

---

# **PyTorch-Lightning-Bolts**

## **Documentation**

***Release 0.2.0***

**PyTorchLightning et al.**

**Mar 29, 2021**



## START HERE

<b>1</b>	<b>Introduction Guide</b>	<b>1</b>
<b>2</b>	<b>Model quality control</b>	<b>11</b>
<b>3</b>	<b>Build a Callback</b>	<b>15</b>
<b>4</b>	<b>Info Callbacks</b>	<b>17</b>
<b>5</b>	<b>Self-supervised Callbacks</b>	<b>19</b>
<b>6</b>	<b>Variational Callbacks</b>	<b>21</b>
<b>7</b>	<b>Vision Callbacks</b>	<b>23</b>
<b>8</b>	<b>DataModules</b>	<b>27</b>
<b>9</b>	<b>Sklearn Datamodule</b>	<b>29</b>
<b>10</b>	<b>Vision DataModules</b>	<b>33</b>
<b>11</b>	<b>AsynchronousLoader</b>	<b>47</b>
<b>12</b>	<b>DummyDataset</b>	<b>49</b>
<b>13</b>	<b>Losses</b>	<b>51</b>
<b>14</b>	<b>Bolts Loggers</b>	<b>53</b>
<b>15</b>	<b>How to use models</b>	<b>55</b>
<b>16</b>	<b>Autoencoders</b>	<b>63</b>
<b>17</b>	<b>Classic ML Models</b>	<b>67</b>
<b>18</b>	<b>Convolutional Architectures</b>	<b>71</b>
<b>19</b>	<b>GANs</b>	<b>75</b>
<b>20</b>	<b>Self-supervised Learning</b>	<b>79</b>
<b>21</b>	<b>Self-supervised learning Transforms</b>	<b>91</b>
<b>22</b>	<b>Self-supervised learning</b>	<b>101</b>

<b>23</b>	<b>Semi-supervised learning</b>	<b>103</b>
<b>24</b>	<b>Self-supervised Learning Contrastive tasks</b>	<b>105</b>
<b>25</b>	<b>Indices and tables</b>	<b>109</b>
	<b>Python Module Index</b>	<b>193</b>
	<b>Index</b>	<b>195</b>

---

CHAPTER  
ONE

---

## INTRODUCTION GUIDE

Welcome to PyTorch Lightning Bolts!

Bolts is a Deep learning research and production toolbox of:

- SOTA pretrained models.
- Model components.
- Callbacks.
- Losses.
- Datasets.

**The Main goal of bolts is to enable trying new ideas as fast as possible!**

All models are tested (daily), benchmarked, documented and work on CPUs, TPUs, GPUs and 16-bit precision.

**some examples!**

```
from pl_bolts.models import VAE, GPT2, ImageGPT, PixelCNN
from pl_bolts.models.self_supervised import AMDIM, CPCV2, SimCLR, MocoV2
from pl_bolts.models import LinearRegression, LogisticRegression
from pl_bolts.models.gans import GAN
from pl_bolts.callbacks import PrintTableMetricsCallback
from pl_bolts.datamodules import FashionMNISTDataModule, CIFAR10DataModule,
    ImagenetDataModule
```

Bolts are built for rapid idea iteration - subclass, override and train!

```
from pl_bolts.models import ImageGPT
from pl_bolts.self_supervised import SimCLR

class VideoGPT(ImageGPT):

    def training_step(self, batch, batch_idx):
        x, y = batch
        x = _shape_input(x)

        logits = self.gpt(x)
        simclr_features = self.simclr(x)

        # -----
        # do something new with GPT logits + simclr_features
        # -----

        loss = self.criterion(logits.view(-1, logits.size(-1)), x.view(-1).long())
```

(continues on next page)

(continued from previous page)

```
    logs = {"loss": loss}
    return {"loss": loss, "log": logs}
```

### Mix and match data, modules and components as you please!

```
model = GAN(datamodule=ImagenetDataModule(PATH))
model = GAN(datamodule=FashionMNISTDataModule(PATH))
model = ImageGPT(datamodule=FashionMNISTDataModule(PATH))
```

### And train on any hardware accelerator

```
import pytorch_lightning as pl

model = ImageGPT(datamodule=FashionMNISTDataModule(PATH))

# cpus
pl.Trainer.fit(model)

# gpus
pl.Trainer(gpus=8).fit(model)

# tpus
pl.Trainer(tpu_cores=8).fit(model)
```

### Or pass in any dataset of your choice

```
model = ImageGPT()
Trainer().fit(
    model,
    train_dataloader=DataLoader(...),
    val_dataloader=DataLoader(...)
)
```

---

## 1.1 Community Built

Bolts are built-by the Lightning community and contributed to bolts. The lightning team guarantees that contributions are:

1. Rigorously Tested (CPUs, GPUs, TPUs).
  2. Rigorously Documented.
  3. Standardized via PyTorch Lightning.
  4. Optimized for speed.
  5. Checked for correctness.
-

### 1.1.1 How to contribute

We accept contributions directly to Bolts or via your own repository.

---

**Note:** We encourage you to have your own repository so we can link to it via our docs!

---

To contribute:

1. Submit a pull request to Bolts (we will help you finish it!).
2. We'll help you add [tests](#).
3. We'll help you refactor models to work on ([GPU](#), [TPU](#), [CPU](#))..
4. We'll help you remove bottlenecks in your model.
5. We'll help you write up [documentation](#).
6. We'll help you pretrain expensive models and host weights for you.
7. We'll create proper attribution for you and link to your repo.
8. Once all of this is ready, we will merge into bolts.

After your model or other contribution is in bolts, our team will make sure it maintains compatibility with the other components of the library!

---

### 1.1.2 Contribution ideas

Don't have something to contribute? Ping us on [Slack](#) or look at our [Github issues](#)!

**We'll help and guide you through the implementation / conversion**

---

## 1.2 When to use Bolts

### 1.2.1 For pretrained models

Most bolts have pretrained weights trained on various datasets or algorithms. This is useful when you don't have enough data, time or money to do your own training.

For example, you could use a pretrained VAE to generate features for an image dataset.

```
from pl_bolts.models.autoencoders import VAE
from pl_bolts.models.self_supervised import CPCV2

model1 = VAE(pretrained='imagenet2012')
encoder = model1.encoder
encoder.freeze()

# bolts are pretrained on different datasets
model2 = CPCV2(encoder='resnet18', pretrained='imagenet128').freeze()
model3 = CPCV2(encoder='resnet18', pretrained='stl10').freeze()
```

(continues on next page)

(continued from previous page)

```
for (x, y) in own_data
    features = encoder(x)
    feat2 = model2(x)
    feat3 = model3(x)

# which is better?
```

## 1.2.2 To finetune on your data

If you have your own data, finetuning can often increase the performance. Since this is pure PyTorch you can use any finetuning protocol you prefer.

### Example 1: Unfrozen finetune

```
# unfrozen finetune
model = CPCV2(encoder='resnet18', pretrained='imagenet128')
resnet18 = model.encoder
# don't call .freeze()

classifier = LogisticRegression()

for (x, y) in own_data:
    feats = resnet18(x)
    y_hat = classifier(feats)
```

### Example 2: Freeze then unfreeze

```
# FREEZE!
model = CPCV2(encoder='resnet18', pretrained='imagenet128')
resnet18 = model.encoder
resnet18.freeze()

classifier = LogisticRegression()

for epoch in epochs:
    for (x, y) in own_data:
        feats = resnet18(x)
        y_hat = classifier(feats)
        loss = cross_entropy_with_logits(y_hat, y)

    # UNFREEZE after 10 epochs
    if epoch == 10:
        resnet18.unfreeze()
```

### 1.2.3 For research

Here is where bolts is very different than other libraries with models. It's not just designed for production, but each module is written to be easily extended for research.

```
from pl_bolts.models import ImageGPT
from pl_bolts.self_supervised import SimCLR

class VideoGPT(ImageGPT):

    def training_step(self, batch, batch_idx):
        x, y = batch
        x = _shape_input(x)

        logits = self.gpt(x)
        simclr_features = self.simclr(x)

        # -----
        # do something new with GPT logits + simclr_features
        # -----

        loss = self.criterion(logits.view(-1, logits.size(-1)), x.view(-1).long())

        logs = {"loss": loss}
        return {"loss": loss, "log": logs}
```

Or perhaps your research is in self\_supervised\_learning and you want to do a new SimCLR. In this case, the only thing you want to change is the loss.

By subclassing you can focus on changing a single piece of a system without worrying that the other parts work (because if they are in Bolts, then they do and we've tested it).

```
# subclass SimCLR and change ONLY what you want to try
class ComplexCLR(SimCLR):

    def init_loss(self):
        return self.new_xent_loss

    def new_xent_loss(self):
        out = torch.cat([out_1, out_2], dim=0) n_samples = len(out)

        # Full similarity matrix
        cov = torch.mm(out, out.t().contiguous())
        sim = torch.exp(cov / temperature)

        # Negative similarity
        mask = ~torch.eye(n_samples, device=sim.device).bool()
        neg = sim.masked_select(mask).view(n_samples, -1).sum(dim=-1)

        # -----
        # some new thing we want to do
        # -----

        # Positive similarity :
        pos = torch.exp(torch.sum(out_1 * out_2, dim=-1) / temperature)
        pos = torch.cat([pos, pos], dim=0)
        loss = -torch.log(pos / neg).mean()
```

(continues on next page)

(continued from previous page)

```
return loss
```

## 1.3 Callbacks

Callbacks are arbitrary programs which can run at any points in time within a training loop in Lightning.

Bolts houses a collection of callbacks that are community contributed and can work in any Lightning Module!

```
from pl_bolts.callbacks import PrintTableMetricsCallback
import pytorch_lightning as pl

trainer = pl.Trainer(callbacks=[PrintTableMetricsCallback()])
```

---

## 1.4 DataModules

In PyTorch, working with data has these major elements.

1. Downloading, saving and preparing the dataset.
2. Splitting into train, val and test.
3. For each split, applying different transforms

A DataModule groups together those actions into a single reproducible *DataModule* that can be shared around to guarantee:

1. Consistent data preprocessing (download, splits, etc...)
2. The same exact splits
3. The same exact transforms

```
from pl_bolts.datamodules import ImagenetDataModule

dm = ImagenetDataModule(data_dir=PATH)

# standard PyTorch!
train_loader = dm.train_dataloader()
val_loader = dm.val_dataloader()
test_loader = dm.test_dataloader()

Trainer().fit(
    model,
    train_loader,
    val_loader
)
```

But when paired with PyTorch LightningModules (all bolts models), you can plug and play full dataset definitions with the same splits, transforms, etc...

```
imagenet = ImagenetDataModule(PATH)
model = VAE(datamodule=imagenet)
model = ImageGPT(datamodule=imagenet)
model = GAN(datamodule=imagenet)
```

We even have prebuilt modules to bridge the gap between Numpy, Sklearn and PyTorch

```
from sklearn.datasets import load_boston
from pl_bolts.datamodules import SklearnDataModule

X, y = load_boston(return_X_y=True)
datamodule = SklearnDataModule(X, y)

model = LitModel(datamodule)
```

## 1.5 Regression Heroes

In case your job or research doesn't need a "hammer", we offer implementations of Classic ML models which benefit from lightning's multi-GPU and TPU support.

So, now you can run huge workloads scalably, without needing to do any engineering. For instance, here we can run Logistic Regression on Imagenet (each epoch takes about 3 minutes)!

```
from pl_bolts.models.regression import LogisticRegression

imagenet = ImagenetDataModule(PATH)

# 224 x 224 x 3
pixels_per_image = 150528
model = LogisticRegression(input_dim=pixels_per_image, num_classes=1000)
model.prepare_data = imagenet.prepare_data

trainer = Trainer(gpus=2)
trainer.fit(
    model,
    imagenet.train_dataloader(batch_size=256),
    imagenet.val_dataloader(batch_size=256)
)
```

### 1.5.1 Linear Regression

Here's an example for Linear regression

```
import pytorch_lightning as pl
from pl_bolts.datamodules import SklearnDataModule
from sklearn.datasets import load_boston

# link the numpy dataset to PyTorch
X, y = load_boston(return_X_y=True)
loaders = SklearnDataModule(X, y)

# training runs training batches while validating against a validation set
```

(continues on next page)

(continued from previous page)

```
model = LinearRegression()
trainer = pl.Trainer(num_gpus=8)
trainer.fit(model, loaders.train_dataloader(), loaders.val_dataloader())
```

Once you're done, you can run the test set if needed.

```
trainer.test(test_dataloaders=loaders.test_dataloader())
```

But more importantly, you can scale up to many GPUs, TPUs or even CPUs

```
# 8 GPUs
trainer = pl.Trainer(num_gpus=8)

# 8 TPU cores
trainer = pl.Trainer(tpu_cores=8)

# 32 GPUs
trainer = pl.Trainer(num_gpus=8, num_nodes=4)

# 128 CPUs
trainer = pl.Trainer(num_processes=128)
```

## 1.5.2 Logistic Regression

Here's an example for Logistic regression

```
from sklearn.datasets import load_iris
from pl_bolts.models.regression import LogisticRegression
from pl_bolts.datamodules import SklearnDataModule
import pytorch_lightning as pl

# use any numpy or sklearn dataset
X, y = load_iris(return_X_y=True)
dm = SklearnDataModule(X, y)

# build model
model = LogisticRegression(input_dim=4, num_classes=3)

# fit
trainer = pl.Trainer(tpu_cores=8, precision=16)
trainer.fit(model, dm.train_dataloader(), dm.val_dataloader())

trainer.test(test_dataloaders=dm.test_dataloader(batch_size=12))
```

Any input will be flattened across all dimensions except the first one (batch). This means images, sound, etc... work out of the box.

```
# create dataset
dm = MNISTDataModule(num_workers=0, data_dir=tmpdir)

model = LogisticRegression(input_dim=28 * 28, num_classes=10, learning_rate=0.001)
model.prepare_data = dm.prepare_data
model.train_dataloader = dm.train_dataloader
model.val_dataloader = dm.val_dataloader
model.test_dataloader = dm.test_dataloader
```

(continues on next page)

(continued from previous page)

```
trainer = pl.Trainer(max_epochs=2)
trainer.fit(model)
trainer.test(model)
# {test_acc: 0.92}
```

But more importantly, you can scale up to many GPUs, TPUs or even CPUs

```
# 8 GPUs
trainer = pl.Trainer(num_gpus=8)

# 8 TPUs
trainer = pl.Trainer(tpu_cores=8)

# 32 GPUs
trainer = pl.Trainer(num_gpus=8, num_nodes=4)

# 128 CPUs
trainer = pl.Trainer(num_processes=128)
```

## 1.6 Regular PyTorch

Everything in bolts also works with regular PyTorch since they are all just nn.Modules!

However, if you train using Lightning you don't have to deal with engineering code :)

## 1.7 Command line support

Any bolt module can also be trained from the command line

```
cd pl_bolts/models/autoencoders/basic_vae
python basic_vae_pl_module.py
```

Each script accepts Argparse arguments for both the lightning trainer and the model

```
python basic_vae_pl_module.py --latent_dim 32 --batch_size 32 --gpus 4 --max_epochs 12
```



## MODEL QUALITY CONTROL

For bolts to be added to the library we have a **rigorous** quality control checklist

### 2.1 Bolts vs my own repo

We hope you keep your own repo still! We want to link to it to let people know. However, by adding your contribution to bolts you get these **additional** benefits!

1. More visibility! (more people all over the world use your code)
2. We test your code on every PR (CPUs, GPUs, TPUs).
3. We host the docs (and test on every PR).
4. We help you build thorough, beautiful documentation.
5. We help you build robust tests.
6. We'll pretrain expensive models for you and host weights.
7. We will improve the speed of your models!
8. Eligible for invited talks to discuss your implementation.
9. Lightning Swag + involvement in the broader contributor community :)

---

**Note:** You still get to keep your attribution and be recognized for your work!

---

---

**Note:** Bolts is a community library built by incredible people like you!

---

### 2.2 Contribution requirements

#### 2.2.1 Benchmarked

Models have known performance results on common baseline datasets.

## 2.2.2 Device agnostic

Models must work on CPUs, GPUs and TPUs without changing code. We help authors with this.

```
# bad
encoder.to(device)
```

## 2.2.3 Fast

We inspect models for computational inefficiencies and help authors meet the bar. Granted, sometimes the approaches are slow for mathematical reasons. But anything related to engineering we help overcome.

```
# bad
mtx = ...
for xi in rows:
    for yi in cols
        mxt[xi, yi] = ...

# good
x = x.item().numpy()
x = np.some_fx(x)
x = torch.tensor(x)
```

## 2.2.4 Tested

Models are tested on every PR (on CPUs, GPUs and soon TPUs).

- Live build
- Tests

## 2.2.5 Modular

Models are modularized to be extended and reused easily.

```
# GOOD!
class LitVAE(pl.LightningModule):

    def init_prior(self, ...):
        # enable users to override interesting parts of each model

    def init_posterior(self, ...):
        # enable users to override interesting parts of each model

# BAD
class LitVAE(pl.LightningModule):

    def __init__(self):
        self.prior = ...
        self.posterior = ...
```

## 2.2.6 Attribution

Any models and weights that are contributed are attributed to you as the author(s).

We request that each contribution have:

- The original paper link
- The list of paper authors
- The link to the original paper code (if available)
- The link to your repo
- Your name and your team's name as the implementation authors.
- Your team's affiliation
- Any generated examples, or result plots.
- Hyperparameters configurations for the results.

Thank you for all your amazing contributions!

---

## 2.3 The bar seems high

If your model doesn't yet meet this bar, no worries! Please open the PR and our team of core contributors will help you get there!

---

## 2.4 Do you have contribution ideas?

Yes! Check the Github issues for requests from the Lightning team and the community! We'll even work with you to finish your implementation! Then we'll help you pretrain it and cover the compute costs when possible.



## BUILD A CALLBACK

This module houses a collection of callbacks that can be passed into the trainer

```
from pl_bolts.callbacks import PrintTableMetricsCallback
import pytorch_lightning as pl

trainer = pl.Trainer(callbacks=[PrintTableMetricsCallback()])

# loss|train_loss|val_loss|epoch
# _____
# 2.2541470527648926|2.2541470527648926|2.2158432006835938|0
```

---

### 3.1 What is a Callback

A callback is a self-contained program that can be intertwined into a training pipeline without polluting the main research logic.

---

### 3.2 Create a Callback

Creating a callback is simple:

```
from pytorch_lightning.callbacks import Callback

class MyCallback(Callback):
    def on_epoch_end(self, trainer, pl_module):
        # do something
```

---

Please refer to [Callback docs](#) for a full list of the 20+ hooks available.



## INFO CALLBACKS

These callbacks give all sorts of useful information during training.

---

### 4.1 Print Table Metrics

This callbacks prints training metrics to a table. It's very bare-bones for speed purposes.

`class pl_bolts.callbacks.printing.PrintTableMetricsCallback`  
Bases: `pytorch_lightning.callbacks.Callback`

Prints a table with the metrics in columns on every epoch end

Example:

```
from pl_bolts.callbacks import PrintTableMetricsCallback
callback = PrintTableMetricsCallback()
```

pass into trainer like so:

```
trainer = pl.Trainer(callbacks=[callback])
trainer.fit(...)

# -----
# at the end of every epoch it will print
# -----

# loss|train_loss|val_loss|epoch
# _____
# 2.2541470527648926|2.2541470527648926|2.2158432006835938|0
```



## SELF-SUPERVISED CALLBACKS

Useful callbacks for self-supervised learning models

---

### 5.1 BYOLMAWeightUpdate

The exponential moving average weight-update rule from Bring Your Own Latent (BYOL).

```
class pl_bolts.callbacks.self_supervised.BYOLMAWeightUpdate(initial_tau=0.996)
Bases: pytorch_lightning.Callback
```

Weight update rule from BYOL.

Your model should have a:

- self.online\_network.
- self.target\_network.

Updates the target\_network params using an exponential moving average update rule weighted by tau. BYOL claims this keeps the online\_network from collapsing.

---

**Note:** Automatically increases tau from *initial\_tau* to 1.0 with every training step

---

Example:

```
from pl_bolts.callbacks.self_supervised import BYOLMAWeightUpdate

# model must have 2 attributes
model = Model()
model.online_network = ...
model.target_network = ...

# make sure to set max_steps in Trainer
trainer = Trainer(callbacks=[BYOLMAWeightUpdate()], max_steps=1000)
```

**Parameters** `initial_tau` – starting tau. Auto-updates with every training step

---

## 5.2 SSLOnlineEvaluator

Appends a MLP for fine-tuning to the given model. Callback has its own mini-inner loop.

```
class pl_bolts.callbacks.self_supervised.SSLOnlineEvaluator(drop_p=0.2,      hid-
                                                          den_dim=1024,
                                                          z_dim=None,
                                                          num_classes=None)
```

Bases: pytorch\_lightning.Callback

Attaches a MLP for finetuning using the standard self-supervised protocol.

Example:

```
from pl_bolts.callbacks.self_supervised import SSLOnlineEvaluator

# your model must have 2 attributes
model = Model()
model.z_dim = ... # the representation dim
model.num_classes = ... # the num of classes in the model
```

### Parameters

- `drop_p` (float) – (0.2) dropout probability
- `hidden_dim` (int) –

(1024) the hidden dimension for the finetune MLP

`get_representations(pl_module, x)`

Override this to customize for the particular model :param\_sphinx\_paramlinks\_pl\_bolts.callbacks.self\_supervised.SSLOnlineEvaluator.get\_representations.x:

## VARIATIONAL CALLBACKS

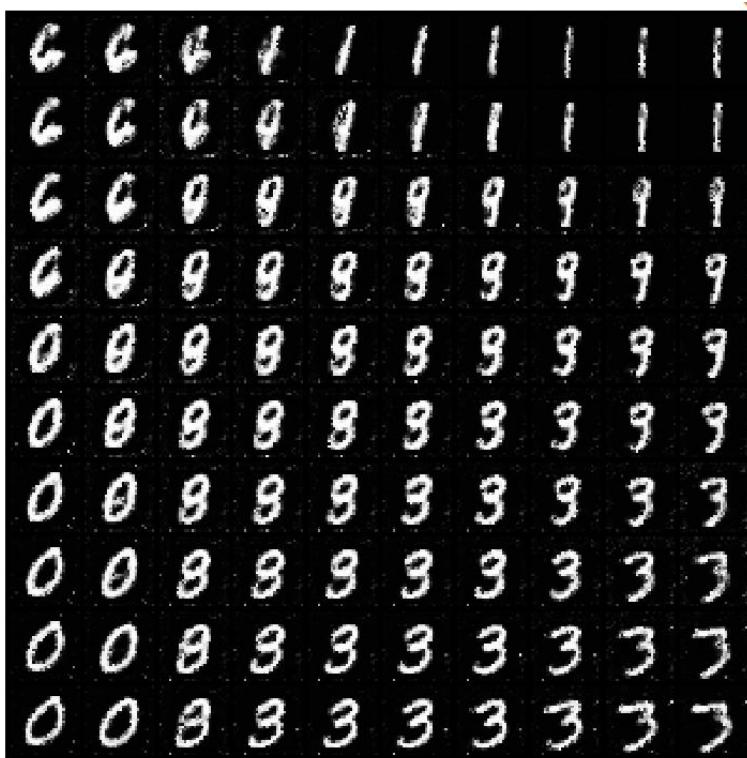
Useful callbacks for GANs, variational-autoencoders or anything with latent spaces.

---

### 6.1 Latent Dim Interpolator

Interpolates latent dims.

Example output:



```
class pl_bolts.callbacks.variational.LatentDimInterpolator(interpolate_epoch_interval=20,  
range_start=-  
5, range_end=5,  
num_samples=2)
```

Bases: pytorch\_lightning.callbacks.Callback

Interpolates the latent space for a model by setting all dims to zero and stepping through the first two dims increasing one unit at a time.

Default interpolates between [-5, 5] (-5, -4, -3, ..., 3, 4, 5)

Example:

```
from pl_bolts.callbacks import LatentDimInterpolator  
Trainer(callbacks=[LatentDimInterpolator()])
```

#### Parameters

- `interpolate_epoch_interval` –
- `range_start` – default -5
- `range_end` – default 5
- `num_samples` – default 2

## VISION CALLBACKS

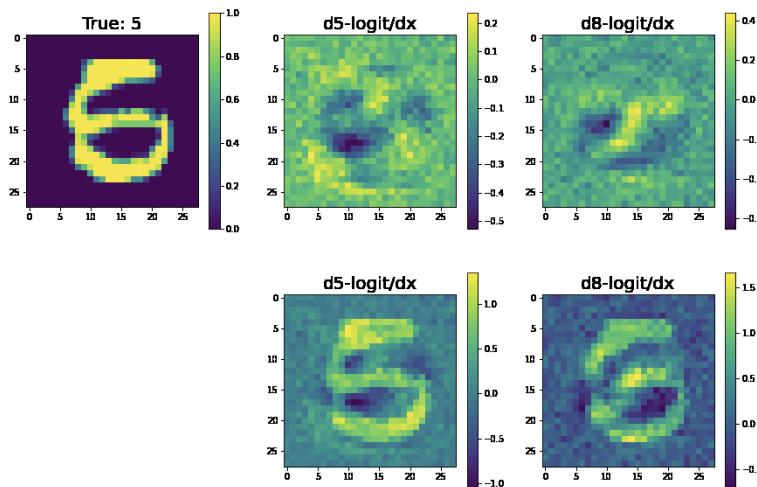
Useful callbacks for vision models

---

### 7.1 Confused Logit

Shows how the input would have to change to move the prediction from one logit to the other

Example outputs:



```
class pl_bolts.callbacks.vision.confused_logit.ConfusedLogitCallback(top_k,
    projec-
    tion_factor=3,
    min_logit_value=5.0,
    log-
    ging_batch_interval=20,
    max_logit_difference=0.1)
```

Bases: `pytorch_lightning.Callback`

Takes the logit predictions of a model and when the probabilities of two classes are very close, the model doesn't have high certainty that it should pick one vs the other class.

This callback shows how the input would have to change to swing the model from one label prediction to the other.

In this case, the network predicts a 5... but gives almost equal probability to an 8. The images show what about the original 5 would have to change to make it more like a 5 or more like an 8.

For each confused logit the confused images are generated by taking the gradient from a logit wrt an input for the top two closest logits.

Example:

```
from pl_bolts.callbacks.vision import ConfusedLogitCallback
trainer = Trainer(callbacks=[ConfusedLogitCallback()])
```

---

**Note:** whenever called, this model will look for self.last\_batch and self.last\_logits in the LightningModule

---

**Note:** this callback supports tensorboard only right now

---

#### Parameters

- `top_k` – How many “offending” images we should plot
- `projection_factor` – How much to multiply the input image to make it look more like this logit label
- `min_logit_value` – Only consider logit values above this threshold
- `logging_batch_interval` – how frequently to inspect/potentially plot something
- `max_logit_difference` – when the top 2 logits are within this threshold we consider them confused

Authored by:

- Alfredo Canziani
- 

## 7.2 Tensorboard Image Generator

Generates images from a generative model and plots to tensorboard

```
class pl_bolts.callbacks.vision.image_generation.TensorboardGenerativeModelImageSampler(num
Bases: pytorch_lightning.Callback
```

Generates images and logs to tensorboard. Your model must implement the forward function for generation

Requirements:

```
# model must have img_dim arg
model.img_dim = (1, 28, 28)

# model forward must work for sampling
z = torch.rand(batch_size, latent_dim)
img_samples = your_model(z)
```

Example:

```
from pl_bolts.callbacks import TensorboardGenerativeModelImageSampler
trainer = Trainer(callbacks=[TensorboardGenerativeModelImageSampler()])
```



---

**CHAPTER  
EIGHT**

---

## **DATAMODULES**

DataModules (introduced in PyTorch Lightning 0.9.0) decouple the data from a model. A DataModule is simply a collection of a training dataloader, val dataloader and test dataloader. In addition, it specifies how to:

- Downloading/preparing data.
- Train/val/test splits.
- Transforms

Then you can use it like this:

Example:

```
dm = MNISTDataModule('path/to/data')
model = LitModel()

trainer = Trainer()
trainer.fit(model, dm)
```

Or use it manually with plain PyTorch

Example:

```
dm = MNISTDataModule('path/to/data')
for batch in dm.train_dataloader():
    ...
for batch in dm.val_dataloader():
    ...
for batch in dm.test_dataloader():
    ...
```

Please visit the PyTorch Lightning documentation for more details on DataModules



## SKLEARN DATAMODULE

Utilities to map sklearn or numpy datasets to PyTorch DataLoaders with automatic data splits and GPU/TPU support.

```
from sklearn.datasets import load_boston
from pl_bolts.datamodules import SklearnDataModule

X, y = load_boston(return_X_y=True)
loaders = SklearnDataModule(X, y)

train_loader = loaders.train_dataloader(batch_size=32)
val_loader = loaders.val_dataloader(batch_size=32)
test_loader = loaders.test_dataloader(batch_size=32)
```

Or build your own torch datasets

```
from sklearn.datasets import load_boston
from pl_bolts.datamodules import SklearnDataset

X, y = load_boston(return_X_y=True)
dataset = SklearnDataset(X, y)
loader = DataLoader(dataset)
```

---

### 9.1 Sklearn Dataset Class

Transforms a sklearn or numpy dataset to a PyTorch Dataset.

```
class pl_bolts.datamodules.sklearn_datamodule.SklearnDataset(X,
                                                               y,
                                                               X_transform=None,
                                                               y_transform=None)
```

Bases: `torch.utils.data.Dataset`

Mapping between numpy (or sklearn) datasets to PyTorch datasets.

#### Parameters

- `X` – Numpy ndarray
- `y` – Numpy ndarray
- `X_transform` – Any transform that works with Numpy arrays
- `y_transform` – Any transform that works with Numpy arrays

### Example

```
>>> from sklearn.datasets import load_boston
>>> from pl_bolts.datamodules import SklearnDataset
...
>>> X, y = load_boston(return_X_y=True)
>>> dataset = SklearnDataset(X, y)
>>> len(dataset)
506
```

---

## 9.2 Sklearn DataModule Class

Automatically generates the train, validation and test splits for a Numpy dataset. They are set up as dataloaders for convenience. Optionally, you can pass in your own validation and test splits.

```
class pl_bolts.datamodules.sklearn_datamodule.SklearnDataModule(X, y,
                                                               x_val=None,
                                                               y_val=None,
                                                               x_test=None,
                                                               y_test=None,
                                                               val_split=0.2,
                                                               test_split=0.1,
                                                               num_workers=2,
                                                               random_state=1234,
                                                               shuffle=True,
                                                               *args,
                                                               **kwargs)
```

Bases: pytorch\_lightning.LightningDataModule

Automatically generates the train, validation and test splits for a Numpy dataset. They are set up as dataloaders for convenience. Optionally, you can pass in your own validation and test splits.

### Example

```
>>> from sklearn.datasets import load_boston
>>> from pl_bolts.datamodules import SklearnDataModule
...
>>> X, y = load_boston(return_X_y=True)
>>> loaders = SklearnDataModule(X, y)
...
>>> # train set
>>> train_loader = loaders.train_dataloader(batch_size=32)
>>> len(train_loader.dataset)
355
>>> len(train_loader)
11
>>> # validation set
>>> val_loader = loaders.val_dataloader(batch_size=32)
>>> len(val_loader.dataset)
100
```

(continues on next page)

(continued from previous page)

```
>>> len(val_loader)
3
>>> # test set
>>> test_loader = loaders.test_dataloader(batch_size=32)
>>> len(test_loader.dataset)
51
>>> len(test_loader)
1
```



## VISION DATAMODULES

The following are pre-built datamodules for computer-vision.

---

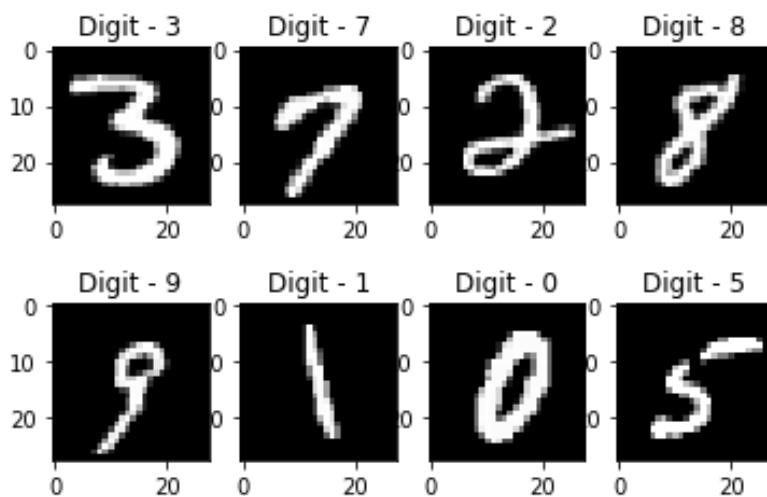
### 10.1 Supervised learning

These are standard vision datasets with the train, test, val splits pre-generated in DataLoaders with the standard transforms (and Normalization) values

#### 10.1.1 BinaryMNIST

```
class pl_bolts.datamodules.binary_mnist_datamodule.BinaryMNISTDataModule(data_dir,  
                           val_split=5000,  
                           num_workers=16,  
                           normalize=False,  
                           seed=42,  
                           *args,  
                           **kwargs)
```

Bases: pytorch\_lightning.LightningDataModule



**Specs:**

- 10 classes (1 per digit)
- Each image is (1 x 28 x 28)

Binary MNIST, train, val, test splits and transforms

Transforms:

```
mnist_transforms = transform_lib.Compose([
    transform_lib.ToTensor()
])
```

Example:

```
from pl_bolts.datamodules import BinaryMNISTDataModule

dm = BinaryMNISTDataModule('.')
model = LitModel()

Trainer().fit(model, dm)
```

**Parameters**

- **data\_dir** (str) – where to save/load the data
- **val\_split** (int) – how many of the training images to use for the validation split
- **num\_workers** (int) – how many workers to use for loading data
- **normalize** (bool) – If true applies image normalize

**prepare\_data()**

Saves MNIST files to data\_dir

**test\_dataloader** (batch\_size=32, transforms=None)

MNIST test set uses the test split

**Parameters**

- **batch\_size** – size of batch
- **transforms** – custom transforms

**train\_dataloader** (batch\_size=32, transforms=None)

MNIST train set removes a subset to use for validation

**Parameters**

- **batch\_size** – size of batch
- **transforms** – custom transforms

**val\_dataloader** (batch\_size=32, transforms=None)

MNIST val set uses a subset of the training set for validation

**Parameters**

- **batch\_size** – size of batch
- **transforms** – custom transforms

**property num\_classes**

Return: 10

### 10.1.2 CityScapes

```
class pl_bolts.datamodules.cityscapes_datamodule.CityscapesDataModule(data_dir,
                                                                      val_split=5000,
                                                                      num_workers=16,
                                                                      batch_size=32,
                                                                      seed=42,
                                                                      *args,
                                                                      **kwargs)
```

Bases: pytorch\_lightning.LightningDataModule



Standard Cityscapes, train, val, test splits and transforms

**Specs:**

- 30 classes (road, person, sidewalk, etc...)
- (image, target) - image dims: (3 x 32 x 32), target dims: (3 x 32 x 32)

Transforms:

```
transforms = transform_lib.Compose([
    transform_lib.ToTensor(),
    transform_lib.Normalize(
        mean=[0.28689554, 0.32513303, 0.28389177],
        std=[0.18696375, 0.19017339, 0.18720214]
    )
])
```

Example:

```
from pl_bolts.datamodules import CityscapesDataModule

dm = CityscapesDataModule(PATH)
model = LitModel()

Trainer().fit(model, dm)
```

Or you can set your own transforms

Example:

```
dm.train_transforms = ...
dm.test_transforms = ...
dm.val_transforms = ...
```

### Parameters

- `data_dir` – where to save/load the data
- `val_split` – how many of the training images to use for the validation split
- `num_workers` – how many workers to use for loading data
- `batch_size` – number of examples per training/eval step

`prepare_data()`

Saves Cityscapes files to data\_dir

`test_dataloader()`

Cityscapes test set uses the test split

`train_dataloader()`

Cityscapes train set with removed subset to use for validation

`val_dataloader()`

Cityscapes val set uses a subset of the training set for validation

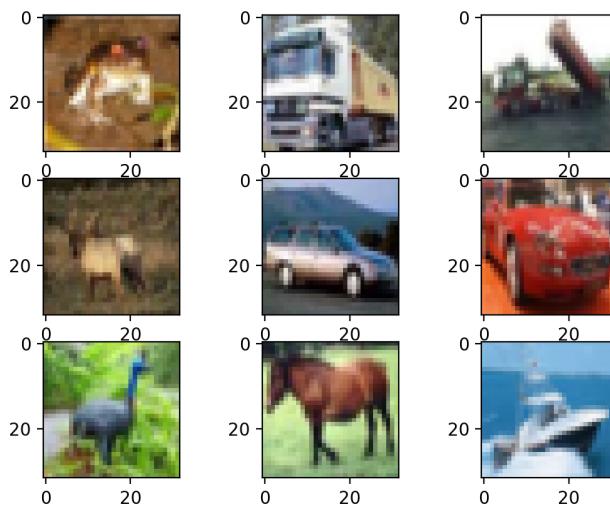
`property num_classes`

Return: 30

### 10.1.3 CIFAR-10

```
class pl_bolts.datamodules.cifar10_datamodule.CIFAR10DataModule(data_dir=None,
                                                               val_split=5000,
                                                               num_workers=16,
                                                               batch_size=32,
                                                               seed=42,
                                                               *args,
                                                               **kwargs)
```

Bases: `pytorch_lightning.LightningDataModule`



**Specs:**

- 10 classes (1 per class)
- Each image is (3 x 32 x 32)

Standard CIFAR10, train, val, test splits and transforms

Transforms:

```
mnist_transforms = transform_lib.Compose([
    transform_lib.ToTensor(),
    transforms.Normalize(
        mean=[x / 255.0 for x in [125.3, 123.0, 113.9]],
        std=[x / 255.0 for x in [63.0, 62.1, 66.7]]
    )
])
```

Example:

```
from pl_bolts.datamodules import CIFAR10DataModule

dm = CIFAR10DataModule(PATH)
model = LitModel()

Trainer().fit(model, dm)
```

Or you can set your own transforms

Example:

```
dm.train_transforms = ...
dm.test_transforms = ...
dm.val_transforms = ...
```

## Parameters

- **data\_dir** (Optional[str]) – where to save/load the data
- **val\_split** (int) – how many of the training images to use for the validation split
- **num\_workers** (int) – how many workers to use for loading data
- **batch\_size** (int) – number of examples per training/eval step

### `prepare_data()`

Saves CIFAR10 files to data\_dir

### `test_dataloader()`

CIFAR10 test set uses the test split

### `train_dataloader()`

CIFAR train set removes a subset to use for validation

### `val_dataloader()`

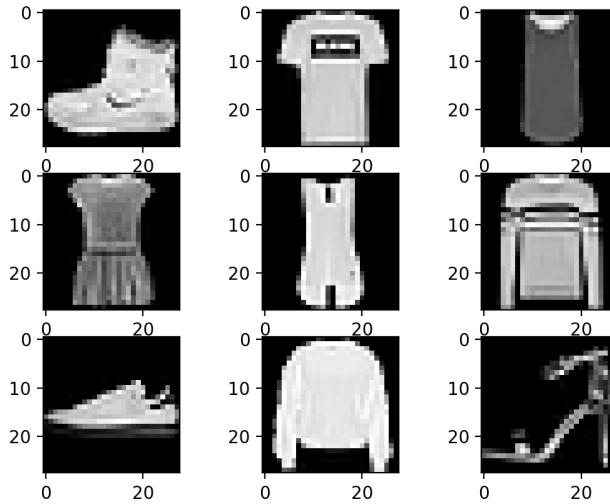
CIFAR10 val set uses a subset of the training set for validation

### `property num_classes`

Return: 10

### 10.1.4 FashionMNIST

```
class pl_bolts.datamodules.fashion_mnist_datamodule.FashionMNISTDataModule(data_dir,
                                                                           val_split=5000,
                                                                           num_workers=16,
                                                                           seed=42,
                                                                           *args,
                                                                           **kwargs)
Bases: pytorch_lightning.LightningDataModule
```



#### Specs:

- 10 classes (1 per type)
- Each image is (1 x 28 x 28)

Standard FashionMNIST, train, val, test splits and transforms

Transforms:

```
mnist_transforms = transform_lib.Compose([
    transform_lib.ToTensor()
])
```

Example:

```
from pl_bolts.datamodules import FashionMNISTDataModule

dm = FashionMNISTDataModule('..')
model = LitModel()

Trainer().fit(model, dm)
```

#### Parameters

- `data_dir` (str) – where to save/load the data

- **val\_split** (int) – how many of the training images to use for the validation split
- **num\_workers** (int) – how many workers to use for loading data

**prepare\_data()**  
Saves FashionMNIST files to data\_dir

**test\_dataloader** (batch\_size=32, transforms=None)  
FashionMNIST test set uses the test split

**Parameters**

- **batch\_size** – size of batch
- **transforms** – custom transforms

**train\_dataloader** (batch\_size=32, transforms=None)  
FashionMNIST train set removes a subset to use for validation

**Parameters**

- **batch\_size** – size of batch
- **transforms** – custom transforms

**val\_dataloader** (batch\_size=32, transforms=None)  
FashionMNIST val set uses a subset of the training set for validation

**Parameters**

- **batch\_size** – size of batch
- **transforms** – custom transforms

**property num\_classes**  
Return: 10

### 10.1.5 Imagenet

```
class pl_bolts.datamodules.imagenet_datamodule.ImagenetDataModule(data_dir,
                                                               meta_dir=None,
                                                               num_imgs_per_val_class=50,
                                                               im-
                                                               age_size=224,
                                                               num_workers=16,
                                                               batch_size=32,
                                                               *args,
                                                               **kwargs)
```

Bases: pytorch\_lightning.LightningDataModule

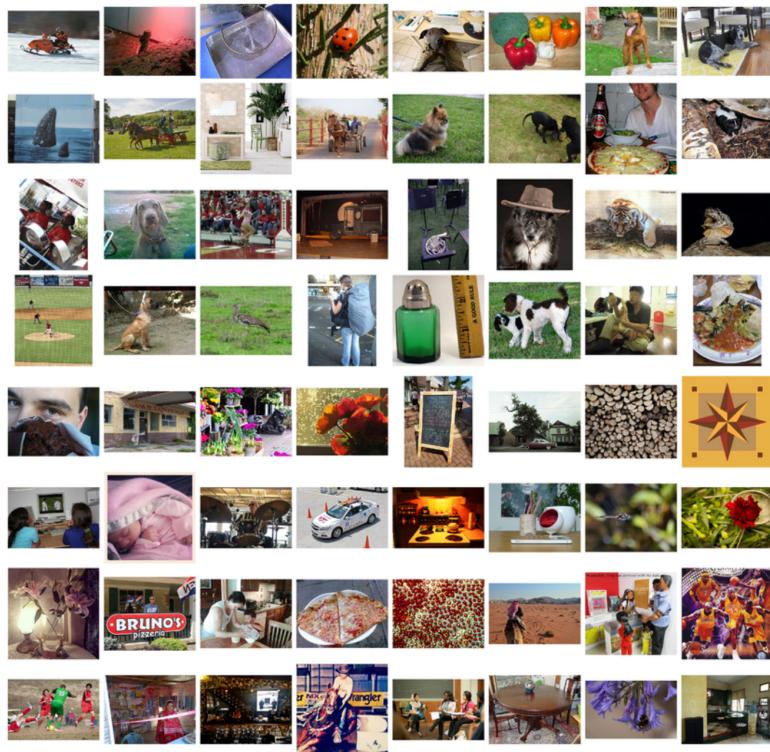
#### Specs:

- 1000 classes
- Each image is (3 x varies x varies) (here we default to 3 x 224 x 224)

Imagenet train, val and test dataloaders.

The train set is the imagenet train.

The val set is taken from the train set with *num\_imgs\_per\_val\_class* images per class. For example if *num\_imgs\_per\_val\_class*=2 then there will be 2,000 images in the validation set.



The test set is the official imangenet validation set.

Example:

```
from pl_bolts.datamodules import ImagenetDataModule

dm = ImagenetDataModule(IMAGENET_PATH)
model = LitModel()

Trainer().fit(model, dm)
```

#### Parameters

- **data\_dir** (str) – path to the imangenet dataset file
- **meta\_dir** (Optional[str]) – path to meta.bin file
- **num\_imgs\_per\_val\_class** (int) – how many images per class for the validation set
- **image\_size** (int) – final image size
- **num\_workers** (int) – how many data workers
- **batch\_size** (int) – batch\_size

#### prepare\_data()

This method already assumes you have imangenet2012 downloaded. It validates the data using the meta.bin.

**Warning:** Please download imangenet on your own first.

**test\_dataloader()**

Uses the validation split of imagenet2012 for testing

**train\_dataloader()**

Uses the train split of imagenet2012 and puts away a portion of it for the validation split

**train\_transform()**

The standard imagenet transforms

```
transform_lib.Compose([
    transform_lib.RandomResizedCrop(self.image_size),
    transform_lib.RandomHorizontalFlip(),
    transform_lib.ToTensor(),
    transform_lib.Normalize(
        mean=[0.485, 0.456, 0.406],
        std=[0.229, 0.224, 0.225]
    ),
])
```

**val\_dataloader()**

Uses the part of the train split of imagenet2012 that was not used for training via *num\_imgs\_per\_val\_class*

**Parameters**

- **batch\_size** – the batch size
- **transforms** – the transforms

**val\_transform()**

The standard imagenet transforms for validation

```
transform_lib.Compose([
    transform_lib.Resize(self.image_size + 32),
    transform_lib.CenterCrop(self.image_size),
    transform_lib.ToTensor(),
    transform_lib.Normalize(
        mean=[0.485, 0.456, 0.406],
        std=[0.229, 0.224, 0.225]
    ),
])
```

**property num\_classes**

Return:

1000

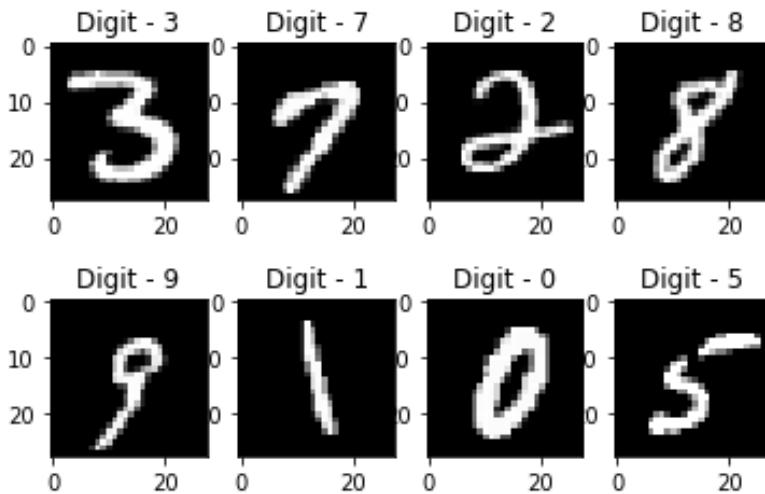
## 10.1.6 MNIST

```
class pl_bolts.datamodules.mnist_datamodule.MNISTDataModule(data_dir,
                                                               val_split=5000,
                                                               num_workers=16,
                                                               normalize=False,
                                                               seed=42,      *args,
                                                               **kwargs)
```

Bases: pytorch\_lightning.LightningDataModule

**Specs:**

- 10 classes (1 per digit)



- Each image is (1 x 28 x 28)

Standard MNIST, train, val, test splits and transforms

Transforms:

```
mnist_transforms = transform_lib.Compose([
    transform_lib.ToTensor()
])
```

Example:

```
from pl_bolts.datamodules import MNISTDataModule

dm = MNISTDataModule('.')
model = LitModel()

Trainer().fit(model, dm)
```

### Parameters

- `data_dir` (str) – where to save/load the data
- `val_split` (int) – how many of the training images to use for the validation split
- `num_workers` (int) – how many workers to use for loading data
- `normalize` (bool) – If true applies image normalize

### `prepare_data()`

Saves MNIST files to `data_dir`

### `test_dataloader(batch_size=32, transforms=None)`

MNIST test set uses the test split

#### Parameters

- `batch_size` – size of batch
- `transforms` – custom transforms

---

```
train_dataloader(batch_size=32, transforms=None)
    MNIST train set removes a subset to use for validation

    Parameters
        • batch_size – size of batch
        • transforms – custom transforms

val_dataloader(batch_size=32, transforms=None)
    MNIST val set uses a subset of the training set for validation

    Parameters
        • batch_size – size of batch
        • transforms – custom transforms

property num_classes
    Return: 10
```

---

## 10.2 Semi-supervised learning

The following datasets have support for unlabeled training and semi-supervised learning where only a few examples are labeled.

### 10.2.1 Imagenet (ssl)

```
class pl_bolts.datamodules.ssl_imagenet_datamodule.SSLImagenetDataModule(data_dir,
    meta_dir=None,
    num_workers=16,
    *args,
    **kwargs)
```

Bases: pytorch\_lightning.LightningDataModule

### 10.2.2 STL-10

```
class pl_bolts.datamodules.stl10_datamodule.STL10DataModule(data_dir=None,
    unla-
    beled_val_split=5000,
    train_val_split=500,
    num_workers=16,
    batch_size=32,
    seed=42,      *args,
    **kwargs)
```

Bases: pytorch\_lightning.LightningDataModule

#### Specs:

- 10 classes (1 per type)
- Each image is (3 x 96 x 96)



Standard STL-10, train, val, test splits and transforms. STL-10 has support for doing validation splits on the labeled or unlabeled splits

Transforms:

```
mnist_transforms = transform_lib.Compose([
    transform_lib.ToTensor(),
    transforms.Normalize(
        mean=(0.43, 0.42, 0.39),
        std=(0.27, 0.26, 0.27)
    )
])
```

Example:

```
from pl_bolts.datamodules import STL10DataModule

dm = STL10DataModule(PATH)
model = LitModel()

Trainer().fit(model, dm)
```

### Parameters

- `data_dir` (Optional[str]) – where to save/load the data
- `unlabeled_val_split` (int) – how many images from the unlabeled training split to use for validation
- `train_val_split` (int) – how many images from the labeled training split to use for validation
- `num_workers` (int) – how many workers to use for loading data
- `batch_size` (int) – the batch size

#### `prepare_data()`

Downloads the unlabeled, train and test split

#### `test_dataloader()`

Loads the test split of STL10

### Parameters

- `batch_size` – the batch size

- **transforms** – the transforms

**train\_dataloader()**

Loads the ‘unlabeled’ split minus a portion set aside for validation via *unlabeled\_val\_split*.

**train\_dataloader\_mixed()**

Loads a portion of the ‘unlabeled’ training data and ‘train’ (labeled) data. both portions have a subset removed for validation via *unlabeled\_val\_split* and *train\_val\_split*

**Parameters**

- **batch\_size** – the batch size
- **transforms** – a sequence of transforms

**val\_dataloader()**

Loads a portion of the ‘unlabeled’ training data set aside for validation The val dataset = (unlabeled - train\_val\_split)

**Parameters**

- **batch\_size** – the batch size
- **transforms** – a sequence of transforms

**val\_dataloader\_mixed()**

Loads a portion of the ‘unlabeled’ training data set aside for validation along with the portion of the ‘train’ dataset to be used for validation

unlabeled\_val = (unlabeled - train\_val\_split)

labeled\_val = (train- train\_val\_split)

full\_val = unlabeled\_val + labeled\_val

**Parameters**

- **batch\_size** – the batch size
- **transforms** – a sequence of transforms



## ASYNCHRONOUSLOADER

This dataloader behaves identically to the standard pytorch dataloader, but will transfer data asynchronously to the GPU with training. You can also use it to wrap an existing dataloader.

**Example::** `dataloader = AsynchronousLoader(DataLoader(ds, batch_size=16), device=device)`

**for b in dataloader:** ...

```
class pl_bolts.datamodules.async_dataloader.AsynchronousLoader(data,           de-  
                                                vice=torch.device,  
                                                q_size=10,  
                                                num_batches=None,  
                                                **kwargs)
```

Bases: `object`

Class for asynchronously loading from CPU memory to device memory with DataLoader.

Note that this only works for single GPU training, multiGPU uses PyTorch's DataParallel or DistributedDataParallel which uses its own code for transferring data across GPUs. This could just break or make things slower with DataParallel or DistributedDataParallel.

### Parameters

- `data` – The PyTorch Dataset or DataLoader we're using to load.
  - `device` – The PyTorch device we are loading to
  - `q_size` – Size of the queue used to store the data loaded to the device
  - `num_batches` – Number of batches to load. This must be set if the dataloader doesn't have a finite `_len_`. It will also override `DataLoader._len_` if set and `DataLoader` has a `_len_`. Otherwise it can be left as None
  - `**kwargs` – Any additional arguments to pass to the dataloader if we're constructing one here
-



---

CHAPTER  
TWELVE

---

## DUMMYDATASET

```
class pl_bolts.datamodules.dummy_dataset.DummyDataset (*shapes,  
                                                    num_samples=10000)
```

Bases: `torch.utils.data.Dataset`

Generate a dummy dataset

### Parameters

- `*shapes` – list of shapes
- `num_samples` – how many samples to use in this dataset

Example:

```
from pl_bolts.datamodules import DummyDataset

# mnist dims
>>> ds = DummyDataset((1, 28, 28), (1,))
>>> dl = DataLoader(ds, batch_size=7)
...
>>> batch = next(iter(dl))
>>> x, y = batch
>>> x.size()
torch.Size([7, 1, 28, 28])
>>> y.size()
torch.Size([7, 1])
```



---

CHAPTER  
**THIRTEEN**

---

## **LOSSES**

This package lists common losses across research domains (This is a work in progress. If you have any losses you want to contribute, please submit a PR!)

---

**Note:** this module is a work in progress

---

---

### **13.1 Your Loss**

We're cleaning up many of our losses, but in the meantime, submit a PR to add your loss here!



---

CHAPTER  
FOURTEEN

---

## BOLTS LOGGERS

The loggers in this package are being considered to be added to the main PyTorch Lightning repository. These loggers may be more unstable, in development, or not fully tested yet.

---

**Note:** This module is a work in progress

---

### 14.1 allegro.ai TRAINS

allegro.ai is a third-party logger. To use *TrainsLogger* as your logger do the following. First, install the package:

```
pip install trains
```

Then configure the logger and pass it to the Trainer:

```
from pl_bolts.loggers import TrainsLogger
trains_logger = TrainsLogger(
    project_name='examples',
    task_name='pytorch lightning test',
)
trainer = Trainer(logger=trains_logger)
```

```
class MyModule(LightningModule):
    def __init__(self):
        some_img = fake_image()
        self.logger.experiment.log_image('debug', 'generated_image_0', some_img, 0)
```

**See also:**

*TrainsLogger* docs.

---

### 14.1.1 Your Logger

Add your loggers here!

## HOW TO USE MODELS

Models are meant to be “bolted” onto your research or production cases.

Bolts are meant to be used in the following ways

---

### 15.1 Predicting on your data

Most bolts have pretrained weights trained on various datasets or algorithms. This is useful when you don’t have enough data, time or money to do your own training.

For example, you could use a pretrained VAE to generate features for an image dataset.

```
from pl_bolts.models.autoencoders import VAE

model = VAE(pretrained='imagenet2012')
encoder = model.encoder
encoder.freeze()

for (x, y) in own_data
    features = encoder(x)
```

The advantage of bolts is that each system can be decomposed and used in interesting ways. For instance, this resnet18 was trained using self-supervised learning (no labels) on Imagenet, and thus might perform better than the same resnet18 trained with labels

```
# trained without labels
from pl_bolts.models.self_supervised import CPCV2

model = CPCV2(encoder='resnet18', pretrained='imagenet128')
resnet18_unsupervised = model.encoder.freeze()

# trained with labels
from torchvision.models import resnet18
resnet18_supervised = resnet18(pretrained=True)

# perhaps the features when trained without labels are much better for classification_
# or other tasks
x = image_sample()
unsup_feats = resnet18_unsupervised(x)
sup_feats = resnet18_supervised(x)

# which one will be better?
```

Bolts are often trained on more than just one dataset.

```
model = CPCV2(encoder='resnet18', pretrained='stl10')
```

---

## 15.2 Finetuning on your data

If you have a little bit of data and can pay for a bit of training, it's often better to finetune on your own data.

To finetune you have two options unfrozen finetuning or unfrozen later.

### 15.2.1 Unfrozen Finetuning

In this approach, we load the pretrained model and unfreeze from the beginning

```
model = CPCV2(encoder='resnet18', pretrained='imagenet128')
resnet18 = model.encoder
# don't call .freeze()

classifier = LogisticRegression()

for (x, y) in own_data:
    feats = resnet18(x)
    y_hat = classifier(feats)
    ...
```

Or as a LightningModule

```
class FineTuner(pl.LightningModule):

    def __init__(self, encoder):
        self.encoder = encoder
        self.classifier = LogisticRegression()

    def training_step(self, batch, batch_idx):
        (x, y) = batch
        feats = self.encoder(x)
        y_hat = self.classifier(feats)
        loss = cross_entropy_with_logits(y_hat, y)
        return loss

trainer = Trainer(gpus=2)
model = FineTuner(resnet18)
trainer.fit(model)
```

Sometimes this works well, but more often it's better to keep the encoder frozen for a while

## 15.2.2 Freeze then unfreeze

The approach that works best most often is to freeze first then unfreeze later

```
# freeze!
model = CPCV2(encoder='resnet18', pretrained='imagenet128')
resnet18 = model.encoder
resnet18.freeze()

classifier = LogisticRegression()

for epoch in epochs:
    for (x, y) in own_data:
        feats = resnet18(x)
        y_hat = classifier(feats)
        loss = cross_entropy_with_logits(y_hat, y)

    # unfreeze after 10 epochs
    if epoch == 10:
        resnet18.unfreeze()
```

---

**Note:** In practice, unfreezing later works MUCH better.

---

Or in Lightning as a Callback so you don't pollute your research code.

```
class UnFreezeCallback(Callback):

    def on_epoch_end(self, trainer, pl_module):
        if trainer.current_epoch == 10:
            encoder.unfreeze()

trainer = Trainer(gpus=2, callbacks=[UnFreezeCallback()])
model = FineTuner(resnet18)
trainer.fit(model)
```

Unless you still need to mix it into your research code.

```
class FineTuner(pl.LightningModule):

    def __init__(self, encoder):
        self.encoder = encoder
        self.classifier = LogisticRegression()

    def training_step(self, batch, batch_idx):

        # option 1 - (not recommended because it's messy)
        if self.trainer.current_epoch == 10:
            self.encoder.unfreeze()

        (x, y) = batch
        feats = self.encoder(x)
        y_hat = self.classifier(feats)
        loss = cross_entropy_with_logits(y_hat, y)
        return loss

    def on_epoch_end(self, trainer, pl_module):
```

(continues on next page)

(continued from previous page)

```
# a hook is cleaner (but a callback is much better)
if self.trainer.current_epoch == 10:
    self.encoder.unfreeze()
```

### 15.2.3 Hyperparameter search

For finetuning to work well, you should try many versions of the model hyperparameters. Otherwise you're unlikely to get the most value out of your data.

```
learning_rates = [0.01, 0.001, 0.0001]
hidden_dim = [128, 256, 512]

for lr in learning_rates:
    for hd in hidden_dim:
        vae = VAE(hidden_dim=hd, learning_rate=lr)
        trainer = Trainer()
        trainer.fit(vae)
```

---

## 15.3 Train from scratch

If you do have enough data and compute resources, then you could try training from scratch.

```
# get data
train_data = DataLoader(YourDataset)
val_data = DataLoader(YourDataset)

# use any bolts model without pretraining
model = VAE()

# fit!
trainer = Trainer(gpus=2)
trainer.fit(model, train_data, val_data)
```

---

**Note:** For this to work well, make sure you have enough data and time to train these models!

---

## 15.4 For research

What separates bolts from all the other libraries out there is that bolts is built by and used by AI researchers. This means every single bolt is modularized so that it can be easily extended or mixed with arbitrary parts of the rest of the code-base.

### 15.4.1 Extending work

Perhaps a research project requires modifying a part of a known approach. In this case, you're better off only changing that part of a system that is already known to perform well. Otherwise, you risk not implementing the work correctly.

#### Example 1: Changing the prior or approx posterior of a VAE

```
from pl_bolts.models.autoencoders import VAE

class MyVAEFlavor(VAE):

    def init_prior(self, z_mu, z_std):
        P = MyPriorDistribution

        # default is standard normal
        # P = distributions.normal.Normal(loc=torch.zeros_like(z_mu), scale=torch.
        ↪ones_like(z_std))
        return P

    def init_posterior(self, z_mu, z_std):
        Q = MyPosteriorDistribution
        # default is normal(z_mu, z_sigma)
        # Q = distributions.normal.Normal(loc=z_mu, scale=z_std)
        return Q
```

And of course train it with lightning.

```
model = MyVAEFlavor()
trainer = Trainer()
trainer.fit(model)
```

In just a few lines of code you changed something fundamental about a VAE... This means you can iterate through ideas much faster knowing that the bolt implementation and the training loop are CORRECT and TESTED.

If your model doesn't work with the new P, Q, then you can discard that research idea much faster than trying to figure out if your VAE implementation was correct, or if your training loop was correct.

#### Example 2: Changing the generator step of a GAN

```
from pl_bolts.models.gans import GAN

class FancyGAN(GAN):

    def generator_step(self, x):
        # sample noise
        z = torch.randn(x.shape[0], self.hparams.latent_dim)
        z = z.type_as(x)

        # generate images
        self.generated_imgs = self(z)

        # ground truth result (ie: all real)
        real = torch.ones(x.size(0), 1)
        real = real.type_as(x)
        g_loss = self.generator_loss(real)

        tqdm_dict = {'g_loss': g_loss}
        output = OrderedDict({
            'loss': g_loss,
```

(continues on next page)

(continued from previous page)

```
        'progress_bar': tqdm_dict,
        'log': tqdm_dict
    })
    return output
```

### Example 3: Changing the way the loss is calculated in a contrastive self-supervised learning approach

```
from pl_bolts.models.self_supervised import AMDIM

class MyDIM(AMDIM):

    def validation_step(self, batch, batch_nb):
        [img_1, img_2], labels = batch

        # generate features
        r1_x1, r5_x1, r7_x1, r1_x2, r5_x2, r7_x2 = self.forward(img_1, img_2)

        # Contrastive task
        loss, lgt_reg = self.contrastive_task((r1_x1, r5_x1, r7_x1), (r1_x2, r5_x2, r7_x2))
        unsupervised_loss = loss.sum() + lgt_reg

        result = {
            'val_nce': unsupervised_loss
        }
    return result
```

---

## 15.4.2 Importing parts

All the bolts are modular. This means you can also arbitrarily mix and match fundamental blocks from across approaches.

### Example 1: Use the VAE encoder for a GAN as a generator

```
from pl_bolts.models.gans import GAN
from pl_bolts.models.autoencoders.basic_vae import Encoder

class FancyGAN(GAN):

    def init_generator(self, img_dim):
        generator = Encoder(...)
        return generator

trainer = Trainer(...)
trainer.fit(FancyGAN())
```

### Example 2: Use the contrastive task of AMDIM in CPC

```
from pl_bolts.models.self_supervised import AMDIM, CPCV2

default_amdim_task = AMDIM().contrastive_task
model = CPCV2(contrastive_task=default_amdim_task, encoder='cpc_default')
# you might need to modify the cpc encoder depending on what you use
```

### 15.4.3 Compose new ideas

You may also be interested in creating completely new approaches that mix and match all sorts of different pieces together

```
# this model is for illustration purposes, it makes no research sense but it's
# intended to show
# that you can be as creative and expressive as you want.
class MyNewContrastiveApproach(pl.LightningModule):

    def __init__(self):
        super().__init__()

        self.gan = GAN()
        self.vae = VAE()
        self.amdim = AMDIM()
        self.cpc = CPCV2

    def training_step(self, batch, batch_idx):
        (x, y) = batch

        feat_a = self.gan.generator(x)
        feat_b = self.vae.encoder(x)

        unsup_loss = self.amdim(feat_a) + self.cpc(feat_b)

        vae_loss = self.vae._step(batch)
        gan_loss = self.gan.generator_loss(x)

    return unsup_loss + vae_loss + gan_loss
```



## AUTOENCODERS

This section houses autoencoders and variational autoencoders.

---

### 16.1 Basic AE

This is the simplest autoencoder. You can use it like so

```
from pl_bolts.models.autoencoders import AE

model = AE()
trainer = Trainer()
trainer.fit(model)
```

You can override any part of this AE to build your own variation.

```
from pl_bolts.models.autoencoders import AE

class MyAEFlavor(AE):

    def init_encoder(self, hidden_dim, latent_dim, input_width, input_height):
        encoder = YourSuperFancyEncoder(...)
        return encoder

class pl_bolts.models.autoencoders.AE(datamodule=None,      input_channels=1,      in-
                                         input_height=28,   input_width=28,   latent_dim=32,
                                         batch_size=32,     hidden_dim=128,     learn-
                                         ing_rate=0.001,   num_workers=8,   data_dir='!',
                                         **kwargs)
Bases: pytorch_lightning.LightningModule
```

#### Parameters

- **datamodule** (Optional[LightningDataModule]) – the datamodule (train, val, test splits)
- **input\_channels** – num of image channels
- **input\_height** – image height
- **input\_width** – image width
- **latent\_dim** – emb dim for encoder
- **batch\_size** – the batch size

- `hidden_dim` – the encoder dim
  - `learning_rate` – the learning rate
  - `num_workers` – num dataloader workers
  - `data_dir` – where to store data
- 

### 16.1.1 Variational Autoencoders

## 16.2 Basic VAE

Use the VAE like so.

```
from pl_bolts.models.autoencoders import VAE

model = VAE()
trainer = Trainer()
trainer.fit(model)
```

You can override any part of this VAE to build your own variation.

```
from pl_bolts.models.autoencoders import VAE

class MyVAEFlavor(VAE):

    def get_posterior(self, mu, std):
        # do something other than the default
        # P = self.get_distribution(self.prior, loc=torch.zeros_like(mu), scale=torch.
        ↵ones_like(std))

    return P

class pl_bolts.models.autoencoders.VAE(hidden_dim=128,          latent_dim=32,          in-
                                         put_channels=3,          input_width=224,         in-
                                         put_height=224,          batch_size=32,          learn-
                                         ing_rate=0.001,          data_dir='.',          datamodule=None,
                                         num_workers=8,          pretrained=None, **kwargs)
Bases: pytorch_lightning.LightningModule
```

Standard VAE with Gaussian Prior and approx posterior.

Model is available pretrained on different datasets:

Example:

```
# not pretrained
vae = VAE()

# pretrained on imagenet
vae = VAE(pretrained='imagenet')

# pretrained on cifar10
vae = VAE(pretrained='cifar10')
```

### Parameters

- `hidden_dim` (int) – encoder and decoder hidden dims
- `latent_dim` (int) – latenet code dim
- `input_channels` (int) – num of channels of the input image.
- `input_width` (int) – image input width
- `input_height` (int) – image input height
- `batch_size` (int) – the batch size
- `the learning rate` (learning\_rate) –
- `data_dir` (str) – the directory to store data
- `datamodule` (Optional[LightningDataModule]) – The Lightning DataModule
- `pretrained` (Optional[str]) – Load weights pretrained on a dataset



---

CHAPTER  
SEVENTEEN

---

## CLASSIC ML MODELS

This module implements classic machine learning models in PyTorch Lightning, including linear regression and logistic regression. Unlike other libraries that implement these models, here we use PyTorch to enable multi-GPU, multi-TPU and half-precision training.

---

### 17.1 Linear Regression

Linear regression fits a linear model between a real-valued target variable ( $y$ ) and one or more features ( $X$ ). We estimate the regression coefficients  $\beta$  that minimizes the mean squared error between the predicted and true target values.

```
from pl_bolts.models.regression import LinearRegression
import pytorch_lightning as pl
from pl_bolts.datamodules import SklearnDataModule
from sklearn.datasets import load_boston

X, y = load_boston(return_X_y=True)
loaders = SklearnDataModule(X, y)

model = LinearRegression(input_dim=13)
trainer = pl.Trainer()
trainer.fit(model, loaders.train_dataloader(), loaders.val_dataloader())
trainer.test(test_dataloaders=loaders.test_dataloader())

class pl_bolts.models.regression.linear_regression.LinearRegression(input_dim,
bias=True,
learning_rate=0.0001,
optimizer=torch.optim.Adam,
l1_strength=None,
l2_strength=None,
**kwargs)
Bases: pytorch_lightning.LightningModule

Linear regression model implementing - with optional L1/L2 regularization  $\min\{W\| (Wx + b) - y \|_2^2$ 
$$
```

**Parameters**

- **input\_dim** (int) – number of dimensions of the input (1+)

- **bias** (bool) – If false, will not use  $+b$
  - **learning\_rate** (float) – learning\_rate for the optimizer
  - **optimizer** (Optimizer) – the optimizer to use (default='Adam')
  - **l1\_strength** (Optional[float]) – L1 regularization strength (default=None)
  - **l2\_strength** (Optional[float]) – L2 regularization strength (default=None)
- 

## 17.2 Logistic Regression

Logistic regression is a non-linear model used for classification, i.e. when we have a categorical target variable. This implementation supports both binary and multi-class classification.

To leverage autograd we think of logistic regression as a one-layer neural network with a sigmoid activation. This allows us to support training on GPUs and TPUs.

```
from sklearn.datasets import load_iris
from pl_bolts.models.regression import LogisticRegression
from pl_bolts.datamodules import SklearnDataModule
import pytorch_lightning as pl

# use any numpy or sklearn dataset
X, y = load_iris(return_X_y=True)
dm = SklearnDataModule(X, y)

# build model
model = LogisticRegression(input_dim=4, num_classes=3)

# fit
trainer = pl.Trainer(tpu_cores=8, precision=16)
trainer.fit(model, dm.train_dataloader(), dm.val_dataloader())

trainer.test(test_dataloaders=dm.test_dataloader(batch_size=12))
```

Any input will be flattened across all dimensions except the first one (batch). This means images, sound, etc... work out of the box.

```
# create dataset
dm = MNISTDataModule(num_workers=0, data_dir=tmpdir)

model = LogisticRegression(input_dim=28 * 28, num_classes=10, learning_rate=0.001)
model.prepare_data = dm.prepare_data
model.train_dataloader = dm.train_dataloader
model.val_dataloader = dm.val_dataloader
model.test_dataloader = dm.test_dataloader

trainer = pl.Trainer(max_epochs=2)
trainer.fit(model)
trainer.test(model)
# {test_acc: 0.92}
```

```
class pl_bolts.models.regression.logistic_regression.LogisticRegression(input_dim,
                           num_classes,
                           bias=True,
                           learning_rate=0.0001,
                           op-
                           ti-
                           mizer=torch.optim.Adam,
                           l1_strength=0.0,
                           l2_strength=0.0,
                           **kwargs)
```

Bases: `pytorch_lightning.LightningModule`

Logistic regression model

#### Parameters

- `input_dim` (int) – number of dimensions of the input (at least 1)
- `num_classes` (int) – number of class labels (binary: 2, multi-class: >2)
- `bias` (bool) – specifies if a constant or intercept should be fitted (equivalent to `fit_intercept` in sklearn)
- `learning_rate` (float) – learning\_rate for the optimizer
- `optimizer` (Optimizer) – the optimizer to use (default='Adam')
- `l1_strength` (float) – L1 regularization strength (default=None)
- `l2_strength` (float) – L2 regularization strength (default=None)



## CONVOLUTIONAL ARCHITECTURES

This package lists contributed convolutional architectures.

---

### 18.1 GPT-2

```
class pl_bolts.models.vision.image_gpt.gpt2.GPT2(embed_dim,      heads,      layers,
                                                num_positions,      vocab_size,
                                                num_classes)
```

Bases: `pytorch_lightning.LightningModule`

GPT-2 from `language Models` are Unsupervised Multitask Learners

Paper by: Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever

Implementation contributed by:

- Teddy Koker

Example:

```
from pl_bolts.models import GPT2

seq_len = 17
batch_size = 32
vocab_size = 16
x = torch.randint(0, vocab_size, (seq_len, batch_size))
model = GPT2(embed_dim=32, heads=2, layers=2, num_positions=seq_len, vocab_
             size=vocab_size, num_classes=4)
results = model(x)
```

`forward(x, classify=False)`

Expect input as shape [sequence len, batch] If classify, return classification logits

---

## 18.2 Image GPT

```
class pl_bolts.models.vision.image_gpt.igpt_module.ImageGPT(datamodule=None,
    embed_dim=16,
    heads=2,      layers=2,      pixels=28,
    vocab_size=16,
    num_classes=10,
    classify=False,
    batch_size=64,
    learning_rate=0.01,
    steps=25000,
    data_dir='.',
    num_workers=8,
    **kwargs)
```

Bases: pytorch\_lightning.LightningModule

**Paper:** Generative Pretraining from Pixels [original paper [code](#)].

**Paper by:** Mark Che, Alec Radford, Rewon Child, Jeff Wu, Heewoo Jun, Prafulla Dhariwal, David Luan, Ilya Sutskever

**Implementation contributed by:**

- Teddy Koker

**Original repo with results and more implementation details:**

- <https://github.com/teddykoker/image-gpt>

**Example Results (Photo credits: Teddy Koker):**



**Default arguments:**

Table 1: Argument Defaults

Argument	Default	iGPT-S (Chen et al.)
<code>-embed_dim</code>	16	512
<code>-heads</code>	2	8
<code>-layers</code>	8	24
<code>-pixels</code>	28	32
<code>-vocab_size</code>	16	512
<code>-num_classes</code>	10	10
<code>-batch_size</code>	64	128
<code>-learning_rate</code>	0.01	0.01
<code>-steps</code>	25000	1000000

Example:

```
import pytorch_lightning as pl
from pl_bolts.models.vision import ImageGPT

dm = MNISTDataModule('.')
model = ImageGPT(dm)

pl.Trainer(gpu=4).fit(model)
```

As script:

```
cd pl_bolts/models/vision/image_gpt
python igpt_module.py --learning_rate 1e-2 --batch_size 32 --gpus 4
```

## Parameters

- `datamodule` (Optional[LightningDataModule]) – LightningDataModule
- `embed_dim` (int) – the embedding dim
- `heads` (int) – number of attention heads
- `layers` (int) – number of layers
- `pixels` (int) – number of input pixels
- `vocab_size` (int) – vocab size
- `num_classes` (int) – number of classes in the input
- `classify` (bool) – true if should classify
- `batch_size` (int) – the batch size
- `learning_rate` (float) – learning rate
- `steps` (int) – number of steps for cosine annealing
- `data_dir` (str) – where to store data
- `num_workers` (int) – num\_data workers

## 18.3 Pixel CNN

```
class pl_bolts.models.vision.pixel_cnn.PixelCNN(input_channels, hid-  
den_channels=256, num_blocks=5)  
Bases: torch.nn.Module
```

Implementation of Pixel CNN.

Paper authors: Aaron van den Oord, Nal Kalchbrenner, Oriol Vinyals, Lasse Espeholt, Alex Graves, Koray Kavukcuoglu

Implemented by:

- William Falcon

Example:

```
>>> from pl_bolts.models.vision import PixelCNN  
>>> import torch  
...  
>>> model = PixelCNN(input_channels=3)  
>>> x = torch.rand(5, 3, 64, 64)  
>>> out = model(x)  
...  
>>> out.shape  
torch.Size([5, 3, 64, 64])
```

---

CHAPTER  
**NINETEEN**

---

**GANS**

Collection of Generative Adversarial Networks

---

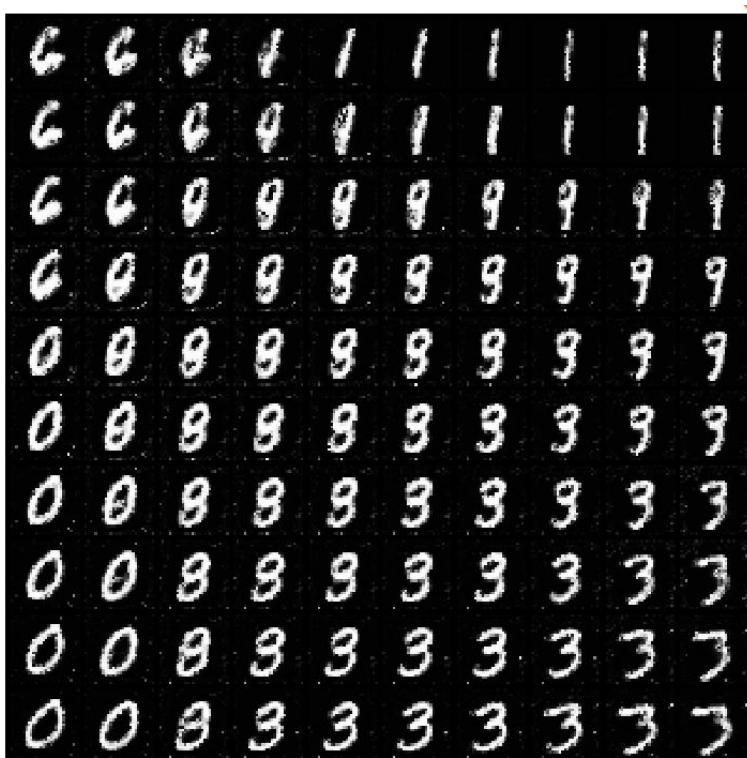
## 19.1 Basic GAN

This is a vanilla GAN. This model can work on any dataset size but results are shown for MNIST. Replace the encoder, decoder or any part of the training loop to build a new method, or simply finetune on your data.

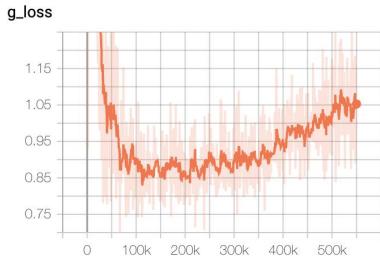
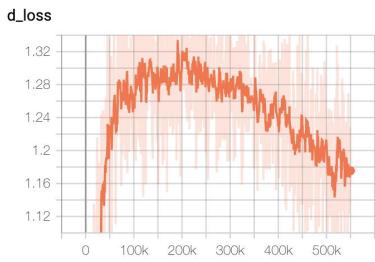
Implemented by:

- William Falcon

Example outputs:



Loss curves:



```
from pl_bolts.models.gans import GAN
...
gan = GAN()
trainer = Trainer()
trainer.fit(gan)
```

**class** pl\_bolts.models.gans.GAN(*datamodule=None*, *latent\_dim=32*, *batch\_size=100*, *learning\_rate=0.0002*, *data\_dir=""*, *num\_workers=8*, *\*\*kwargs*)  
Bases: pytorch\_lightning.LightningModule

Vanilla GAN implementation.

Example:

```
from pl_bolts.models.gan import GAN

m = GAN()
Trainer(gpus=2).fit(m)
```

Example CLI:

```
# mnist
python basic_gan_module.py --gpus 1

# imagenet
python basic_gan_module.py --gpus 1 --dataset 'imagenet2012'
--data_dir /path/to/imagenet/folder/ --meta_dir ~/path/to/meta/bin/folder
--batch_size 256 --learning_rate 0.0001
```

## Parameters

- **datamodule** (Optional[LightningDataModule]) – the datamodule (train, val, test splits)
- **latent\_dim** (int) – emb dim for encoder
- **batch\_size** (int) – the batch size
- **learning\_rate** (float) – the learning rate
- **data\_dir** (str) – where to store data

- **num\_workers** `(int)` – data workers

**forward** (`z`)

Generates an image given input noise `z`

Example:

```
z = torch.rand(batch_size, latent_dim)
gan = GAN.load_from_checkpoint(PATH)
img = gan(z)
```



## SELF-SUPERVISED LEARNING

This bolts module houses a collection of all self-supervised learning models.

Self-supervised learning extracts representations of an input by solving a pretext task. In this package, we implement many of the current state-of-the-art self-supervised algorithms.

Self-supervised models are trained with unlabeled datasets

---

### 20.1 Use cases

Here are some use cases for the self-supervised package.

#### 20.1.1 Extracting image features

The models in this module are trained unsupervised and thus can capture better image representations (features).

In this example, we'll load a resnet 18 which was pretrained on imagenet using CPC as the pretext task.

Example:

```
from pl_bolts.models.self_supervised import CPCV2

# load resnet18 pretrained using CPC on imagenet
model = CPCV2(pretrained='resnet18')
cpc_resnet18 = model.encoder
cpc_resnet18.freeze()

# it supports any torchvision resnet
model = CPCV2(pretrained='resnet50')
```

This means you can now extract image representations that were pretrained via unsupervised learning.

Example:

```
my_dataset = SomeDataset()
for batch in my_dataset:
    x, y = batch
    out = cpc_resnet18(x)
```

---

## 20.1.2 Train with unlabeled data

These models are perfect for training from scratch when you have a huge set of unlabeled images

```
from pl_bolts.models.self_supervised import SimCLR
from pl_bolts.models.self_supervised.simclr import SimCLREvalDataTransform, ▶
    SimCLRTrainDataTransform

train_dataset = MyDataset(transforms=SimCLRTrainDataTransform())
val_dataset = MyDataset(transforms=SimCLREvalDataTransform())

# simclr needs a lot of compute!
model = SimCLR()
trainer = Trainer(tpu_cores=128)
trainer.fit(
    model,
    DataLoader(train_dataset),
    DataLoader(val_dataset),
)
```

## 20.1.3 Research

Mix and match any part, or subclass to create your own new method

```
from pl_bolts.models.self_supervised import CPCV2
from pl_bolts.losses.self_supervised_learning import FeatureMapContrastiveTask

amdim_task = FeatureMapContrastiveTask(comparisons='01, 11, 02', bidirectional=True)
model = CPCV2(contrastive_task=amdim_task)
```

---

## 20.2 Contrastive Learning Models

Contrastive self-supervised learning (CSL) is a self-supervised learning approach where we generate representations of instances such that similar instances are near each other and far from dissimilar ones. This is often done by comparing triplets of positive, anchor and negative representations.

In this section, we list Lightning implementations of popular contrastive learning approaches.

### 20.2.1 AMDIM

```
class pl_bolts.models.self_supervised.AMDIM(datamodule='cifar10',
                                              encoder='amdim_encoder',
                                              contrastive_task=torch.nn.Module,
                                              image_channels=3,
                                              image_height=32,
                                              encoder_feature_dim=320,
                                              embedding_fx_dim=1280,
                                              conv_block_depth=10,
                                              use_bn=False,
                                              tclip=20.0,
                                              learning_rate=0.0002,
                                              num_classes=10,
                                              data_dir='',
                                              batch_size=200,
                                              **kwargs)
```

Bases: `pytorch_lightning.LightningModule`

PyTorch Lightning implementation of Augmented Multiscale Deep InfoMax (AMDIM)

Paper authors: Philip Bachman, R Devon Hjelm, William Buchwalter.

Model implemented by: [William Falcon](#)

This code is adapted to Lightning using the original author repo ([the original repo](#)).

## Example

```
>>> from pl_bolts.models.self_supervised import AMDIM
...
>>> model = AMDIM(encoder='resnet18')
```

Train:

```
trainer = Trainer()
trainer.fit(model)
```

## Parameters

- **datamodule** (Union[str, LightningDataModule]) – A LightningDatamodule
- **encoder** (Union[str, Module, LightningModule]) – an encoder string or model
- **image\_channels** (int) – 3
- **image\_height** (int) – pixels
- **encoder\_feature\_dim** (int) – Called *ndf* in the paper, this is the representation size for the encoder.
- **embedding\_fx\_dim** (int) – Output dim of the embedding function (*nrkhs* in the paper) (Reproducing Kernel Hilbert Spaces).
- **conv\_block\_depth** (int) – Depth of each encoder block,
- **use\_bn** (bool) – If true will use batchnorm.
- **tclip** (int) – soft clipping non-linearity to the scores after computing the regularization term and before computing the log-softmax. This is the ‘second trick’ used in the paper
- **learning\_rate** (int) – The learning rate
- **data\_dir** (str) – Where to store data
- **num\_classes** (int) – How many classes in the dataset
- **batch\_size** (int) – The batch size

## 20.2.2 BYOL

```
class pl_bolts.models.self_supervised.BYOL(num_classes, learning_rate=0.2,
                                             weight_decay=1.5e-05, input_height=32,
                                             batch_size=32, num_workers=0,
                                             warmup_epochs=10, max_epochs=1000,
                                             **kwargs)
```

Bases: pytorch\_lightning.LightningModule

PyTorch Lightning implementation of Bring Your Own Latent (BYOL)

Paper authors: Jean-Bastien Grill ,Florian Strub, Florent Altché, Corentin Tallec, Pierre H. Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Daniel Guo, Mohammad Gheshlaghi Azar, Bilal Piot, Koray Kavukcuoglu, Rémi Munos, Michal Valko.

**Model implemented by:**

- Annika Brundyn

**Warning:** Work in progress. This implementation is still being verified.

**TODOs:**

- verify on CIFAR-10
- verify on STL-10
- pre-train on imagenet

Example:

```
import pytorch_lightning as pl
from pl_bolts.models.self_supervised import BYOL
from pl_bolts.datamodules import CIFAR10DataModule
from pl_bolts.models.self_supervised.simclr.simclr_transforms import (
    SimCLREvalDataTransform, SimCLRTrainDataTransform)

# model
model = BYOL(num_classes=10)

# data
dm = CIFAR10DataModule(num_workers=0)
dm.train_transforms = SimCLRTrainDataTransform(32)
dm.val_transforms = SimCLREvalDataTransform(32)

trainer = pl.Trainer()
trainer.fit(model, dm)
```

Train:

```
trainer = Trainer()
trainer.fit(model)
```

CLI command:

```
# cifar10
python byol_module.py --gpus 1
```

(continues on next page)

(continued from previous page)

```
# imagenet
python byol_module.py
--gpus 8
--dataset imagenet2012
--data_dir /path/to/imagenet/
--meta_dir /path/to/folder/with/meta.bin/
--batch_size 32
```

### Parameters

- **datamodule** – The datamodule
- **learning\_rate** – the learning rate
- **weight\_decay** – optimizer weight decay
- **input\_height** – image input height
- **batch\_size** – the batch size
- **num\_workers** – number of workers
- **warmup\_epochs** – num of epochs for scheduler warm up
- **max\_epochs** – max epochs for scheduler

### 20.2.3 CPC (V2)

```
class pl_bolts.models.self_supervised.CPCV2(datamodule=None, encoder='cpc_encoder',
                                             patch_size=8, patch_overlap=4, online_ft=True,
                                             task='cpc', num_workers=4, learning_rate=0.0001,
                                             data_dir='', batch_size=32, pretrained=None,
                                             **kwargs)
```

Bases: `pytorch_lightning.LightningModule`

PyTorch Lightning implementation of Data-Efficient Image Recognition with Contrastive Predictive Coding

Paper authors: (Olivier J. Hénaff, Aravind Srinivas, Jeffrey De Fauw, Ali Razavi, Carl Doersch, S. M. Ali Eslami, Aaron van den Oord).

Model implemented by:

- William Falcon
- Tullie Murrell

## Example

```
>>> from pl_bolts.models.self_supervised import CPCV2
...
>>> model = CPCV2()
```

Train:

```
trainer = Trainer()
trainer.fit(model)
```

CLI command:

```
# cifar10
python cpc_module.py --gpus 1

# imagenet
python cpc_module.py
--gpus 8
--dataset imagenet2012
--data_dir /path/to/imagenet/
--meta_dir /path/to/folder/with/meta.bin/
--batch_size 32
```

To Finetune:

```
python cpc_finetuner.py --ckpt_path path/to/checkpoint.ckpt --dataset cifar10 --
--gpus x
```

Some uses:

```
# load resnet18 pretrained using CPC on imagenet
model = CPCV2(encoder='resnet18', pretrained=True)
resnet18 = model.encoder
resnet18.freeze()

# it supports any torchvision resnet
model = CPCV2(encoder='resnet50', pretrained=True)

# use it as a feature extractor
x = torch.rand(2, 3, 224, 224)
out = model(x)
```

## Parameters

- **datamodule** (Optional[LightningDataModule]) – A Datamodule (optional). Otherwise set the dataloaders directly
- **encoder** (Union[str, Module, LightningModule]) – A string for any of the resnets in torchvision, or the original CPC encoder, or a custom nn.Module encoder
- **patch\_size** (int) – How big to make the image patches
- **patch\_overlap** (int) – How much overlap should each patch have.
- **online\_ft** (int) – Enable a 1024-unit MLP to fine-tune online
- **task** (str) – Which self-supervised task to use ('cpc', 'amdim', etc...)
- **num\_workers** (int) – num dataloader workers

- `learning_rate` (int) – what learning rate to use
- `data_dir` (str) – where to store data
- `batch_size` (int) – batch size
- `pretrained` (Optional[str]) – If true, will use the weights pretrained (using CPC) on Imagenet

## 20.2.4 Moco (V2)

```
class pl_bolts.models.self_supervised.MoCoV2(base_encoder='resnet18', emb_dim=128,
                                              num_negatives=65536, encoder_momentum=0.999,
                                              max_temperature=0.07, learning_rate=0.03,
                                              weight_decay=0.0001, momentum=0.9,
                                              datamodule=None, data_dir='./', batch_size=256,
                                              use_mlp=False, num_workers=8, *args,
                                              **kwargs)
```

Bases: `pytorch_lightning.LightningModule`

PyTorch Lightning implementation of `Moco`

Paper authors: Xinlei Chen, Haoqi Fan, Ross Girshick, Kaiming He.

Code adapted from [facebookresearch/moco](#) to Lightning by:

- William Falcon

### Example

```
>>> from pl_bolts.models.self_supervised import MoCoV2
...
>>> model = MoCoV2()
```

Train:

```
trainer = Trainer()
trainer.fit(model)
```

CLI command:

```
# cifar10
python moco2_module.py --gpus 1

# imagenet
python moco2_module.py
--gpus 8
--dataset imagenet2012
--data_dir /path/to/imagenet/
--meta_dir /path/to/folder/with/meta.bin/
--batch_size 32
```

### Parameters

- `base_encoder` (Union[str, Module]) – torchvision model name or torch.nn.Module
- `emb_dim` (int) – feature dimension (default: 128)
- `num_negatives` (int) – queue size; number of negative keys (default: 65536)
- `encoder_momentum` (float) – moco momentum of updating key encoder (default: 0.999)
- `softmax_temperature` (float) – softmax temperature (default: 0.07)
- `learning_rate` (float) – the learning rate
- `momentum` (float) – optimizer momentum
- `weight_decay` (float) – optimizer weight decay
- `datamodule` (Optional[LightningDataModule]) – the DataModule (train, val, test dataloaders)
- `data_dir` (str) – the directory to store data
- `batch_size` (int) – batch size
- `use_mlp` (bool) – add an mlp to the encoders
- `num_workers` (int) – workers for the loaders

`_batch_shuffle_ddp (x)`  
Batch shuffle, for making use of BatchNorm. \* Only support DistributedDataParallel (DDP) model. \*

`_batch_unshuffle_ddp (x, idx_unshuffle)`  
Undo batch shuffle. \* Only support DistributedDataParallel (DDP) model. \*

`_momentum_update_key_encoder ()`  
Momentum update of the key encoder

`forward (img_q, img_k)`

**Input:** im\_q: a batch of query images im\_k: a batch of key images

**Output:** logits, targets

`init_encoders (base_encoder)`  
Override to add your own encoders

---

## 20.2.5 SimCLR

PyTorch Lightning implementation of SimCLR

Paper authors: Ting Chen, Simon Kornblith, Mohammad Norouzi, Geoffrey Hinton.

Model implemented by:

- William Falcon
- Tullie Murrell

To Train:

```

import pytorch_lightning as pl
from pl_bolts.models.self_supervised import SimCLR
from pl_bolts.datamodules import CIFAR10DataModule
from pl_bolts.models.self_supervised.simclr.simclr_transforms import (
    SimCLREvalDataTransform, SimCLRTrainDataTransform)

# data
dm = CIFAR10DataModule(num_workers=0)
dm.train_transforms = SimCLRTrainDataTransform(32)
dm.val_transforms = SimCLREvalDataTransform(32)

# model
model = SimCLR(num_samples=dm.num_samples, batch_size=dm.batch_size)

# fit
trainer = pl.Trainer()
trainer.fit(model, dm)

```

## CIFAR-10 baseline

Table 1: Cifar-10 implementation results

Implementation	test acc	Encoder	Optimizer	Batch	Epochs	Hardware	LR
Original	92.00?	resnet50	LARS	512	1000	1 V100 (32GB)	1.0
Ours	85.68	resnet50	LARS	512	960 (12 hr)	1 V100 (32GB)	1e-6

CIFAR-10 pretrained model:

```

from pl_bolts.models.self_supervised import SimCLR

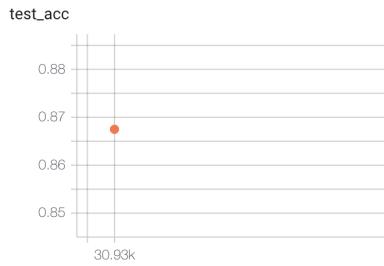
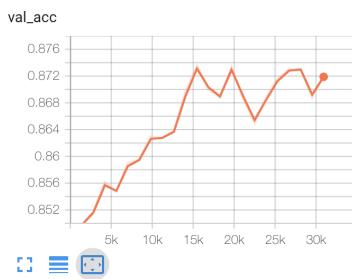
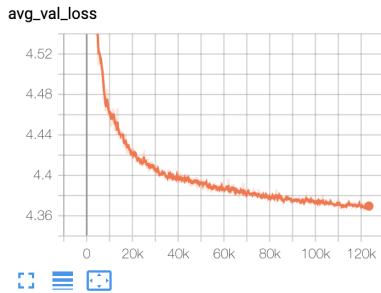
weight_path = 'https://pl-bolts-weights.s3.us-east-2.amazonaws.com/simclr/simclr-'
            'cifar10-v1-exp12_87_52/epoch%3D960.ckpt'
simclr = SimCLR.load_from_checkpoint(weight_path, strict=False)

simclr.freeze()

```

Pre-training:

Fine-tuning (Single layer MLP, 1024 hidden units):



To reproduce:

```
# pretrain
python simclr_module.py
--gpus 1
--dataset cifar10
--batch_size 512
--learning_rate 1e-06
--num_workers 8

# finetune
python simclr_finetuner.py
--ckpt_path path/to/epoch=xyz.ckpt
--gpus 1
```

## SimCLR API

```
class pl_bolts.models.self_supervised.SimCLR(batch_size, num_samples,
                                              warmup_epochs=10, lr=0.0001,
                                              opt_weight_decay=1e-06,
                                              loss_temperature=0.5, **kwargs)
```

Bases: pytorch\_lightning.LightningModule

### Parameters

- **batch\_size** – the batch size
- **num\_samples** – num samples in the dataset
- **warmup\_epochs** – epochs to warmup the lr for
- **lr** – the optimizer learning rate
- **opt\_weight\_decay** – the optimizer weight decay
- **loss\_temperature** – the loss temperature



## SELF-SUPERVISED LEARNING TRANSFORMS

These transforms are used in various self-supervised learning approaches.

---

### 21.1 CPC transforms

Transforms used for CPC

#### 21.1.1 CIFAR-10 Train (c)

```
class pl_bolts.models.self_supervised.cpc.transforms.CPCTrainTransformsCIFAR10(patch_size=8,  
                                over-  
                                lap=4)  
Bases: object
```

Transforms used for CPC:

##### Parameters

- **patch\_size** – size of patches when cutting up the image into overlapping patches
- **overlap** – how much to overlap patches

Transforms:

```
random_flip  
img_jitter  
col_jitter  
rnd_gray  
transforms.ToTensor()  
normalize  
Patchify(patch_size=patch_size, overlap_size=patch_size // 2)
```

Example:

```
# in a regular dataset  
CIFAR10(..., transforms=CPCTrainTransformsCIFAR10())  
  
# in a DataModule  
module = CIFAR10DataModule(PATH)  
train_loader = module.train_dataloader(batch_size=32,   
                                       transforms=CPCTrainTransformsCIFAR10())
```

```
__call__(inp)
Call self as a function.
```

### 21.1.2 CIFAR-10 Eval (c)

```
class pl_bolts.models.self_supervised.cpc.transforms.CPCEvalTransformsCIFAR10(patch_size=8,
over-
lap=4)
```

Bases: `object`

Transforms used for CPC:

#### Parameters

- `patch_size` – size of patches when cutting up the image into overlapping patches
- `overlap` – how much to overlap patches

Transforms:

```
random_flip
transforms.ToTensor()
normalize
Patchify(patch_size=patch_size, overlap_size=overlap)
```

Example:

```
# in a regular dataset
CIFAR10(..., transforms=CPCEvalTransformsCIFAR10())

# in a DataModule
module = CIFAR10DataModule(PATH)
train_loader = module.train_dataloader(batch_size=32,_
transforms=CPCEvalTransformsCIFAR10())
```

```
__call__(inp)
```

Call self as a function.

### 21.1.3 Imagenet Train (c)

```
class pl_bolts.models.self_supervised.cpc.transforms.CPCTrainTransformsImageNet128(patch_size-
over-
lap=16)
```

Bases: `object`

Transforms used for CPC:

#### Parameters

- `patch_size` – size of patches when cutting up the image into overlapping patches
- `overlap` – how much to overlap patches

Transforms:

```
random_flip
transforms.ToTensor()
normalize
Patchify(patch_size=patch_size, overlap_size=patch_size // 2)
```

Example:

```
# in a regular dataset
Imagenet(..., transforms=CPCTrainTransformsImageNet128())

# in a DataModule
module = ImagenetDataModule(PATH)
train_loader = module.train_dataloader(batch_size=32, ↴
transforms=CPCTrainTransformsImageNet128())
```

### `__call__(inp)`

Call self as a function.

## 21.1.4 Imagenet Eval (c)

```
class pl_bolts.models.self_supervised.cpc.transforms.CPCEvalTransformsImageNet128(patch_size=..., over-
lap=16)
```

Bases: `object`

Transforms used for CPC:

### Parameters

- `patch_size` – size of patches when cutting up the image into overlapping patches
- `overlap` – how much to overlap patches

Transforms:

```
random_flip
transforms.ToTensor()
normalize
Patchify(patch_size=patch_size, overlap_size=patch_size // 2)
```

Example:

```
# in a regular dataset
Imagenet(..., transforms=CPCEvalTransformsImageNet128())

# in a DataModule
module = ImagenetDataModule(PATH)
train_loader = module.train_dataloader(batch_size=32, ↴
transforms=CPCEvalTransformsImageNet128())
```

### `__call__(inp)`

Call self as a function.

## 21.1.5 STL-10 Train (c)

```
class pl_bolts.models.self_supervised.cpc.transforms.CPCTrainTransformsSTL10(patch_size=16,  
                                                               over-  
                                                               lap=8)
```

Bases: `object`

Transforms used for CPC:

### Parameters

- `patch_size` – size of patches when cutting up the image into overlapping patches
- `overlap` – how much to overlap patches

Transforms:

```
random_flip  
img_jitter  
col_jitter  
rnd_gray  
transforms.ToTensor()  
normalize  
Patchify(patch_size=patch_size, overlap_size=patch_size // 2)
```

Example:

```
# in a regular dataset  
STL10(..., transforms=CPCTrainTransformsSTL10())  
  
# in a DataModule  
module = STL10DataModule(PATH)  
train_loader = module.train_dataloader(batch_size=32, ▾  
                                         transforms=CPCTrainTransformsSTL10())
```

### `__call__(inp)`

Call self as a function.

## 21.1.6 STL-10 Eval (c)

```
class pl_bolts.models.self_supervised.cpc.transforms.CPCEvalTransformsSTL10(patch_size=16,  
                                                               over-  
                                                               lap=8)
```

Bases: `object`

Transforms used for CPC:

### Parameters

- `patch_size` – size of patches when cutting up the image into overlapping patches
- `overlap` – how much to overlap patches

Transforms:

```
random_flip  
transforms.ToTensor()  
normalize  
Patchify(patch_size=patch_size, overlap_size=patch_size // 2)
```

Example:

```
# in a regular dataset
STL10(..., transforms=CPCEvalTransformsSTL10())

# in a DataModule
module = STL10DataModule(PATH)
train_loader = module.train_dataloader(batch_size=32, ↴
                                         transforms=CPCEvalTransformsSTL10())
```

**\_\_call\_\_(inp)**  
Call self as a function.

## 21.2 AMDIM transforms

Transforms used for AMDIM

### 21.2.1 CIFAR-10 Train (a)

```
class pl_bolts.models.self_supervised.amdim.transforms.AMDIMTrainTransformsCIFAR10
Bases: object
```

Transforms applied to AMDIM

Transforms:

```
img_jitter,
col_jitter,
rnd_gray,
transforms.ToTensor(),
normalize
```

Example:

```
x = torch.rand(5, 3, 32, 32)

transform = AMDIMTrainTransformsCIFAR10()
(view1, view2) = transform(x)
```

**\_\_call\_\_(inp)**  
Call self as a function.

### 21.2.2 CIFAR-10 Eval (a)

```
class pl_bolts.models.self_supervised.amdim.transforms.AMDIMEvalTransformsCIFAR10
Bases: object
```

Transforms applied to AMDIM

Transforms:

```
transforms.ToTensor(),
normalize
```

Example:

```
x = torch.rand(5, 3, 32, 32)

transform = AMDIMEvalTransformsCIFAR10()
(view1, view2) = transform(x)
```

**\_\_call\_\_(inp)**  
Call self as a function.

### 21.2.3 Imagenet Train (a)

**class** pl\_bolts.models.self\_supervised.amdim.transforms.**AMDIMTrainTransformsImageNet128** (*height*)  
Bases: **object**

Transforms applied to AMDIM

Transforms:

```
img_jitter,
col_jitter,
rnd_gray,
transforms.ToTensor(),
normalize
```

Example:

```
x = torch.rand(5, 3, 128, 128)

transform = AMDIMTrainTransformsSTL10()
(view1, view2) = transform(x)
```

**\_\_call\_\_(inp)**  
Call self as a function.

### 21.2.4 Imagenet Eval (a)

**class** pl\_bolts.models.self\_supervised.amdim.transforms.**AMDIMEvalTransformsImageNet128** (*height*)  
Bases: **object**

Transforms applied to AMDIM

Transforms:

```
transforms.Resize(height + 6, interpolation=3),
transforms.CenterCrop(height),
transforms.ToTensor(),
normalize
```

Example:

```
x = torch.rand(5, 3, 128, 128)

transform = AMDIMEvalTransformsImageNet128()
view1 = transform(x)
```

**\_\_call\_\_(inp)**  
Call self as a function.

## 21.2.5 STL-10 Train (a)

```
class pl_bolts.models.self_supervised.amdim.transforms.AMDIMTrainTransformsSTL10 (height=64)
Bases: object
```

Transforms applied to AMDIM

Transforms:

```
img_jitter,
col_jitter,
rnd_gray,
transforms.ToTensor(),
normalize
```

Example:

```
x = torch.rand(5, 3, 64, 64)

transform = AMDIMTrainTransformsSTL10()
(view1, view2) = transform(x)
```

**\_\_call\_\_(*inp*)**

Call self as a function.

## 21.2.6 STL-10 Eval (a)

```
class pl_bolts.models.self_supervised.amdim.transforms.AMDIMEvalTransformsSTL10 (height=64)
Bases: object
```

Transforms applied to AMDIM

Transforms:

```
transforms.Resize(height + 6, interpolation=3),
transforms.CenterCrop(height),
transforms.ToTensor(),
normalize
```

Example:

```
x = torch.rand(5, 3, 64, 64)

transform = AMDIMTrainTransformsSTL10()
view1 = transform(x)
```

**\_\_call\_\_(*inp*)**

Call self as a function.

## 21.3 MOCO V2 transforms

Transforms used for MOCO V2

### 21.3.1 CIFAR-10 Train (m2)

```
class pl_bolts.models.self_supervised.moco.transforms.Moco2TrainCIFAR10Transforms (height=32)
Bases: object
```

Moco 2 augmentation: <https://arxiv.org/pdf/2003.04297.pdf>

```
__call__(inp)
```

Call self as a function.

### 21.3.2 CIFAR-10 Eval (m2)

```
class pl_bolts.models.self_supervised.moco.transforms.Moco2EvalCIFAR10Transforms (height=32)
Bases: object
```

Moco 2 augmentation: <https://arxiv.org/pdf/2003.04297.pdf>

```
__call__(inp)
```

Call self as a function.

### 21.3.3 Imagenet Train (m2)

```
class pl_bolts.models.self_supervised.moco.transforms.Moco2TrainSTL10Transforms (height=64)
Bases: object
```

Moco 2 augmentation: <https://arxiv.org/pdf/2003.04297.pdf>

```
__call__(inp)
```

Call self as a function.

### 21.3.4 Imagenet Eval (m2)

```
class pl_bolts.models.self_supervised.moco.transforms.Moco2EvalSTL10Transforms (height=64)
Bases: object
```

Moco 2 augmentation: <https://arxiv.org/pdf/2003.04297.pdf>

```
__call__(inp)
```

Call self as a function.

### 21.3.5 STL-10 Train (m2)

```
class pl_bolts.models.self_supervised.moco.transforms.Moco2TrainImagenetTransforms (height=128)
    Bases: object

    Moco 2 augmentation: https://arxiv.org/pdf/2003.04297.pdf

    __call__ (inp)
        Call self as a function.
```

### 21.3.6 STL-10 Eval (m2)

```
class pl_bolts.models.self_supervised.moco.transforms.Moco2EvalImagenetTransforms (height=128)
    Bases: object

    Moco 2 augmentation: https://arxiv.org/pdf/2003.04297.pdf

    __call__ (inp)
        Call self as a function.
```

---

## 21.4 SimCLR transforms

Transforms used for SimCLR

### 21.4.1 Train (sc)

```
class pl_bolts.models.self_supervised.simclr.simclr_transforms.SimCLRTrainDataTransform (input_height=32, input_width=32)
    Bases: object
```

Transforms for SimCLR

Transform:

```
RandomResizedCrop(size=self.input_height)
RandomHorizontalFlip()
RandomApply([color_jitter], p=0.8)
RandomGrayscale(p=0.2)
GaussianBlur(kernel_size=int(0.1 * self.input_height))
transforms.ToTensor()
```

Example:

```
from pl_bolts.models.self_supervised.simclr.transforms import SimCLRTrainDataTransform

transform = SimCLRTrainDataTransform(input_height=32)
x = sample()
(xi, xj) = transform(x)
```

```
__call__ (sample)
    Call self as a function.
```

## 21.4.2 Eval (sc)

```
class pl_bolts.models.self_supervised.simclr.simclr_transforms.SimCLREvalDataTransform(input_
s=1)

Bases: object
```

Transforms for SimCLR

Transform:

```
Resize(input_height + 10, interpolation=3)
transforms.CenterCrop(input_height),
transforms.ToTensor()
```

Example:

```
from pl_bolts.models.self_supervised.simclr.transforms import_
SimCLREvalDataTransform

transform = SimCLREvalDataTransform(input_height=32)
x = sample()
(xi, xj) = transform(x)
```

```
__call__(sample)
Call self as a function.
```

## SELF-SUPERVISED LEARNING

Collection of useful functions for self-supervised learning

---

### 22.1 Identity class

Example:

```
from pl_bolts.utils import Identity

class pl_bolts.utils.self_supervised.Identity
    Bases: torch.nn.Module

An identity class to replace arbitrary layers in pretrained models
```

Example:

```
from pl_bolts.utils import Identity

model = resnet18()
model.fc = Identity()
```

---

### 22.2 SSL-ready resnets

Torchvision resnets with the fc layers removed and with the ability to return all feature maps instead of just the last one.

Example:

```
from pl_bolts.utils.self_supervised import torchvision_ssl_encoder

resnet = torchvision_ssl_encoder('resnet18', pretrained=False, return_all_feature_
    ↴maps=True)
x = torch.rand(3, 3, 32, 32)

feat_maps = resnet(x)
```

pl\_bolts.utils.self\_supervised.**torchvision\_ssl\_encoder**(*name*, *pretrained=False*, *re-*  
*turn\_all\_feature\_maps=False*)

---

## 22.3 SSL backbone finetuner

```
class pl_bolts.models.self_supervised.ssl_finetuner.SSLFineTuner(backbone,
                                                               in_features,
                                                               num_classes,
                                                               hid-
                                                               den_dim=1024)
```

Bases: `pytorch_lightning.LightningModule`

Finetunes a self-supervised learning backbone using the standard evaluation protocol of a singler layer MLP with 1024 units

Example:

```
from pl_bolts.utils.self_supervised import SSLFineTuner
from pl_bolts.models.self_supervised import CPCV2
from pl_bolts.datamodules import CIFAR10DataModule
from pl_bolts.models.self_supervised.cpc.transforms import_
    CPCEvalTransformsCIFAR10,
    CPCTrainTransformsCIFAR10

# pretrained model
backbone = CPCV2.load_from_checkpoint(PATH, strict=False)

# dataset + transforms
dm = CIFAR10DataModule(data_dir='.')
dm.train_transforms = CPCTrainTransformsCIFAR10()
dm.val_transforms = CPCEvalTransformsCIFAR10()

# finetuner
finetuner = SSLFineTuner(backbone, in_features=backbone.z_dim, num_
    ↪classes=backbone.num_classes)

# train
trainer = pl.Trainer()
trainer.fit(finetuner, dm)

# test
trainer.test(datamodule=dm)
```

### Parameters

- `backbone` – a pretrained model
- `in_features` – feature dim of backbone outputs
- `num_classes` – classes of the dataset
- `hidden_dim` – dim of the MLP (1024 default used in self-supervised literature)

## SEMI-SUPERVISED LEARNING

Collection of utilities for semi-supervised learning where some part of the data is labeled and the other part is not.

---

### 23.1 Balanced classes

Example:

```
from pl_bolts.utils.semi_supervised import balance_classes
```

pl\_bolts.utils.semi\_supervised.**balance\_classes**(*X, Y, batch\_size*)

Makes sure each batch has an equal amount of data from each class. Perfect balance

#### Parameters

- **X** (ndarray) – input features
- **Y** (list) – mixed labels (ints)
- **batch\_size** (int) – the ultimate batch size

### 23.2 half labeled batches

Example:

```
from pl_bolts.utils.semi_supervised import balance_classes
```

pl\_bolts.utils.semi\_supervised.**generate\_half\_labeled\_batches**(*smaller\_set\_X,  
smaller\_set\_Y,  
larger\_set\_X,  
larger\_set\_Y,  
batch\_size*)

Given a labeled dataset and an unlabeled dataset, this function generates a joint pair where half the batches are labeled and the other half is not



---

CHAPTER  
**TWENTYFOUR**

---

## SELF-SUPERVISED LEARNING CONTRASTIVE TASKS

This section implements popular contrastive learning tasks used in self-supervised learning.

---

### 24.1 FeatureMapContrastiveTask

This task compares sets of feature maps.

In general the feature map comparison pretext task uses triplets of features. Here are the abstract steps of comparison.

Generate multiple views of the same image

```
x1_view_1 = data_augmentation(x1)
x1_view_2 = data_augmentation(x1)
```

Use a different example to generate additional views (usually within the same batch or a pool of candidates)

```
x2_view_1 = data_augmentation(x2)
x2_view_2 = data_augmentation(x2)
```

Pick 3 views to compare, these are the anchor, positive and negative features

```
anchor = x1_view_1
positive = x1_view_2
negative = x2_view_1
```

Generate feature maps for each view

```
(a0, a1, a2) = encoder(anchor)
(p0, p1, p2) = encoder(positive)
```

Make a comparison for a set of feature maps

```
phi = some_score_function()

# the '01' comparison
score = phi(a0, p1)

# and can be bidirectional
score = phi(p0, a1)
```

In practice the contrastive task creates a BxB matrix where B is the batch size. The diagonals for set 1 of feature maps are the anchors, the diagonals of set 2 of the feature maps are the positives, the non-diagonals of set 1 are the negatives.

```
class pl_bolts.losses.self_supervised_learning.FeatureMapContrastiveTask(comparisons='00,
11',
tclip=10.0,
bidi-
rec-
tional=True)
```

Bases: `torch.nn.Module`

Performs an anchor, positive negative pair comparison for each each tuple of feature maps passed.

```
# extract feature maps
pos_0, pos_1, pos_2 = encoder(x_pos)
anc_0, anc_1, anc_2 = encoder(x_anchor)

# compare only the 0th feature maps
task = FeatureMapContrastiveTask('00')
loss, regularizer = task((pos_0), (anc_0))

# compare (pos_0 to anc_1) and (pos_0, anc_2)
task = FeatureMapContrastiveTask('01, 02')
losses, regularizer = task((pos_0, pos_1, pos_2), (anc_0, anc_1, anc_2))
loss = losses.sum()

# compare (pos_1 vs a anc_random)
task = FeatureMapContrastiveTask('0r')
loss, regularizer = task((pos_0, pos_1, pos_2), (anc_0, anc_1, anc_2))
```

### Parameters

- `comparisons` (`str`) – groupings of feature map indices to compare (zero indexed, ‘r’ means random) ex: ‘00, 1r’
- `tclip` (`float`) – stability clipping value
- `bidirectional` (`bool`) – if true, does the comparison both ways

```
# with bidirectional the comparisons are done both ways
task = FeatureMapContrastiveTask('01, 02')

# will compare the following:
# 01: (pos_0, anc_1), (anc_0, pos_1)
# 02: (pos_0, anc_2), (anc_0, pos_2)
```

### forward(anchor\_maps, positive\_maps)

Takes in a set of tuples, each tuple has two feature maps with all matching dimensions

### Example

```
>>> import torch
>>> from pytorch_lightning import seed_everything
>>> seed_everything(0)
0
>>> a1 = torch.rand(3, 5, 2, 2)
>>> a2 = torch.rand(3, 5, 2, 2)
>>> b1 = torch.rand(3, 5, 2, 2)
>>> b2 = torch.rand(3, 5, 2, 2)
```

(continues on next page)

(continued from previous page)

```

...
>>> task = FeatureMapContrastiveTask('01, 11')
...
>>> losses, regularizer = task((a1, a2), (b1, b2))
>>> losses
tensor([2.2351, 2.1902])
>>> regularizer
tensor(0.0324)

```

**static parse\_map\_indexes**(comparisons)

Example:

```

>>> FeatureMapContrastiveTask.parse_map_indexes('11')
[(1, 1)]
>>> FeatureMapContrastiveTask.parse_map_indexes('11,59')
[(1, 1), (5, 9)]
>>> FeatureMapContrastiveTask.parse_map_indexes('11,59, 2r')
[(1, 1), (5, 9), (2, -1)]

```

## 24.2 Context prediction tasks

The following tasks aim to predict a target using a context representation.

### 24.2.1 CPCContrastiveTask

This is the predictive task from CPC (v2).

```

task = CPCTask(num_input_channels=32)

# (batch, channels, rows, cols)
# this should be thought of as 49 feature vectors, each with 32 dims
Z = torch.random.rand(3, 32, 7, 7)

loss = task(Z)

class pl_bolts.losses.self_supervised_learning.CPCTask(num_input_channels,
                                                       target_dim=64,
                                                       embed_scale=0.1)
Bases: torch.nn.Module
Loss used in CPC

```



---

CHAPTER  
TWENTYFIVE

---

## INDICES AND TABLES

- genindex
- modindex
- search

Logo

## 25.1 PyTorch Lightning Bolts

Pretrained SOTA Deep Learning models, callbacks and more for research and production with PyTorch Lightning and PyTorch

---

### 25.1.1 Trending contributors

### 25.1.2 Continuous Integration

| System / PyTorch ver. | 1.4 (min. req.) | 1.6 (latest) | | :—: | :—: | :—: | | Linux py3.6 / py3.7 / py3.8 | CI testing | CI testing | | OSX py3.6 / py3.7 / py3.8 | CI testing | CI testing | | Windows py3.6 / py3.7 / py3.8 | wip | wip |

### 25.1.3 Install

Simple installation from PyPI

```
pip install pytorch-lightning
```

Install bleeding-edge (no guarantees)

```
pip install git+https://github.com/PytorchLightning/pytorch-lightning-bolts.  
→git@master --upgrade
```

## 25.1.4 Docs

- master
- stable
- 0.2.0
- 0.1.1

## 25.1.5 What is Bolts

Bolts is a Deep learning research and production toolbox of:

- SOTA pretrained models.
- Model components.
- Callbacks.
- Losses.
- Datasets.

## 25.1.6 Main Goals of Bolts

The main goal of Bolts is to enable rapid model idea iteration.

### Example 1: Finetuning on data

```
from pl_bolts.models.self_supervised import SimCLR
from pl_bolts.models.self_supervised.simclr.transforms import_ 
    ↪SimCLRTrainDataTransform, SimCLREvalDataTransform
import pytorch_lightning as pl

# data
train_data = DataLoader(MyDataset(transforms=SimCLRTrainDataTransform(input_
    ↪height=32)))
val_data = DataLoader(MyDataset(transforms=SimCLREvalDataTransform(input_height=32)))

# model
model = SimCLR(pretrained='imagenet2012')

# train!
trainer = pl.Trainer(gpus=8)
trainer.fit(model, train_data, val_data)
```

## Example 2: Subclass and ideate

```
from pl_bolts.models import ImageGPT
from pl_bolts.self_supervised import SimCLR

class VideoGPT(ImageGPT):

    def training_step(self, batch, batch_idx):
        x, y = batch
        x = _shape_input(x)

        logits = self.gpt(x)
        simclr_features = self.simclr(x)

        # -----
        # do something new with GPT logits + simclr_features
        # -----

        loss = self.criterion(logits.view(-1, logits.size(-1)), y.view(-1).long())

        logs = {"loss": loss}
        return {"loss": loss, "log": logs}
```

### 25.1.7 Who is Bolts for?

- Corporate production teams
- Professional researchers
- Ph.D. students
- Linear + Logistic regression heroes

### 25.1.8 I don't need deep learning

Great! We have LinearRegression and LogisticRegression implementations with numpy and sklearn bridges for datasets! But our implementations work on multiple GPUs, TPUs and scale dramatically...

Check out our Linear Regression on TPU demo

```
from pl_bolts.models.regression import LinearRegression
from pl_bolts.datamodules import SklearnDataModule

# sklearn dataset
X, y = load_boston(return_X_y=True)
loaders = SklearnDataModule(X, y)

model = LinearRegression(input_dim=13)
trainer = pl.Trainer(num_tpu_cores=1)
trainer.fit(model, loaders.train_dataloader(), loaders.val_dataloader())
trainer.test(test_dataloaders=loaders.test_dataloader())
```

### 25.1.9 Is this another model zoo?

No!

Bolts is unique because models are implemented using PyTorch Lightning and structured so that they can be easily subclassed and iterated on.

For example, you can override the elbo loss of a VAE, or the generator\_step of a GAN to quickly try out a new idea. The best part is that all the models are benchmarked so you won't waste time trying to "reproduce" or find the bugs with your implementation.

### 25.1.10 Team

Bolts is supported by the PyTorch Lightning team and the PyTorch Lightning community!

## 25.2 pl\_bolts.callbacks package

Collection of PyTorchLightning callbacks

### 25.2.1 Subpackages

#### pl\_bolts.callbacks.vision package

##### Submodules

#### pl\_bolts.callbacks.vision.confused\_logit module

```
class pl_bolts.callbacks.vision.confused_logit.ConfusedLogitCallback(top_k,
                                                                    projec-
                                                                    tion_factor=3,
                                                                    min_logit_value=5.0,
                                                                    log-
                                                                    ging_batch_interval=20,
                                                                    max_logit_difference=0.1)
```

Bases: pytorch\_lightning.Callback

Takes the logit predictions of a model and when the probabilities of two classes are very close, the model doesn't have high certainty that it should pick one vs the other class.

This callback shows how the input would have to change to swing the model from one label prediction to the other.

In this case, the network predicts a 5... but gives almost equal probability to an 8. The images show what about the original 5 would have to change to make it more like a 5 or more like an 8.

For each confused logit the confused images are generated by taking the gradient from a logit wrt an input for the top two closest logits.

Example:

```
from pl_bolts.callbacks.vision import ConfusedLogitCallback
trainer = Trainer(callbacks=[ConfusedLogitCallback()])
```

---

**Note:** whenever called, this model will look for self.last\_batch and self.last\_logits in the LightningModule

---



---

**Note:** this callback supports tensorboard only right now

---

### Parameters

- **top\_k** – How many “offending” images we should plot
- **projection\_factor** – How much to multiply the input image to make it look more like this logit label
- **min\_logit\_value** – Only consider logit values above this threshold
- **logging\_batch\_interval** – how frequently to inspect/potentially plot something
- **max\_logit\_difference** – when the top 2 logits are within this threshold we consider them confused

Authored by:

- Alfredo Canziani

```
static _ConfusedLogitCallback__draw_sample (fig, axarr, row_idx, col_idx, img, title)
    _plot (confusing_x, confusing_y, trainer, model, mask_idxs)
on_train_batch_end (trainer, pl_module, batch, batch_idx, dataloader_idx)
```

## pl\_bolts.callbacks.vision.image\_generation module

```
class pl_bolts.callbacks.vision.image_generation.TensorboardGenerativeModelImageSampler (num
Bases: pytorch_lightning.Callback
```

Generates images and logs to tensorboard. Your model must implement the forward function for generation

Requirements:

```
# model must have img_dim arg
model.img_dim = (1, 28, 28)

# model forward must work for sampling
z = torch.rand(batch_size, latent_dim)
img_samples = your_model(z)
```

Example:

```
from pl_bolts.callbacks import TensorboardGenerativeModelImageSampler

trainer = Trainer(callbacks=[TensorboardGenerativeModelImageSampler()])

on_epoch_end (trainer, pl_module)
```

## 25.2.2 Submodules

### pl\_bolts.callbacks.printing module

```
class pl_bolts.callbacks.printing.PrintTableMetricsCallback  
Bases: pytorch_lightning.callbacks.Callback
```

Prints a table with the metrics in columns on every epoch end

Example:

```
from pl_bolts.callbacks import PrintTableMetricsCallback  
  
callback = PrintTableMetricsCallback()
```

pass into trainer like so:

```
trainer = pl.Trainer(callbacks=[callback])  
trainer.fit(...)  
  
# -----  
# at the end of every epoch it will print  
# -----  
  
# loss|train_loss|val_loss|epoch  
# _____  
# 2.2541470527648926|2.2541470527648926|2.2158432006835938|0
```

`on_epoch_end(trainer, pl_module)`

```
pl_bolts.callbacks.printing.dicts_to_table(dicts, keys=None, pads=None,  
                                            fcodes=None, convert_headers=None,  
                                            header_names=None, skip_none_lines=False,  
                                            replace_values=None)
```

Generate ascii table from dictionary Taken from (<https://stackoverflow.com/questions/40056747/print-a-list-of-dictionaries-in-table-form>)

#### Parameters

- **dicts** (`List[Dict]`) – input dictionary list; empty lists make keys OR header\_names mandatory
- **keys** (`Optional[List[str]]`) – order list of keys to generate columns for; no key/dict-key should suffix with ‘\_\_’ else adjust code-suffix
- **pads** (`Optional[List[str]]`) – indicate padding direction and size, eg <10 to right pad alias left-align
- **fcodes** (`Optional[List[str]]`) – formating codes for respective column type, eg .3f
- **convert\_headers** (`Optional[Dict[str, Callable]]`) – apply converters(dict) on column keys k, eg timestamps
- **header\_names** (`Optional[List[str]]`) – supply for custom column headers instead of keys
- **skip\_none\_lines** (`bool`) – skip line if contains None
- **replace\_values** (`Optional[Dict[str, Any]]`) – specify per column keys k a map from seen value to new value; new value must comply with the columns fcode; CAUTION: modifies input (due speed)

## Example

```
>>> a = {'a': 1, 'b': 2}
>>> b = {'a': 3, 'b': 4}
>>> print(dicts_to_table([a, b]))
a|b
-----
1|2
3|4
```

## pl\_bolts.callbacks.self\_supervised module

**class** pl\_bolts.callbacks.self\_supervised.BYOLMAWeightUpdate (*initial\_tau=0.996*)  
Bases: pytorch\_lightning.Callback

Weight update rule from BYOL.

Your model should have a:

- self.online\_network.
- self.target\_network.

Updates the target\_network params using an exponential moving average update rule weighted by tau. BYOL claims this keeps the online\_network from collapsing.

---

**Note:** Automatically increases tau from *initial\_tau* to 1.0 with every training step

---

Example:

```
from pl_bolts.callbacks.self_supervised import BYOLMAWeightUpdate

# model must have 2 attributes
model = Model()
model.online_network = ...
model.target_network = ...

# make sure to set max_steps in Trainer
trainer = Trainer(callbacks=[BYOLMAWeightUpdate()], max_steps=1000)
```

**Parameters** **initial\_tau** – starting tau. Auto-updates with every training step

**on\_train\_batch\_end**(*trainer, pl\_module, batch, batch\_idx, dataloader\_idx*)  
**update\_tau**(*pl\_module, trainer*)  
**update\_weights**(*online\_net, target\_net*)

**class** pl\_bolts.callbacks.self\_supervised.SSLOnlineEvaluator (*drop\_p=0.2, hidden\_dim=1024, z\_dim=None, num\_classes=None*)

Bases: pytorch\_lightning.Callback

Attaches a MLP for finetuning using the standard self-supervised protocol.

Example:

```
from pl_bolts.callbacks.self_supervised import SSLOnlineEvaluator

# your model must have 2 attributes
model = Model()
model.z_dim = ... # the representation dim
model.num_classes = ... # the num of classes in the model
```

#### Parameters

- **drop\_p** (`float`) – (0.2) dropout probability

- **hidden\_dim** (`int`) –

(1024) the hidden dimension for the finetune MLP

**get\_representations** (*pl\_module, x*)

Override this to customize for the particular model :param\_sphinx\_paramlinks\_pl\_bolts.callbacks.self\_supervised.SSLOnlineEvaluator.get\_representations.x:

**on\_pretrain\_routine\_start** (*trainer, pl\_module*)

**on\_train\_batch\_end** (*trainer, pl\_module, batch, batch\_idx, dataloader\_idx*)

**to\_device** (*batch, device*)

## pl\_bolts.callbacks.variational module

```
class pl_bolts.callbacks.variational.LatentDimInterpolator(interpolate_epoch_interval=20,
                                                               range_start=-5,
                                                               range_end=5,
                                                               num_samples=2)
```

Bases: `pytorch_lightning.callbacks.Callback`

Interpolates the latent space for a model by setting all dims to zero and stepping through the first two dims increasing one unit at a time.

Default interpolates between [-5, 5] (-5, -4, -3, ..., 3, 4, 5)

Example:

```
from pl_bolts.callbacks import LatentDimInterpolator

Trainer(callbacks=[LatentDimInterpolator()])
```

#### Parameters

- **interpolate\_epoch\_interval** –
- **range\_start** – default -5
- **range\_end** – default 5
- **num\_samples** – default 2

**interpolate\_latent\_space** (*pl\_module, latent\_dim*)

**on\_epoch\_end** (*trainer, pl\_module*)

## 25.3 pl\_bolts.datamodules package

### 25.3.1 Submodules

#### pl\_bolts.datamodules.async\_dataloader module

```
class pl_bolts.datamodules.async_dataloader.AsyncDataLoader(data,      de-
                                                               vice=torch.device,
                                                               q_size=10,
                                                               num_batches=None,
                                                               **kwargs)
```

Bases: `object`

Class for asynchronously loading from CPU memory to device memory with DataLoader.

Note that this only works for single GPU training, multiGPU uses PyTorch's DataParallel or DistributedDataParallel which uses its own code for transferring data across GPUs. This could just break or make things slower with DataParallel or DistributedDataParallel.

#### Parameters

- `data` – The PyTorch Dataset or DataLoader we're using to load.
- `device` – The PyTorch device we are loading to
- `q_size` – Size of the queue used to store the data loaded to the device
- `num_batches` – Number of batches to load. This must be set if the dataloader doesn't have a finite `__len__`. It will also override `DataLoader.__len__` if set and `DataLoader` has a `__len__`. Otherwise it can be left as None
- `**kwargs` – Any additional arguments to pass to the dataloader if we're constructing one here

`load_instance(sample)`

`load_loop()`

#### pl\_bolts.datamodules.base\_dataset module

```
class pl_bolts.datamodules.base_dataset.LightDataset(*args, **kwargs)
```

Bases: `abc.ABC, torch.utils.data.Dataset`

`_download_from_url(base_url, data_folder, file_name)`

`static _prepare_subset(full_data, full_targets, num_samples, labels)`

Prepare a subset of a common dataset.

`Return type Tuple[Tensor, Tensor]`

`DATASET_NAME = 'light'`

`cache_folder_name: str = None`

`property cached_folder_path`

`Return type str`

`data: torch.Tensor = None`

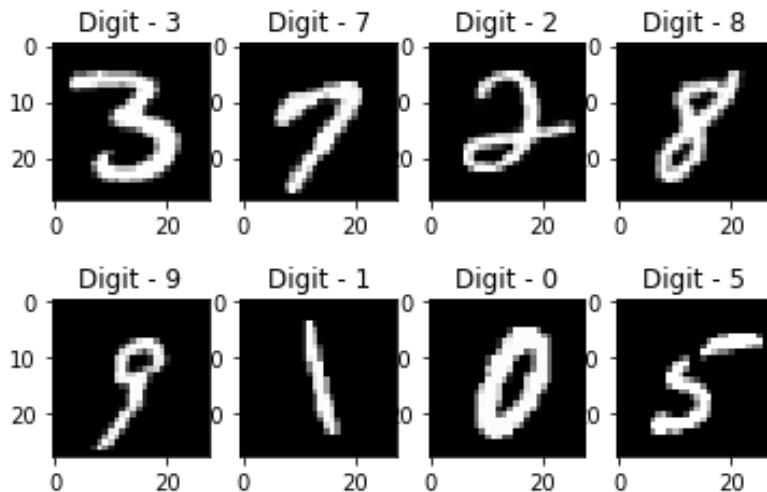
`dir_path: str = None`

```
normalize: tuple = None
targets: torch.Tensor = None
```

### pl\_bolts.datamodules.binary\_mnist\_datamodule module

```
class pl_bolts.datamodules.binary_mnist_datamodule.BinaryMNIST(*args,
                                                               **kwargs)
Bases: torchvision.datasets.MNIST

class pl_bolts.datamodules.binary_mnist_datamodule.BinaryMNISTDataModule(data_dir,
                                                                           val_split=5000,
                                                                           num_workers=16,
                                                                           nor-
                                                                           mal-
                                                                           ize=False,
                                                                           seed=42,
                                                                           *args,
                                                                           **kwargs)
Bases: pytorch_lightning.LightningDataModule
```



#### Specs:

- 10 classes (1 per digit)
- Each image is (1 x 28 x 28)

Binary MNIST, train, val, test splits and transforms

Transforms:

```
mnist_transforms = transform_lib.Compose([
    transform_lib.ToTensor()
])
```

Example:

```
from pl_bolts.datamodules import BinaryMNISTDataModule

dm = BinaryMNISTDataModule('.')
model = LitModel()

Trainer().fit(model, dm)
```

**Parameters**

- **data\_dir** (`str`) – where to save/load the data
- **val\_split** (`int`) – how many of the training images to use for the validation split
- **num\_workers** (`int`) – how many workers to use for loading data
- **normalize** (`bool`) – If true applies image normalize

`_default_transforms()`

`prepare_data()`

Saves MNIST files to `data_dir`

`test_dataloader(batch_size=32, transforms=None)`

MNIST test set uses the test split

**Parameters**

- **batch\_size** – size of batch
- **transforms** – custom transforms

`train_dataloader(batch_size=32, transforms=None)`

MNIST train set removes a subset to use for validation

**Parameters**

- **batch\_size** – size of batch
- **transforms** – custom transforms

`val_dataloader(batch_size=32, transforms=None)`

MNIST val set uses a subset of the training set for validation

**Parameters**

- **batch\_size** – size of batch
- **transforms** – custom transforms

`name = 'mnist'`

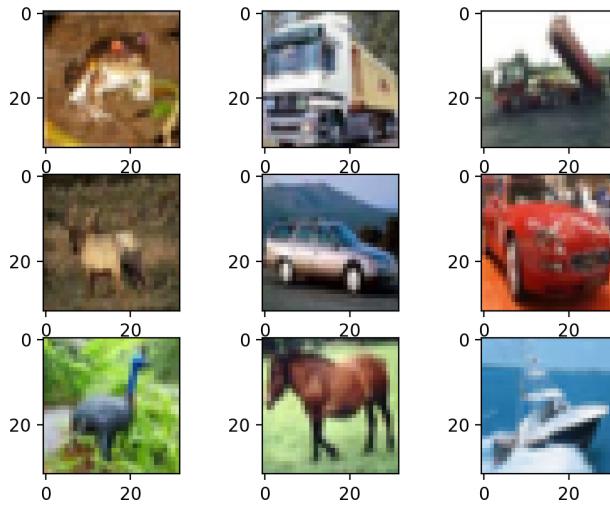
`property num_classes`

Return: 10

## pl\_bolts.datamodules.cifar10\_datamodule module

```
class pl_bolts.datamodules.cifar10_datamodule.CIFAR10DataModule(data_dir=None,
                                                               val_split=5000,
                                                               num_workers=16,
                                                               batch_size=32,
                                                               seed=42,
                                                               *args,
                                                               **kwargs)
```

Bases: pytorch\_lightning.LightningDataModule



### Specs:

- 10 classes (1 per class)
- Each image is (3 x 32 x 32)

Standard CIFAR10, train, val, test splits and transforms

Transforms:

```
mnist_transforms = transform_lib.Compose([
    transform_lib.ToTensor(),
    transforms.Normalize(
        mean=[x / 255.0 for x in [125.3, 123.0, 113.9]],
        std=[x / 255.0 for x in [63.0, 62.1, 66.7]]
    )
])
```

Example:

```
from pl_bolts.datamodules import CIFAR10DataModule

dm = CIFAR10DataModule(PATH)
model = LitModel()

Trainer().fit(model, dm)
```

Or you can set your own transforms

Example:

```
dm.train_transforms = ...
dm.test_transforms = ...
dm.val_transforms = ...
```

### Parameters

- **data\_dir** (`Optional[str]`) – where to save/load the data
- **val\_split** (`int`) – how many of the training images to use for the validation split
- **num\_workers** (`int`) – how many workers to use for loading data
- **batch\_size** (`int`) – number of examples per training/eval step

```
default_transforms()

prepare_data()
    Saves CIFAR10 files to data_dir

test_dataloader()
    CIFAR10 test set uses the test split

train_dataloader()
    CIFAR train set removes a subset to use for validation

val_dataloader()
    CIFAR10 val set uses a subset of the training set for validation

extra_args = {}

name = 'cifar10'

property num_classes
    Return: 10

class pl_bolts.datamodules.cifar10_datamodule.TinyCIFAR10DataModule(data_dir,
    val_split=50,
    num_workers=16,
    num_samples=100,
    la-
    bels=(1,
    5,     8),
    *args,
    **kwargs)
Bases: pl_bolts.datamodules.cifar10_datamodule.CIFAR10DataModule
```

Standard CIFAR10, train, val, test splits and transforms

Transforms:

```
mnist_transforms = transform_lib.Compose([
    transform_lib.ToTensor(),
    transforms.Normalize(mean=[x / 255.0 for x in [125.3, 123.0, 113.9]],
                        std=[x / 255.0 for x in [63.0, 62.1, 66.7]])
])
```

Example:

```
from pl_bolts.datamodules import CIFAR10DataModule  
  
dm = CIFAR10DataModule(PATH)  
model = LitModel(datamodule=dm)
```

### Parameters

- **data\_dir** (str) – where to save/load the data
- **val\_split** (int) – how many of the training images to use for the validation split
- **num\_workers** (int) – how many workers to use for loading data
- **num\_samples** (int) – number of examples per selected class/label
- **labels** (Optional[Sequence]) – list selected CIFAR10 classes/labels

#### property num\_classes

Return number of classes.

Return type int

## pl\_bolts.datamodules.cifar10\_dataset module

```
class pl_bolts.datamodules.cifar10_dataset.CIFAR10(data_dir='.', train=True, transform=None, download=True)
```

Bases: *pl\_bolts.datamodules.base\_dataset.LightDataset*

Customized CIFAR10 dataset for testing Pytorch Lightning without the torchvision dependency.

Part of the code was copied from <https://github.com/pytorch/vision/blob/v0.5.0/torchvision/datasets/>

### Parameters

- **data\_dir** (str) – Root directory of dataset where CIFAR10/processed/training.pt and CIFAR10/processed/test.pt exist.
- **train** (bool) – If True, creates dataset from training.pt, otherwise from test.pt.
- **download** (bool) – If true, downloads the dataset from the internet and puts it in root directory. If dataset is already downloaded, it is not downloaded again.

## Examples

```
>>> from torchvision import transforms  
>>> from pl_bolts.transforms.dataset_normalizations import cifar10_normalization  
>>> cf10_transforms = transforms.Compose([transforms.ToTensor(), cifar10_normalization()])  
>>> dataset = CIFAR10(download=True, transform=cf10_transforms)  
>>> len(dataset)  
50000  
>>> torch.bincount(dataset.targets)  
tensor([5000, 5000, 5000, 5000, 5000, 5000, 5000, 5000, 5000, 5000])  
>>> data, label = dataset[0]  
>>> data.shape  
torch.Size([3, 32, 32])
```

(continues on next page)

(continued from previous page)

```
>>> label
6
```

Labels:

```
airplane: 0
automobile: 1
bird: 2
cat: 3
deer: 4
dog: 5
frog: 6
horse: 7
ship: 8
truck: 9
```

**classmethod** `_check_exists`(*data\_folder*, *file\_names*)

**Return type** `bool`

`_extract_archive_save_torch`(*download\_path*)

`_unpickle`(*path\_folder*, *file\_name*)

**Return type** `Tuple[Tensor, Tensor]`

`download`(*data\_folder*)

Download the data if it doesn't exist in `cached_folder_path` already.

**Return type** `None`

`prepare_data`(*download*)

```
BASE_URL = 'https://www.cs.toronto.edu/~kriz/'
```

```
DATASET_NAME = 'CIFAR10'
```

```
FILE_NAME = 'cifar-10-python.tar.gz'
```

```
TEST_FILE_NAME = 'test.pt'
```

```
TRAIN_FILE_NAME = 'training.pt'
```

```
cache_folder_name: str = 'complete'
```

```
data = None
```

```
dir_path = None
```

```
labels = {0, 1, 2, 3, 4, 5, 6, 7, 8, 9}
```

```
normalize = None
```

```
relabel = False
```

```
targets = None
```

```
class pl_bolts.datamodules.cifar10_dataset.TrialCIFAR10 (data_dir='.', train=True,
                                                          transform=None,
                                                          download=False,
                                                          num_samples=100,
                                                          labels=(1, 5, 8), relabel=True)
```

Bases: `pl_bolts.datamodules.cifar10_dataset.CIFAR10`

Customized CIFAR10 dataset for testing Pytorch Lightning without the torchvision dependency.

### Parameters

- **data\_dir** (str) – Root directory of dataset where CIFAR10/processed/training.pt and CIFAR10/processed/test.pt exist.
- **train** (bool) – If True, creates dataset from training.pt, otherwise from test.pt.
- **download** (bool) – If true, downloads the dataset from the internet and puts it in root directory. If dataset is already downloaded, it is not downloaded again.
- **num\_samples** (int) – number of examples per selected class/digit
- **labels** (Optional[Sequence]) – list selected CIFAR10 digits/classes

### Examples

```
>>> dataset = TrialCIFAR10(download=True, num_samples=150, labels=(1, 5, 8))
>>> len(dataset)
450
>>> sorted(set([d.item() for d in dataset.targets]))
[1, 5, 8]
>>> torch.bincount(dataset.targets)
tensor([ 0, 150,   0,   0,   0, 150,   0,   0, 150])
>>> data, label = dataset[0]
>>> data.shape
torch.Size([3, 32, 32])
```

**prepare\_data(download)**

**Return type** None

```
data = None
dir_path = None
normalize = None
targets = None
```

## pl\_bolts.datamodules.cityscapes\_datamodule module

```
class pl_bolts.datamodules.cityscapes_datamodule.CityscapesDataModule(data_dir,
                                                                      val_split=5000,
                                                                      num_workers=16,
                                                                      batch_size=32,
                                                                      seed=42,
                                                                      *args,
                                                                      **kwargs)
```

Bases: pytorch\_lightning.LightningDataModule

Standard Cityscapes, train, val, test splits and transforms

### Specs:

- 30 classes (road, person, sidewalk, etc...)
- (image, target) - image dims: (3 x 32 x 32), target dims: (3 x 32 x 32)



Transforms:

```
transforms = transform_lib.Compose([
    transform_lib.ToTensor(),
    transform_lib.Normalize(
        mean=[0.28689554, 0.32513303, 0.28389177],
        std=[0.18696375, 0.19017339, 0.18720214]
    )
])
```

Example:

```
from pl_bolts.datamodules import CityscapesDataModule

dm = CityscapesDataModule(PATH)
model = LitModel()

Trainer().fit(model, dm)
```

Or you can set your own transforms

Example:

```
dm.train_transforms = ...
dm.test_transforms = ...
dm.val_transforms = ...
```

### Parameters

- **data\_dir** – where to save/load the data
- **val\_split** – how many of the training images to use for the validation split
- **num\_workers** – how many workers to use for loading data
- **batch\_size** – number of examples per training/eval step

```
default_transforms()
prepare_data()
Saves Cityscapes files to data_dir
test_dataloader()
Cityscapes test set uses the test split
```

```
train_dataloader()
    Cityscapes train set with removed subset to use for validation

val_dataloader()
    Cityscapes val set uses a subset of the training set for validation

extra_args = {}

name = 'Cityscapes'

property num_classes
    Return: 30
```

## pl\_bolts.datamodules.concat\_dataset module

```
class pl_bolts.datamodules.concat_dataset.ConcatDataset(*datasets)
Bases: torch.utils.data.Dataset
```

## pl\_bolts.datamodules.dummy\_dataset module

```
class pl_bolts.datamodules.dummy_dataset.DummyDataset(*shapes,
                                                       num_samples=10000)
Bases: torch.utils.data.Dataset

Generate a dummy dataset
```

### Parameters

- **\*shapes** – list of shapes
- **num\_samples** – how many samples to use in this dataset

Example:

```
from pl_bolts.datamodules import DummyDataset

# mnist dims
>>> ds = DummyDataset((1, 28, 28), (1,))
>>> dl = DataLoader(ds, batch_size=7)
...
>>> batch = next(iter(dl))
>>> x, y = batch
>>> x.size()
torch.Size([7, 1, 28, 28])
>>> y.size()
torch.Size([7, 1])
```

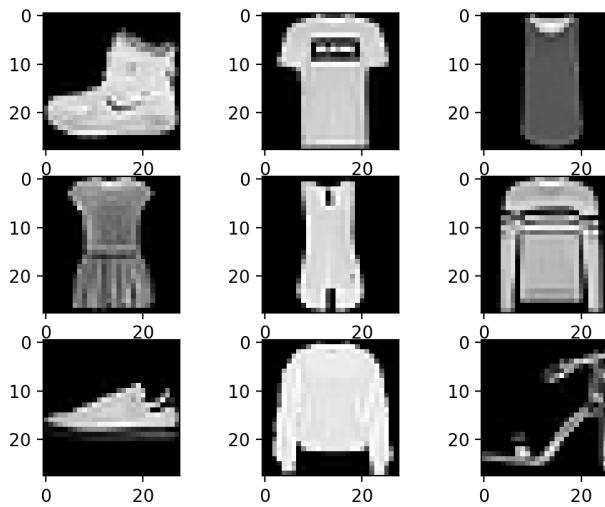
```
class pl_bolts.datamodules.dummy_dataset.DummyDetectionDataset(img_shape=(3,
                                                               256, 256),
                                                               num_boxes=1,
                                                               num_classes=2,
                                                               num_samples=10000)
Bases: torch.utils.data.Dataset

_random_bbox()
```

## pl\_bolts.datamodules.fashion\_mnist\_datamodule module

```
class pl_bolts.datamodules.fashion_mnist_datamodule.FashionMNISTDataModule(data_dir,
                                                                           val_split=5000,
                                                                           num_workers=16,
                                                                           seed=42,
                                                                           *args,
                                                                           **kwargs)
```

Bases: `pytorch_lightning.LightningDataModule`



### Specs:

- 10 classes (1 per type)
- Each image is (1 x 28 x 28)

Standard FashionMNIST, train, val, test splits and transforms

Transforms:

```
mnist_transforms = transform_lib.Compose([
    transform_lib.ToTensor()
])
```

Example:

```
from pl_bolts.datamodules import FashionMNISTDataModule

dm = FashionMNISTDataModule('.')
model = LitModel()

Trainer().fit(model, dm)
```

### Parameters

- `data_dir (str)` – where to save/load the data

- **val\_split** (`int`) – how many of the training images to use for the validation split
- **num\_workers** (`int`) – how many workers to use for loading data

```
_default_transforms()
prepare_data()
    Saves FashionMNIST files to data_dir
test_dataloader(batch_size=32, transforms=None)
    FashionMNIST test set uses the test split
    Parameters
        • batch_size – size of batch
        • transforms – custom transforms
train_dataloader(batch_size=32, transforms=None)
    FashionMNIST train set removes a subset to use for validation
    Parameters
        • batch_size – size of batch
        • transforms – custom transforms
val_dataloader(batch_size=32, transforms=None)
    FashionMNIST val set uses a subset of the training set for validation
    Parameters
        • batch_size – size of batch
        • transforms – custom transforms
name = 'fashion_mnist'
property num_classes
    Return: 10
```

## pl\_bolts.datamodules.imagenet\_datamodule module

```
class pl_bolts.datamodules.imagenet_datamodule.ImagenetDataModule(data_dir,
    meta_dir=None,
    num_imgs_per_val_class=50,
    im-
    age_size=224,
    num_workers=16,
    batch_size=32,
    *args,
    **kwargs)
```

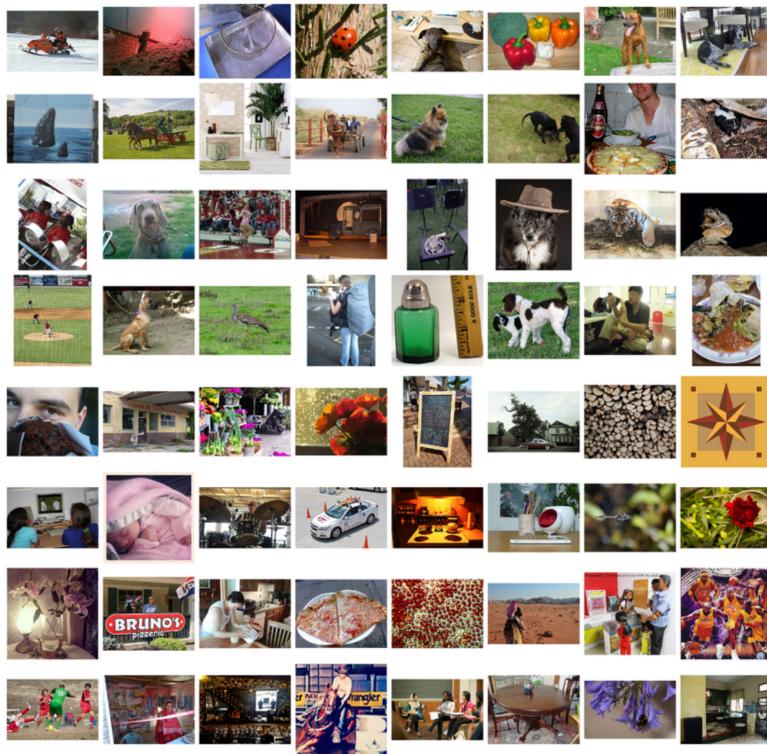
Bases: `pytorch_lightning.LightningDataModule`

### Specs:

- 1000 classes
- Each image is (3 x varies x varies) (here we default to 3 x 224 x 224)

Imagenet train, val and test dataloaders.

The train set is the imagenet train.



The val set is taken from the train set with `num_imgs_per_val_class` images per class. For example if `num_imgs_per_val_class=2` then there will be 2,000 images in the validation set.

The test set is the official imagenet validation set.

Example:

```
from pl_bolts.datamodules import ImagenetDataModule

dm = ImagenetDataModule(IMAGENET_PATH)
model = LitModel()

Trainer().fit(model, dm)
```

### Parameters

- `data_dir` (`str`) – path to the imagenet dataset file
- `meta_dir` (`Optional[str]`) – path to meta.bin file
- `num_imgs_per_val_class` (`int`) – how many images per class for the validation set
- `image_size` (`int`) – final image size
- `num_workers` (`int`) – how many data workers
- `batch_size` (`int`) – batch\_size

`_verify_splits(data_dir, split)`

`prepare_data()`

This method already assumes you have imagenet2012 downloaded. It validates the data using the meta.bin.

**Warning:** Please download imagenet on your own first.

**test\_dataloader()**

Uses the validation split of imagenet2012 for testing

**train\_dataloader()**

Uses the train split of imagenet2012 and puts away a portion of it for the validation split

**train\_transform()**

The standard imagenet transforms

```
transform_lib.Compose([
    transform_lib.RandomResizedCrop(self.image_size),
    transform_lib.RandomHorizontalFlip(),
    transform_lib.ToTensor(),
    transform_lib.Normalize(
        mean=[0.485, 0.456, 0.406],
        std=[0.229, 0.224, 0.225]
    ),
])
```

**val\_dataloader()**

Uses the part of the train split of imagenet2012 that was not used for training via *num\_imgs\_per\_val\_class*

**Parameters**

- **batch\_size** – the batch size
- **transforms** – the transforms

**val\_transform()**

The standard imagenet transforms for validation

```
transform_lib.Compose([
    transform_lib.Resize(self.image_size + 32),
    transform_lib.CenterCrop(self.image_size),
    transform_lib.ToTensor(),
    transform_lib.Normalize(
        mean=[0.485, 0.456, 0.406],
        std=[0.229, 0.224, 0.225]
    ),
])
```

**name = 'imagenet'****property num\_classes**

Return:

1000

## pl\_bolts.datamodules.imagenet\_dataset module

```
class pl_bolts.datamodules.imagenet_dataset.UnlabeledImagenet(root, split='train',
                                                               num_classes=-1,
                                                               num_imgs_per_class=1,
                                                               num_imgs_per_class_val_split=50,
                                                               meta_dir=None,
                                                               **kwargs)
```

Bases: `torchvision.datasets.ImageNet`

Official train set gets split into train, val. (using `nb_imgs_per_val_class` for each class). Official validation becomes test set

Within each class, we further allow limiting the number of samples per class (for semi-sup Ing)

### Parameters

- `root` – path of dataset
- `split (str)` –
- `num_classes (int)` – Sets the limit of classes
- `num_imgs_per_class (int)` – Limits the number of images per class
- `num_imgs_per_class_val_split (int)` – How many images per class to generate the val split
- `download` –
- `kwargs` –

```
classmethod generate_meta_bins(devkit_dir)
```

```
partition_train_set(imgs, nb_imgs_in_val)
```

```
pl_bolts.datamodules.imagenet_dataset._calculate_md5(fpather, chunk_size=1048576)
```

```
pl_bolts.datamodules.imagenet_dataset._check_integrity(fpather, md5=None)
```

```
pl_bolts.datamodules.imagenet_dataset._check_md5(fpather, md5, **kwargs)
```

```
pl_bolts.datamodules.imagenet_dataset._is_gzip(filename)
```

```
pl_bolts.datamodules.imagenet_dataset._is_tar(filename)
```

```
pl_bolts.datamodules.imagenet_dataset._is_targz(filename)
```

```
pl_bolts.datamodules.imagenet_dataset._is_tarxz(filename)
```

```
pl_bolts.datamodules.imagenet_dataset._is_zip(filename)
```

```
pl_bolts.datamodules.imagenet_dataset._verify_archive(root, file, md5)
```

```
pl_bolts.datamodules.imagenet_dataset.extract_archive(from_path, to_path=None, remove_finished=False)
```

```
pl_bolts.datamodules.imagenet_dataset.parse_devkit_archive(root, file=None)
```

Parse the devkit archive of the ImageNet2012 classification dataset and save the meta information in a binary file.

### Parameters

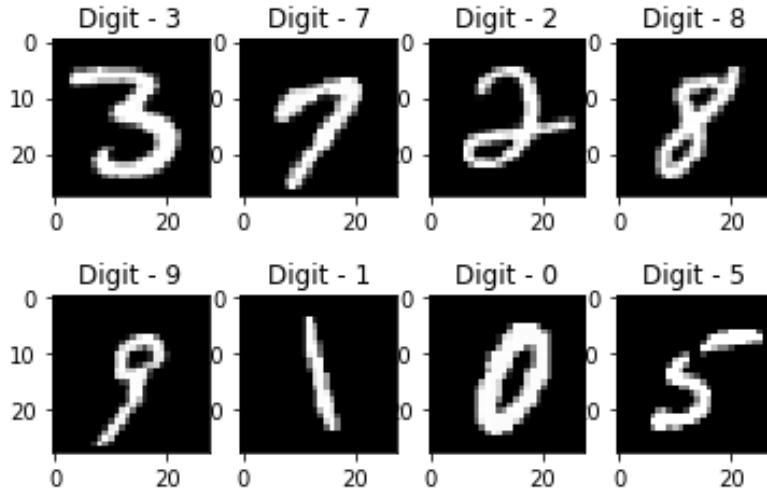
- `root (str)` – Root directory containing the devkit archive

- **file** (*str, optional*) – Name of devkit archive. Defaults to ‘ILSVRC2012\_devkit\_t12.tar.gz’

### pl\_bolts.datamodules.mnist\_datamodule module

```
class pl_bolts.datamodules.mnist_datamodule.MNISTDataModule(data_dir,
                                                               val_split=5000,
                                                               num_workers=16,
                                                               normalize=False,
                                                               seed=42,      *args,
                                                               **kwargs)
```

Bases: pytorch\_lightning.LightningDataModule



#### Specs:

- 10 classes (1 per digit)
- Each image is (1 x 28 x 28)

Standard MNIST, train, val, test splits and transforms

Transforms:

```
mnist_transforms = transform_lib.Compose([
    transform_lib.ToTensor()
])
```

Example:

```
from pl_bolts.datamodules import MNISTDataModule

dm = MNISTDataModule('.')
model = LitModel()

Trainer().fit(model, dm)
```

#### Parameters

- **data\_dir** (*str*) – where to save/load the data

- **val\_split** (`int`) – how many of the training images to use for the validation split
- **num\_workers** (`int`) – how many workers to use for loading data
- **normalize** (`bool`) – If true applies image normalize

```
_default_transforms()
prepare_data()
    Saves MNIST files to data_dir
test_dataloader(batch_size=32, transforms=None)
    MNIST test set uses the test split

Parameters
    • batch_size – size of batch
    • transforms – custom transforms

train_dataloader(batch_size=32, transforms=None)
    MNIST train set removes a subset to use for validation

Parameters
    • batch_size – size of batch
    • transforms – custom transforms

val_dataloader(batch_size=32, transforms=None)
    MNIST val set uses a subset of the training set for validation

Parameters
    • batch_size – size of batch
    • transforms – custom transforms

name = 'mnist'
property num_classes
    Return: 10
```

## pl\_bolts.datamodules.sklearn\_datamodule module

```
class pl_bolts.datamodules.sklearn_datamodule.SklearnDataModule(X, y,
    x_val=None,
    y_val=None,
    x_test=None,
    y_test=None,
    val_split=0.2,
    test_split=0.1,
    num_workers=2,
    random_state=1234,
    shuffle=True,
    *args,
    **kwargs)
```

Bases: `pytorch_lightning.LightningDataModule`

Automatically generates the train, validation and test splits for a Numpy dataset. They are set up as dataloaders for convenience. Optionally, you can pass in your own validation and test splits.

## Example

```
>>> from sklearn.datasets import load_boston
>>> from pl_bolts.datamodules import SklearnDataModule
...
>>> X, y = load_boston(return_X_y=True)
>>> loaders = SklearnDataModule(X, y)
...
>>> # train set
>>> train_loader = loaders.train_dataloader(batch_size=32)
>>> len(train_loader.dataset)
355
>>> len(train_loader)
11
>>> # validation set
>>> val_loader = loaders.val_dataloader(batch_size=32)
>>> len(val_loader.dataset)
100
>>> len(val_loader)
3
>>> # test set
>>> test_loader = loaders.test_dataloader(batch_size=32)
>>> len(test_loader.dataset)
51
>>> len(test_loader)
1
```

```
_init_datasets(X, y, x_val, y_val, x_test, y_test)

test_dataloader(batch_size=16)
train_dataloader(batch_size=16)
val_dataloader(batch_size=16)
name = 'sklearn'

class pl_bolts.datamodules.sklearn_datamodule.SklearnDataset(X, y,
                                                               X_transform=None,
                                                               y_transform=None)
```

Bases: `torch.utils.data.Dataset`

Mapping between numpy (or sklearn) datasets to PyTorch datasets.

### Parameters

- `x` (ndarray) – Numpy ndarray
- `y` (ndarray) – Numpy ndarray
- `x_transform` (Optional[Any]) – Any transform that works with Numpy arrays
- `y_transform` (Optional[Any]) – Any transform that works with Numpy arrays

## Example

```
>>> from sklearn.datasets import load_boston
>>> from pl_bolts.datamodules import SklearnDataset
...
>>> X, y = load_boston(return_X_y=True)
>>> dataset = SklearnDataset(X, y)
>>> len(dataset)
506
```

```
class pl_bolts.datamodules.sklearn_datamodule.TensorDataModule(X, y,
                                                               x_val=None,
                                                               y_val=None,
                                                               x_test=None,
                                                               y_test=None,
                                                               val_split=0.2,
                                                               test_split=0.1,
                                                               num_workers=2,
                                                               random_state=1234,
                                                               shuffle=True,
                                                               *args,
                                                               **kwargs)
Bases: pl_bolts.datamodules.sklearn_datamodule.SklearnDataModule
```

Automatically generates the train, validation and test splits for a PyTorch tensor dataset. They are set up as dataloaders for convenience. Optionally, you can pass in your own validation and test splits.

## Example

```
>>> from pl_bolts.datamodules import TensorDataModule
>>> import torch
...
>>> # create dataset
>>> X = torch.rand(100, 3)
>>> y = torch.rand(100)
>>> loaders = TensorDataModule(X, y)
...
>>> # train set
>>> train_loader = loaders.train_dataloader(batch_size=10)
>>> len(train_loader.dataset)
70
>>> len(train_loader)
7
>>> # validation set
>>> val_loader = loaders.val_dataloader(batch_size=10)
>>> len(val_loader.dataset)
20
>>> len(val_loader)
2
>>> # test set
>>> test_loader = loaders.test_dataloader(batch_size=10)
>>> len(test_loader.dataset)
10
>>> len(test_loader)
1
```

Automatically generates the train, validation and test splits for a Numpy dataset. They are set up as dataloaders for convenience. Optionally, you can pass in your own validation and test splits.

## Example

```
>>> from sklearn.datasets import load_boston
>>> from pl_bolts.datamodules import SklearnDataModule
...
>>> X, y = load_boston(return_X_y=True)
>>> loaders = SklearnDataModule(X, y)
...
>>> # train set
>>> train_loader = loaders.train_dataloader(batch_size=32)
>>> len(train_loader.dataset)
355
>>> len(train_loader)
11
>>> # validation set
>>> val_loader = loaders.val_dataloader(batch_size=32)
>>> len(val_loader.dataset)
100
>>> len(val_loader)
3
>>> # test set
>>> test_loader = loaders.test_dataloader(batch_size=32)
>>> len(test_loader.dataset)
51
>>> len(test_loader)
1
```

```
class pl_bolts.datamodules.sklearn_datamodule.TensorDataset(X, y,
                                                               X_transform=None,
                                                               y_transform=None)
```

Bases: `torch.utils.data.Dataset`

Prepare PyTorch tensor dataset for data loaders.

### Parameters

- `X (Tensor)` – PyTorch tensor
- `y (Tensor)` – PyTorch tensor
- `X_transform (Optional[Any])` – Any transform that works with PyTorch tensors
- `y_transform (Optional[Any])` – Any transform that works with PyTorch tensors

## Example

```
>>> from pl_bolts.datamodules import TensorDataset
...
>>> X = torch.rand(10, 3)
>>> y = torch.rand(10)
>>> dataset = TensorDataset(X, y)
>>> len(dataset)
10
```

## pl\_bolts.datamodules.ssl\_imagenet\_datamodule module

```
class pl_bolts.datamodules.ssl_imagenet_datamodule.SSLImagenetDataModule(data_dir,  

meta_dir=None,  

num_workers=16,  

*args,  

**kwargs)  
Bases: pytorch_lightning.LightningDataModule  
_default_transforms()  
_verify_splits(data_dir, split)  
prepare_data()  
test_dataloader(batch_size, num_images_per_class, add_normalize=False)  
train_dataloader(batch_size, num_images_per_class=-1, add_normalize=False)  
val_dataloader(batch_size, num_images_per_class=50, add_normalize=False)  
name = 'imagenet'  
property num_classes
```

## pl\_bolts.datamodules.stl10\_datamodule module

```
class pl_bolts.datamodules.stl10_datamodule.STL10DataModule(data_dir=None,  

unla-  
beled_val_split=5000,  
train_val_split=500,  
num_workers=16,  
batch_size=32,  
seed=42, *args,  
**kwargs)  
Bases: pytorch_lightning.LightningDataModule
```



### Specs:

- 10 classes (1 per type)
- Each image is (3 x 96 x 96)

Standard STL-10, train, val, test splits and transforms. STL-10 has support for doing validation splits on the labeled or unlabeled splits

Transforms:

```
mnist_transforms = transform_lib.Compose([
    transform_lib.ToTensor(),
    transforms.Normalize(
        mean=(0.43, 0.42, 0.39),
        std=(0.27, 0.26, 0.27)
    )
])
```

Example:

```
from pl_bolts.datamodules import STL10DataModule

dm = STL10DataModule(PATH)
model = LitModel()

Trainer().fit(model, dm)
```

### Parameters

- **data\_dir** (Optional[str]) – where to save/load the data
- **unlabeled\_val\_split** (int) – how many images from the unlabeled training split to use for validation
- **train\_val\_split** (int) – how many images from the labeled training split to use for validation
- **num\_workers** (int) – how many workers to use for loading data
- **batch\_size** (int) – the batch size

**default\_transforms()**

**prepare\_data()**

Downloads the unlabeled, train and test split

**test\_dataloader()**

Loads the test split of STL10

### Parameters

- **batch\_size** – the batch size
- **transforms** – the transforms

**train\_dataloader()**

Loads the ‘unlabeled’ split minus a portion set aside for validation via *unlabeled\_val\_split*.

**train\_dataloader\_labeled()**

**train\_dataloader\_mixed()**

Loads a portion of the ‘unlabeled’ training data and ‘train’ (labeled) data. both portions have a subset removed for validation via *unlabeled\_val\_split* and *train\_val\_split*

### Parameters

- **batch\_size** – the batch size
- **transforms** – a sequence of transforms

---

**val\_dataloader()**  
Loads a portion of the ‘unlabeled’ training data set aside for validation The val dataset = (unlabeled - train\_val\_split)

**Parameters**

- **batch\_size** – the batch size
- **transforms** – a sequence of transforms

**val\_dataloader\_labeled()****val\_dataloader\_mixed()**

Loads a portion of the ‘unlabeled’ training data set aside for validation along with the portion of the ‘train’ dataset to be used for validation

unlabeled\_val = (unlabeled - train\_val\_split)  
labeled\_val = (train- train\_val\_split)  
full\_val = unlabeled\_val + labeled\_val

**Parameters**

- **batch\_size** – the batch size
- **transforms** – a sequence of transforms

**name = 'stl10'****property num\_classes****pl\_bolts.datamodules.vocdetection\_datamodule module**

**class** pl\_bolts.datamodules.vocdetection\_datamodule.**Compose**(transforms)  
Bases: `object`

Like `torchvision.transforms.compose` but works for (image, target)

**\_\_call\_\_(image, target)**  
Call self as a function.

**class** pl\_bolts.datamodules.vocdetection\_datamodule.**VOCDetectionDataModule**(data\_dir,  
year='2012',  
num\_workers=16,  
nor-  
mal-  
ize=False,  
\*args,  
\*\*kwargs)

Bases: `pytorch_lightning.LightningDataModule`

TODO(teddykoker) docstring

**\_default\_transforms()****prepare\_data()**

Saves VOCDetection files to data\_dir

**train\_dataloader(batch\_size=1, transforms=None)**

VOCDetection train set uses the *train* subset

**Parameters**

- **batch\_size** – size of batch
- **transforms** – custom transforms

**val\_dataloader** (*batch\_size=1, transforms=None*)  
VOCDetection val set uses the *val* subset

**Parameters**

- **batch\_size** – size of batch
- **transforms** – custom transforms

**name = 'vocdetection'**

**property num\_classes**  
Return: 21

pl\_bolts.datamodules.vocdetection\_datamodule.\_**collate\_fn**(*batch*)

pl\_bolts.datamodules.vocdetection\_datamodule.\_**prepare\_voc\_instance**(*image, target*)  
Prepares VOC dataset into appropriate target for fasterrcnn  
<https://github.com/pytorch/vision/issues/1097#issuecomment-508917489>

## 25.4 pl\_bolts.metrics package

### 25.4.1 Submodules

#### pl\_bolts.metrics.aggregation module

pl\_bolts.metrics.aggregation.**accuracy** (*preds, labels*)  
pl\_bolts.metrics.aggregation.**mean** (*res, key*)  
pl\_bolts.metrics.aggregation.**precision\_at\_k** (*output, target, top\_k=(1, )*)  
Computes the accuracy over the k top predictions for the specified values of k

## 25.5 pl\_bolts.models package

Collection of PyTorchLightning models

### 25.5.1 Subpackages

#### pl\_bolts.models.autoencoders package

Here are a VAE and GAN

## Subpackages

### pl\_bolts.models.autoencoders.basic\_ae package

#### AE Template

This is a basic template for implementing an Autoencoder in PyTorch Lightning.

A default encoder and decoder have been provided but can easily be replaced by custom models.

**This template uses the MNIST dataset but image data of any dimension can be fed in as long as the image width and image height are even values.** For other types of data, such as sound, it will be necessary to change the Encoder and Decoder.

**The default encoder and decoder are both convolutional with a 128-dimensional hidden layer and a 32-dimensional latent space.** The model accepts arguments for these dimensions (see example below) if you want to use the default encoder + decoder but with different hidden layer and latent layer dimensions. The model also assumes a Gaussian prior and a Gaussian approximate posterior distribution.

```
from pl_bolts.models.autoencoders import AE

model = AE()
trainer = pl.Trainer()
trainer.fit(model)
```

## Submodules

### pl\_bolts.models.autoencoders.basic\_ae.basic\_ae\_module module

```
class pl_bolts.models.autoencoders.basic_ae.basic_ae_module.AE(datamodule=None,
                                                               in-
                                                               put_channels=1,
                                                               in-
                                                               put_height=28,
                                                               in-
                                                               put_width=28,
                                                               latent_dim=32,
                                                               batch_size=32,
                                                               hid-
                                                               den_dim=128,
                                                               learn-
                                                               ing_rate=0.001,
                                                               num_workers=8,
                                                               data_dir='.',
                                                               **kwargs)
```

Bases: pytorch\_lightning.LightningModule

#### Parameters

- **datamodule** (`Optional[LightningDataModule]`) – the datamodule (train, val, test splits)
- **input\_channels** – num of image channels
- **input\_height** – image height

- `input_width` – image width
- `latent_dim` – emb dim for encoder
- `batch_size` – the batch size
- `hidden_dim` – the encoder dim
- `learning_rate` – the learning rate
- `num_workers` – num dataloader workers
- `data_dir` – where to store data

`_run_step(batch)`

`static add_model_specific_args(parent_parser)`

`configure_optimizers()`

`forward(z)`

`init_decoder(hidden_dim, latent_dim)`

`init_encoder(hidden_dim, latent_dim, input_width, input_height)`

`test_epoch_end(outputs)`

`test_step(batch, batch_idx)`

`training_step(batch, batch_idx)`

`validation_epoch_end(outputs)`

`validation_step(batch, batch_idx)`

`pl_bolts.models.autoencoders.basic_ae.basic_ae_module.cli_main()`

## pl\_bolts.models.autoencoders.basic\_ae.components module

`class pl_bolts.models.autoencoders.basic_ae.components.AEEncoder(hidden_dim,  
                         latent_dim,  
                         input_width,  
                         in-  
                         put_height)`

Bases: `torch.nn.Module`

Takes as input an image, uses a CNN to extract features which get split into a mu and sigma vector

`_calculate_output_dim(input_width, input_height)`

`forward(x)`

`class pl_bolts.models.autoencoders.basic_ae.components.DenseBlock(in_dim,  
                         out_dim,  
                         drop_p=0.2)`

Bases: `torch.nn.Module`

`forward(x)`

## pl\_bolts.models.autoencoders.basic\_vae package

### VAE Template

This is a basic template for implementing a Variational Autoencoder in PyTorch Lightning.

A default encoder and decoder have been provided but can easily be replaced by custom models.

**This template uses the MNIST dataset but image data of any dimension can be fed in as long as the image width and image height are even values.** For other types of data, such as sound, it will be necessary to change the Encoder and Decoder.

**The default encoder and decoder are both convolutional with a 128-dimensional hidden layer and a 32-dimensional latent space.** The model accepts arguments for these dimensions (see example below) if you want to use the default encoder + decoder but with different hidden layer and latent layer dimensions. The model also assumes a Gaussian prior and a Gaussian approximate posterior distribution.

### Submodules

#### pl\_bolts.models.autoencoders.basic\_vae.basic\_vae module

```
class pl_bolts.models.autoencoders.basic_vae.basic_vae_module.VAE(hidden_dim=128,
la-
tent_dim=32,
in-
put_channels=3,
in-
put_width=224,
in-
put_height=224,
batch_size=32,
learn-
ing_rate=0.001,
data_dir='.',
datamod-
ule=None,
num_workers=8,
pre-
trained=None,
**kwargs)
```

Bases: `pytorch_lightning.LightningModule`

Standard VAE with Gaussian Prior and approx posterior.

Model is available pretrained on different datasets:

Example:

```
# not pretrained
vae = VAE()

# pretrained on imagenet
vae = VAE(pretrained='imagenet')

# pretrained on cifar10
vae = VAE(pretrained='cifar10')
```

## Parameters

- `hidden_dim`(`int`) – encoder and decoder hidden dims
- `latent_dim`(`int`) – latenet code dim
- `input_channels`(`int`) – num of channels of the input image.
- `input_width`(`int`) – image input width
- `input_height`(`int`) – image input height
- `batch_size`(`int`) – the batch size
- `the learning rate`(`learning_rate`) –
- `data_dir`(`str`) – the directory to store data
- `datamodule`(`Optional[LightningDataModule]`) – The Lightning DataModule
- `pretrained`(`Optional[str]`) – Load weights pretrained on a dataset

```
_VAE__init_system()
_VAE__set_pretrained_dims(pretrained)
_run_step(batch)
_set_default_datamodule(datamodule)
static add_model_specific_args(parent_parser)
configure_optimizers()
elbo_loss(x, P, Q, num_samples)
forward(z)
get_approx_posterior(z_mu, z_std)
get_prior(z_mu, z_std)
init_decoder()
init_encoder()
load_pretrained(pretrained)
test_step(batch, batch_idx)
training_step(batch, batch_idx)
validation_step(batch, batch_idx)
pl_bolts.models.autoencoders.basic_vae.basic_vae_module.cli_main()
```

## pl\_bolts.models.autoencoders.basic\_vae.components module

```
class pl_bolts.models.autoencoders.basic_vae.components.Decoder(hidden_dim,
latent_dim, in-
put_width, in-
put_height, in-
put_channels)
```

Bases: `torch.nn.Module`

Takes in latent vars and reconstructs an image

```

    _calculate_output_size(input_width, input_height)
    forward(z)

class pl_bolts.models.autoencoders.basic_vae.components.DenseBlock(in_dim,
                                                               out_dim,
                                                               drop_p=0.2)
Bases: torch.nn.Module

forward(x)

class pl_bolts.models.autoencoders.basic_vae.components.Encoder(hidden_dim,
                                                               latent_dim, in-
                                                               put_channels,
                                                               input_width,
                                                               input_height)
Bases: torch.nn.Module

Takes as input an image, uses a CNN to extract features which get split into a mu and sigma vector

    _calculate_output_dim(input_width, input_height)
    forward(x)

```

## pl\_bolts.models.detection package

### Submodules

#### pl\_bolts.models.detection.faster\_rcnn module

```

class pl_bolts.models.detection.faster_rcnn.FasterRCNN(learning_rate=0.0001,
                                                       num_classes=91, pre-
                                                       trained=False, pre-
                                                       trained_backbone=True,
                                                       train-
                                                       able_backbone_layers=3,
                                                       replace_head=True,
                                                       **kwargs)

```

Bases: pytorch\_lightning.LightningModule

PyTorch Lightning implementation of Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks.

Paper authors: Shaoqing Ren, Kaiming He, Ross Girshick, Jian Sun

#### Model implemented by:

- *Teddy Koker <https://github.com/teddykoker>*

**During training, the model expects both the input tensors, as well as targets (list of dictionary), containing:**

- boxes (*FloatTensor[N, 4]*): the ground truth boxes in  $[x1, y1, x2, y2]$  format.
- labels (*Int64Tensor[N]*): the class label for each ground truh box

CLI command:

```
# PascalVOC
python faster_rcnn.py --gpus 1 --pretrained True
```

## Parameters

- **learning\_rate** (`float`) – the learning rate
- **num\_classes** (`int`) – number of detection classes (including background)
- **pretrained** (`bool`) – if true, returns a model pre-trained on COCO train2017
- **pretrained\_backbone** (`bool`) – if true, returns a model with backbone pre-trained on Imagenet
- **trainable\_backbone\_layers** (`int`) – number of trainable resnet layers starting from final block

```
static add_model_specific_args(parent_parser)
configure_optimizers()
forward(x)
training_step(batch, batch_idx)
validation_epoch_end(outs)
validation_step(batch, batch_idx)

pl_bolts.models.detection.faster_rcnn._evaluate_iou(target, pred)
    Evaluate intersection over union (IOU) for target from dataset and output prediction from model
pl_bolts.models.detection.faster_rcnn.run_cli()
```

## pl\_bolts.models.gans package

### Subpackages

#### pl\_bolts.models.gans.basic package

### Submodules

#### pl\_bolts.models.gans.basic.basic\_gan\_module module

```
class pl_bolts.models.gans.basic.basic_gan_module.GAN(datamodule=None,
                                                       latent_dim=32,
                                                       batch_size=100,      learning_rate=0.0002, data_dir="",
                                                       num_workers=8, **kwargs)
```

Bases: `pytorch_lightning.LightningModule`

Vanilla GAN implementation.

Example:

```
from pl_bolts.models.gan import GAN

m = GAN()
Trainer(gpus=2).fit(m)
```

Example CLI:

```
# mnist
python basic_gan_module.py --gpus 1

# imagenet
python basic_gan_module.py --gpus 1 --dataset 'imagenet2012'
--data_dir /path/to/imagenet/folder/ --meta_dir ~/path/to/meta/bin/folder
--batch_size 256 --learning_rate 0.0001
```

## Parameters

- **datamodule** (`Optional[LightningDataModule]`) – the datamodule (train, val, test splits)
- **latent\_dim** (`int`) – emb dim for encoder
- **batch\_size** (`int`) – the batch size
- **learning\_rate** (`float`) – the learning rate
- **data\_dir** (`str`) – where to store data
- **num\_workers** (`int`) – data workers

```
_set_default_datamodule(datamodule)
static add_model_specific_args(parent_parser)
configure_optimizers()
discriminator_loss(x)
discriminator_step(x)
forward(z)
Generates an image given input noise z
```

Example:

```
z = torch.rand(batch_size, latent_dim)
gan = GAN.load_from_checkpoint(PATH)
img = gan(z)
```

```
generator_loss(x)
generator_step(x)
init_discriminator(img_dim)
init_generator(img_dim)
training_step(batch, batch_idx, optimizer_idx)

pl_bolts.models.gans.basic_gan_module.cli_main()
```

**pl\_bolts.models.gans.basic.components module**

```
class pl_bolts.models.gans.basic.components.Discriminator(img_shape, hid-  
den_dim=1024)  
    Bases: torch.nn.Module  
  
    forward(img)  
  
class pl_bolts.models.gans.basic.components.Generator(latent_dim, img_shape, hid-  
den_dim=256)  
    Bases: torch.nn.Module  
  
    forward(z)
```

**pl\_bolts.models.regression package****Submodules****pl\_bolts.models.regression.linear\_regression module**

```
class pl_bolts.models.regression.linear_regression.LinearRegression(input_dim,  
bias=True,  
learn-  
ing_rate=0.0001,  
opti-  
mizer=torch.optim.Adam,  
l1_strength=None,  
l2_strength=None,  
**kwargs)  
Bases: pytorch_lightning.LightningModule
```

Linear regression model implementing - with optional L1/L2 regularization  $\min_{\{W\}} \|Wx + b - y\|_2^2$

**Parameters**

- **input\_dim** (`int`) – number of dimensions of the input (1+)
- **bias** (`bool`) – If false, will not use  $+b$
- **learning\_rate** (`float`) – learning\_rate for the optimizer
- **optimizer** (`Optimizer`) – the optimizer to use (default='Adam')
- **l1\_strength** (`Optional[float]`) – L1 regularization strength (default=None)
- **l2\_strength** (`Optional[float]`) – L2 regularization strength (default=None)

```
static add_model_specific_args(parent_parser)  
configure_optimizers()  
forward(x)  
test_epoch_end(outputs)  
test_step(batch, batch_idx)  
training_step(batch, batch_idx)  
validation_epoch_end(outputs)
```

```
validation_step(batch, batch_idx)
pl_bolts.models.regression.linear_regression.cli_main()
```

## pl\_bolts.models.regression.logistic\_regression module

```
class pl_bolts.models.regression.logistic_regression.LogisticRegression(input_dim,
num_classes,
bias=True,
learn-
ing_rate=0.0001,
op-
ti-
mizer=torch.optim.Adam,
l1_strength=0.0,
l2_strength=0.0,
**kwargs)
```

Bases: pytorch\_lightning.LightningModule

Logistic regression model

### Parameters

- **input\_dim** (*int*) – number of dimensions of the input (at least 1)
- **num\_classes** (*int*) – number of class labels (binary: 2, multi-class: >2)
- **bias** (*bool*) – specifies if a constant or intercept should be fitted (equivalent to `fit_intercept` in sklearn)
- **learning\_rate** (*float*) – learning\_rate for the optimizer
- **optimizer** (*Optimizer*) – the optimizer to use (default='Adam')
- **l1\_strength** (*float*) – L1 regularization strength (default=None)
- **l2\_strength** (*float*) – L2 regularization strength (default=None)

```
static add_model_specific_args(parent_parser)
```

```
configure_optimizers()
```

```
forward(x)
```

```
test_epoch_end(outputs)
```

```
test_step(batch, batch_idx)
```

```
training_step(batch, batch_idx)
```

```
validation_epoch_end(outputs)
```

```
validation_step(batch, batch_idx)
```

```
pl_bolts.models.regression.logistic_regression.cli_main()
```

## pl\_bolts.models.self\_supervised package

These models have been pre-trained using self-supervised learning. The models can also be used without pre-training and overwritten for your own research.

Here's an example for using these as pretrained models.

```
from pl_bolts.models.self_supervised import CPCV2

images = get_imagenet_batch()

# extract unsupervised representations
pretrained = CPCV2(pretrained=True)
representations = pretrained(images)

# use these in classification or any downstream task
classifications = classifier(representations)
```

## Subpackages

### pl\_bolts.models.self\_supervised.amdim package

#### Submodules

##### pl\_bolts.models.self\_supervised.amdim.amdim\_module module

```
class pl_bolts.models.self_supervised.amdim.amdim_module.AMDIM(datamodule='cifar10',
                                                               en-
                                                               coder='amdim_encoder',
                                                               con-
                                                               trastive_task=torch.nn.Module,
                                                               im-
                                                               age_channels=3,
                                                               im-
                                                               age_height=32,
                                                               en-
                                                               coder_feature_dim=320,
                                                               embed-
                                                               ding_fx_dim=1280,
                                                               conv_block_depth=10,
                                                               use_bn=False,
                                                               tclip=20.0,
                                                               learn-
                                                               ing_rate=0.0002,
                                                               data_dir='',
                                                               num_classes=10,
                                                               batch_size=200,
                                                               **kwargs)
```

Bases: `pytorch_lightning.LightningModule`

PyTorch Lightning implementation of Augmented Multiscale Deep InfoMax (AMDIM)

Paper authors: Philip Bachman, R Devon Hjelm, William Buchwalter.

Model implemented by: [William Falcon](#)

This code is adapted to Lightning using the original author repo ([the original repo](#)).

## Example

```
>>> from pl_bolts.models.self_supervised import AMDIM
...
>>> model = AMDIM(encoder='resnet18')
```

Train:

```
trainer = Trainer()
trainer.fit(model)
```

## Parameters

- **datamodule** (`Union[str, LightningDataModule]`) – A LightningDatamodule
- **encoder** (`Union[str, Module, LightningModule]`) – an encoder string or model
- **image\_channels** (`int`) – 3
- **image\_height** (`int`) – pixels
- **encoder\_feature\_dim** (`int`) – Called *ndf* in the paper, this is the representation size for the encoder.
- **embedding\_fx\_dim** (`int`) – Output dim of the embedding function (*nrkhs* in the paper) (Reproducing Kernel Hilbert Spaces).
- **conv\_block\_depth** (`int`) – Depth of each encoder block,
- **use\_bn** (`bool`) – If true will use batchnorm.
- **tclip** (`int`) – soft clipping non-linearity to the scores after computing the regularization term and before computing the log-softmax. This is the ‘second trick’ used in the paper
- **learning\_rate** (`int`) – The learning rate
- **data\_dir** (`str`) – Where to store data
- **num\_classes** (`int`) – How many classes in the dataset
- **batch\_size** (`int`) – The batch size

```
static add_model_specific_args(parent_parser)
configure_optimizers()
forward(img_1, img_2)
init_encoder()
train_dataloader()
training_step(batch, batch_nb)
training_step_end(outputs)
val_dataloader()
validation_epoch_end(outputs)
validation_step(batch, batch_nb)

pl_bolts.models.self_supervised.amdim.amdim_module.cli_main()
```

## pl\_bolts.models.self\_supervised.amdim.datasets module

```
class pl_bolts.models.self_supervised.amdim.datasets.AMDIMPatchesPretraining
Bases: object

    " For pretraining we use the train transform for both train and val.

    static cifar10(dataset_root, patch_size, patch_overlap, split='train')
    static imagenet(dataset_root, nb_classes, patch_size, patch_overlap, split='train')
    static stl(dataset_root, patch_size, patch_overlap, split=None)

class pl_bolts.models.self_supervised.amdim.datasets.AMDIMPretraining
Bases: object

    " For pretraining we use the train transform for both train and val.

    static cifar10(dataset_root, split='train')
    static cifar10_tiny(dataset_root, split='train')
    static get_dataset(datamodule, data_dir, split='train', **kwargs)
    static imagenet(dataset_root, nb_classes, split='train')
    static stl(dataset_root, split=None)
```

## pl\_bolts.models.self\_supervised.amdim.networks module

```
class pl_bolts.models.self_supervised.amdim.networks.AMDIMEncoder(dummy_batch,
    num_channels=3,
    en-
    coder_feature_dim=64,
    embed-
    ding_fx_dim=512,
    conv_block_depth=3,
    en-
    coder_size=32,
    use_bn=False)

Bases: torch.nn.Module

_config_modules(x, output_widths, n_rkhs, use_bn)
    Configure the modules for extracting fake rkhs embeddings for infomax.

_forward_acts(x)
    Return activations from all layers.

forward(x)

init_weights(init_scale=1.0)
    Run custom weight init for modules...

class pl_bolts.models.self_supervised.amdim.networks.Conv3x3(n_in,           n_out,
    n_kern,
    n_stride,   n_pad,
    use_bn=True,
    pad_mode='constant')

Bases: torch.nn.Module

forward(x)
```

```

class pl_bolts.models.self_supervised.amdim.networks.ConvResBlock (n_in, n_out,
width,
stride,
pad, depth,
use_bn)
Bases: torch.nn.Module

forward(x)
init_weights(init_scale=1.0)
    Do a fixup-ish init for each ConvResNxN in this block.

class pl_bolts.models.self_supervised.amdim.networks.ConvResNxN (n_in, n_out,
width,
stride, pad,
use_bn=False)
Bases: torch.nn.Module

forward(x)
init_weights(init_scale=1.0)

class pl_bolts.models.self_supervised.amdim.networks.FakeRKHSCovNet (n_input,
n_output,
use_bn=False)
Bases: torch.nn.Module

forward(x)
init_weights(init_scale=1.0)

class pl_bolts.models.self_supervised.amdim.networks.MaybeBatchNorm2d (n_ftr,
affine,
use_bn)
Bases: torch.nn.Module

forward(x)
forward(x)
```

**class** pl\_bolts.models.self\_supervised.amdim.networks.NopNet (*norm\_dim*=None)

Bases: torch.nn.Module

**forward**(*x*)

## pl\_bolts.models.self\_supervised.amdim.ssl\_datasets module

```

class pl_bolts.models.self_supervised.amdim.ssl_datasets.CIFAR10Mixed (root,
split='val',
trans-
form=None,
tar-
get_transform=None,
down-
load=False,
nb_labeled_per_class=None,
val_pct=0.1)
Bases: pl_bolts.models.self_supervised.amdim.ssl_datasets.SSLDatasetMixin,
torchvision.datasets.CIFAR10

class pl_bolts.models.self_supervised.amdim.ssl_datasets.SSLDatasetMixin
Bases: abc.ABC
```

```
classmethod deterministic_shuffle(x, y)
classmethod generate_train_val_split(examples, labels, pct_val)
    Splits dataset uniformly across classes :param _sphinx_paramlinks_pl_bolts.models.self_supervised.amdim.ssl_datasets.SSL
    :param _sphinx_paramlinks_pl_bolts.models.self_supervised.amdim.ssl_datasets.SSLDatasetMixin.generate_train_val_split
    :param _sphinx_paramlinks_pl_bolts.models.self_supervised.amdim.ssl_datasets.SSLDatasetMixin.generate_train_val_split
    :return:
classmethod select_nb_imgs_per_class(examples, labels, nb_imgs_in_val)
    Splits a dataset into two parts. The labeled split has nb_imgs_in_val per class :param
    _sphinx_paramlinks_pl_bolts.models.self_supervised.amdim.ssl_datasets.SSLDatasetMixin.select_nb_imgs_per_class.exam
    :param _sphinx_paramlinks_pl_bolts.models.self_supervised.amdim.ssl_datasets.SSLDatasetMixin.select_nb_imgs_per_cla
    :param _sphinx_paramlinks_pl_bolts.models.self_supervised.amdim.ssl_datasets.SSLDatasetMixin.select_nb_imgs_per_cla
    :return:
```

## pl\_bolts.models.self\_supervised.amdim.transforms module

**class** pl\_bolts.models.self\_supervised.amdim.transforms.**AMDIMEvalTransformsCIFAR10**  
Bases: `object`

Transforms applied to AMDIM

Transforms:

```
transforms.ToTensor(),
normalize
```

Example:

```
x = torch.rand(5, 3, 32, 32)

transform = AMDIMEvalTransformsCIFAR10()
(view1, view2) = transform(x)
```

**\_\_call\_\_(inp)**

Call self as a function.

**class** pl\_bolts.models.self\_supervised.amdim.transforms.**AMDIMEvalTransformsImageNet128** (*height*)  
Bases: `object`

Transforms applied to AMDIM

Transforms:

```
transforms.Resize(height + 6, interpolation=3),
transforms.CenterCrop(height),
transforms.ToTensor(),
normalize
```

Example:

```
x = torch.rand(5, 3, 128, 128)

transform = AMDIMEvalTransformsImageNet128()
view1 = transform(x)
```

**\_\_call\_\_(inp)**

Call self as a function.

---

```
class pl_bolts.models.self_supervised.amdim.transforms.AMDIMEvalTransformsSTL10 (height=64)
Bases: object
```

Transforms applied to AMDIM

Transforms:

```
transforms.Resize(height + 6, interpolation=3),
transforms.CenterCrop(height),
transforms.ToTensor(),
normalize
```

Example:

```
x = torch.rand(5, 3, 64, 64)

transform = AMDIMTrainTransformsSTL10()
view1 = transform(x)
```

**\_\_call\_\_(*inp*)**

Call self as a function.

```
class pl_bolts.models.self_supervised.amdim.transforms.AMDIMTrainTransformsCIFAR10
Bases: object
```

Transforms applied to AMDIM

Transforms:

```
img_jitter,
col_jitter,
rnd_gray,
transforms.ToTensor(),
normalize
```

Example:

```
x = torch.rand(5, 3, 32, 32)

transform = AMDIMTrainTransformsCIFAR10()
(view1, view2) = transform(x)
```

**\_\_call\_\_(*inp*)**

Call self as a function.

```
class pl_bolts.models.self_supervised.amdim.transforms.AMDIMTrainTransformsImageNet128 (height
Bases: object
```

Transforms applied to AMDIM

Transforms:

```
img_jitter,
col_jitter,
rnd_gray,
transforms.ToTensor(),
normalize
```

Example:

```
x = torch.rand(5, 3, 128, 128)

transform = AMDIMTrainTransformsSTL10()
(view1, view2) = transform(x)
```

### `__call__(inp)`

Call self as a function.

**class** `pl_bolts.models.self_supervised.amdim.transforms.AMDIMTrainTransformsSTL10 (height=64)`  
Bases: `object`

Transforms applied to AMDIM

Transforms:

```
img_jitter,
col_jitter,
rnd_gray,
transforms.ToTensor(),
normalize
```

Example:

```
x = torch.rand(5, 3, 64, 64)

transform = AMDIMTrainTransformsSTL10()
(view1, view2) = transform(x)
```

### `__call__(inp)`

Call self as a function.

## `pl_bolts.models.self_supervised.biol package`

### Submodules

#### `pl_bolts.models.self_supervised.biol.biol_module module`

**class** `pl_bolts.models.self_supervised.biol.biol_module.BYOL (num_classes, learning_rate=0.2, weight_decay=1.5e-05, input_height=32, batch_size=32, num_workers=0, warmup_epochs=10, max_epochs=1000, **kwargs)`  
Bases: `pytorch_lightning.LightningModule`

PyTorch Lightning implementation of Bring Your Own Latent (BYOL)

Paper authors: Jean-Bastien Grill ,Florian Strub, Florent Altché, Corentin Tallec, Pierre H. Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Daniel Guo, Mohammad Gheshlaghi Azar, Bilal Piot, Koray Kavukcuoglu, Rémi Munos, Michal Valko.

**Model implemented by:**

- Annika Brundyn

**Warning:** Work in progress. This implementation is still being verified.

### TODOs:

- verify on CIFAR-10
- verify on STL-10
- pre-train on imagenet

Example:

```
import pytorch_lightning as pl
from pl_bolts.models.self_supervised import BYOL
from pl_bolts.datamodules import CIFAR10DataModule
from pl_bolts.models.self_supervised.simclr.simclr_transforms import (
    SimCLREvalDataTransform, SimCLRTrainDataTransform)

# model
model = BYOL(num_classes=10)

# data
dm = CIFAR10DataModule(num_workers=0)
dm.train_transforms = SimCLRTrainDataTransform(32)
dm.val_transforms = SimCLREvalDataTransform(32)

trainer = pl.Trainer()
trainer.fit(model, dm)
```

Train:

```
trainer = Trainer()
trainer.fit(model)
```

CLI command:

```
# cifar10
python byol_module.py --gpus 1

# imagenet
python byol_module.py
--gpus 8
--dataset imagenet2012
--data_dir /path/to/imagenet/
--meta_dir /path/to/folder/with/meta.bin/
--batch_size 32
```

### Parameters

- **datamodule** – The datamodule
- **learning\_rate** (`float`) – the learning rate
- **weight\_decay** (`float`) – optimizer weight decay
- **input\_height** (`int`) – image input height
- **batch\_size** (`int`) – the batch size

- **num\_workers** (`int`) – number of workers
- **warmup\_epochs** (`int`) – num of epochs for scheduler warm up
- **max\_epochs** (`int`) – max epochs for scheduler

```
static add_model_specific_args(parent_parser)
configure_optimizers()
forward(x)
on_train_batch_end(batch, batch_idx, dataloader_idx)

Return type None

shared_step(batch, batch_idx)
training_step(batch, batch_idx)
validation_step(batch, batch_idx)
```

## pl\_bolts.models.self\_supervised.byol.models module

```
class pl_bolts.models.self_supervised.byol.models.MLP(input_dim=2048, hid-
                                                     den_size=4096, out-
                                                     put_dim=256)
Bases: torch.nn.Module

forward(x)

class pl_bolts.models.self_supervised.byol.models.SiameseArm(encoder=None)
Bases: torch.nn.Module

forward(x)
```

## pl\_bolts.models.self\_supervised.cpc package

### Submodules

#### pl\_bolts.models.self\_supervised.cpc.cpc\_finetuner module

```
pl_bolts.models.self_supervised.cpc.cpc_finetuner.cli_main()
```

#### pl\_bolts.models.self\_supervised.cpc.cpc\_module module

## CPC V2

```
class pl_bolts.models.self_supervised.cpc.cpc_module.CPCV2 (datamodule=None,
    en-
    coder='cpc_encoder',
    patch_size=8,
    patch_overlap=4,
    online_ft=True,
    task='cpc',
    num_workers=4,
    learn-
    ing_rate=0.0001,
    data_dir='',
    batch_size=32,
    pretrained=None,
    **kwargs)
```

Bases: `pytorch_lightning.LightningModule`

PyTorch Lightning implementation of [Data-Efficient Image Recognition with Contrastive Predictive Coding](#)

Paper authors: (Olivier J. Hénaff, Aravind Srinivas, Jeffrey De Fauw, Ali Razavi, Carl Doersch, S. M. Ali Eslami, Aaron van den Oord).

Model implemented by:

- [William Falcon](#)
- [Tullie Murrell](#)

## Example

```
>>> from pl_bolts.models.self_supervised import CPCV2
...
>>> model = CPCV2()
```

Train:

```
trainer = Trainer()
trainer.fit(model)
```

CLI command:

```
# cifar10
python cpc_module.py --gpus 1

# imagenet
python cpc_module.py
    --gpus 8
    --dataset imagenet2012
    --data_dir /path/to/imagenet/
    --meta_dir /path/to/folder/with/meta.bin/
    --batch_size 32
```

To Finetune:

```
python cpc_finetuner.py --ckpt_path path/to/checkpoint.ckpt --dataset cifar10 --
    ↵ gpus x
```

Some uses:

```
# load resnet18 pretrained using CPC on imagenet
model = CPCV2(encoder='resnet18', pretrained=True)
resnet18 = model.encoder
resnet18.freeze()

# it supports any torchvision resnet
model = CPCV2(encoder='resnet50', pretrained=True)

# use it as a feature extractor
x = torch.rand(2, 3, 224, 224)
out = model(x)
```

## Parameters

- **datamodule** (`Optional[LightningDataModule]`) – A Datamodule (optional). Otherwise set the dataloaders directly
- **encoder** (`Union[str, Module, LightningModule]`) – A string for any of the resnets in torchvision, or the original CPC encoder, or a custom nn.Module encoder
- **patch\_size** (`int`) – How big to make the image patches
- **patch\_overlap** (`int`) – How much overlap should each patch have.
- **online\_ft** (`int`) – Enable a 1024-unit MLP to fine-tune online
- **task** (`str`) – Which self-supervised task to use ('cpc', 'amdim', etc...)
- **num\_workers** (`int`) – num dataloader workers
- **learning\_rate** (`int`) – what learning rate to use
- **data\_dir** (`str`) – where to store data
- **batch\_size** (`int`) – batch size
- **pretrained** (`Optional[str]`) – If true, will use the weights pretrained (using CPC) on Imagenet

```
_CPCV2__compute_final_nb_c(patch_size)
_CPCV2__recover_z_shape(Z, b)
static add_model_specific_args(parent_parser)
configure_optimizers()
forward(img_I)
init_encoder()
load_pretrained(encoder)
shared_step(batch)
training_step(batch, batch_nb)
validation_step(batch, batch_nb)
```

**pl\_bolts.models.self\_supervised.cpc.networks module**

```
class pl_bolts.models.self_supervised.cpc.networks.CPCResNet101 (sample_batch,  

                                                               zero_init_residual=False,  

                                                               groups=1,  

                                                               width_per_group=64,  

                                                               re-  

                                                               place_stride_with_dilation=None,  

                                                               norm_layer=None)  
  
Bases: torch.nn.Module  
  
_make_layer (sample_batch, block, planes, blocks, stride=1, dilate=False, expansion=4)  
  
flatten (x)  
  
forward (x)  
  
class pl_bolts.models.self_supervised.cpc.networks.LNBottleneck (sample_batch,  

                                                               inplanes,  

                                                               planes,  

                                                               stride=1,  

                                                               downsam-  

                                                               ple_conv=None,  

                                                               groups=1,  

                                                               base_width=64,  

                                                               dilation=1,  

                                                               norm_layer=None,  

                                                               expansion=4)  
  
Bases: torch.nn.Module  
  
_LNBottleneck_init_layer_norms (x, conv1, conv2, conv3, downsample_conv)  
  
forward (x)  
  
pl_bolts.models.self_supervised.cpc.networks.conv1x1 (in_planes, out_planes,  

                                                       stride=1)  
1x1 convolution  
  
pl_bolts.models.self_supervised.cpc.networks.conv3x3 (in_planes, out_planes,  

                                                       stride=1, groups=1, dila-  

                                                       tion=1)  
3x3 convolution with padding
```

**pl\_bolts.models.self\_supervised.cpc.transforms module**

```
class pl_bolts.models.self_supervised.cpc.transforms.CPCEvalTransformsCIFAR10 (patch_size=8,  

                                                               over-  

                                                               lap=4)
```

Bases: object

Transforms used for CPC:

**Parameters**

- **patch\_size** – size of patches when cutting up the image into overlapping patches
- **overlap** – how much to overlap patches

Transforms:

```
random_flip
transforms.ToTensor()
normalize
Patchify(patch_size=patch_size, overlap_size=overlap)
```

Example:

```
# in a regular dataset
CIFAR10(..., transforms=CPCEvalTransformsCIFAR10())

# in a DataModule
module = CIFAR10DataModule(PATH)
train_loader = module.train_dataloader(batch_size=32,_
    ↴transforms=CPCEvalTransformsCIFAR10())
```

### `__call__(inp)`

Call self as a function.

```
class pl_bolts.models.self_supervised.cpc.transforms.CPCEvalTransformsImageNet128(patch_size=_
    ↴over-
    ↴lap=16)
```

Bases: `object`

Transforms used for CPC:

#### Parameters

- **patch\_size** – size of patches when cutting up the image into overlapping patches
- **overlap** – how much to overlap patches

Transforms:

```
random_flip
transforms.ToTensor()
normalize
Patchify(patch_size=patch_size, overlap_size=patch_size // 2)
```

Example:

```
# in a regular dataset
Imagenet(..., transforms=CPCEvalTransformsImageNet128())

# in a DataModule
module = ImagenetDataModule(PATH)
train_loader = module.train_dataloader(batch_size=32,_
    ↴transforms=CPCEvalTransformsImageNet128())
```

### `__call__(inp)`

Call self as a function.

```
class pl_bolts.models.self_supervised.cpc.transforms.CPCEvalTransformsSTL10(patch_size=16,
    ↴over-
    ↴lap=8)
```

Bases: `object`

Transforms used for CPC:

#### Parameters

- **patch\_size** – size of patches when cutting up the image into overlapping patches

- **overlap** – how much to overlap patches

Transforms:

```
random_flip
transforms.ToTensor()
normalize
Patchify(patch_size=patch_size, overlap_size=patch_size // 2)
```

Example:

```
# in a regular dataset
STL10(..., transforms=CPCEvalTransformsSTL10())

# in a DataModule
module = STL10DataModule(PATH)
train_loader = module.train_dataloader(batch_size=32, ▶
    transforms=CPCEvalTransformsSTL10())
```

### `__call__(inp)`

Call self as a function.

```
class pl_bolts.models.self_supervised.cpc.transforms.CPCTrainTransformsCIFAR10(patch_size=8,
    over-
    lap=4)
```

Bases: `object`

Transforms used for CPC:

#### Parameters

- **patch\_size** – size of patches when cutting up the image into overlapping patches
- **overlap** – how much to overlap patches

Transforms:

```
random_flip
img_jitter
col_jitter
rnd_gray
transforms.ToTensor()
normalize
Patchify(patch_size=patch_size, overlap_size=patch_size // 2)
```

Example:

```
# in a regular dataset
CIFAR10(..., transforms=CPCTrainTransformsCIFAR10())

# in a DataModule
module = CIFAR10DataModule(PATH)
train_loader = module.train_dataloader(batch_size=32, ▶
    transforms=CPCTrainTransformsCIFAR10())
```

### `__call__(inp)`

Call self as a function.

```
class pl_bolts.models.self_supervised.cpc.transforms.CPCTrainTransformsImageNet128(patch_size-
    over-
    lap=16)
```

Bases: `object`

Transforms used for CPC:

#### Parameters

- **patch\_size** – size of patches when cutting up the image into overlapping patches
- **overlap** – how much to overlap patches

Transforms:

```
random_flip
transforms.ToTensor()
normalize
Patchify(patch_size=patch_size, overlap_size=patch_size // 2)
```

Example:

```
# in a regular dataset
Imagenet(..., transforms=CPCTrainTransformsImageNet128())

# in a DataModule
module = ImagenetDataModule(PATH)
train_loader = module.train_dataloader(batch_size=32, ↴
transforms=CPCTrainTransformsImageNet128())
```

#### `__call__(inp)`

Call self as a function.

```
class pl_bolts.models.self_supervised.cpc.transforms.CPCTrainTransformsSTL10(patch_size=16,
over-
lap=8)
```

Bases: `object`

Transforms used for CPC:

#### Parameters

- **patch\_size** – size of patches when cutting up the image into overlapping patches
- **overlap** – how much to overlap patches

Transforms:

```
random_flip
img_jitter
col_jitter
rnd_gray
transforms.ToTensor()
normalize
Patchify(patch_size=patch_size, overlap_size=patch_size // 2)
```

Example:

```
# in a regular dataset
STL10(..., transforms=CPCTrainTransformsSTL10())

# in a DataModule
module = STL10DataModule(PATH)
train_loader = module.train_dataloader(batch_size=32, ↴
transforms=CPCTrainTransformsSTL10())
```

---

**\_\_call\_\_(inp)**  
Call self as a function.

## pl\_bolts.models.self\_supervised.moco package

### Submodules

#### pl\_bolts.models.self\_supervised.moco.callbacks module

```
class pl_bolts.models.self_supervised.moco.callbacks.MocoLRScheduler (initial_lr=0.03,
use_cosine_scheduler=False,
sched-
ule=(120,
160),
max_epochs=200)
```

Bases: pytorch\_lightning.Callback

**on\_epoch\_start** (*trainer, pl\_module*)

#### pl\_bolts.models.self\_supervised.moco.moco2\_module module

Adapted from: <https://github.com/facebookresearch/moco>

Original work is: Copyright (c) Facebook, Inc. and its affiliates. All Rights Reserved This implementation is: Copyright (c) PyTorch Lightning, Inc. and its affiliates. All Rights Reserved

```
class pl_bolts.models.self_supervised.moco.moco2_module.MocoV2 (base_encoder='resnet18',
emb_dim=128,
num_negatives=65536,
en-
coder_momentum=0.999,
soft-
max_temperature=0.07,
learn-
ing_rate=0.03,
momentum=0.9,
weight_decay=0.0001,
datamod-
ule=None,
data_dir='./',
batch_size=256,
use_mlp=False,
num_workers=8,
*args,
**kwargs)
```

Bases: pytorch\_lightning.LightningModule

PyTorch Lightning implementation of Moco

Paper authors: Xinlei Chen, Haoqi Fan, Ross Girshick, Kaiming He.

Code adapted from [facebookresearch/moco](https://github.com/facebookresearch/moco) to Lightning by:

- William Falcon

## Example

```
>>> from pl_bolts.models.self_supervised import MocoV2
...
>>> model = MocoV2()
```

Train:

```
trainer = Trainer()
trainer.fit(model)
```

CLI command:

```
# cifar10
python moco2_module.py --gpus 1

# imagenet
python moco2_module.py
--gpus 8
--dataset imagenet2012
--data_dir /path/to/imagenet/
--meta_dir /path/to/folder/with/meta.bin/
--batch_size 32
```

## Parameters

- **base\_encoder** (`Union[str, Module]`) – torchvision model name or `torch.nn.Module`
- **emb\_dim** (`int`) – feature dimension (default: 128)
- **num\_negatives** (`int`) – queue size; number of negative keys (default: 65536)
- **encoder\_momentum** (`float`) – moco momentum of updating key encoder (default: 0.999)
- **softmax\_temperature** (`float`) – softmax temperature (default: 0.07)
- **learning\_rate** (`float`) – the learning rate
- **momentum** (`float`) – optimizer momentum
- **weight\_decay** (`float`) – optimizer weight decay
- **datamodule** (`Optional[LightningDataModule]`) – the DataModule (train, val, test dataloaders)
- **data\_dir** (`str`) – the directory to store data
- **batch\_size** (`int`) – batch size
- **use\_mlp** (`bool`) – add an mlp to the encoders
- **num\_workers** (`int`) – workers for the loaders

\_batch\_shuffle\_ddp (`x`)

Batch shuffle, for making use of BatchNorm. \* **Only support DistributedDataParallel (DDP) model.** \*

\_batch\_unshuffle\_ddp (`x, idx_unshuffle`)

Undo batch shuffle. \* **Only support DistributedDataParallel (DDP) model.** \*

\_dequeue\_and\_enqueue (`keys`)

```

_momentum_update_key_encoder()
    Momentum update of the key encoder

static add_model_specific_args(parent_parser)
configure_optimizers()
forward(img_q, img_k)

Input: im_q: a batch of query images im_k: a batch of key images
Output: logits, targets

init_encoders(base_encoder)
Override to add your own encoders

training_step(batch, batch_idx)
validation_epoch_end(outputs)
validation_step(batch, batch_idx)

pl_bolts.models.self_supervised.moco.moco2_module.cli_main()

pl_bolts.models.self_supervised.moco.moco2_module.concat_all_gather(tensor)
    Performs all_gather operation on the provided tensors. * Warning: torch.distributed.all_gather has no gradient.

```

## pl\_bolts.models.self\_supervised.moco.transforms module

```

class pl_bolts.models.self_supervised.moco.transforms.GaussianBlur(sigma=(0.1,
                                                                    2.0))
Bases: object
Gaussian blur augmentation in SimCLR https://arxiv.org/abs/2002.05709

__call__(x)
    Call self as a function.

class pl_bolts.models.self_supervised.moco.transforms.Moco2EvalCIFAR10Transforms(height=32)
Bases: object
Moco 2 augmentation: https://arxiv.org/pdf/2003.04297.pdf

__call__(inp)
    Call self as a function.

class pl_bolts.models.self_supervised.moco.transforms.Moco2EvalImagenetTransforms(height=128)
Bases: object
Moco 2 augmentation: https://arxiv.org/pdf/2003.04297.pdf

__call__(inp)
    Call self as a function.

class pl_bolts.models.self_supervised.moco.transforms.Moco2EvalSTL10Transforms(height=64)
Bases: object
Moco 2 augmentation: https://arxiv.org/pdf/2003.04297.pdf

__call__(inp)
    Call self as a function.

```

```
class pl_bolts.models.self_supervised.moco.transforms.Moco2TrainCIFAR10Transforms (height=32)
Bases: object
Moco 2 augmentation: https://arxiv.org/pdf/2003.04297.pdf
__call__(inp)
    Call self as a function.

class pl_bolts.models.self_supervised.moco.transforms.Moco2TrainImagenetTransforms (height=128)
Bases: object
Moco 2 augmentation: https://arxiv.org/pdf/2003.04297.pdf
__call__(inp)
    Call self as a function.

class pl_bolts.models.self_supervised.moco.transforms.Moco2TrainSTL10Transforms (height=64)
Bases: object
Moco 2 augmentation: https://arxiv.org/pdf/2003.04297.pdf
__call__(inp)
    Call self as a function.
```

## pl\_bolts.models.self\_supervised.simclr package

### Submodules

#### pl\_bolts.models.self\_supervised.simclr.simclr\_finetuner module

```
pl_bolts.models.self_supervised.simclr.simclr_finetuner.cli_main()
```

#### pl\_bolts.models.self\_supervised.simclr.simclr\_module module

```
class pl_bolts.models.self_supervised.simclr.simclr_module.DenseNetEncoder
Bases: torch.nn.Module
forward(x)

class pl_bolts.models.self_supervised.simclr.simclr_module.Projection (input_dim=2048,
hid-
den_dim=2048,
out-
put_dim=128)
Bases: torch.nn.Module
forward(x)

class pl_bolts.models.self_supervised.simclr.simclr_module.SimCLR (batch_size,
num_samples,
warmup_epochs=10,
lr=0.0001,
opt_weight_decay=1e-
06,
loss_temperature=0.5,
**kwargs)
Bases: pytorch_lightning.LightningModule
```

## Parameters

- **batch\_size** – the batch size
- **num\_samples** – num samples in the dataset
- **warmup\_epochs** – epochs to warmup the lr for
- **lr** – the optimizer learning rate
- **opt\_weight\_decay** – the optimizer weight decay
- **loss\_temperature** – the loss temperature

```
static add_model_specific_args(parent_parser)
configure_optimizers()
exclude_from_wt_decay(named_params, weight_decay, skip_list=['bias', 'bn'])
forward(x)
init_encoder()
setup(stage)
shared_step(batch, batch_idx)
training_step(batch, batch_idx)
validation_step(batch, batch_idx)
pl_bolts.models.self_supervised.simclr.simclr_module.cli_main()
```

## pl\_bolts.models.self\_supervised.simclr.simclr\_transforms module

```
class pl_bolts.models.self_supervised.simclr.simclr_transforms.GaussianBlur(kernel_size,
min=0.1,
max=2.0)
Bases: object
__call__(sample)
    Call self as a function.
```

**class** pl\_bolts.models.self\_supervised.simclr.simclr\_transforms.**SimCLREvalDataTransform**(input\_size=1)

Bases: **object**

Transforms for SimCLR

Transform:

```
Resize(input_height + 10, interpolation=3)
transforms.CenterCrop(input_height),
transforms.ToTensor()
```

Example:

```
from pl_bolts.models.self_supervised.simclr.transforms import_
SimCLREvalDataTransform

transform = SimCLREvalDataTransform(input_height=32)
x = sample()
(xi, xj) = transform(x)
```

```
__call__(sample)
Call self as a function.
```

```
class pl_bolts.models.self_supervised.simclr.simclr_transforms.SimCLRTrainDataTransform(input_height=32, input_width=32, output_size=224, zero_centered=False, s=1.0)
Bases: object
```

Transforms for SimCLR

Transform:

```
RandomResizedCrop(size=self.input_height)
RandomHorizontalFlip()
RandomApply([color_jitter], p=0.8)
RandomGrayscale(p=0.2)
GaussianBlur(kernel_size=int(0.1 * self.input_height))
transforms.ToTensor()
```

Example:

```
from pl_bolts.models.self_supervised.simclr.transforms import_
SimCLRTrainDataTransform

transform = SimCLRTrainDataTransform(input_height=32)
x = sample()
(xi, xj) = transform(x)
```

```
__call__(sample)
Call self as a function.
```

## Submodules

### pl\_bolts.models.self\_supervised.evaluator module

```
class pl_bolts.models.self_supervised.evaluator.Flatten
Bases: torch.nn.Module

forward(input_tensor)

class pl_bolts.models.self_supervised.evaluator.SSLEvaluator(n_input, n_classes,
                                                               n_hidden=512,
                                                               p=0.1)
Bases: torch.nn.Module

forward(x)
```

### pl\_bolts.models.self\_supervised.resnets module

```
class pl_bolts.models.self_supervised.resnets.ResNet(block,
                                                       layers,
                                                       num_classes=1000,
                                                       zero_init_residual=False,
                                                       groups=1,
                                                       width_per_group=64,
                                                       replace_stride_with_dilation=None,
                                                       norm_layer=None,
                                                       return_all_feature_maps=False)
Bases: torch.nn.Module
```

```

_make_layer(block, planes, blocks, stride=1, dilate=False)

forward(x)

pl_bolts.models.self_supervised.resnets.resnet18(pretrained=False, progress=True,
                                                **kwargs)
    ResNet-18 model from “Deep Residual Learning for Image Recognition” :param
    _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet18.pretrained: If True, returns a model pre-
    trained on ImageNet :type _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet18.pretrained:
    bool :param _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet18.progress:
    If True, displays a progress bar of the download to stderr :type
    _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet18.progress: bool

pl_bolts.models.self_supervised.resnets.resnet34(pretrained=False, progress=True,
                                                **kwargs)
    ResNet-34 model from “Deep Residual Learning for Image Recognition” :param
    _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet34.pretrained: If True, returns a model pre-
    trained on ImageNet :type _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet34.pretrained:
    bool :param _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet34.progress:
    If True, displays a progress bar of the download to stderr :type
    _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet34.progress: bool

pl_bolts.models.self_supervised.resnets.resnet50(pretrained=False, progress=True,
                                                **kwargs)
    ResNet-50 model from “Deep Residual Learning for Image Recognition” :param
    _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet50.pretrained: If True, returns a model pre-
    trained on ImageNet :type _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet50.pretrained:
    bool :param _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet50.progress:
    If True, displays a progress bar of the download to stderr :type
    _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet50.progress: bool

pl_bolts.models.self_supervised.resnets.resnet50_bn(pretrained=False,
                                                    progress=True, **kwargs)
    ResNet-50 model from “Deep Residual Learning for Image Recognition” :param
    _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet50_bn.pretrained: If True, returns a model pre-
    trained on ImageNet :type _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet50_bn.pretrained:
    bool :param _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet50_bn.progress:
    If True, displays a progress bar of the download to stderr :type
    _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet50_bn.progress: bool

pl_bolts.models.self_supervised.resnets.resnet101(pretrained=False, progress=True,
                                                **kwargs)
    ResNet-101 model from “Deep Residual Learning for Image Recognition” :param
    _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet101.pretrained: If True, returns a model pre-
    trained on ImageNet :type _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet101.pretrained:
    bool :param _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet101.progress:
    If True, displays a progress bar of the download to stderr :type
    _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet101.progress: bool

pl_bolts.models.self_supervised.resnets.resnet152(pretrained=False, progress=True,
                                                **kwargs)
    ResNet-152 model from “Deep Residual Learning for Image Recognition” :param
    _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet152.pretrained: If True, returns a model pre-
    trained on ImageNet :type _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet152.pretrained:
    bool :param _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet152.progress:
    If True, displays a progress bar of the download to stderr :type
    _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet152.progress: bool

```

```
pl_bolts.models.self_supervised.resnets.resnext50_32x4d(pretrained=False,  
                                         progress=True, **kwargs)  
    ResNeXt-50 32x4d model from “Aggregated Residual Transformation for Deep Neural Networks” :param  
    _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnext50_32x4d.pretrained: If True, returns a  
    model pre-trained on ImageNet :type _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnext50_32x4d.pretrained:  
    bool :param _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnext50_32x4d.progress:  
    If True, displays a progress bar of the download to stderr :type  
    _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnext50_32x4d.progress: bool  
  
pl_bolts.models.self_supervised.resnets.resnext101_32x8d(pretrained=False,  
                                         progress=True,  
                                         **kwargs)  
    ResNeXt-101 32x8d model from “Aggregated Residual Transformation for Deep Neural Networks” :param  
    _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnext101_32x8d.pretrained: If True, returns a  
    model pre-trained on ImageNet :type _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnext101_32x8d.pretrained:  
    bool :param _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnext101_32x8d.progress:  
    If True, displays a progress bar of the download to stderr :type  
    _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnext101_32x8d.progress: bool  
  
pl_bolts.models.self_supervised.resnets.wide_resnet50_2(pretrained=False,  
                                         progress=True, **kwargs)  
    Wide ResNet-50-2 model from “Wide Residual Networks” The model is the same as ResNet  
    except for the bottleneck number of channels which is twice larger in every block. The  
    number of channels in outer 1x1 convolutions is the same, e.g. last block in ResNet-  
    50 has 2048-512-2048 channels, and in Wide ResNet-50-2 has 2048-1024-2048. :param  
    _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.wide_resnet50_2.pretrained: If True, returns a  
    model pre-trained on ImageNet :type _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.wide_resnet50_2.pretrained:  
    bool :param _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.wide_resnet50_2.progress:  
    If True, displays a progress bar of the download to stderr :type  
    _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.wide_resnet50_2.progress: bool  
  
pl_bolts.models.self_supervised.resnets.wide_resnet101_2(pretrained=False,  
                                         progress=True,  
                                         **kwargs)  
    Wide ResNet-101-2 model from “Wide Residual Networks” The model is the same as ResNet  
    except for the bottleneck number of channels which is twice larger in every block. The  
    number of channels in outer 1x1 convolutions is the same, e.g. last block in ResNet-  
    50 has 2048-512-2048 channels, and in Wide ResNet-50-2 has 2048-1024-2048. :param  
    _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.wide_resnet101_2.pretrained: If True, returns a  
    model pre-trained on ImageNet :type _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.wide_resnet101_2.pretrained:  
    bool :param _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.wide_resnet101_2.progress:  
    If True, displays a progress bar of the download to stderr :type  
    _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.wide_resnet101_2.progress: bool
```

## pl\_bolts.models.self\_supervised.ssl\_finetuner module

```
class pl_bolts.models.self_supervised.ssl_finetuner.SSLFineTuner(backbone,  
                                         in_features,  
                                         num_classes,  
                                         hid-  
                                         den_dim=1024)
```

Bases: `pytorch_lightning.LightningModule`

Finetunes a self-supervised learning backbone using the standard evaluation protocol of a singler layer MLP  
with 1024 units

Example:

```
from pl_bolts.utils.self_supervised import SSLFineTuner
from pl_bolts.models.self_supervised import CPCV2
from pl_bolts.datamodules import CIFAR10DataModule
from pl_bolts.models.self_supervised.cpc.transforms import_
    CPCEvalTransformsCIFAR10,
    CPCTrainTransformsCIFAR10

# pretrained model
backbone = CPCV2.load_from_checkpoint(PATH, strict=False)

# dataset + transforms
dm = CIFAR10DataModule(data_dir='.')
dm.train_transforms = CPCTrainTransformsCIFAR10()
dm.val_transforms = CPCEvalTransformsCIFAR10()

# finetuner
finetuner = SSLFineTuner(backbone, in_features=backbone.z_dim, num_
    classes=backbone.num_classes)

# train
trainer = pl.Trainer()
trainer.fit(finetuner, dm)

# test
trainer.test(datamodule=dm)
```

## Parameters

- **backbone** – a pretrained model
- **in\_features** – feature dim of backbone outputs
- **num\_classes** – classes of the dataset
- **hidden\_dim** – dim of the MLP (1024 default used in self-supervised literature)

**configure\_optimizers()**

**on\_train\_epoch\_start()**

**Return type** None

**shared\_step**(batch)

**test\_step**(batch, batch\_idx)

**training\_step**(batch, batch\_idx)

**validation\_step**(batch, batch\_idx)

## pl\_bolts.models.vision package

### Subpackages

#### pl\_bolts.models.vision.image\_gpt package

### Submodules

#### pl\_bolts.models.vision.image\_gpt.gpt2 module

**class** pl\_bolts.models.vision.image\_gpt.gpt2.**Block** (*embed\_dim, heads*)

Bases: torch.nn.Module

**forward** (*x*)

**class** pl\_bolts.models.vision.image\_gpt.gpt2.**GPT2** (*embed\_dim, heads, layers, num\_positions, vocab\_size, num\_classes*)

Bases: pytorch\_lightning.LightningModule

GPT-2 from language Models are Unsupervised Multitask Learners

Paper by: Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever

Implementation contributed by:

- Teddy Koker

Example:

```
from pl_bolts.models import GPT2

seq_len = 17
batch_size = 32
vocab_size = 16
x = torch.randint(0, vocab_size, (seq_len, batch_size))
model = GPT2(embed_dim=32, heads=2, layers=2, num_positions=seq_len, vocab_
    ↴size=vocab_size, num_classes=4)
results = model(x)
```

**\_init\_embeddings()**

**\_init\_layers()**

**\_init\_sos\_token()**

**forward** (*x, classify=False*)

Expect input as shape [sequence len, batch] If classify, return classification logits

## pl\_bolts.models.vision.image\_gpt.igpt\_module module

```
class pl_bolts.models.vision.image_gpt.igpt_module.ImageGPT(datamodule=None,
                                                               embed_dim=16,
                                                               heads=2,      layers=2,      pixels=28,
                                                               vocab_size=16,
                                                               num_classes=10,
                                                               classify=False,
                                                               batch_size=64,
                                                               learning_rate=0.01,
                                                               steps=25000,
                                                               data_dir='.',
                                                               num_workers=8,
                                                               **kwargs)
```

Bases: `pytorch_lightning.LightningModule`

**Paper:** Generative Pretraining from Pixels [original paper [code](#)].

**Paper by:** Mark Che, Alec Radford, Rewon Child, Jeff Wu, Heewoo Jun, Prafulla Dhariwal, David Luan, Ilya Sutskever

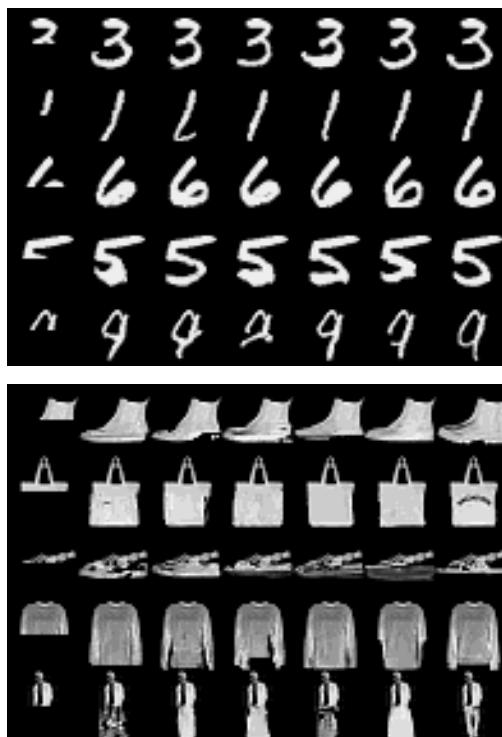
**Implementation contributed by:**

- Teddy Koker

**Original repo with results and more implementation details:**

- <https://github.com/teddykoker/image-gpt>

**Example Results (Photo credits: Teddy Koker):**



**Default arguments:**

Table 1: Argument Defaults

Argument	Default	iGPT-S (Chen et al.)
<code>-embed_dim</code>	16	512
<code>-heads</code>	2	8
<code>-layers</code>	8	24
<code>-pixels</code>	28	32
<code>-vocab_size</code>	16	512
<code>-num_classes</code>	10	10
<code>-batch_size</code>	64	128
<code>-learning_rate</code>	0.01	0.01
<code>-steps</code>	25000	1000000

Example:

```
import pytorch_lightning as pl
from pl_bolts.models.vision import ImageGPT

dm = MNISTDataModule('.')
model = ImageGPT(dm)

pl.Trainer(gpu=4).fit(model)
```

As script:

```
cd pl_bolts/models/vision/image_gpt
python igpt_module.py --learning_rate 1e-2 --batch_size 32 --gpus 4
```

## Parameters

- `datamodule` (`Optional[LightningDataModule]`) – LightningDataModule
- `embed_dim` (`int`) – the embedding dim
- `heads` (`int`) – number of attention heads
- `layers` (`int`) – number of layers
- `pixels` (`int`) – number of input pixels
- `vocab_size` (`int`) – vocab size
- `num_classes` (`int`) – number of classes in the input
- `classify` (`bool`) – true if should classify
- `batch_size` (`int`) – the batch size
- `learning_rate` (`float`) – learning rate
- `steps` (`int`) – number of steps for cosine annealing
- `data_dir` (`str`) – where to store data
- `num_workers` (`int`) – num\_data workers

```
static add_model_specific_args(parent_parser)
configure_optimizers()
forward(x, classify=False)
```

```

test_epoch_end(outs)
test_step(batch, batch_idx)
training_step(batch, batch_idx)
validation_epoch_end(outs)
validation_step(batch, batch_idx)

pl_bolts.models.vision.image_gpt.igpt_module._shape_input(x)
    shape batch of images for input into GPT2 model

pl_bolts.models.vision.image_gpt.igpt_module.cli_main()

```

## Submodules

### pl\_bolts.models.vision.pixel\_cnn module

PixelCNN Implemented by: William Falcon Reference: <https://arxiv.org/pdf/1905.09272.pdf> (page 15) Accessed: May 14, 2020

```

class pl_bolts.models.vision.pixel_cnn.PixelCNN(input_channels, hid-
den_channels=256, num_blocks=5)

```

Bases: `torch.nn.Module`

Implementation of Pixel CNN.

Paper authors: Aaron van den Oord, Nal Kalchbrenner, Oriol Vinyals, Lasse Espeholt, Alex Graves, Koray Kavukcuoglu

Implemented by:

- William Falcon

Example:

```

>>> from pl_bolts.models.vision import PixelCNN
>>> import torch
...
>>> model = PixelCNN(input_channels=3)
>>> x = torch.rand(5, 3, 64, 64)
>>> out = model(x)
...
>>> out.shape
torch.Size([5, 3, 64, 64])

```

`conv_block`(*input\_channels*)

`forward`(*z*)

## 25.5.2 Submodules

### pl\_bolts.models.mnist\_module module

```
class pl_bolts.models.mnist_module.LitMNIST(hidden_dim=128,      learning_rate=0.001,
                                             batch_size=32,           num_workers=4,
                                             data_dir='', **kwargs)
Bases: pytorch_lightning.LightningModule
static add_model_specific_args(parent_parser)
configure_optimizers()
forward(x)
prepare_data()
test_dataloader()
test_epoch_end(outputs)
test_step(batch, batch_idx)
train_dataloader()
training_step(batch, batch_idx)
val_dataloader()
validation_epoch_end(outputs)
validation_step(batch, batch_idx)
pl_bolts.models.mnist_module.cli_main()
```

## 25.6 pl\_bolts.losses package

### 25.6.1 Submodules

#### pl\_bolts.losses.self\_supervised\_learning module

```
class pl_bolts.losses.self_supervised_learning.AmdimNCELoss(tclip)
Bases: torch.nn.Module
forward(anchor_representations, positive_representations, mask_mat)
    Compute the NCE scores for predicting r_src->r_trg. :param
    _sphinx_paramlinks_pl_bolts.losses.self_supervised_learning.AmdimNCELoss.forward.anchor_representations:
    (batch_size, emb_dim) :param _sphinx_paramlinks_pl_bolts.losses.self_supervised_learning.AmdimNCELoss.forward.posi-
    (emb_dim, n_batch * w* h) (ie: nb_feat_vectors x embedding_dim) :param
    _sphinx_paramlinks_pl_bolts.losses.self_supervised_learning.AmdimNCELoss.forward.mask_mat:
    (n_batch_gpu, n_batch)

    Output: raw_scores : (n_batch_gpu, n_locs) nce_scores : (n_batch_gpu, n_locs) lgt_reg : scalar
class pl_bolts.losses.self_supervised_learning.CPCTask(num_input_channels,
                                                       target_dim=64,          em-
                                                       bed_scale=0.1)
Bases: torch.nn.Module
Loss used in CPC
```

```
compute_loss_h(targets, preds, i)
forward(Z)

class pl_bolts.losses.self_supervised_learning.FeatureMapContrastiveTask(comparisons='00,
    11',
    tclip=10.0,
    bidi-
    rec-
    tional=True)

Bases: torch.nn.Module
```

Performs an anchor, positive negative pair comparison for each each tuple of feature maps passed.

```
# extract feature maps
pos_0, pos_1, pos_2 = encoder(x_pos)
anc_0, anc_1, anc_2 = encoder(x_anchor)

# compare only the 0th feature maps
task = FeatureMapContrastiveTask('00')
loss, regularizer = task((pos_0), (anc_0))

# compare (pos_0 to anc_1) and (pos_0, anc_2)
task = FeatureMapContrastiveTask('01, 02')
losses, regularizer = task((pos_0, pos_1, pos_2), (anc_0, anc_1, anc_2))
loss = losses.sum()

# compare (pos_1 vs a anc_random)
task = FeatureMapContrastiveTask('0r')
loss, regularizer = task((pos_0, pos_1, pos_2), (anc_0, anc_1, anc_2))
```

### Parameters

- **comparisons** (`str`) – groupings of feature map indices to compare (zero indexed, ‘r’ means random) ex: ‘00, 1r’
- **tclip** (`float`) – stability clipping value
- **bidirectional** (`bool`) – if true, does the comparison both ways

```
# with bidirectional the comparisons are done both ways
task = FeatureMapContrastiveTask('01, 02')

# will compare the following:
# 01: (pos_0, anc_1), (anc_0, pos_1)
# 02: (pos_0, anc_2), (anc_0, pos_2)
```

```
_FeatureMapContrastiveTask__cache_dimension_masks(*args)
_FeatureMapContrastiveTask__compare_maps(m1, m2)
_sample_src_ftr(r_cnv, masks)
feat_size_w_mask(w, feature_map)
forward(anchor_maps, positive_maps)
```

Takes in a set of tuples, each tuple has two feature maps with all matching dimensions

## Example

```
>>> import torch
>>> from pytorch_lightning import seed_everything
>>> seed_everything(0)
0
>>> a1 = torch.rand(3, 5, 2, 2)
>>> a2 = torch.rand(3, 5, 2, 2)
>>> b1 = torch.rand(3, 5, 2, 2)
>>> b2 = torch.rand(3, 5, 2, 2)
...
>>> task = FeatureMapContrastiveTask('01', '11')
...
>>> losses, regularizer = task((a1, a2), (b1, b2))
>>> losses
tensor([2.2351, 2.1902])
>>> regularizer
tensor(0.0324)
```

**static parse\_map\_indexes**(comparisons)

Example:

```
>>> FeatureMapContrastiveTask.parse_map_indexes('11')
[(1, 1)]
>>> FeatureMapContrastiveTask.parse_map_indexes('11,59')
[(1, 1), (5, 9)]
>>> FeatureMapContrastiveTask.parse_map_indexes('11,59, 2r')
[(1, 1), (5, 9), (2, -1)]
```

pl\_bolts.losses.self\_supervised\_learning.**nt\_xent\_loss**(out\_1, out\_2, temperature)  
Loss used in SimCLR

pl\_bolts.losses.self\_supervised\_learning.**tanh\_clip**(x, clip\_val=10.0)  
soft clip values to the range [-clip\_val, +clip\_val]

## 25.7 pl\_bolts.loggers package

Collection of PyTorchLightning loggers

```
class pl_bolts.loggers.TrainsLogger(project_name=None, task_name=None,
                                      task_type='training', reuse_last_task_id=True, output_uri=None,
                                      auto_connect_arg_parser=True, auto_connect_frameworks=True,
                                      auto_resource_monitoring=True)
```

Bases: pytorch\_lightning.loggers.base.LightningLoggerBase

Log using [allegro.ai TRAINS](#). Install it with pip:

```
pip install trains
```

## Example

```
>>> from pytorch_lightning import Trainer
>>> trains_logger = TrainsLogger(
...     project_name='pytorch lightning',
...     task_name='default',
...     output_uri='.',
... )
TRAINNS Task: ...
TRAINNS results page: ...
>>> trainer = Trainer(logger=trains_logger)
```

Use the logger anywhere in your LightningModule as follows:

```
>>> from pytorch_lightning import LightningModule
>>> class LitModel(LightningModule):
...     def training_step(self, batch, batch_idx):
...         # example
...         self.logger.experiment.whatever_trains_supports(...)
...
...     def any_lightning_module_function_or_hook(self):
...         self.logger.experiment.whatever_trains_supports(...)
```

## Parameters

- **project\_name** (`Optional[str]`) – The name of the experiment's project. Defaults to `None`.
- **task\_name** (`Optional[str]`) – The name of the experiment. Defaults to `None`.
- **task\_type** (`str`) – The name of the experiment. Defaults to '`training`'.
- **reuse\_last\_task\_id** (`bool`) – Start with the previously used task id. Defaults to `True`.
- **output\_uri** (`Optional[str]`) – Default location for output models. Defaults to `None`.
- **auto\_connect\_arg\_parser** (`bool`) – Automatically grab the `ArgumentParser` and connect it with the task. Defaults to `True`.
- **auto\_connect\_frameworks** (`bool`) – If `True`, automatically patch to trains back-end. Defaults to `True`.
- **auto\_resource\_monitoring** (`bool`) – If `True`, machine vitals will be sent along side the task scalars. Defaults to `True`.

## Examples

```
>>> logger = TrainsLogger("pytorch lightning", "default", output_uri=".")
TRAINNS Task: ...
TRAINNS results page: ...
>>> logger.log_metrics({"val_loss": 1.23}, step=0)
>>> logger.log_text("sample test")
sample test
>>> import numpy as np
>>> logger.log_artifact("confusion matrix", np.ones((2, 3)))
>>> logger.log_image("passed", "Image 1", np.random.randint(0, 255, (200, 150, 3),
...     dtype=np.uint8))
```

```
classmethod bypass_mode()  
    Returns the bypass mode state.
```

**Note:** `GITHUB_ACTIONS` env will automatically set `bypass_mode` to `True` unless overridden specifically with `TrainsLogger.set_bypass_mode(False)`.

---

**Return type** `bool`

**Returns** If `True`, all outside communication is skipped.

```
finalize(status=None)
```

**Return type** `None`

```
log_artifact(name, artifact, metadata=None, delete_after_upload=False)
```

Save an artifact (file/object) in TRAINS experiment storage.

#### Parameters

- **name** (`str`) – Artifact name. Notice! it will override the previous artifact if the name already exists.
- **artifact** (`Union[str, Path, Dict[str, Any], ndarray, Image]`) – Artifact object to upload. Currently supports:
  - string / `pathlib.Path` are treated as path to artifact file to upload If a wildcard or a folder is passed, a zip file containing the local files will be created and uploaded.
  - dict will be stored as .json file and uploaded
  - `pandas.DataFrame` will be stored as .csv.gz (compressed CSV file) and uploaded
  - `numpy.ndarray` will be stored as .npz and uploaded
  - `PIL.Image.Image` will be stored to .png file and uploaded
- **metadata** (`Optional[Dict[str, Any]]`) – Simple key/value dictionary to store on the artifact. Defaults to `None`.
- **delete\_after\_upload** (`bool`) – If `True`, the local artifact will be deleted (only applies if artifact is a local file). Defaults to `False`.

**Return type** `None`

```
log_hyperparams(params)
```

Log hyperparameters (numeric values) in TRAINS experiments.

**Parameters** **params** (`Union[Dict[str, Any], Namespace]`) – The hyperparameters that passed through the model.

**Return type** `None`

```
log_image(title, series, image, step=None)
```

Log Debug image in TRAINS experiment

#### Parameters

- **title** (`str`) – The title of the debug image, i.e. “failed”, “passed”.
- **series** (`str`) – The series name of the debug image, i.e. “Image 0”, “Image 1”.
- **image** (`Union[str, ndarray, Image, Tensor]`) – Debug image to log. If `numpy.ndarray` or `torch.Tensor`, the image is assumed to be the following:

- shape: CHW
- color space: RGB
- value range: [0., 1.] (float) or [0, 255] (uint8)
- **step** (`Optional[int]`) – Step number at which the metrics should be recorded. Defaults to None.

**Return type** `None`

**log\_metric** (`title, series, value, step=None`)

Log metrics (numeric values) in TRAINS experiments. This method will be called by the users.

**Parameters**

- **title** (`str`) – The title of the graph to log, e.g. loss, accuracy.
- **series** (`str`) – The series name in the graph, e.g. classification, localization.
- **value** (`float`) – The value to log.
- **step** (`Optional[int]`) – Step number at which the metrics should be recorded. Defaults to None.

**Return type** `None`

**log\_metrics** (`metrics, step=None`)

Log metrics (numeric values) in TRAINS experiments. This method will be called by Trainer.

**Parameters**

- **metrics** (`Dict[str, float]`) – The dictionary of the metrics. If the key contains “/”, it will be split by the delimiter, then the elements will be logged as “title” and “series” respectively.
- **step** (`Optional[int]`) – Step number at which the metrics should be recorded. Defaults to None.

**Return type** `None`

**log\_text** (`text`)

Log console text data in TRAINS experiment.

**Parameters** `text` (`str`) – The value of the log (data-point).

**Return type** `None`

**classmethod set\_bypass\_mode** (`bypass`)

Will bypass all outside communication, and will drop all logs. Should only be used in “standalone mode”, when there is no access to the `trains-server`.

**Parameters** `bypass` (`bool`) – If `True`, all outside communication is skipped.

**Return type** `None`

**classmethod set\_credentials** (`api_host=None, web_host=None, files_host=None, key=None, secret=None`)

Set new default TRAINS-server host and credentials. These configurations could be overridden by either OS environment variables or `trains.conf` configuration file.

---

**Note:** Credentials need to be set *prior* to Logger initialization.

---

**Parameters**

- **api\_host** (`Optional[str]`) – Trains API server url, example: `host='http://localhost:8008'`
- **web\_host** (`Optional[str]`) – Trains WEB server url, example: `host='http://localhost:8080'`
- **files\_host** (`Optional[str]`) – Trains Files server url, example: `host='http://localhost:8081'`
- **key** (`Optional[str]`) – user key/secret pair, example: `key='thisisakey123'`
- **secret** (`Optional[str]`) – user key/secret pair, example: `secret='thisisseceret123'`

**Return type** `None`

`_bypass = None`

**property experiment**

Actual TRAINS object. To use TRAINS features in your `LightningModule` do the following.

Example:

```
self.logger.experiment.some_trains_function()
```

**Return type** `Task`

**property id**

ID is a uuid (string) representing this specific experiment in the entire system.

**Return type** `Optional[str]`

**property name**

Name is a human readable non-unique name (str) of the experiment.

**Return type** `Optional[str]`

**property version**

**Return type** `Optional[str]`

## 25.7.1 Submodules

### pl\_bolts.loggers.trains module

#### TRAINSM

```
class pl_bolts.loggers.trains.TrainsLogger(project_name=None,
                                             task_name=None,          task_type='training',
                                             reuse_last_task_id=True, output_uri=None,
                                             auto_connect_arg_parser=True,
                                             auto_connect_frameworks=True,
                                             auto_resource_monitoring=True)
```

Bases: `pytorch_lightning.loggers.base.LightningLoggerBase`

Log using [allegro.ai TRAINS](#). Install it with pip:

```
pip install trains
```

## Example

```
>>> from pytorch_lightning import Trainer
>>> trains_logger = TrainsLogger(
...     project_name='pytorch lightning',
...     task_name='default',
...     output_uri='.',
... )
TRAINNS Task: ...
TRAINNS results page: ...
>>> trainer = Trainer(logger=trains_logger)
```

Use the logger anywhere in your LightningModule as follows:

```
>>> from pytorch_lightning import LightningModule
>>> class LitModel(LightningModule):
...     def training_step(self, batch, batch_idx):
...         # example
...         self.logger.experiment.whatever_trains_supports(...)
...
...     def any_lightning_module_function_or_hook(self):
...         self.logger.experiment.whatever_trains_supports(...)
```

## Parameters

- **project\_name** (`Optional[str]`) – The name of the experiment's project. Defaults to `None`.
- **task\_name** (`Optional[str]`) – The name of the experiment. Defaults to `None`.
- **task\_type** (`str`) – The name of the experiment. Defaults to '`training`'.
- **reuse\_last\_task\_id** (`bool`) – Start with the previously used task id. Defaults to `True`.
- **output\_uri** (`Optional[str]`) – Default location for output models. Defaults to `None`.
- **auto\_connect\_arg\_parser** (`bool`) – Automatically grab the `ArgumentParser` and connect it with the task. Defaults to `True`.
- **auto\_connect\_frameworks** (`bool`) – If `True`, automatically patch to trains back-end. Defaults to `True`.
- **auto\_resource\_monitoring** (`bool`) – If `True`, machine vitals will be sent along side the task scalars. Defaults to `True`.

## Examples

```
>>> logger = TrainsLogger("pytorch lightning", "default", output_uri=".")
TRAINNS Task: ...
TRAINNS results page: ...
>>> logger.log_metrics({"val_loss": 1.23}, step=0)
>>> logger.log_text("sample test")
sample test
>>> import numpy as np
>>> logger.log_artifact("confusion matrix", np.ones((2, 3)))
>>> logger.log_image("passed", "Image 1", np.random.randint(0, 255, (200, 150, 3),
...     dtype=np.uint8))
```

```
classmethod bypass_mode()  
    Returns the bypass mode state.
```

**Note:** `GITHUB_ACTIONS` env will automatically set `bypass_mode` to `True` unless overridden specifically with `TrainsLogger.set_bypass_mode(False)`.

---

**Return type** `bool`

**Returns** If `True`, all outside communication is skipped.

```
finalize(status=None)
```

**Return type** `None`

```
log_artifact(name, artifact, metadata=None, delete_after_upload=False)
```

Save an artifact (file/object) in TRAINS experiment storage.

#### Parameters

- **name** (`str`) – Artifact name. Notice! it will override the previous artifact if the name already exists.
- **artifact** (`Union[str, Path, Dict[str, Any], ndarray, Image]`) – Artifact object to upload. Currently supports:
  - string / `pathlib.Path` are treated as path to artifact file to upload If a wildcard or a folder is passed, a zip file containing the local files will be created and uploaded.
  - dict will be stored as .json file and uploaded
  - `pandas.DataFrame` will be stored as .csv.gz (compressed CSV file) and uploaded
  - `numpy.ndarray` will be stored as .npz and uploaded
  - `PIL.Image.Image` will be stored to .png file and uploaded
- **metadata** (`Optional[Dict[str, Any]]`) – Simple key/value dictionary to store on the artifact. Defaults to `None`.
- **delete\_after\_upload** (`bool`) – If `True`, the local artifact will be deleted (only applies if artifact is a local file). Defaults to `False`.

**Return type** `None`

```
log_hyperparams(params)
```

Log hyperparameters (numeric values) in TRAINS experiments.

**Parameters** **params** (`Union[Dict[str, Any], Namespace]`) – The hyperparameters that passed through the model.

**Return type** `None`

```
log_image(title, series, image, step=None)
```

Log Debug image in TRAINS experiment

#### Parameters

- **title** (`str`) – The title of the debug image, i.e. “failed”, “passed”.
- **series** (`str`) – The series name of the debug image, i.e. “Image 0”, “Image 1”.
- **image** (`Union[str, ndarray, Image, Tensor]`) – Debug image to log. If `numpy.ndarray` or `torch.Tensor`, the image is assumed to be the following:

- shape: CHW
- color space: RGB
- value range: [0., 1.] (float) or [0, 255] (uint8)
- **step** (`Optional[int]`) – Step number at which the metrics should be recorded. Defaults to None.

**Return type** `None`

**log\_metric** (`title, series, value, step=None`)

Log metrics (numeric values) in TRAINS experiments. This method will be called by the users.

**Parameters**

- **title** (`str`) – The title of the graph to log, e.g. loss, accuracy.
- **series** (`str`) – The series name in the graph, e.g. classification, localization.
- **value** (`float`) – The value to log.
- **step** (`Optional[int]`) – Step number at which the metrics should be recorded. Defaults to None.

**Return type** `None`

**log\_metrics** (`metrics, step=None`)

Log metrics (numeric values) in TRAINS experiments. This method will be called by Trainer.

**Parameters**

- **metrics** (`Dict[str, float]`) – The dictionary of the metrics. If the key contains “/”, it will be split by the delimiter, then the elements will be logged as “title” and “series” respectively.
- **step** (`Optional[int]`) – Step number at which the metrics should be recorded. Defaults to None.

**Return type** `None`

**log\_text** (`text`)

Log console text data in TRAINS experiment.

**Parameters** `text` (`str`) – The value of the log (data-point).

**Return type** `None`

**classmethod set\_bypass\_mode** (`bypass`)

Will bypass all outside communication, and will drop all logs. Should only be used in “standalone mode”, when there is no access to the `trains-server`.

**Parameters** `bypass` (`bool`) – If `True`, all outside communication is skipped.

**Return type** `None`

**classmethod set\_credentials** (`api_host=None, web_host=None, files_host=None, key=None, secret=None`)

Set new default TRAINS-server host and credentials. These configurations could be overridden by either OS environment variables or `trains.conf` configuration file.

---

**Note:** Credentials need to be set *prior* to Logger initialization.

---

**Parameters**

- **api\_host** (`Optional[str]`) – Trains API server url, example: `host='http://localhost:8008'`
- **web\_host** (`Optional[str]`) – Trains WEB server url, example: `host='http://localhost:8080'`
- **files\_host** (`Optional[str]`) – Trains Files server url, example: `host='http://localhost:8081'`
- **key** (`Optional[str]`) – user key/secret pair, example: `key='thisisakey123'`
- **secret** (`Optional[str]`) – user key/secret pair, example: `secret='thisisseceret123'`

**Return type** `None`

`_bypass = None`

**property experiment**

Actual TRAINS object. To use TRAINS features in your `LightningModule` do the following.

Example:

```
self.logger.experiment.some_trains_function()
```

**Return type** `Task`

**property id**

ID is a uuid (string) representing this specific experiment in the entire system.

**Return type** `Optional[str]`

**property name**

Name is a human readable non-unique name (str) of the experiment.

**Return type** `Optional[str]`

**property version**

**Return type** `Optional[str]`

## 25.8 pl\_bolts.optimizers package

### 25.8.1 Submodules

#### pl\_bolts.optimizers.lars\_scheduling module

#### References

- <https://github.com/NVIDIA/apex/blob/master/apex/parallel/LARC.py>
- <https://arxiv.org/pdf/1708.03888.pdf>
- <https://github.com/noahgolmant/pytorch-lars/blob/master/lars.py>

**class** `pl_bolts.optimizers.lars_scheduling.LARSWrapper(optimizer, eta=0.02, clip=True, eps=1e-08)`

Bases: `object`

Wrapper that adds LARS scheduling to any optimizer. This helps stability with huge batch sizes.

**Parameters**

- **optimizer** – torch optimizer
- **eta** – LARS coefficient (trust)
- **clip** – True to clip LR
- **eps** – adaptive\_lr stability coefficient

```
step()
update_p(p, group, weight_decay)
property param_groups
property state
```

**pl\_bolts.optimizers.lr\_scheduler module**

```
class pl_bolts.optimizers.lr_scheduler.LinearWarmupCosineAnnealingLR(optimizer,
warmup_epochs,
max_epochs,
warmup_start_lr=0.0,
eta_min=0.0,
last_epoch=-1)
```

Bases: torch.optim.lr\_scheduler.\_LRScheduler

Sets the learning rate of each parameter group to follow a linear warmup schedule between `warmup_start_lr` and `base_lr` followed by a cosine annealing schedule between `base_lr` and `eta_min`.

**Warning:** It is recommended to call `step()` for `LinearWarmupCosineAnnealingLR` after each iteration as calling it after each epoch will keep the starting lr at `warmup_start_lr` for the first epoch which is 0 in most cases.

**Warning:** passing epoch to `step()` is being deprecated and comes with an EPOCH\_DEPRECATED\_WARNING. It calls the `_get_closed_form_lr()` method for this scheduler instead of `get_lr()`. Though this does not change the behavior of the scheduler, when passing epoch param to `step()`, the user should call the `step()` function before calling train and validation methods.

**Parameters**

- **optimizer** (*Optimizer*) – Wrapped optimizer.
- **warmup\_epochs** (*int*) – Maximum number of iterations for linear warmup
- **max\_epochs** (*int*) – Maximum number of iterations
- **warmup\_start\_lr** (*float*) – Learning rate to start the linear warmup. Default: 0.
- **eta\_min** (*float*) – Minimum learning rate. Default: 0.
- **last\_epoch** (*int*) – The index of last epoch. Default: -1.

## Example

```
>>> layer = nn.Linear(10, 1)
>>> optimizer = Adam(layer.parameters(), lr=0.02)
>>> scheduler = LinearWarmupCosineAnnealingLR(optimizer, warmup_epochs=10, max_
->epochs=40)
>>> #
>>> # the default case
>>> for epoch in range(40):
...     # train(...)
...     # validate(...)
...     scheduler.step()
>>> #
>>> # passing epoch param case
>>> for epoch in range(40):
...     scheduler.step(epoch)
...     # train(...)
...     # validate(...)
```

### `_get_closed_form_lr()`

Called when epoch is passed as a param to the `step` function of the scheduler.

**Return type** `List[float]`

### `get_lr()`

Compute learning rate using chainable form of the scheduler

**Return type** `List[float]`

## 25.9 `pl_bolts.transforms` package

### 25.9.1 Subpackages

#### `pl_bolts.transforms.self_supervised` package

##### Submodules

##### `pl_bolts.transforms.self_supervised.ssl_transforms` module

```
class pl_bolts.transforms.self_supervised.ssl_transforms.Patchify(patch_size,
                                                               over-
                                                               lap_size)
```

Bases: `object`

##### `__call__(x)`

Call self as a function.

```
class pl_bolts.transforms.self_supervised.ssl_transforms.RandomTranslateWithReflect(max_trans_
Bases: object
```

Translate image randomly Translate vertically and horizontally by n pixels where n is integer drawn uniformly independently for each axis from [-max\_translation, max\_translation]. Fill the uncovered blank area with reflect padding.

##### `__call__(old_image)`

Call self as a function.

## 25.9.2 Submodules

### pl\_bolts.transforms.dataset\_normalizations module

```
pl_bolts.transforms.dataset_normalizations.cifar10_normalization()
pl_bolts.transforms.dataset_normalizations.imagenet_normalization()
pl_bolts.transforms.dataset_normalizations.stl10_normalization()
```

## 25.10 pl\_bolts.utils package

### 25.10.1 Submodules

#### pl\_bolts.utils.pretrained\_weights module

```
pl_bolts.utils.pretrained_weights.load_pretrained(model, class_name=None)
```

#### pl\_bolts.utils.self\_supervised module

```
pl_bolts.utils.self_supervised.torchvision_ssl_encoder(name, pretrained=False, return_all_feature_maps=False)
```

#### pl\_bolts.utils.semi\_supervised module

**class** pl\_bolts.utils.semi\_supervised.Identity

Bases: torch.nn.Module

An identity class to replace arbitrary layers in pretrained models

Example:

```
from pl_bolts.utils import Identity

model = resnet18()
model.fc = Identity()
```

**forward**(x)

pl\_bolts.utils.semi\_supervised.balance\_classes(X, Y, batch\_size)

Makes sure each batch has an equal amount of data from each class. Perfect balance

#### Parameters

- **X** (ndarray) – input features
- **Y** (list) – mixed labels (ints)
- **batch\_size** (int) – the ultimate batch size

```
pl_bolts.utils.semi_supervised.generate_half_labeled_batches(smaller_set_X,
                                                               smaller_set_Y,
                                                               larger_set_X,
                                                               larger_set_Y,
                                                               batch_size)
```

Given a labeled dataset and an unlabeled dataset, this function generates a joint pair where half the batches are

labeled and the other half is not

### pl\_bolts.utils.shaping module

`pl_bolts.utils.shaping.tile(a, dim, n_tile)`

## PYTHON MODULE INDEX

### p

pl\_bolts.callbacks, 112  
pl\_bolts.callbacks.printing, 114  
pl\_bolts.callbacks.self\_supervised, 115  
pl\_bolts.callbacks.variational, 116  
pl\_bolts.callbacks.vision, 112  
pl\_bolts.callbacks.vision.confused\_logit, 112  
pl\_bolts.callbacks.vision.image\_generation, 113  
pl\_bolts.datamodules, 117  
pl\_bolts.datamodules.async\_dataloader, 117  
pl\_bolts.datamodules.base\_dataset, 117  
pl\_bolts.datamodules.binary\_mnist\_datamodule, 118  
pl\_bolts.datamodules.cifar10\_datamodule, 120  
pl\_bolts.datamodules.cifar10\_dataset, 122  
pl\_bolts.datamodules.cityscapes\_datamodule, 124  
pl\_bolts.datamodules.concat\_dataset, 126  
pl\_bolts.datamodules.dummy\_dataset, 126  
pl\_bolts.datamodules.fashion\_mnist\_datamodule, 127  
pl\_bolts.datamodules.imagenet\_datamodule, 128  
pl\_bolts.datamodules.imagenet\_dataset, 131  
pl\_bolts.datamodules.mnist\_datamodule, 132  
pl\_bolts.datamodules.sklearn\_datamodule, 133  
pl\_bolts.datamodules.ssl\_imagenet\_datamodule, 137  
pl\_bolts.datamodules.stl10\_datamodule, 137  
pl\_bolts.datamodules.vocdetection\_datamodule, 139  
pl\_bolts.loggers, 180  
pl\_bolts.loggers.trains, 184  
pl\_bolts.losses, 178  
pl\_bolts.losses.self\_supervised\_learning, 178  
pl\_bolts.metrics, 140  
pl\_bolts.metrics.aggregation, 140  
pl\_bolts.models, 140  
pl\_bolts.models.autoencoders, 140  
pl\_bolts.models.autoencoders.basic\_ae, 140  
pl\_bolts.models.autoencoders.basic\_ae.components, 141  
pl\_bolts.models.autoencoders.basic\_ae.basic\_ae\_module, 141  
pl\_bolts.models.autoencoders.basic\_vae, 143  
pl\_bolts.models.autoencoders.basic\_vae.basic\_vae\_module, 143  
pl\_bolts.models.autoencoders.basic\_vae.components, 144  
pl\_bolts.models.detection, 145  
pl\_bolts.models.detection.faster\_rcnn, 145  
pl\_bolts.models.gans, 146  
pl\_bolts.models.gans.basic, 146  
pl\_bolts.models.gans.basic.basic\_gan\_module, 146  
pl\_bolts.models.gans.basic.components, 148  
pl\_bolts.models.mnist\_module, 178  
pl\_bolts.models.regression, 148  
pl\_bolts.models.regression.linear\_regression, 148  
pl\_bolts.models.regression.logistic\_regression, 149  
pl\_bolts.models.self\_supervised, 150  
pl\_bolts.models.self\_supervised.amdim, 150  
pl\_bolts.models.self\_supervised.amdim.amdim\_module, 150  
pl\_bolts.models.self\_supervised.amdim.datasets, 152  
pl\_bolts.models.self\_supervised.amdim.networks, 152

```
    152
pl_bolts.models.self_supervised.amdim.ssl_transforms, 190
    153           pl_bolts.transforms.self_supervised.ssl_transforms,
pl_bolts.models.self_supervised.amdim.transforms, 190
    154           pl_bolts.utils, 191
pl_bolts.models.self_supervised.byol,      pl_bolts.utils.pretrained_weights, 191
    156           pl_bolts.utils.self_supervised, 191
pl_bolts.models.self_supervised.byol.byol,  pl_bolts.utils.semi_supervised, 191
    156           pl_bolts.utils.shaping, 192
pl_bolts.models.self_supervised.byol.models,
    158
pl_bolts.models.self_supervised.cpc, 158
pl_bolts.models.self_supervised.cpc.cpc_finetuner,
    158
pl_bolts.models.self_supervised.cpc.cpc_module,
    158
pl_bolts.models.self_supervised.cpc.networks,
    161
pl_bolts.models.self_supervised.cpc.transforms,
    161
pl_bolts.models.self_supervised.evaluator,
    170
pl_bolts.models.self_supervised.moco,
    165
pl_bolts.models.self_supervised.moco.callbacks,
    165
pl_bolts.models.self_supervised.moco.moco2_module,
    165
pl_bolts.models.self_supervised.moco.transforms,
    167
pl_bolts.models.self_supervised.resnets,
    170
pl_bolts.models.self_supervised.simclr,
    168
pl_bolts.models.self_supervised.simclr.simclr_finetuner,
    168
pl_bolts.models.self_supervised.simclr.simclr_module,
    168
pl_bolts.models.self_supervised.simclr.simclr_transforms,
    169
pl_bolts.models.self_supervised.ssl_finetuner,
    172
pl_bolts.models.vision, 174
pl_bolts.models.vision.image_gpt, 174
pl_bolts.models.vision.image_gpt.gpt2,
    174
pl_bolts.models.vision.image_gpt.igpt_module,
    175
pl_bolts.models.vision.pixel_cnn, 177
pl_bolts.optimizers, 188
pl_bolts.optimizers.lars_scheduling, 188
pl_bolts.optimizers.lr_scheduler, 189
pl_bolts.transforms, 190
pl_bolts.transforms.dataset_normalizations,
```

# INDEX

## Symbols

\_CPCV2\_\_compute\_final\_nb\_c ()  
    (*pl\_bolts.models.self\_supervised.cpc.cpc\_module.CPCV2 method*), 160

\_CPCV2\_\_recover\_z\_shape ()  
    (*pl\_bolts.models.self\_supervised.cpc.cpc\_module.CPCV2 method*), 160

\_ConfusedLogitCallback\_\_draw\_sample ()  
    (*pl\_bolts.callbacks.vision.confused\_logit.ConfusedLogitCallback static method*), 113

\_FeatureMapContrastiveTask\_\_cache\_dimension\_masks ()  
    (*pl\_bolts.losses.self\_supervised\_learning.FeatureMapContrastiveTask method*), 179

\_FeatureMapContrastiveTask\_\_compare\_maps ()  
    (*pl\_bolts.losses.self\_supervised\_learning.FeatureMapContrastiveTask method*), 179

\_LNBottleneck\_\_init\_layer\_norms ()  
    (*pl\_bolts.models.self\_supervised.cpc.networks.LNBottleneck method*), 161

\_VAE\_\_init\_system ()  
    (*pl\_bolts.models.autoencoders.basic\_vae.basic\_vae\_module.VAE method*), 144

\_VAE\_\_set\_pretrained\_dims ()  
    (*pl\_bolts.models.autoencoders.basic\_vae.basic\_vae\_module.VAE method*), 144

\_\_call\_\_ () (*pl\_bolts.datamodules.vocdetection\_datamodule.Compose method*), 139

\_\_call\_\_ () (*pl\_bolts.models.self\_supervised.amdim.transforms.AMDIMTrainTransformsCIFAR10 method*), 154

\_\_call\_\_ () (*pl\_bolts.models.self\_supervised.amdim.transforms.AMDIMTrainTransformsImageNet128 method*), 154

\_\_call\_\_ () (*pl\_bolts.models.self\_supervised.amdim.transforms.AMDIMTrainTransformsImageNet128 method*), 155

\_\_call\_\_ () (*pl\_bolts.models.self\_supervised.amdim.transforms.AMDIMTrainTransformsImageNet128 method*), 155

\_\_call\_\_ () (*pl\_bolts.models.self\_supervised.amdim.transforms.AMDIMTrainTransformsImageNet128 method*), 156

\_\_call\_\_ () (*pl\_bolts.models.self\_supervised.amdim.transforms.AMDIMTrainTransformsImageNet128 method*), 156

\_\_call\_\_ () (*pl\_bolts.models.self\_supervised.cpc.transforms.CPCEvalTransformsCIFAR10 method*), 162

\_\_call\_\_ () (*pl\_bolts.models.self\_supervised.cpc.transforms.CPCEvalTransformsImageNet128 method*), 188

*method*), 162

    \_\_call\_\_ () (*pl\_bolts.models.self\_supervised.cpc.transforms.CPCEvalTrainLogger method*), 163

    \_\_call\_\_ () (*pl\_bolts.models.self\_supervised.cpc.transforms.CPCTrainLogger method*), 163

*method*), 164

    \_\_call\_\_ () (*pl\_bolts.models.self\_supervised.cpc.transforms.CPCTrainLogger static method*), 164

*method*), 167

    \_\_call\_\_ () (*pl\_bolts.models.self\_supervised.moco.transforms.Moco2EvalTransformsImageNet128 method*), 167

    \_\_call\_\_ () (*pl\_bolts.models.self\_supervised.moco.transforms.Moco2EvalTransformsImageNet128 method*), 167

*method*), 167

    \_\_call\_\_ () (*pl\_bolts.models.self\_supervised.moco.transforms.Moco2EvalTransformsImageNet128 method*), 167

*method*), 167

*method*), 168

    \_\_call\_\_ () (*pl\_bolts.models.self\_supervised.moco.transforms.Moco2EvalTransformsImageNet128 method*), 168

*method*), 168

*method*), 168

*method*), 169

*method*), 170

*method*), 190

*method*), 190

*method*), 166

*batch\_unshuffle\_ddp* ()

*method*), 166

*method*), 166

*bypass* (*pl\_bolts.loggers.TrainsLogger attribute*), 184

*TrainsLogger attribute*), 188

```
_calculate_md5() (in module _get_closed_form_lr())
    pl_bolts.datamodules.imagenet_dataset), (pl_bolts.optimizers.lr_scheduler.LinearWarmupCosineAnnealing
    131 method), 190
_calculate_output_dim() _init_datasets() (pl_bolts.datamodules.sklearn_datamodule.Sklear
    (pl_bolts.models.autoencoders.basic_ae.components.AEEncoderMethod), 134
    method), 142 _init_embeddings()
    _init_layers() (pl_bolts.models.vision.image_gpt.gpt2.GPT2
    (pl_bolts.models.autoencoders.basic_vae.components.EncoderMethod), 174
    method), 145 _method), 174
_calculate_output_size() _init_layers() (pl_bolts.models.vision.image_gpt.gpt2.GPT2
    (pl_bolts.models.autoencoders.basic_vae.components.DecoderToken)
    method), 144 _method), 174
_check_exists() (pl_bolts.datamodules.cifar10_dataset.CIFAR10Method), 174
    class method), 123 _is_gzip() (in module
    _check_integrity() (in module pl_bolts.datamodules.imagenet_dataset),
    pl_bolts.datamodules.imagenet_dataset), 131 _is_tar()
    131 _is_targz() (in module
    _check_md5() (in module pl_bolts.datamodules.imagenet_dataset),
    pl_bolts.datamodules.imagenet_dataset), 131 _is_tarxz()
    131 _is_zip() (in module
    _collate_fn() (in module pl_bolts.datamodules.vocdetection_datamodule),
    pl_bolts.datamodules.vocdetection_datamodule), 140 _make_layer() (pl_bolts.models.self_supervised.cpc.networks.CPCRes
    _config_modules() (pl_bolts.models.self_supervised.amdim.networks.AMDIMEncoder
    method), 152 _method), 161
    _default_transforms() (pl_bolts.datamodules.binary_mnist_datamodule.BinaryMNISTDataModule
    method), 119 _momentum_update_key_encoder()
    _default_transforms() (pl_bolts.datamodules.fashion_mnist_datamodule.FashionMNISTDataModule
    method), 128 _plot() (pl_bolts.callbacks.vision.confused_logit.ConfusedLogitCallback
    _default_transforms() (pl_bolts.datamodules.mnist_datamodule.MNISTDataModule
    method), 133 _SSLImageDatasetModule
    _default_transforms() (pl_bolts.datamodules.ssl_imagenet_datamodule.SSLImageDatasetModule
    method), 137 _prepare_subset()
    _default_transforms() (pl_bolts.datamodules.base_dataset.LightDataset
    (pl_bolts.datamodules.vocdetection_datamodule.VOCDetectionDatasetModule), 139 _prepare_voc_instance()
    method), 139 _random_bbox() (pl_bolts.datamodules.dummy_dataset.DummyDetectio
    _dequeue_and_enqueue() (pl_bolts.models.self_supervised.moco.moco2_module.MocoV2
    method), 166 _method), 171
    _download_from_url() (pl_bolts.datamodules.base_dataset.LightDataset _run_step() (pl_bolts.models.autoencoders.basic_ae.basic_ae_module
    method), 117 _method), 142
    _evaluate_iou() (in module _run_step() (pl_bolts.models.autoencoders.basic_vae.basic_vae_modu
    pl_bolts.models.detection.faster_rcnn), 146 _method), 144
    _extract_archive_save_torch() _sample_src_ftr()
    (pl_bolts.datamodules.cifar10_dataset.CIFAR10 _pl_bolts.losses.self_supervised_learning.FeatureMapContrastive
    method), 123 _method), 179
    _forwardActs() (pl_bolts.models.self_supervised.amdim.workers.AMDIMFitter
    method), 152 _pl_bolts.models.autoencoders.basic_vae.basic_vae_module.VAE
```

```

        method), 144
_set_default_datamodule()
    (pl_bolts.models.gans.basic.basic_gan_module.GAN
     method), 147
_shape_input() (in module pl_bolts.models.vision.image_gpt.igpt_module), AE (class in pl_bolts.models.autoencoders.basic_ae.basic_ae_module),
    177
_unpickle() (pl_bolts.datamodules.cifar10_dataset.CIFAR10
    method), 123
_verify_archive() (in module pl_bolts.datamodules.imagenet_dataset),
    131
_verify_splits() (pl_bolts.datamodules.imagenet_dataModule
    method), 129
_verify_splits() (pl_bolts.datamodules.ssl_imagenet_datamodul
    method), 137

```

**A**

```

accuracy() (in module pl_bolts.metrics.aggregation), AMDIMEvalTransformsImageNet128 (class in
    140
add_model_specific_args()
    (pl_bolts.models.autoencoders.basic_ae.basic_ae
     static method), 142
add_model_specific_args()
    (pl_bolts.models.autoencoders.basic_vae.basic_vae
     static method), 144
add_model_specific_args()
    (pl_bolts.models.detection.faster_rcnn.FasterRCNN
     static method), 146
add_model_specific_args()
    (pl_bolts.models.gans.basic.basic_gan_module.GAN
     static method), 147
add_model_specific_args()
    (pl_bolts.models.mnist_module.LitMNIST
     static method), 178
add_model_specific_args()
    (pl_bolts.models.regression.linear_regression.LinearRegressio
     static method), 148
add_model_specific_args()
    (pl_bolts.models.regression.logistic_regression.LogisticRegres
     static method), 149
add_model_specific_args()
    (pl_bolts.models.self_supervised.amdim.amdim_
     static method), 151
add_model_specific_args()
    (pl_bolts.models.self_supervised.byol.byol_module.BYOL
     static method), 158
add_model_specific_args()
    (pl_bolts.models.self_supervised.cpc.cpc_module.CPCV2
     static method), 160
add_model_specific_args()
    (pl_bolts.models.self_supervised.moco.moco2_
     static method), 167
add_model_specific_args()

```

**B**

```

        (pl_bolts.models.self_supervised.simclr.simclr_module.SimCLR
         static method), 169
GAN_model_specific_args()
    (pl_bolts.models.vision.image_gpt.igpt_module.ImageGPT
     static method), 176
AE (class in pl_bolts.models.autoencoders.basic_ae.basic_ae_module),
    141
(class in
    pl_bolts.models.autoencoders.basic_ae.components),
    142
AMDIM (class in pl_bolts.models.self_supervised.amdim.amdim_module),
    150
AMDIImagenetDataModule (class in
    pl_bolts.models.self_supervised.amdim.networks),
    152
SSLImagenetDataModule
    (class in
    pl_bolts.models.self_supervised.amdim.transforms),
    154
AMDIEvalTransformsCIFAR10 (class in
    pl_bolts.models.self_supervised.amdim.transforms),
    154
AMDIEvalTransformsSTL10 (class in
    pl_bolts.models.self_supervised.amdim.transforms),
    154
AMDIInCELoss (class in
    pl_bolts.losses.self_supervised_learning),
    178
AMDIMPatchesPretraining (class in
    pl_bolts.models.self_supervised.amdim.datasets),
    152
AMDIMPretraining (class in
    pl_bolts.models.self_supervised.amdim.datasets),
    152
AMDIMTrainTransformsCIFAR10 (class in
    pl_bolts.models.self_supervised.amdim.transforms),
    155
AMDIMTrainTransformsImageNet128 (class in
    pl_bolts.models.self_supervised.amdim.transforms),
    155
AMDIMTrainTransformsSTL10 (class in
    pl_bolts.models.self_supervised.amdim.transforms),
    156
AMDIMTrainTransformsSTL10 (class in
    pl_bolts.datamodules.async_dataloader),
    117
balance_classes() (in module
    pl_bolts.utils.semi_supervised), 191
BASE_URL (pl_bolts.datamodules.cifar10_dataset.CIFAR10
    attribute), 123
FileMNIST (class in
    pl_bolts.datamodules.binary_mnist_datamodule),
    118

```

```

BinaryMNISTDataModule      (class      in  cli_main()          (in      module
                           pl_bolts.datamodules.binary_mnist_datamodule),      pl_bolts.models.mnist_module), 178
                           118
                           cli_main()          (in      module
Block (class in pl_bolts.models.vision.image_gpt.gpt2),           pl_bolts.models.regression.linear_regression),
                           174
                           149
BYOL (class in pl_bolts.models.self_supervised.byol.byol_module) main ()          (in      module
                           156
                           149
BYOLMAWeightUpdate        (class      in  cli_main()          (in      module
                           pl_bolts.callbacks.self_supervised), 115
                           149
                           cli_main()          (in      module
bypass_mode () (pl_bolts.loggers.TrainsLogger           pl_bolts.models.self_supervised.amdim.amdim_module),
                           class method), 185
                           151
bypass_mode () (pl_bolts.loggers.TrainsLogger class   cli_main()          (in      module
                           method), 181
                           158
                           cli_main()          (in      module
C
cache_folder_name          (pl_bolts.datamodules.base_dataset.LightDataset
                           attribute), 117
                           167
                           cli_main()          (in      module
cache_folder_name          (pl_bolts.datamodules.cifar10_dataset.CIFAR10
                           attribute), 123
                           168
                           cli_main()          (in      module
cached_folder_path ()      (pl_bolts.datamodules.base_dataset.LightDataset
                           property), 117
                           169
                           cli_main()          (in      module
CIFAR10 (class in pl_bolts.datamodules.cifar10_dataset),           pl_bolts.models.vision.image_gpt.igpt_module),
                           122
                           177
cifar10 () (pl_bolts.models.self_supervised.amdim.datasets.AMDIM (class in pl_bolts.datamodules.vocdetection_datamodule),
                           static method), 152
                           139
cifar10 () (pl_bolts.models.self_supervised.amdim.datasets.AMDIM (class in pl_bolts.losses.self_supervised_learning.CPCTask
                           static method), 152
                           178
cifar10_normalization ()    (in      module  concat_all_gather ()          (in      module
                           pl_bolts.transforms.dataset_normalizations),
                           167
                           pl_bolts.models.self_supervised.moco.moco2_module),
                           191
cifar10_tiny () (pl_bolts.models.self_supervised.amdim.datasets.AMDIM (class in
                           static method), 152
                           pl_bolts.datamodules.concat_dataset), 126
CIFAR10DataModule          (class      in  configure_optimizers ()
                           pl_bolts.datamodules.cifar10_datamodule),
                           (pl_bolts.models.autoencoders.basic_ae.basic_ae_module.AE
                           method), 142
                           120
CIFAR10Mixed               (class      in  configure_optimizers ()
                           pl_bolts.models.self_supervised.amdim.ssl_datasets),
                           (pl_bolts.models.autoencoders.basic_vae.basic_vae_module.VAE
                           method), 144
                           153
CityscapesDataModule        (class      in  configure_optimizers ()
                           pl_bolts.datamodules.cityscapes_datamodule),
                           (pl_bolts.models.detection.faster_rcnn.FasterRCNN
                           method), 146
                           124
cli_main ()                 (in      module  configure_optimizers ()
                           pl_bolts.models.autoencoders.basic_ae.basic_ae_module),
                           (pl_bolts.models.gans.basic.basic_gan_module.GAN
                           method), 147
                           142
cli_main ()                 (in      module  configure_optimizers ()
                           pl_bolts.models.autoencoders.basic_vae.basic_vae_module),
                           (pl_bolts.models.mnist_module.LitMNIST
                           method), 178
                           144
cli_main ()                 (in      module  configure_optimizers ()
                           pl_bolts.models.gans.basic.basic_gan_module),
                           (pl_bolts.models.regression.linear_regression.LinearRegression
                           method), 148
                           147
                           configure_optimizers ()
```

(*pl\_bolts.models.regression.logistic\_regression.LogisticRegression* in *pl\_bolts.losses.self\_supervised\_learning*),  
*method*), 149  
*configure\_optimizers()* CPCTrainTransformsCIFAR10 (class in  
(*pl\_bolts.models.self\_supervised.amdim.amdim\_module.AMDIM* in *pl\_bolts.models.self\_supervised.cpc.transforms*),  
*method*), 151 163  
*configure\_optimizers()* CPCTrainTransformsImageNet128 (class in  
(*pl\_bolts.models.self\_supervised.byol.byol\_module.BYOL* in *pl\_bolts.models.self\_supervised.cpc.transforms*),  
*method*), 158 163  
*configure\_optimizers()* CPCTrainTransformsSTL10 (class in  
(*pl\_bolts.models.self\_supervised.cpc.cpc\_module.CPCV2* in *pl\_bolts.models.self\_supervised.cpc.transforms*),  
*method*), 160 164  
*configure\_optimizers()* CPCV2 (class in *pl\_bolts.models.self\_supervised.cpc.cpc\_module*),  
(*pl\_bolts.models.self\_supervised.moco.moco2\_module.Moco2*  
*method*), 167  
*configure\_optimizers()* D  
(*pl\_bolts.models.self\_supervised.simclr.simclr\_module.SimCLR* in *pl\_bolts.datamodules.base\_dataset.LightDataset*  
*method*), 169 attribute), 117  
*configure\_optimizers()* data (*pl\_bolts.datamodules.cifar10\_dataset.CIFAR10*  
(*pl\_bolts.models.self\_supervised.ssl\_finetuner.SSLFineTuner* attribute), 123  
*method*), 173 data (*pl\_bolts.datamodules.cifar10\_dataset.TrialCIFAR10*  
attribute), 124  
*configure\_optimizers()* (*pl\_bolts.models.vision.image\_gpt.igpt\_module.ImageGPT* NAME (*pl\_bolts.datamodules.base\_dataset.LightDataset*  
*method*), 176 attribute), 117  
ConfusedLogitCallback (class in *pl\_bolts.callbacks.vision.confused\_logit*), DATASET\_NAME (*pl\_bolts.datamodules.cifar10\_dataset.CIFAR10*  
*method*), 112 attribute), 123  
conv1x1 () (in module Decoder (class in *pl\_bolts.models.autoencoders.basic\_vae.components*),  
*pl\_bolts.models.self\_supervised.cpc.networks*), 151 default\_transforms ()  
161 (*pl\_bolts.datamodules.cifar10\_datamodule.CIFAR10DataModule*  
152 default\_transforms ())  
Conv3x3 (class in *pl\_bolts.models.self\_supervised.amdim.networks*), 121 default\_transforms ()  
152  
conv3x3 () (in module (*pl\_bolts.datamodules.cityscapes\_datamodule.CityscapesDataModule*  
*pl\_bolts.models.self\_supervised.cpc.networks*), 161 method), 125  
default\_transforms ()  
conv\_block () (*pl\_bolts.models.vision.pixel\_cnn.PixelCNN* (*pl\_bolts.datamodules.stl10\_datamodule.STL10DataModule*  
*method*), 177 method), 138  
ConvResBlock (class in DenseBlock (class in  
*pl\_bolts.models.self\_supervised.amdim.networks*), *pl\_bolts.models.autoencoders.basic\_ae.components*),  
152 142  
ConvResNxN (class in DenseBlock (class in  
*pl\_bolts.models.self\_supervised.amdim.networks*), *pl\_bolts.models.autoencoders.basic\_vae.components*),  
153 145  
CPCEvalTransformsCIFAR10 (class in DensenetEncoder (class in  
*pl\_bolts.models.self\_supervised.cpc.transforms*), *pl\_bolts.models.self\_supervised.simclr.simclr\_module*),  
161 168  
CPCEvalTransformsImageNet128 (class in deterministic\_shuffle ()  
*pl\_bolts.models.self\_supervised.cpc.transforms*, (*pl\_bolts.models.self\_supervised.amdim.ssl\_datasets.SSLDataset*  
162 class method), 153  
CPCEvalTransformsSTL10 (class in dicts\_to\_table () (in module  
*pl\_bolts.callbacks.printing*), 114  
*pl\_bolts.models.self\_supervised.cpc.transforms*), dir\_path (*pl\_bolts.datamodules.base\_dataset.LightDataset*  
162 attribute), 117  
CPCResNet101 (class in dir\_path (*pl\_bolts.datamodules.cifar10\_dataset.CIFAR10*  
*pl\_bolts.models.self\_supervised.cpc.networks*), attribute), 123  
161

dir\_path(*pl\_bolts.datamodules.cifar10\_dataset.TrialCIFAR10\_NAME*(*pl\_bolts.datamodules.cifar10\_dataset.CIFAR10 attribute*), 124  
Discriminator (class *in finalize()* (*pl\_bolts.loggers.trains.TrainsLogger method*), 186  
*pl\_bolts.models.gans.basic.components*), 148  
discriminator\_loss () (*pl\_bolts.models.gans.basic.basic\_gan\_module.GANNatten* (*class in pl\_bolts.models.self\_supervised.evaluator*), 147  
*method*), 147  
discriminator\_step () (*pl\_bolts.models.gans.basic.basic\_gan\_module.GAN* (*method*), 147  
*method*), 147  
download () (*pl\_bolts.datamodules.cifar10\_dataset.CIFAR10* (*method*), 123  
*method*), 123  
DummyDataset (class *in* (*pl\_bolts.datamodules.dummy\_dataset*), 126  
DummyDetectionDataset (class *in* (*pl\_bolts.datamodules.dummy\_dataset*), 126  
**E**  
elbo\_loss () (*pl\_bolts.models.autoencoders.basic\_vae.basic\_vae\_module.VAE* (*method*), 142  
*method*), 144  
Encoder (*class in pl\_bolts.models.autoencoders.basic\_vae.components*) (*list* (*method*), 142  
145  
*forward()* (*pl\_bolts.models.autoencoders.basic\_vae.basic\_vae\_module.VAE* (*method*), 142  
*method*), 144  
exclude\_from\_wt\_decay () (*pl\_bolts.models.self\_supervised.simclr.simclr\_module.SIMCLR* (*pl\_bolts.models.autoencoders.basic\_vae.components.Decode* (*method*), 145  
*method*), 169  
experiment () (*pl\_bolts.loggers.trains.TrainsLogger* (*property*), 188  
experiment () (*pl\_bolts.loggers.TrainsLogger* (*property*), 184  
extra\_args (*pl\_bolts.datamodules.cifar10\_datamodule.CIFAR10DataModule* (*attribute*), 121  
attribute), 121  
extra\_args (*pl\_bolts.datamodules.cityscapes\_datamodule.CityscapesDataModule* (*attribute*), 126  
attribute), 126  
extract\_archive () (*in module* (*pl\_bolts.datamodules.imagenet\_dataset*), 131  
**F**  
FakeRKHSCovNet (class *in* (*pl\_bolts.models.self\_supervised.amdim.networks*), 153  
forward () (*pl\_bolts.models.regression.linear\_regression.LinearRegression* (*method*), 148  
method), 153  
FashionMNISTDataModule (class *in* (*pl\_bolts.datamodules.fashion\_mnist\_datamodule*), 127  
forward () (*pl\_bolts.models.regression.logistic\_regression.LogisticRegression* (*method*), 149  
method), 127  
FasterRCNN (class *in* (*pl\_bolts.models.detection.faster\_rcnn*), 145  
forward () (*pl\_bolts.models.self\_supervised.amdim.networks.AMDIMEN* (*method*), 151  
method), 145  
feat\_size\_w\_mask () (*pl\_bolts.losses.self\_supervised\_learning.FeatureMapContrastiveTask* (*method*), 179  
method), 179  
FeatureMapContrastiveTask (class *in* (*pl\_bolts.losses.self\_supervised\_learning*), 179  
forward () (*pl\_bolts.models.self\_supervised.amdim.networks.Conv3x3* (*method*), 152  
method), 179  
forward () (*pl\_bolts.models.self\_supervised.amdim.networks.ConvResBl* (*method*), 153  
method), 179  
forward () (*pl\_bolts.models.self\_supervised.amdim.networks.ConvResNx* (*method*), 153  
method), 179

```

forward() (pl_bolts.models.self_supervised.amdim.networks.EdarkHSGrayNets())
    method), 153
forward() (pl_bolts.models.self_supervised.amdim.networks.MaybeBatchNthOrder131
    method), 153
                                                generate_train_val_split()
forward() (pl_bolts.models.self_supervised.amdim.networks.NopNpl_bolts.models.self_supervised.amdim.ssl_datasets.SSLDataset)
    method), 153
                                                class method), 154
forward() (pl_bolts.models.self_supervised.biol.biol_modules.YOLOr
    method), 158
                                                (class in
                                                pl_bolts.models.gans.basic.components),
forward() (pl_bolts.models.self_supervised.biol.biol_modules.MLP 148
    method), 158
                                                generator_loss() (pl_bolts.models.gans.basic.basic_gan_module.GAN
forward() (pl_bolts.models.self_supervised.biol.biol_modules.SiameseAmethod), 147
    method), 158
                                                generator_step() (pl_bolts.models.gans.basic.basic_gan_module.GAN
forward() (pl_bolts.models.self_supervised.cpc.cpc_module.CPCV2method), 147
    method), 160
                                                get_approx_posterior()
forward() (pl_bolts.models.self_supervised.cpc.networks.CPCResNet100lts.models.autoencoders.basic_vae.basic_vae_module.VAE
    method), 161
                                                method), 144
forward() (pl_bolts.models.self_supervised.cpc.networks.ILNBottleneck () (pl_bolts.models.self_supervised.amdim.datasets.AMDIM
    method), 161
                                                static method), 152
forward() (pl_bolts.models.self_supervised.evaluator.Flatten_lr () (pl_bolts.optimizers.lr_scheduler.LinearWarmupCosineAnnealing
    method), 170
                                                method), 190
forward() (pl_bolts.models.self_supervised.evaluator.SSLEvaluator ()) (pl_bolts.models.autoencoders.basic_vae.basic_vae_modu
    method), 170
                                                method), 144
forward() (pl_bolts.models.self_supervised.moco.moco2gpt2module.MocoV2
    method), 167
                                                (pl_bolts.callbacks.self_supervised.SSLOnlineEvaluator
forward() (pl_bolts.models.self_supervised.resnets.ResNet 116
    method), 171
                                                GPT2 (class in pl_bolts.models.vision.image_gpt.gpt2),
forward() (pl_bolts.models.self_supervised.simclr.simclr_module.DensenetEncoder
    method), 168
forward() (pl_bolts.models.self_supervised.simclr.simclr1module.Projection
    method), 168
                                                id() (pl_bolts.loggers.trains.TrainsLogger property),
forward() (pl_bolts.models.self_supervised.simclr.simclr_module.SimCLR
    method), 169
                                                id() (pl_bolts.loggers.TrainsLogger property), 184
forward() (pl_bolts.models.vision.image_gpt.gpt2.Block Identity (class in pl_bolts.utils.semi_supervised), 191
    method), 174
                                                ImageGPT (class in pl_bolts.models.vision.image_gpt.igpt_module),
forward() (pl_bolts.models.vision.image_gpt.gpt2.GPT2 175
    method), 174
                                                imagenet () (pl_bolts.models.self_supervised.amdim.datasets.AMDIMP
forward() (pl_bolts.models.vision.image_gpt.igpt_module.ImageGPTstatic method), 152
    method), 176
                                                imagenet () (pl_bolts.models.self_supervised.amdim.datasets.AMDIMP
forward() (pl_bolts.models.vision.pixel_cnn.PixelCNN 152
    method), 177
                                                static method), 152
forward() (pl_bolts.utils.semi_supervised.Identity 191
    method), 191
                                                imangenet_normalization() (in module
                                                pl_bolts.transforms.dataset_normalizations),
                                                191
                                                ImagenetDataModule (class in
                                                pl_bolts.datamodules.imagenet_datamodule),
GAN (class in pl_bolts.models.gans.basic.basic_gan_module), 128
    146
                                                init_decoder () (pl_bolts.models.autoencoders.basic_ae.basic_ae_mo
GaussianBlur (class in pl_bolts.models.self_supervised.moco.transforms), 142
    167
                                                init_decoder () (pl_bolts.models.autoencoders.basic_vae.basic_vae_m
GaussianBlur (class in pl_bolts.models.self_supervised.simclr.simclr_transforms), 144
    169
                                                init_discriminator()
                                                (pl_bolts.models.gans.basic.basic_gan_module.GAN
                                                method), 147
generate_half_labeled_batches() (in mod- 191
    ule pl_bolts.utils.semi_supervised), 191
                                                init_encoder () (pl_bolts.models.autoencoders.basic_ae.basic_ae_mo
                                                method), 142

```

```

init_encoder() (pl_bolts.models.autoencoders.basic_vae.basic_vae_module.MAEbolts.loggers.trains.TrainsLogger
    method), 144
    method), 186
init_encoder() (pl_bolts.models.self_supervised.amdim_and_dim_mde.AMDIM(pl_bolts.loggers.TrainsLogger
    method), 151
    method), 182
init_encoder() (pl_bolts.models.self_supervised.cpc.cpc_module.CPCV2bolts.loggers.trains.TrainsLogger
    method), 160
    (pl_bolts.loggers.trains.TrainsLogger method),
init_encoder() (pl_bolts.models.self_supervised.simclr.simclr_module.SimCLR
    method), 169
    log_hyperparams() (pl_bolts.loggers.TrainsLogger
    method), 180
init_encoders() (pl_bolts.models.self_supervised.moco.moco2_method).MocoV2
    method), 167
    log_image() (pl_bolts.loggers.trains.TrainsLogger
    method), 181
init_generator() (pl_bolts.models.gans.basic.basic_gan_module.GAN), 186
    method), 147
    log_image() (pl_bolts.loggers.TrainsLogger method),
init_weights() (pl_bolts.models.self_supervised.amdim.networks.AMDIMEncoder
    method), 152
    log_metric() (pl_bolts.loggers.trains.TrainsLogger
    method), 182
init_weights() (pl_bolts.models.self_supervised.amdim.networks.ChainResBlock
    method), 153
    log_metric() (pl_bolts.loggers.TrainsLogger
    method), 183
init_weights() (pl_bolts.models.self_supervised.amdim.networks.ChainResNN
    method), 153
    log_metrics() (pl_bolts.loggers.trains.TrainsLogger
    method), 184
init_weights() (pl_bolts.models.self_supervised.amdim.networks.EfficientKHSConvNet
    method), 153
    log_metrics() (pl_bolts.loggers.TrainsLogger
    method), 185
interpolate_latent_space()
    (pl_bolts.callbacks.variational.LatentDimInterpolator).log_text() (pl_bolts.loggers.trains.TrainsLogger
    method), 116
    log_text() (pl_bolts.loggers.TrainsLogger method),
    186
L
labels (pl_bolts.datamodules.cifar10_dataset.CIFAR10
    attribute), 123
LogisticRegression (class
    in pl_bolts.models.regression.logistic_regression),
    149
LARSWrapper (class
    in pl_bolts.optimizers.lars_scheduling), 188
LatentDimInterpolator (class
    in pl_bolts.callbacks.variational), 116
LightDataset (class
    in pl_bolts.datamodules.base_dataset), 117
LinearRegression (class
    in pl_bolts.models.regression.linear_regression),
    148
LinearWarmupCosineAnnealingLR (class
    in pl_bolts.optimizers.lr_scheduler), 189
LitMNIST (class in pl_bolts.models.mnist_module), 178
LNBottleneck (class
    in pl_bolts.models.self_supervised.cpc.networks),
    161
load_instance() (pl_bolts.datamodules.async_dataloader.AsyncDataLoader).Transforms (class
    in pl_bolts.models.self_supervised.moco.transforms),
    162
load_loop() (pl_bolts.datamodules.async_dataloader.AsyncDataLoader
    method), 117
    Moco2EvalCIFAR10Transforms (class
    in pl_bolts.models.self_supervised.moco.transforms),
    163
load_pretrained() (in module pl_bolts.utils.pretrained_weights), 191
load_pretrained() (in module pl_bolts.models.autoencoders.basic_vae.basic_vae_module).VAEbolts.loggers.trains.TrainsLogger
    method), 144
    Moco2TrainCIFAR10Transforms (class
    in pl_bolts.models.self_supervised.moco.transforms),
    167
load_pretrained() (pl_bolts.models.self_supervised.cpc.cpc_module.CPCV2
    method), 160
    Moco2TrainImagenetTransforms (class
    in pl_bolts.models.self_supervised.moco.transforms),
    168

```

Moco2TrainSTL10Transforms (class in `num_classes()` (`pl_bolts.datamodules.imagenet_datamodule.ImagenetDataModule`), `property`), 130  
`pl_bolts.models.self_supervised.moco.transforms`, `property`), 130  
`168`  
`num_classes()` (`pl_bolts.datamodules.mnist_datamodule.MNISTDataModule`), `property`), 133  
`168`  
MocoLRScheduler (class in `num_classes()` (`pl_bolts.datamodules.ssl_imagenet_datamodule.SSLImageNetDataModule`), `property`), 137  
`pl_bolts.models.self_supervised.moco.callbacks`, `num_classes()` (`pl_bolts.datamodules.ssl_imagenet_datamodule.SSLImageNetDataModule`), `property`), 137  
`165`  
MocoV2 (class in `pl_bolts.models.self_supervised.moco.moco2_moco_v2_modules`), `(pl_bolts.datamodules.stl10_datamodule.STL10DataModule)`, `property`), 139  
`165`  
`num_classes()` (`pl_bolts.datamodules.vocdetection_datamodule.VOCDataModule`), `property`), 140  
**N**  
name (`pl_bolts.datamodules.binary_mnist_datamodule.BinaryMNISTDataModule`)  
`attribute`), 119  
**O**  
name (`pl_bolts.datamodules.cifar10_datamodule.CIFAR10DataModule`)  
`on_epoch_end()` (`pl_bolts.callbacks.printing.PrintTableMetricsCallback`), `method`), 114  
`attribute`), 121  
name (`pl_bolts.datamodules.cityscapes_datamodule.CityscapesDataModule`)  
`on_epoch_end()` (`pl_bolts.callbacks.variational.LatentDimInterpolator`), `method`), 116  
`attribute`), 126  
name (`pl_bolts.datamodules.fashion_mnist_datamodule.FashionMNISTDataModule`)  
`on_epoch_end()` (`pl_bolts.callbacks.vision.image_generation.TensorboardCallback`), `method`), 113  
`attribute`), 128  
name (`pl_bolts.datamodules.imagenet_datamodule.ImagenetDataModule`)  
`on_epoch_start()` (`pl_bolts.models.self_supervised.moco.callbacks.Moco2DataModule`), `method`), 165  
`attribute`), 130  
name (`pl_bolts.datamodules.mnist_datamodule.MNISTDataModule`)  
`on_train_routine_start()` (`pl_bolts.callbacks.self_supervised.SSLOnlineEvaluator`), `method`), 133  
`attribute`), 133  
name (`pl_bolts.datamodules.sklearn_datamodule.SklearnDataModule`)  
`on_train_routine_start()` (`pl_bolts.callbacks.self_supervised.SSLOnlineEvaluator`), `method`), 116  
`attribute`), 134  
`on_train_batch_end()` (`pl_bolts.callbacks.vision.confused_logit.ConfusedLogitCallback`), `method`), 116  
name (`pl_bolts.datamodules.ssl_imagenet_datamodule.SSLImageNetDataModule`)  
`on_train_batch_end()` (`pl_bolts.callbacks.self_supervised.BYOLMAWeightUpdateCallback`), `method`), 115  
`attribute`), 137  
name (`pl_bolts.datamodules.stl10_datamodule.STL10DataModule`)  
`on_train_batch_end()` (`pl_bolts.callbacks.self_supervised.SSLOnlineEvaluator`), `method`), 139  
`attribute`), 139  
name (`pl_bolts.datamodules.vocdetection_datamodule.VOCDataModule`)  
`on_train_batch_end()` (`pl_bolts.callbacks.vision.confused_logit.ConfusedLogitCallback`), `method`), 140  
`attribute`), 140  
name () (`pl_bolts.loggers.TrainsLogger`), `property`), 188  
name () (`pl_bolts.loggers.TrainsLogger`), `property`), 184  
NopNet (class in `pl_bolts.models.self_supervised.amdim.networks`), `(pl_bolts.models.self_supervised.byol.byol_module.BYOLModule)`, `method`), 153  
normalize (`pl_bolts.datamodules.base_dataset.LightDataset`)  
`on_train_epoch_start()` (`pl_bolts.callbacks.vision.confused_logit.ConfusedLogitCallback`), `method`), 117  
normalize (`pl_bolts.datamodules.cifar10_dataset.CIFAR10Dataset`)  
`on_train_epoch_start()` (`pl_bolts.callbacks.vision.confused_logit.ConfusedLogitCallback`), `method`), 173  
normalize (`pl_bolts.datamodules.cifar10_dataset.TrialCIFAR10Dataset`)  
`on_train_epoch_start()` (`pl_bolts.callbacks.vision.confused_logit.ConfusedLogitCallback`), `method`), 124  
nt\_xent\_loss () (in `module` `param_groups()` (`pl_bolts.optimizers.lars_scheduling.LARSWrapper`), `property`), 189  
`pl_bolts.losses.self_supervised_learning`), `parse_devkit_archive()` (in `module` `pl_bolts.datamodules.imagenet_dataset`), `180`  
`180`  
num\_classes () (`pl_bolts.datamodules.binary_mnist_datamodule.BinaryMNISTDataModule`)  
`property`), 119  
parse\_map\_indexes ()  
num\_classes () (`pl_bolts.datamodules.cifar10_datamodule.CIFAR10DataModule`)  
`pl_bolts.losses.self_supervised_learning`.`FeatureMapContrastive` (`static method`), 180  
`property`), 121  
num\_classes () (`pl_bolts.datamodules.cifar10_datamodule.TinyCIFAR10DataModule`)  
`partition_train_set()` (`pl_bolts.datamodules.imagenet_dataset.UnlabeledImagenetDataset`), `method`), 151  
`property`), 122  
num\_classes () (`pl_bolts.datamodules.cityscapes_datamodule.CityscapesDataModule`)  
`property`), 126  
Patchify (class in `pl_bolts.transforms.self_supervised.ssl_transforms`), `Patchify` (`class in pl_bolts.transforms.self_supervised.ssl_transforms`), `190`  
num\_classes () (`pl_bolts.datamodules.fashion_mnist_datamodule.FashionMNISTDataModule`)  
`property`), 128

PixelCNN (*class in pl\_bolts.models.vision.pixel\_cnn*), 177  
pl\_bolts.callbacks (*module*), 112  
pl\_bolts.callbacks.printing (*module*), 114  
pl\_bolts.callbacks.self\_supervised (*module*), 115  
pl\_bolts.callbacks.variational (*module*), 116  
pl\_bolts.callbacks.vision (*module*), 112  
pl\_bolts.callbacks.vision.confused\_logitpl\_bolts.models.autoencoders.basic\_vae.basic\_vae\_m  
(*module*), 112  
pl\_bolts.callbacks.vision.image\_generatipl\_bolts.models.autoencoders.basic\_vae.components  
(*module*), 113  
pl\_bolts.datamodules (*module*), 117  
pl\_bolts.datamodules.async\_dataloader (*module*), 117  
pl\_bolts.datamodules.base\_dataset (*mod-  
ule*), 117  
pl\_bolts.datamodules.binary\_mnist\_datamodulpl\_bolts.models.gans.basic.basic\_gan\_module  
(*module*), 118  
pl\_bolts.datamodules.cifar10\_datamodulpl\_bolts.models.gans.basic.components  
(*module*), 120  
pl\_bolts.datamodules.cifar10\_dataset (*mod-  
ule*), 122  
pl\_bolts.datamodules.cityscapes\_datamodulpl\_bolts.models.regression.linear\_regression  
(*module*), 124  
pl\_bolts.datamodules.concat\_dataset (*mod-  
ule*), 126  
pl\_bolts.datamodules.dummy\_dataset (*mod-  
ule*), 126  
pl\_bolts.datamodules.fashion\_mnist\_datamodulpl\_bolts.models.self\_supervised.amdim  
(*module*), 127  
pl\_bolts.datamodules.imagenet\_datamodulpl\_bolts.models.self\_supervised.amdim.amdim\_module  
(*module*), 128  
pl\_bolts.datamodules.imagenet\_dataset (*mod-  
ule*), 131  
pl\_bolts.datamodules.mnist\_datamodulpl\_bolts.models.self\_supervised.amdim.networks  
(*module*), 132  
pl\_bolts.datamodules.sklearn\_datamodulpl\_bolts.models.self\_supervised.amdim.ssl\_datasets  
(*module*), 133  
pl\_bolts.datamodules.ssl\_imagenet\_datamodulpl\_bolts.models.self\_supervised.amdim.transforms  
(*module*), 137  
pl\_bolts.datamodules.stl10\_datamodulpl\_bolts.models.self\_supervised.byol  
(*module*), 137  
pl\_bolts.datamodules.vocdetection\_datamodulpl\_bolts.models.self\_supervised.byol.byol\_module  
(*module*), 139  
pl\_bolts.loggers (*module*), 180  
pl\_bolts.loggers.trains (*module*), 184  
pl\_bolts.losses (*module*), 178  
pl\_bolts.losses.self\_supervised\_learning (*mod-  
ule*), 178  
pl\_bolts.metrics (*module*), 140  
pl\_bolts.metrics.aggregation (*module*), 140  
pl\_bolts.models (*module*), 140  
pl\_bolts.models.autoencoders (*module*), 140  
pl\_bolts.models.autoencoders.basic\_ae  
(*module*), 141  
pl\_bolts.models.autoencoders.basic\_ae.basic\_ae\_modu  
(*module*), 141  
pl\_bolts.models.autoencoders.basic\_ae.components  
(*module*), 142  
pl\_bolts.models.autoencoders.basic\_vae  
(*module*), 143  
pl\_bolts.models.autoencoders.basic\_vae.basic\_vae\_m  
(*module*), 143  
pl\_bolts.models.detection (*module*), 145  
pl\_bolts.models.detection.faster\_rcnn  
(*module*), 145  
pl\_bolts.models.gans (*module*), 146  
pl\_bolts.models.gans.basic (*module*), 146  
pl\_bolts.models.gans.basic.basic\_gan\_module  
(*module*), 146  
pl\_bolts.models.gans.basic.components  
(*module*), 148  
pl\_bolts.models.mnist\_module (*module*), 178  
pl\_bolts.models.regression (*module*), 148  
pl\_bolts.models.regression.linear\_regression  
(*module*), 148  
pl\_bolts.models.regression.logistic\_regression  
(*module*), 149  
pl\_bolts.models.self\_supervised (*module*),  
150  
pl\_bolts.models.self\_supervised.amdim  
(*module*), 150  
pl\_bolts.models.self\_supervised.amdim.amdim\_module  
(*module*), 150  
pl\_bolts.models.self\_supervised.amdim.datasets  
(*module*), 152  
pl\_bolts.models.self\_supervised.amdim.networks  
(*module*), 152  
pl\_bolts.models.self\_supervised.amdim.ssl\_datasets  
(*module*), 153  
pl\_bolts.models.self\_supervised.amdim.transforms  
(*module*), 154  
pl\_bolts.models.self\_supervised.byol  
(*module*), 156  
pl\_bolts.models.self\_supervised.byol\_byol\_module  
(*module*), 156  
pl\_bolts.models.self\_supervised.byol.models  
(*module*), 158  
pl\_bolts.models.self\_supervised.cpc  
(*module*), 158  
pl\_bolts.models.self\_supervised.cpc.cpc\_finetuner  
(*module*), 158  
pl\_bolts.models.self\_supervised.cpc.cpc\_module  
(*module*), 158

```

pl_bolts.models.self_supervised.cpc.netwrbolts.utils.shaping (module), 192
    (module), 161                               precision_at_k ()           (in      module
pl_bolts.models.self_supervised.cpc.transforms pl_bolts.metrics.aggregation), 140
    (module), 161                               prepare_data () (pl_bolts.datamodules.binary_mnist_datamodule.Binary
pl_bolts.models.self_supervised.evaluator     method), 119
    (module), 170                               prepare_data () (pl_bolts.datamodules.cifar10_datamodule.CIFAR10D
pl_bolts.models.self_supervised.moco          method), 121
    (module), 165                               prepare_data () (pl_bolts.datamodules.cifar10_dataset.CIFAR10
pl_bolts.models.self_supervised.moco.callbacks method), 123
    (module), 165                               prepare_data () (pl_bolts.datamodules.cifar10_dataset.TrialCIFAR10
pl_bolts.models.self_supervised.moco2_module method), 124
    (module), 165                               prepare_data () (pl_bolts.datamodules.cityscapes_datamodule.Cityscap
pl_bolts.models.self_supervised.moco.transforms method), 125
    (module), 167                               prepare_data () (pl_bolts.datamodules.fashion_mnist_datamodule.Fash
pl_bolts.models.self_supervised.resnets        method), 128
    (module), 170                               prepare_data () (pl_bolts.datamodules.imagenet_datamodule.Imagenet
pl_bolts.models.self_supervised.simclr         method), 129
    (module), 168                               prepare_data () (pl_bolts.datamodules.mnist_datamodule.MNISTData
pl_bolts.models.self_supervised.simclr.simclr method), 133
    (module), 168                               prepare_data () (pl_bolts.datamodules.ssl_imagenet_datamodule.SSLI
pl_bolts.models.self_supervised.simclr.simclr method), 137
    (module), 168                               prepare_data () (pl_bolts.datamodules.stl10_datamodule.STL10DataM
pl_bolts.models.self_supervised.simclr.simclr method), 138
    (module), 169                               prepare_data () (pl_bolts.datamodules.vocdetection_datamodule.VOC
pl_bolts.models.self_supervised.ssl_finetuner method), 139
    (module), 172                               prepare_data () (pl_bolts.models.mnist_module.LitMNIST
pl_bolts.models.vision (module), 174             method), 178
pl_bolts.models.vision.image_gpt   (mod- PrintTableMetricsCallback (class      in
    ule), 174                               pl_bolts.callbacks.printing), 114
pl_bolts.models.vision.image_gpt.gpt2       Projection      (class      in
    (module), 174                           pl_bolts.models.self_supervised.simclr.simclr_module),
pl_bolts.models.vision.image_gpt.igpt_module (module), 168
    (module), 175
pl_bolts.models.vision.pixel_cnn   (mod- R
    ule), 177
pl_bolts.optimizers (module), 188
pl_bolts.optimizers.lars_scheduling (module), 188 RandomTranslateWithReflect (class      in
pl_bolts.optimizers.lr_scheduler  (mod- pl_bolts.transforms.self_supervised.ssl_transforms),
    ule), 189                               190
pl_bolts.transforms (module), 190
pl_bolts.transforms.dataset_normalization (module), 191 relabel (pl_bolts.datamodules.cifar10_dataset.CIFAR10
pl_bolts.transforms.self_supervised          attribute), 123
    (module), 190
pl_bolts.transforms.self_supervised          ResNet (class in pl_bolts.models.self_supervised.resnets),
    (module), 190                               170
pl_bolts.transforms.ssl_transforms          resnet101 ()           (in      module
    (module), 190                           pl_bolts.models.self_supervised.resnets),
    171
pl_bolts.utils (module), 191
pl_bolts.utils.pretrained_weights  (mod- resnet152 ()           (in      module
    ule), 191                           pl_bolts.models.self_supervised.resnets),
    171
pl_bolts.utils.self_supervised   (module), 191 resnet18 ()           (in      module
pl_bolts.utils.semi_supervised   (module), 191                           pl_bolts.models.self_supervised.resnets),
    171
pl_bolts.utils.semi_supervised   (module), 191 resnet34 ()           (in      module
pl_bolts.utils.semi_supervised   (module), 191                           pl_bolts.models.self_supervised.resnets),
    171
pl_bolts.utils.semi_supervised   (module), 191 resnet50 ()           (in      module
pl_bolts.utils.semi_supervised   (module), 191

```

```

    pl_bolts.models.self_supervised.resnets),
171
resnet50_bn()           (in      module
    pl_bolts.models.self_supervised.resnets),
171
resnext101_32x8d()     (in      module
    pl_bolts.models.self_supervised.resnets),
172
resnext50_32x4d()      (in      module
    pl_bolts.models.self_supervised.resnets),
171
run_cli()              (in      module
    pl_bolts.models.detection.faster_rcnn), 146

S
select_nb_imgs_per_class()
    (pl_bolts.models.self_supervised.amdim.ssl_datasets.SSLDataModule), 189
    class method), 154
set_bypass_mode()
    (pl_bolts.loggers.trains.TrainsLogger      class
        method), 187
set_bypass_mode() (pl_bolts.loggers.TrainsLogger
    class method), 183
set_credentials()
    (pl_bolts.loggers.trains.TrainsLogger      class
        method), 187
set_credentials() (pl_bolts.loggers.TrainsLogger
    class method), 183
setup() (pl_bolts.models.self_supervised.simclr.simclr_module.SimCLR
    method), 169
shared_step() (pl_bolts.models.self_supervised.biol.biol_module.BYOL
    method), 158
shared_step() (pl_bolts.models.self_supervised.cpc.cpc_module.CPCModule,
    pl_bolts.losses.self_supervised_learning),
    180
shared_step() (pl_bolts.models.self_supervised.simclr.simclr_module.SimCLR
    method), 160
shared_step() (pl_bolts.models.self_supervised.simclr.simclr_module.SimCLR
    method), 169
shared_step() (pl_bolts.models.self_supervised.ssl_finetuner.SSLFineTuner,
    pl_bolts.datamodules.base_dataset.LightDataset
    attribute), 118
shared_step() (pl_bolts.models.self_supervised.ssl_finetuner.SSLFineTuner,
    pl_bolts.datamodules.cifar10_dataset.CIFAR10
    attribute), 123
SiameseArm          (class      in  targets (pl_bolts.datamodules.cifar10_dataset.TrialCIFAR10
    pl_bolts.models.self_supervised.biol.models),
    158
    TensorboardGenerativeModelImageSampler
SimCLR(class in pl_bolts.models.self_supervised.simclr.simclr_module,
    168
    class in pl_bolts.callbacks.vision.image_generation),
    113
SimCLREvalDataTransform (class      in  TensorDataModule          (class      in
    pl_bolts.models.self_supervised.simclr.simclr_transforms), pl_bolts.datamodules.sklearn_datamodule),
    169
    135
SimCLRTrainDataTransform (class      in  TensorDataset           (class      in
    pl_bolts.models.self_supervised.simclr.simclr_transforms), pl_bolts.datamodules.sklearn_datamodule),
    170
    136
SklearnDataModule     (class      in  test_dataloader()          (pl_bolts.datamodules.binary_mnist_datamodule.BinaryMNISTDataModule
    pl_bolts.datamodules.sklearn_datamodule),
    133
    method), 119
SklearnDataset        (class      in  test_dataloader()          (pl_bolts.datamodules.cifar10_datamodule.CIFAR10DataModule
    pl_bolts.datamodules.sklearn_datamodule),
    133
    method), 119
SSLDataModule         (class      in
    pl_bolts.models.self_supervised.amdim.ssl_datasets),
    153
SSLEvaluator          (class      in
    pl_bolts.models.self_supervised.evaluator),
    170
SSLFineTuner          (class      in
    pl_bolts.models.self_supervised.ssl_finetuner),
    172
SSLImagenetDataModule (class      in
    pl_bolts.datamodules.ssl_imagenet_datamodule),
    137
SSLOnlineEvaluator    (class      in
    pl_bolts.callbacks.self_supervised), 115
state() (pl_bolts.optimizers.lars_scheduling.LARSWrapper
    static method), 152
step() (pl_bolts.optimizers.lars_scheduling.LARSWrapper
    method), 189
stl() (pl_bolts.models.self_supervised.amdim.datasets.AMDIMPatchesPretrainin
    static method), 152
stl() (pl_bolts.models.self_supervised.amdim.datasets.AMDIMPretrainin
    static method), 152
stl10_normalization() (in      module
    pl_bolts.transforms.dataset_normalizations),
    191
STL10DataModule        (class      in
    pl_bolts.datamodules.stl10_datamodule),
    127
T
tanh_clip()           (in      module
    pl_bolts.losses.self_supervised_learning),
    180
targets(pl_bolts.datamodules.cifar10_dataset.TrialCIFAR10
    attribute), 124
TensorboardGenerativeModelImageSampler

```

```

        method), 121
test_dataloader()
    (pl_bolts.datamodules.cityscapes_datamodule.CityscapesDataModule)
        method), 125
test_dataloader()
    (pl_bolts.datamodules.fashion_mnist_datamodule.FashionMNISTDataModule)
        method), 128
test_dataloader()
    (pl_bolts.datamodules.imagenet_datamodule.ImagenetDataModule)
        method), 130
test_dataloader()
    (pl_bolts.datamodules.mnist_datamodule.MNISTDataModule)
        method), 133
test_dataloader()
    (pl_bolts.datamodules.sklearn_datamodule.SklearnDataModule)
        method), 134
test_dataloader()
    (pl_bolts.datamodules.ssl_imagenet_datamodule.SSLImagenetDataModule)
        method), 137
test_dataloader()
    (pl_bolts.datamodules.stl10_datamodule.STL10DataModule)
        method), 138
test_dataloader()
    (pl_bolts.models.mnist_module.LitMNIST
        method), 178
test_epoch_end() (pl_bolts.models.autoencoders.basic_ae.basic_ae_module)
    (method), 142
test_epoch_end() (pl_bolts.models.mnist_module.LitMNIST
    method), 178
test_epoch_end() (pl_bolts.models.regression.linear_regression)
    (method), 148
test_epoch_end() (pl_bolts.models.regression.logistic_regression)
    (method), 149
test_epoch_end() (pl_bolts.models.vision.image_gpt_igpt_module)
    (method), 176
TEST_FILE_NAME (pl_bolts.datamodules.cifar10_dataset)
    attribute), 123
test_step() (pl_bolts.models.autoencoders.basic_ae.basic_ae_module)
    (method), 142
test_step() (pl_bolts.models.autoencoders.basic_vae.basic_vae
    (method), 144
test_step() (pl_bolts.models.mnist_module.LitMNIST
    method), 178
test_step() (pl_bolts.models.regression.linear_regression.LinearRegression)
    (method), 148
test_step() (pl_bolts.models.regression.logistic_regression.LogisticRegression)
    (method), 149
test_step() (pl_bolts.models.self_supervised.ssl_finetuner)
    (method), 173
test_step() (pl_bolts.models.vision.image_gpt_igpt_module)
    (method), 177
tile() (in module pl_bolts.utils.shaping), 192
TinyCIFAR10DataModule (class
    pl_bolts.datamodules.cifar10_datamodule),
        121
        to_device() (pl_bolts.callbacks.self_supervised.SSLOnlineEvaluator
            pl_bolts.utils.self_supervised), 191
        torchvision_ssl_encoder() (in module
            pl_bolts.utils.self_supervised), 191
        (pl_bolts.datamodules.binary_mnist_datamodule.BinaryMNISTDataModule)
            method), 119
        (pl_bolts.datamodules.cityscapes_datamodule.CityscapesDataModule)
            method), 125
        (pl_bolts.datamodules.fashion_mnist_datamodule.FashionMNISTDataModule)
            method), 128
        (pl_bolts.datamodules.imagenet_datamodule.ImagenetDataModule)
            method), 130
        (pl_bolts.datamodules.mnist_datamodule.MNISTDataModule)
            method), 133
        (pl_bolts.datamodules.sklearn_datamodule.SklearnDataModule)
            method), 134
        (pl_bolts.datamodules.ssl_imagenet_datamodule.SSLImagenetDataModule)
            method), 137
        (pl_bolts.datamodules.stl10_datamodule.STL10DataModule)
            method), 138
        (pl_bolts.models.mnist_module.LitMNIST
            method), 178
        train_dataloader()
            (pl_bolts.datamodules.sklearn_datamodule.SklearnDataModule)
                method), 133
        train_dataloader()
            (pl_bolts.datamodules.ssl_imagenet_datamodule.SSLImagenetDataModule)
                method), 137
        train_dataloader()
            (pl_bolts.datamodules.stl10_datamodule.STL10DataModule)
                method), 138
        train_dataloader()
            (pl_bolts.datamodules.voc_detection_datamodule.VOCDetectionDataModule)
                method), 139
        train_dataloader()
            (pl_bolts.models.mnist_module.LitMNIST
                method), 123
        (pl_bolts.models.mnist_module.LitMNIST
            method), 178
        train_dataloader()
            (pl_bolts.datamodules.stl10_datamodule.STL10DataModule)
                method), 151
        train_dataloader_labeled()
            (pl_bolts.datamodules.stl10_datamodule.STL10DataModule)
                method), 138
        train_dataloader_mixed()
            (pl_bolts.datamodules.stl10_datamodule.STL10DataModule)
                method), 138
        train_dataloader()
            (pl_bolts.datamodules.voc_detection_datamodule.VOCDetectionDataModule)
                method), 139
        train_dataloader()
            (pl_bolts.models.mnist_module.LitMNIST
                method), 123
        (pl_bolts.models.mnist_module.LitMNIST
            method), 173
        ImageGPE form()
            (pl_bolts.datamodules.imagenet_datamodule.ImagenetDataModule)
                method), 130
        in training_step() (pl_bolts.models.autoencoders.basic_ae.basic_ae_module)
            (method), 142

```

```

training_step() (pl_bolts.models.autoencoders.basic_vae.basic_vae_module.VAEbolts.datamodules.fashion_mnist_datamodule.
    method), 144
    method), 128
training_step() (pl_bolts.models.detection.faster_rcnn.FasterRCNNader () (pl_bolts.datamodules.imagenet_datamodule.Image
    method), 146
    method), 130
training_step() (pl_bolts.models.gans.basic.basic_gan_datagenerator () (pl_bolts.datamodules.mnist_datamodule.MNISTD
    method), 147
    method), 133
training_step() (pl_bolts.models.mnist_module.LitMNIST_dataloader () (pl_bolts.datamodules.sklearn_datamodule.Sklear
    method), 178
    method), 134
training_step() (pl_bolts.models.regression.linear_regression.LinearRegressionpl_bolts.datamodules.ssl_imagenet_datamodule.S
    method), 148
    method), 137
training_step() (pl_bolts.models.regression.logistic_regression.LogisticRegressionpl_bolts.datamodules.stl10_datamodule.STL10Da
    method), 149
    method), 138
training_step() (pl_bolts.models.self_supervised.amdim_amdim_module.AMDIMbolts.datamodules.vocdetection_datamodule.V
    method), 151
    method), 140
training_step() (pl_bolts.models.self_supervised.byzol_byol_module.BYOL () (pl_bolts.models.mnist_module.LitMNIST
    method), 158
    method), 178
training_step() (pl_bolts.models.self_supervised.cpc.cpc_module.CPCV2 () (pl_bolts.models.self_supervised.amdim.amdim_ma
    method), 160
    method), 151
training_step() (pl_bolts.models.self_supervised.moco.moco2_module.MocoV2lede ()
    method), 167
    (pl_bolts.datamodules.stl10_datamodule.STL10DataModule
training_step() (pl_bolts.models.self_supervised.simclr.simclr_module.SimCLR
    method), 169
    val_dataloader_mixed()
training_step() (pl_bolts.models.self_supervised.ssl_finetuner.SSLFinetunerpl_bolts.datamodules.stl10_datamodule.STL10DataModule
    method), 173
    method), 139
training_step() (pl_bolts.models.vision.image_gpt.igpt_module.ImageGPT (pl_bolts.datamodules.imagenet_datamodule.Imager
    method), 177
    method), 130
training_step_end()
    validation_epoch_end()
    (pl_bolts.models.self_supervised.amdim.amdim_module.AMDIMbolts.models.autoencoders.basic_ae.basic_ae_module.AE
        method), 151
        method), 142
TrainsLogger (class in pl_bolts.loggers), 180
    validation_epoch_end()
TrainsLogger (class in pl_bolts.loggers.trains), 184
    (pl_bolts.models.detection.faster_rcnn.FasterRCNN
TrialCIFAR10 (class
    in
    method), 146
    (pl_bolts.datamodules.cifar10_dataset), 123
    validation_epoch_end()
    (pl_bolts.models.mnist_module.LitMNIST
    method), 178
U
UnlabeledImagenet (class
    in
    validation_epoch_end()
    pl_bolts.datamodules.imagenet_dataset),
    (pl_bolts.models.regression.linear_regression.LinearRegression
    method), 131
    method), 148
update_p() (pl_bolts.optimizers.lars_scheduling.LARSWrapper)
    validation_epoch_end()
    (pl_bolts.models.regression.logistic_regression.LogisticRegression
    method), 189
    method), 149
update_tau() (pl_bolts.callbacks.self_supervised.BYOLMAWeightUpdate), 149
    validation_epoch_end()
    method), 115
update_weights() (pl_bolts.callbacks.self_supervised.BYOLMAWeightUpdate)
    validation_epoch_end()
    (pl_bolts.models.self_supervised.amdim.amdim_module.AMDIM
    method), 115
    method), 151
    validation_epoch_end()
V
VAE (class in pl_bolts.models.autoencoders.basic_vae.basic_vae_module),
    validation_epoch_end()
    143
    method), 167
val_dataloader() (pl_bolts.datamodules.binary_mnist_datamodule.BinaryMNISTDataModule)
    (pl_bolts.models.vision.image_gpt.igpt_module.ImageGPT
    method), 119
    method), 177
val_dataloader() (pl_bolts.datamodules.cifar10_datamodule.CIFAR10DataModule)
    validation_step()
    (pl_bolts.models.autoencoders.basic_ae.basic_ae_module.AE
    method), 121
    method), 142
val_dataloader() (pl_bolts.datamodules.cityscapes_datamodule.CityscapesDataModule
    validation_step()
    method), 126
    method), 142

```

```

(pl_bolts.models.autoencoders.basic_vae.basic_vae_module.VAE
method), 144
validation_step()
(pl_bolts.models.detection.faster_rcnn.FasterRCNN
method), 146
validation_step()
(pl_bolts.models.mnist_module.LitMNIST
method), 178
validation_step()
(pl_bolts.models.regression.linear_regression.LinearRegression
method), 148
validation_step()
(pl_bolts.models.regression.logistic_regression.LogisticRegression
method), 149
validation_step()
(pl_bolts.models.self_supervised.amdim.amdim_module.AMDIM
method), 151
validation_step()
(pl_bolts.models.self_supervised.byol.byol_module.BYOL
method), 158
validation_step()
(pl_bolts.models.self_supervised.cpc.cpc_module.CPCV2
method), 160
validation_step()
(pl_bolts.models.self_supervised.moco.moco2_module.MocoV2
method), 167
validation_step()
(pl_bolts.models.self_supervised.simclr.simclr_module.SimCLR
method), 169
validation_step()
(pl_bolts.models.self_supervised.ssl_finetuner.SSLFineTuner
method), 173
validation_step()
(pl_bolts.models.vision.image_gpt.igpt_module.ImageGPT
method), 177
version() (pl_bolts.loggers.trains.TrainsLogger prop-
erty), 188
version() (pl_bolts.loggers.TrainsLogger property),
184
VOCDetectionDataModule      (class      in
pl_bolts.datamodules.vocdetection_datamodule),
139

```

## W

```

wide_resnet101_2()      (in      module
pl_bolts.models.self_supervised.resnets),
172
wide_resnet50_2()      (in      module
pl_bolts.models.self_supervised.resnets),
172

```