSIM_Multiscale*
method), 127, 154, 155

W

`weight()` (*ashpy.losses.executor.Executor* *property*),
64, 86, 88