

---

# Flash

## PyTorch Lightning

Mar 10, 2021



## GET STARTED:

<b>1</b>	<b>Quick Start</b>	<b>1</b>
<b>2</b>	<b>Installation</b>	<b>7</b>
<b>3</b>	<b>Tutorial: Creating a Custom Task</b>	<b>9</b>
<b>4</b>	<b>From Flash to Lightning</b>	<b>13</b>
<b>5</b>	<b>General Task</b>	<b>17</b>
<b>6</b>	<b>Image Classification</b>	<b>19</b>
<b>7</b>	<b>Image Embedder</b>	<b>25</b>
<b>8</b>	<b>Summarization</b>	<b>29</b>
<b>9</b>	<b>Text Classification</b>	<b>35</b>
<b>10</b>	<b>Tabular Classification</b>	<b>39</b>
<b>11</b>	<b>Translation</b>	<b>45</b>
<b>12</b>	<b>Model</b>	<b>51</b>
<b>13</b>	<b>Data</b>	<b>53</b>
<b>14</b>	<b>Training from scratch</b>	<b>55</b>
<b>15</b>	<b>Finetuning</b>	<b>59</b>
<b>16</b>	<b>Predictions (inference)</b>	<b>63</b>
<b>17</b>	<b>Indices and tables</b>	<b>65</b>
	<b>Index</b>	<b>67</b>



## **QUICK START**

Flash is a high-level deep learning framework for fast prototyping, baselining, finetuning and solving deep learning problems. It features a set of tasks for you to use for inference and finetuning out of the box, and an easy to implement API to customize every step of the process.

Flash is excellent for:

- data scientists
- kagglers
- applied corporate researchers
- applied academic researchers

### **1.1 Why Flash?**

Flash is built on top of [Pytorch Lightning](#), a powerful deep learning research framework for training models at scale. With the power of Lightning, you can train your flash tasks on any hardware: CPUs, GPUs or TPUs without any code changes.

#### **1.1.1 End-to-end deep learning**

Flash is built for 3 major use-cases:

- Inference (predictions)
- Finetuning
- Training

#### **1.1.2 Scalability**

Flash is built on top of [Pytorch Lightning](#), a powerful deep learning research framework for training models at scale. With the power of Lightning, you can train your flash tasks on any hardware: CPUs, GPUs or TPUs without any code changes.

### 1.1.3 Flexibility

Unlike other high-level frameworks, it's easy to customize the Flash tasks with [Pytorch Lightning](#) components to get all the flexibility you need. Lightning is just organized PyTorch with the unnecessary engineering details abstracted away.

- Flash (high level)
- Lightning (mid-level)
- PyTorch (low-level)

When you need more flexibility you can build your own tasks or simply use Lightning directly.

---

**Note:** Read [here](#) to understand when to use Flash vs Lightning.

---

## 1.2 Install

You can install flash using pip or conda:

```
pip install pytorch-lightning-flash -U -pre
```

---

## 1.3 Tasks

Flash is comprised of a collection of Tasks. The Flash tasks are opinionated and laser-focused objects designed to solve a specific type of problem, using state-of-the-art methods.

The Flash tasks contain all the relevant information to solve the task at hand- the number of class labels you want to predict, number of columns in your dataset, as well as details on the model architecture used such as loss function, optimizers, etc.

Here are examples of tasks:

```
from flash.text import TextClassifier
from flash.vision import ImageClassifier
from flash.tabular import TabularClassifier
```

---

**Note:** Tasks are inflexible by definition! To get more flexibility, you can simply use `LightningModule` directly or modify an existing task in just a few lines.

---

## 1.4 Inference

Inference is the process of generating predictions from trained models. To use a task for inference:

1. Init your task with pretrained weights using a checkpoint (a checkpoint is simply a file that capture the exact value of all parameters used by a model). Local file or URL works.
2. Pass in the data to `flash.core.model.Task.predict()`.

Here's an example of inference.

```
# import our libraries
from flash.text import TextClassifier

# 1. Init the finetuned task from URL
model = TextClassifier.load_from_checkpoint("https://flash-weights.s3.amazonaws.com/
↪text_classification_model.pt")

# 2. Perform inference from list of sequences
predictions = model.predict([
    "Turgid dialogue, feeble characterization - Harvey Keitel a judge?.",
    "The worst movie in the history of cinema.",
    "This guy has done a great job with this movie!",
])

# Expect [0,0, 1] which means [negative, negative, positive]
print(predictions)
```

## 1.5 Finetune

Finetuning (or transfer-learning) is the process of tweaking a model trained on a large dataset, to your particular (likely much smaller) dataset. To use a Task for finetuning:

1. Download and set up your own data (`DataLoader` or `LightningModule` work).
2. Init your task.
3. Init a `flash.core.trainer.Trainer` (or a `Lightning Trainer`).
4. Call `flash.core.trainer.Trainer.finetune()` with your data set.
5. Use your finetuned model for predictions

Here's an example of finetuning.

```
import flash
from flash import download_data
from flash.vision import ImageClassificationData, ImageClassifier

# 1. Download the data
download_data("https://pl-flash-data.s3.amazonaws.com/hymenoptera_data.zip", 'data/')

# 2. Load the data from folders
datamodule = ImageClassificationData.from_folders(
    backbone="resnet18",
    train_folder="data/hymenoptera_data/train/",
    valid_folder="data/hymenoptera_data/val/",
    test_folder="data/hymenoptera_data/test/",
)

# 3. Build the model using desired Task
model = ImageClassifier(num_classes=datamodule.num_classes)

# 4. Create the trainer (run one epoch for demo)
trainer = flash.Trainer(max_epochs=1)

# 5. Finetune the model
trainer.finetune(model, datamodule=datamodule, strategy="freeze")

# 6. Use the model for predictions
predictions = model.predict('data/hymenoptera_data/val/bees/65038344_52a45d090d.jpg')
# Expect 1 -> bee
print(predictions)

predictions = model.predict('data/hymenoptera_data/val/ants/2255445811_dabcdf7258.jpg
↪')
# Expect 0 -> ant
print(predictions)

# 7. Save the new model!
trainer.save_checkpoint("image_classification_model.pt")
```

Once your model is finetuned, use it for prediction anywhere you want!

```
from flash.vision import ImageClassifier

# load finetuned checkpoint
model = ImageClassifier.load_from_checkpoint("image_classification_model.pt")

predictions = model.predict('path/to/your/own/image.png')
```



## 1.6 Train

When you have enough data, you're likely better off training from scratch instead of finetuning. Steps here are similar to finetune:

1. Download and set up your own data (`DataLoader` or `LightningModule` work).
  2. Init your task.
  3. Init a `flash.core.trainer.Trainer` (or a `Lightning Trainer`).
  4. Call `flash.core.trainer.Trainer.fit()` with your data set.
  5. Use your finetuned model for predictions
- 

## 1.7 A few Built-in Tasks

- *Task*
  - *ImageClassification*
  - *TextClassification*
  - *TabularClassification*
- 

## 1.8 Contribute a task

The lightning + Flash team is hard at work building more tasks for common deep-learning use cases. But we're looking for incredible contributors like you to submit new tasks!

Join our [Slack](#) to get help becoming a contributor!



## INSTALLATION

Flash is tested on Python 3.6+, and PyTorch 1.6

### 2.1 Install with pip/conda

```
pip install lightning-flash
```

### 2.2 Install from source

```
git clone https://github.com/PyTorchLightning/flash.git
cd flash
pip install -e .
```



## TUTORIAL: CREATING A CUSTOM TASK

In this tutorial we will go over the process of creating a custom task, along with a custom data module.

```
import flash

import torch
from torch.utils.data import TensorDataset, DataLoader
from torch import nn
from sklearn import datasets
from sklearn.model_selection import train_test_split
```

### 3.1 The Task: Linear regression

Here we create a basic linear regression task by subclassing `flash.Task`. For the majority of tasks, you will likely only need to override the `__init__` and `forward` methods.

```
class LinearRegression(flash.Task):
    def __init__(self, num_inputs, learning_rate=0.001, metrics=None):
        # what kind of model do we want?
        model = nn.Linear(num_inputs, 1)

        # what loss function do we want?
        loss_fn = torch.nn.functional.mse_loss

        # what optimizer do we want?
        optimizer = torch.optim.SGD

        super().__init__(
            model=model,
            loss_fn=loss_fn,
            optimizer=optimizer,
            metrics=metrics,
            learning_rate=learning_rate,
        )

    def forward(self, x):
        # we don't actually need to override this method for this example
        return self.model(x)
```

### 3.1.1 Where is the training step?

Most models can be trained simply by passing the output of `forward` to the supplied `loss_fn`, and then passing the resulting loss to the supplied `optimizer`. If you need a more custom configuration, you can override `step` (which is called for training, validation, and testing) or override `training_step`, `validation_step`, and `test_step` individually. These methods behave identically to PyTorch Lightning's [methods](#).

## 3.2 The Data

For a task you will likely need a specific way of loading data. For this example, let's say we want a `flash.DataModule` to be used explicitly for the prediction of diabetes disease progression. We can create this `DataModule` below, wrapping the scikit-learn [Diabetes dataset](#).

```
class DiabetesPipeline(flash.core.data.TaskDataPipeline):
    def after_uncollate(self, samples):
        return [f"disease progression: {float(s):.2f}" for s in samples]

class DiabetesData(flash.DataModule):
    def __init__(self, batch_size=64, num_workers=0):
        x, y = datasets.load_diabetes(return_X_y=True)
        x = torch.from_numpy(x).float()
        y = torch.from_numpy(y).float().unsqueeze(1)
        x_train, x_test, y_train, y_test = train_test_split(x, y, test_size=.20,
        ↪random_state=0)

        train_ds = TensorDataset(x_train, y_train)
        test_ds = TensorDataset(x_test, y_test)

        super().__init__(
            train_ds=train_ds,
            test_ds=test_ds,
            batch_size=batch_size,
            num_workers=num_workers
        )
        self.num_inputs = x.shape[1]

    @staticmethod
    def default_pipeline():
        return DiabetesPipeline()
```

You'll notice we added a `DataPipeline`, which will be used when we call `.predict()` on our model. In this case we want to nicely format our output from the model with the string "disease progression", but you could do any sort of post processing you want (see [DataPipeline](#)).

### 3.3 Fit

Like any Flash Task, we can fit our model using the `flash.Trainer` by supplying the task itself, and the associated data:

```
data = DiabetesData()
model = LinearRegression(num_inputs=data.num_inputs)

trainer = flash.Trainer(max_epochs=1000)
trainer.fit(model, data)
```

With a trained model we can now perform inference. Here we will use a few examples from the test set of our data:

```
predict_data = torch.tensor([
    [ 0.0199,  0.0507,  0.1048,  0.0701, -0.0360, -0.0267, -0.0250, -0.0026,  0.0037,  ↵
↪0.0403],
    [-0.0128, -0.0446,  0.0606,  0.0529,  0.0480,  0.0294, -0.0176,  0.0343,  0.0702,  ↵
↪0.0072],
    [ 0.0381,  0.0507,  0.0089,  0.0425, -0.0428, -0.0210, -0.0397, -0.0026, -0.0181, ↵
↪ 0.0072],
    [-0.0128, -0.0446, -0.0235, -0.0401, -0.0167,  0.0046, -0.0176, -0.0026, -0.0385, ↵
↪-0.0384],
    [-0.0237, -0.0446,  0.0455,  0.0907, -0.0181, -0.0354,  0.0707, -0.0395, -0.0345, ↵
↪-0.0094]])

model.predict(predict_data)
```

Because of our custom data pipeline's `after_uncollate` method, we will get a nicely formatted output like the following:

```
['disease progression: 155.90',
 'disease progression: 156.59',
 'disease progression: 152.69',
 'disease progression: 149.05',
 'disease progression: 150.90']
```





## FROM FLASH TO LIGHTNING

Flash is built on top of [Pytorch Lightning](#) to abstract away the unnecessary boilerplate for:

- Data science
- Kaggle
- Business use cases
- Applied research

Flash is a HIGH level library and Lightning is a LOW level library.

- Flash (High-level)
- Lightning (medium-level)
- PyTorch (low-level)

As the complexity increases or decreases, users can move between Flash and Lightning seamlessly to find the level of abstraction that works for them.

Table 1: Abstraction levels

Approach	Flexibility	Minimum DL Expertise level	PyTorch Knowledge	Use cases
Using an out-of-the-box task	Low	Novice+	Low+	Fast baseline, Data Science, Analysis, Applied Research
Using the Generic Task	Medium	Intermediate+	Intermediate+	Fast baseline, data science
Building a custom task	High	Intermediate+	Intermediate+	Fast baseline, custom business context, applied research
Building a LightningModule	Ultimate (organized PyTorch)	Expert+	Expert+	For anything you can do with PyTorch, AI research (academic and corporate)

## 4.1 Using an out-of-the-box task

Tasks can come from a variety of places:

- Flash
- Other Lightning-based libraries
- Your own library

Using a task requires almost zero knowledge of deep learning and PyTorch. The focus is on solving a problem as quickly as possible. This is great for:

- data science
  - analysis
  - applied research
- 

## 4.2 Using the Generic Task

If you encounter a problem that does not have a matching task, you can use the generic task. However, this does require a bit of PyTorch knowledge but not a lot of knowledge over all the details of deep learning.

This is great for:

- data science
  - kaggle baselines
  - a quick baseline
  - applied research
  - learning about deep learning
- 

**Note:** If you've used something like Keras, this is the most similar level of abstraction.

---

## 4.3 Building a custom task

If you're feeling adventurous and there isn't an out-of-the-box task for a particular applied problem, consider building your own task. This requires a decent amount of PyTorch knowledge, but not too much because tasks are Lightning-Modules that already abstract a lot of the details for you.

This is great for:

- data science
- researchers building for corporate data science teams
- applied research
- custom business context

---

**Note:** In a company setting, a good setup here is to have your own Flash-like library with tasks contextualized with your business problems.

---

## 4.4 Building a LightningModule

Once you've reached the threshold of flexibility offered by Flash, it's time to move to a LightningModule directly. LightningModule is organized PyTorch but gives you the same flexibility. However, you must already know PyTorch fairly well and be comfortable with at least basic deep learning concepts.

This is great for:

- experts
- academic AI research
- corporate AI research
- advanced applied research
- publishing papers



## GENERAL TASK

A majority of data science problems that involve machine learning can be tackled using Task. With Task you can:

- Pass an arbitrary model
- Pass an arbitrary loss
- Pass an arbitrary optimizer

### 5.1 Example: Image Classification

```
from flash import Task
from torch import nn, optim
from torch.utils.data import DataLoader, random_split
from torchvision import transforms, datasets
import pytorch_lightning as pl

# model
model = nn.Sequential(
    nn.Flatten(),
    nn.Linear(28 * 28, 128),
    nn.ReLU(),
    nn.Linear(128, 10)
)

# data
dataset = datasets.MNIST('./data_folder', download=True, transform=transforms.
    ↳ ToTensor())
train, val = random_split(dataset, [55000, 5000])

# task
classifier = Task(model, loss_fn=nn.functional.cross_entropy, optimizer=optim.Adam)

# train
pl.Trainer().fit(classifier, DataLoader(train), DataLoader(val))
```

## 5.2 API reference

### 5.2.1 Task

```
class flash.core.Task(model=None, loss_fn=None, optimizer=torch.optim.Adam, metrics=None,  
                    learning_rate=5e-05)
```

A general Task.

#### Parameters

- **model** `Optional[Module]` – Model to use for the task.
- **loss\_fn** `Union[Callable, Mapping, Sequence, None]` – Loss function for training
- **optimizer** `Type[Optimizer]` – Optimizer to use for training, defaults to `torch.optim.SGD`.
- **metrics** `Union[Metric, Mapping, Sequence, None]` – Metrics to compute for training and evaluation.
- **learning\_rate** `float` – Learning rate to use for training, defaults to `5e-5`

```
static default_pipeline()
```

Pipeline to use when there is no datamodule or it has not defined its pipeline

**Return type** `DataPipeline`

```
step(batch, batch_idx)
```

The training/validation/test step. Override for custom behavior.

**Return type** `Any`

## IMAGE CLASSIFICATION

### 6.1 The task

The task of identifying what is in an image is called **image classification**. Typically, Image Classification is used to identify images containing a single object. The task predicts which ‘class’ the image most likely belongs to with a degree of certainty. A class is a label that describes what is in an image, such as ‘car’, ‘house’, ‘cat’ etc. For example, we can train the image classifier task on images of ants and it will learn to predict the probability that an image contains an ant.

---

### 6.2 Inference

The *ImageClassifier* is already pre-trained on **ImageNet**, a dataset of over 14 million images.

Use the *ImageClassifier* pretrained model for inference on any string sequence using `predict()`:

```
# import our libraries
from flash import Trainer
from flash import download_data
from flash.vision import ImageClassificationData, ImageClassifier

# 1. Download the data
download_data("https://pl-flash-data.s3.amazonaws.com/hymenoptera_data.zip", 'data/')

# 2. Load the model from a checkpoint
model = ImageClassifier.load_from_checkpoint("image_classification_model.pt")

# 3a. Predict what's on a few images! ants or bees?
predictions = model.predict([
    "data/hymenoptera_data/val/bees/65038344_52a45d090d.jpg",
    "data/hymenoptera_data/val/bees/590318879_68cf112861.jpg",
    "data/hymenoptera_data/val/ants/540543309_ddbb193ee5.jpg",
])
print(predictions)

# 3b. Or generate predictions with a whole folder!
datamodule = ImageClassificationData.from_folder(folder="data/hymenoptera_data/
↳predict/")
predictions = Trainer().predict(model, datamodule=datamodule)
print(predictions)
```

For more advanced inference options, see *Predictions (inference)*.

---

## 6.3 Finetuning

Lets say you wanted to develop a model that could determine whether an image contains **ants** or **bees**, using the hymenoptera dataset. Once we download the data using `download_data()`, all we need is the train data and validation data folders to create the *ImageClassificationData*.

**Note:** The dataset contains `train` and `validation` folders, and then each folder contains a **bees** folder, with pictures of bees, and an **ants** folder with images of, you guessed it, ants.

---

```
hymenoptera_data
├── train
│   ├── ants
│   │   ├── 0013035.jpg
│   │   ├── 1030023514_aad5c608f9.jpg
│   │   └── ...
│   └── bees
│       ├── 1092977343_cb42b38d62.jpg
│       ├── 1093831624_fb5fbe2308.jpg
│       └── ...
└── val
    ├── ants
    │   ├── 10308379_1b6c72e180.jpg
    │   ├── 1053149811_f62a3410d3.jpg
    │   └── ...
    └── bees
        ├── 1032546534_06907fe3b3.jpg
        ├── 10870992_eebeeb3a12.jpg
        └── ...
```

Now all we need is three lines of code to build to train our task!

```
import flash
from flash import download_data
from flash.vision import ImageClassificationData, ImageClassifier

# 1. Download the data
download_data("https://pl-flash-data.s3.amazonaws.com/hymenoptera_data.zip", 'data/')

# 2. Load the data
datamodule = ImageClassificationData.from_folders(
    backbone="resnet18",
    train_folder="data/hymenoptera_data/train/",
    valid_folder="data/hymenoptera_data/val/",
    test_folder="data/hymenoptera_data/test/",
)

# 3. Build the model
model = ImageClassifier(num_classes=datamodule.num_classes)

# 4. Create the trainer. Run once on data
```

(continues on next page)



(continued from previous page)

```

trainer = flash.Trainer(max_epochs=1)

# 5. Train the model
trainer.finetune(model, datamodule=datamodule, strategy="freeze_unfreeze")

# 6. Test the model
trainer.test()

# 7. Save it!
trainer.save_checkpoint("image_classification_model.pt")

```

## 6.4 Changing the backbone

By default, we use a [ResNet-18](#) for image classification. You can change the model run by the task by passing in a different backbone.

**Note:** When changing the backbone, make sure you pass in the same backbone to the Task and the Data object!

```

# 1. organize the data
data = ImageClassificationData.from_folders(
    backbone="resnet34",
    train_folder="data/hymenoptera_data/train/",
    valid_folder="data/hymenoptera_data/val/"
)

# 2. build the task
task = ImageClassifier(num_classes=2, backbone="resnet34")

```

Available backbones:

- resnet34
- resnet50
- resnet101
- resnet152

## 6.5 API reference

### 6.5.1 ImageClassifier

```

class flash.vision.ImageClassifier(num_classes,                backbone='resnet18',
                                   num_features=None,          pretrained=True,
                                   loss_fn=torch.nn.functional.cross_entropy,
                                   optimizer=torch.optim.SGD,   met-
                                   rics=pytorch_lightning.metrics.Accuracy,  learn-
                                   ing_rate=0.001)

```

Task that classifies images.

**Parameters**

- **num\_classes** – Number of classes to classify.
- **backbone** – A model to use to compute image features.
- **pretrained** – Use a pretrained backbone.
- **loss\_fn** (Callable) – Loss function for training, defaults to cross entropy.
- **optimizer** (Type[Optimizer]) – Optimizer to use for training, defaults to *torch.optim.SGD*.
- **metrics** (Union[Callable, Mapping, Sequence, None]) – Metrics to compute for training and evaluation.
- **learning\_rate** (float) – Learning rate to use for training, defaults to *1e-3*

**static default\_pipeline()**

Pipeline to use when there is no datamodule or it has not defined its pipeline

**Return type** ImageClassificationDataPipeline

**6.5.2 ImageClassificationData**

```
class flash.vision.ImageClassificationData (train_ds=None,          valid_ds=None,
                                           test_ds=None,           batch_size=1,
                                           num_workers=None)
```

Data module for image classification tasks.

```
classmethod ImageClassificationData.from_filepaths (train_filepaths=None,
                                                    train_labels=None,
                                                    train_transform=torchvision.transforms.Compose,
                                                    valid_filepaths=None,
                                                    valid_labels=None,
                                                    valid_transform=torchvision.transforms.Compose,
                                                    test_filepaths=None,
                                                    test_labels=None,
                                                    loader=<function
                                                    _pil_loader>,    batch_size=64,
                                                    num_workers=None, **kwargs)
```

Creates a ImageClassificationData object from lists of image filepaths and labels

**Parameters**

- **train\_filepaths** (Optional[Sequence[Union[str, Path]]]) – sequence of file paths for training dataset. Defaults to None.
- **train\_labels** (Optional[Sequence]) – sequence of labels for training dataset. Defaults to None.
- **train\_transform** (Optional[Callable]) – transforms for training dataset. Defaults to None.
- **valid\_filepaths** (Optional[Sequence[Union[str, Path]]]) – sequence of file paths for validation dataset. Defaults to None.
- **valid\_labels** (Optional[Sequence]) – sequence of labels for validation dataset. Defaults to None.
- **valid\_transform** (Optional[Callable]) – transforms for validation and testing dataset. Defaults to None.

- **test\_filepaths** (Optional[Sequence[Union[str, Path]]]) – sequence of file paths for test dataset. Defaults to None.
- **test\_labels** (Optional[Sequence]) – sequence of labels for test dataset. Defaults to None.
- **loader** (Callable) – function to load an image file. Defaults to None.
- **batch\_size** (int) – the batchsize to use for parallel loading. Defaults to 64.
- **num\_workers** (Optional[int]) – The number of workers to use for parallelized loading. Defaults to None which equals the number of available CPU threads.

**Returns** The constructed data module.

**Return type** *ImageClassificationData*

### Examples

```
>>> img_data = ImageClassificationData.from_filepaths(["a.png", "b.png"], [0, 1])
```

```
classmethod ImageClassificationData.from_folders(train_folder,
                                                  train_transform=torchvision.transforms.Compose,
                                                  valid_folder=None,
                                                  valid_transform=torchvision.transforms.Compose,
                                                  test_folder=None, loader=<function
                                                  _pil_loader>, batch_size=4,
                                                  num_workers=None, **kwargs)
```

Creates a ImageClassificationData object from folders of images arranged in this way:

```
train/dog/xxx.png
train/dog/xyx.png
train/dog/xxz.png
train/cat/123.png
train/cat/nsdf3.png
train/cat/asd932.png
```

### Parameters

- **train\_folder** (Union[str, Path, None]) – Path to training folder.
- **train\_transform** (Optional[Callable]) – Image transform to use for training set.
- **valid\_folder** (Union[str, Path, None]) – Path to validation folder.
- **valid\_transform** (Optional[Callable]) – Image transform to use for validation and test set.
- **test\_folder** (Union[str, Path, None]) – Path to test folder.
- **loader** (Callable) – A function to load an image given its path.
- **batch\_size** (int) – Batch size for data loading.
- **num\_workers** (Optional[int]) – The number of workers to use for parallelized loading. Defaults to None which equals the number of available CPU threads.

**Returns** the constructed data module

**Return type** *ImageClassificationData*

**Examples**

```
>>> img_data = ImageClassificationData.from_folders("train/")
```

## IMAGE EMBEDDER

### 7.1 The task

Image embedding encodes an image into a vector of image features which can be used for anything like clustering, similarity search or classification.

---

### 7.2 Inference

The *ImageEmbedder* is already pre-trained on [ImageNet](<http://www.image-net.org/>), a dataset of over 14 million images.

Use the *ImageEmbedder* pretrained model for inference on any image tensor or image path using `predict()`:

```
from flash.vision import ImageEmbedder

# Load finetuned task
embedder = ImageEmbedder(backbone='resnet18')

# 2. Perform inference on an image file
embeddings = model.predict('path/to/image.png')
print(predictions)
```

Or on a random image tensor

```
# 2. Perform inference on an image file
import torch
images = torch.rand(32, 3, 224, 224)
embeddings = model.predict(images)
print(predictions)
```

For more advanced inference options, see *Predictions (inference)*.

---

## 7.3 Finetuning

To tailor this image embedder to your dataset, finetune first.

```
import flash
from flash import download_data
from flash.vision import ImageClassificationData, ImageClassifier

# 1. Download the data
download_data("https://pl-flash-data.s3.amazonaws.com/hymenoptera_data.zip", 'data/')

# 2. Load the data
datamodule = ImageClassificationData.from_folders(
    train_folder="data/hymenoptera_data/train/",
    valid_folder="data/hymenoptera_data/val/",
    test_folder="data/hymenoptera_data/test/",
)

# 3. Build the model
embedder = ImageEmbedder(backbone="resnet18")

# 4. Create the trainer. Run once on data
trainer = flash.Trainer(max_epochs=1)

# 5. Train the model
trainer.finetune(model, datamodule=datamodule)

# 6. Test the model
trainer.test()

# 7. Save it!
trainer.save_checkpoint("image_embedder_model.pt")
```

---

## 7.4 Changing the backbone

By default, we use the encoder from [Swav](#) pretrained on Imagenet via contrastive learning. You can change the model run by the task by passing in a different backbone.

---

**Note:** When changing the backbone, make sure you pass in the same backbone to the Task and the Data object!

---

```
# 1. organize the data
data = ImageClassificationData.from_folders(
    backbone="resnet34",
    train_folder="data/hymenoptera_data/train/",
    valid_folder="data/hymenoptera_data/val/"
)

# 2. build the task
task = ImageClassifier(num_classes=2, backbone="resnet34")
```

Backbones available

Table 1: Backbones

backbone	dataset	training method
resnet18	Imagenet	supervised
resnet34	Imagenet	supervised
resnet50	Imagenet	supervised
resnet101	Imagenet	supervised
resnet152	Imagenet	supervised
swav-imagenet	Imagenet	self-supervised (clustering)

## 7.5 API reference

### 7.5.1 ImageEmbedder

**class** `flash.vision.ImageEmbedder` (*embedding\_dim=None*, *backbone='swav-imagenet'*, *pre-trained=True*, *loss\_fn=torch.nn.functional.cross\_entropy*, *optimizer=torch.optim.SGD*, *metrics=pytorch\_lightning.metrics.Accuracy*, *learning\_rate=0.001*, *pooling\_fn=torch.max*)

Task that classifies images.

#### Parameters

- **embedding\_dim** (Optional[int]) – Dimension of the embedded vector. None uses the default from the backbone
- **backbone** (str) – A model to use to extract image features.
- **pretrained** (bool) – Use a pretrained backbone.
- **loss\_fn** (Callable) – Loss function for training and finetuning, defaults to cross entropy.
- **optimizer** (Type[Optimizer]) – Optimizer to use for training and finetuning, defaults to `torch.optim.SGD`.
- **metrics** (Union[Callable, Mapping, Sequence, None]) – Metrics to compute for training and evaluation.
- **learning\_rate** (float) – Learning rate to use for training, defaults to `1e-3`
- **pooling\_fn** (Callable) – Function used to pool image to generate embeddings. (Default: `torch.max`)

Example:

```
from flash.vision import ImageEmbedder

embedder = ImageEmbedder(backbone='swav-imagenet')
image = torch.rand(32, 3, 32, 32)
embeddings = embedder(image)
```

**static default\_pipeline()**

Pipeline to use when there is no datamodule or it has not defined its pipeline

**Return type** ImageEmbedderDataPipeline



## SUMMARIZATION

### 8.1 The task

Summarization is the task of summarizing text from a larger document/article into a short sentence/description. For example, taking a web article and describing the topic in a short sentence. This task is a subset of Sequence to Sequence tasks, which requires the model to generate a variable length sequence given an input sequence. In our case the article would be our input sequence, and the short description/sentence would be the output sequence from the model.

---

### 8.2 Inference

The SummarizationTask is already pre-trained on [XSUM](<https://arxiv.org/abs/1808.08745>), a dataset of online British Broadcasting Corporation articles.

Use the SummarizationTask pretrained model for inference on any string sequence using `predict()`:

```
# import our libraries
from flash.text import SummarizationTask

# 2. Load the model from a checkpoint
model = SummarizationTask.load_from_checkpoint("https://flash-weights.s3.amazonaws.
↳com/summarization_model_xsum.pt")

# 2. Perform inference from a sequence
predictions = model.predict([
    """
    Camilla bought a box of mangoes with a Brixton Â£10 note, introduced last year to
    ↳try to keep the money of local
    people within the community.The couple were surrounded by shoppers as they walked
    ↳along Electric Avenue.
    They came to Brixton to see work which has started to revitalise the borough.
    It was Charles' first visit to the area since 1996, when he was accompanied by
    ↳the former
    South African president Nelson Mandela.Green grocer Derek Chong, who has run a
    ↳stall on Electric Avenue
    for 20 years, said Camilla had been ""nice and pleasant"" when she purchased the
    ↳fruit.
    ""She asked me what was nice, what would I recommend, and I said we've got some
    ↳nice mangoes.
    She asked me were they ripe and I said yes - they're from the Dominican Republic."
    ↳""
```

(continues on next page)

(continued from previous page)

```

    Mr Chong is one of 170 local retailers who accept the Brixton Pound.
    Customers exchange traditional pound coins for Brixton Pounds and then spend them_
    ↪at the market
    or in participating shops.
    During the visit, Prince Charles spent time talking to youth worker Marcus West,_
    ↪who works with children
    nearby on an estate off Coldharbour Lane. Mr West said:
    ""He's on the level, really down-to-earth. They were very cheery. The prince is a_
    ↪lovely man.""
    He added: ""I told him I was working with young kids and he said, 'Keep up all_
    ↪the good work.'""
    Prince Charles also visited the Railway Hotel, at the invitation of his charity_
    ↪The Prince's Regeneration Trust.
    The trust hopes to restore and refurbish the building,
    where once Jimi Hendrix and The Clash played, as a new community and business_
    ↪centre."
    ""
])
print(predictions)

```

Or on a given dataset:

```

# import our libraries
from pytorch_lightning import Trainer
from flash import download_data
from flash.text import SummarizationData, SummarizationTask

# 2. Load the model from a checkpoint
model = SummarizationTask.load_from_checkpoint("https://flash-weights.s3.amazonaws._
    ↪com/summarization_model_xsum.pt")

# 3. Create dataset from file
datamodule = SummarizationData.from_file(
    predict_file="data/xsum/predict.csv",
    input="input",
)

# 4. generate summaries
predictions = Trainer().predict(model, datamodule=datamodule)
print(predictions)

```

For more advanced inference options, see *Predictions (inference)*.

## 8.3 Finetuning

Say you want to finetune to your own summarization data. We use the XSUM dataset as an example which contains a `train.csv` and `valid.csv`, structured like so:

```

input,target
"The researchers have sequenced the genome of a strain of bacterium that causes the_
    ↪virulent infection...", "A team of UK scientists hopes to shed light on the_
    ↪mysteries of bleeding canker, a disease that is threatening the nation's horse_
    ↪chestnut trees."

```

(continues on next page)

(continued from previous page)

```
"Knight was shot in the leg by an unknown gunman at Miami's Shore Club where West was
↳holding a pre-MTV Awards..." Hip hop star Kanye West is being sued by Death Row
↳Records founder Suge Knight over a shooting at a beach party in August 2005.
...
```

In the above the input column represents the long articles/documents, and the target is the short description used as the target.

All we need is three lines of code to train our model!

```
# import our libraries
import flash
from flash import download_data
from flash.text import SummarizationData, SummarizationTask

# 1. Download data
download_data("https://pl-flash-data.s3.amazonaws.com/xsum.zip", 'data/')

# Organize the data
datamodule = SummarizationData.from_files(
    train_file="data/xsum/train.csv",
    valid_file="data/xsum/valid.csv",
    test_file="data/xsum/test.csv",
    input="input",
    target="target"
)

# 2. Build the task
model = SummarizationTask()

# 4. Create trainer
trainer = flash.Trainer(max_epochs=1, gpus=1)

# 5. Finetune the task
trainer.finetune(model, datamodule=datamodule)

# 6. Save trainer task
trainer.save_checkpoint("summarization_model_xsum.pt")
```

To run the example:

```
python flash_examples/finetuning/summarization.py
```

## 8.4 Changing the backbone

By default, we use the `t5` model for summarization. You can change the model run by passing in the backbone parameter.

**Note:** When changing the backbone, make sure you pass in the same backbone to the Task and the Data object! Since this is a Seq2Seq task, make sure you use a Seq2Seq model.

```
datamodule = SummarizationData.from_files(  
    train_file="data/wmt_en_ro/train.csv",  
    valid_file="data/wmt_en_ro/valid.csv",  
    test_file="data/wmt_en_ro/test.csv",  
    input="input",  
    target="target",  
    backbone="google/mt5-small",  
)  
  
model = SummarizationTask(backbone="google/mt5-small")
```

---

## 8.5 API reference

### 8.5.1 SummarizationTask

```
class flash.text.seq2seq.summarization.model.SummarizationTask(backbone='t5-  
small',  
loss_fn=None,  
optimizer=torch.optim.Adam,  
metrics=None,  
learning_rate=5e-05,  
val_target_max_length=None,  
num_beams=4,  
use_stemmer=True,  
rouge_newline_sep=True)
```

**property task**

Override to define AutoConfig task specific parameters stored within the model.

**Return type** `str`

### 8.5.2 SummarizationData

```
class flash.text.seq2seq.summarization.data.SummarizationData(train_ds=None,  
valid_ds=None,  
test_ds=None,  
batch_size=1,  
num_workers=None)
```

```
classmethod SummarizationData.from_files(train_file, input='input', target=  
get=None, filetype='csv',  
backbone='t5-small', valid_file=None,  
test_file=None, max_source_length=512,  
max_target_length=128, padding='max_length',  
batch_size=16, num_workers=None)
```

Creates a SummarizationData object from files.

**Parameters**

- **train\_file** `(str)` – Path to training data.

- **input** (str) – The field storing the source translation text.
- **target** (Optional[str]) – The field storing the target translation text.
- **filetype** (str) – .csv or .json
- **backbone** (str) – tokenizer to use, can use any HuggingFace tokenizer.
- **valid\_file** (Optional[str]) – Path to validation data.
- **test\_file** (Optional[str]) – Path to test data.
- **max\_source\_length** (int) – Maximum length of the source text. Any text longer will be truncated.
- **max\_target\_length** (int) – Maximum length of the target text. Any text longer will be truncated.
- **padding** (Union[str, bool]) – Padding strategy for batches. Default is pad to maximum length.
- **batch\_size** (int) – the batchsize to use for parallel loading. Defaults to 16.
- **num\_workers** (Optional[int]) – The number of workers to use for parallelized loading. Defaults to None which equals the number of available CPU threads.

**Returns** The constructed data module.

**Return type** *SummarizationData*

Examples:

```
train_df = pd.read_csv("train_data.csv")
tab_data = TabularData.from_df(train_df, target="fraud",
                               numerical_input=["account_value"],
                               categorical_input=["account_type"])
```



## TEXT CLASSIFICATION

### 9.1 The task

Text classification is the task of assigning a piece of text (word, sentence or document) an appropriate class, or category. The categories depend on the chosen dataset and can range from topics. For example, we can use text classification to understand the sentiment of a given sentence- if it is positive or negative.

---

### 9.2 Inference

The *TextClassifier* is already pre-trained on **IMDB**, a dataset of highly polarized movie reviews, trained for binary classification- to predict if a given review has a positive or negative sentiment.

Use the *TextClassifier* pretrained model for inference on any string sequence using `predict()`:

```
from pytorch_lightning import Trainer

from flash import download_data
from flash.text import TextClassificationData, TextClassifier

# 1. Download the data
download_data("https://pl-flash-data.s3.amazonaws.com/imdb.zip", 'data/')

# 2. Load the model from a checkpoint
model = TextClassifier.load_from_checkpoint("https://flash-weights.s3.amazonaws.com/
↪text_classification_model.pt")

# 2a. Classify a few sentences! How was the movie?
predictions = model.predict([
    "Turgid dialogue, feeble characterization - Harvey Keitel a judge?.",
    "The worst movie in the history of cinema.",
    "I come from Bulgaria where it 's almost impossible to have a tornado."
    "Very, very afraid"
    "This guy has done a great job with this movie!",
])
print(predictions)

# 2b. Or generate predictions from a sheet file!
datamodule = TextClassificationData.from_file(
    predict_file="data/imdb/predict.csv",
```

(continues on next page)

(continued from previous page)

```
        input="review",
    )
    predictions = Trainer().predict(model, datamodule=datamodule)
    print(predictions)
```

For more advanced inference options, see *Predictions (inference)*.

## 9.3 Finetuning

Say you wanted to create a model that can predict whether a movie review is **positive** or **negative**. We will be using the IMDB dataset, that contains a `train.csv` and `valid.csv`, structured like so:

```
review,sentiment
"Japanese indie film with humor ... ",positive
"Isaac Florentine has made some ...",negative
"After seeing the low-budget ...",negative
"I've seen the original English version ...",positive
"Hunters chase what they think is a man through ...",negative
...
```

All we need is three lines of code to train our model!

```
import flash
from flash.core.data import download_data
from flash.text import TextClassificationData, TextClassifier

# 1. Download the data
download_data("https://pl-flash-data.s3.amazonaws.com/imdb.zip", 'data/')

# 2. Load the data
datamodule = TextClassificationData.from_files(
    train_file="data/imdb/train.csv",
    valid_file="data/imdb/valid.csv",
    test_file="data/imdb/test.csv",
    input="review",
    target="sentiment",
    batch_size=512
)

# 3. Build the task (using the default backbone="bert-base-cased")
model = TextClassifier(num_classes=datamodule.num_classes)

# 4. Create the trainer. Run once on data
trainer = flash.Trainer(max_epochs=1)

# 5. Finetune the task
trainer.finetune(model, datamodule=datamodule, strategy="freeze_unfreeze")

# 6. Test model
trainer.test()

# 7. Save it!
trainer.save_checkpoint("text_classification_model.pt")
```



To run the example:

```
python flash_examples/finetuning/text_classification.py
```

## 9.4 Changing the backbone

By default, we use the `bert-base-uncased` model for text classification. You can change the model run by the task to any BERT model from [HuggingFace/transformers](#) by passing in a different backbone.

**Note:** When changing the backbone, make sure you pass in the same backbone to the Task and the Data object!

```
datamodule = TextClassificationData.from_files(
    backbone="bert-base-chinese",
    train_file="data/imdb/train.csv",
    valid_file="data/imdb/valid.csv",
    input="review",
    target="sentiment",
    batch_size=512
)

task = TextClassifier(backbone="bert-base-chinese", num_classes=datamodule.num_
    ↪ classes)
```

## 9.5 API reference

### 9.5.1 TextClassifier

```
class flash.text.classification.model.TextClassifier(num_classes,
                                                    backbone='prajjwal1/bert-
                                                    tiny',
                                                    optimizer=torch.optim.Adam, met-
                                                    rics=[pytorch_lightning.metrics.classification.Accuracy
                                                    learning_rate=0.001])
```

Task that classifies text.

#### Parameters

- **num\_classes** (int) – Number of classes to classify.
- **backbone** (str) – A model to use to compute text features can be any BERT model from HuggingFace/transformersimage .
- **optimizer** (Type[Optimizer]) – Optimizer to use for training, defaults to `torch.optim.Adam`.
- **metrics** (Union[Callable, Mapping, Sequence, None]) – Metrics to compute for training and evaluation.

- **learning\_rate** (float) – Learning rate to use for training, defaults to  $1e-3$

**static default\_pipeline()**

Pipeline to use when there is no datamodule or it has not defined its pipeline

**Return type** `ClassificationDataPipeline`

**step** (*batch*, *batch\_idx*)

The training/validation/test step. Override for custom behavior.

**Return type** `dict`

## 9.5.2 TextClassificationData

```
class flash.text.classification.data.TextClassificationData (train_ds=None,
                                                         valid_ds=None,
                                                         test_ds=None,
                                                         batch_size=1,
                                                         num_workers=None)
```

Data module for text classification tasks.

```
classmethod TextClassificationData.from_files (train_file, input, target, filetype='csv',
                                                backbone='prajjwal1/bert-tiny',
                                                valid_file=None, test_file=None,
                                                max_length=128, batch_size=16,
                                                num_workers=None)
```

Creates a TextClassificationData object from files.

### Parameters

- **train\_file** – Path to training data.
- **input** – The field storing the text to be classified.
- **target** – The field storing the class id of the associated text.
- **filetype** – .csv or .json
- **backbone** – tokenizer to use, can use any HuggingFace tokenizer.
- **valid\_file** – Path to validation data.
- **test\_file** – Path to test data.
- **batch\_size** (int) – the batchsize to use for parallel loading. Defaults to 64.
- **num\_workers** (Optional[int]) – The number of workers to use for parallelized loading. Defaults to None which equals the number of available CPU threads.

**Returns** The constructed data module.

**Return type** `TextClassificationData`

Examples:

```
train_df = pd.read_csv("train_data.csv")
tab_data = TabularData.from_df(train_df, target="fraud",
                               numerical_input=["account_value"],
                               categorical_input=["account_type"])
```

## TABULAR CLASSIFICATION

### 10.1 The task

Tabular classification is the task of assigning a class to samples of structured or relational data. The Flash Tabular Classification task can be used for multi-class classification, or classification of samples in more than two classes. In the following example, the Tabular data is structured into rows and columns, where columns represent properties or features. The task will learn to predict a single target column.

---

### 10.2 Finetuning

Say we want to build a model to predict if a passenger survived on the Titanic. We can organize our data in .csv files (exportable from Excel, but you can find the kaggle dataset [here](#)):

```
PassengerId,Survived,Pclass,Name,Sex,Age,SibSp,Parch,Ticket,Fare,Cabin,Embarked
1,0,3,"Braund, Mr. Owen Harris",male,22,1,0,A/5 21171,7.25,,S
3,1,3,"Heikkinen, Miss. Laina",female,26,0,0,STON/O2. 3101282,7.925,,S
5,0,3,"Allen, Mr. William Henry",male,35,0,0,373450,8.05,,S
6,0,3,"Moran, Mr. James",male,,0,0,330877,8.4583,,Q
...
```

We can use the Flash Tabular classification task to predict the probability a passenger survived (1 means survived, 0 otherwise), using the feature columns.

We can create *TabularData* from csv files using the *from\_csv()* method. We will pass in:

- **train\_csv**- csv file containing the training data converted to a Pandas DataFrame
- **categorical\_input**- a list of the names of columns that contain categorical data (strings or integers)
- **numerical\_input**- a list of the names of columns that contain numerical continuous data (floats)
- **target**- the name of the column we want to predict

---

**Tip:** you can pass in *val\_size* and *test\_size* to partition your training data into a separate validation and test set like so:

---

Next, we create the *TabularClassifier* task, using the Data module we created.

```

import flash
from flash import download_data
from flash.tabular import TabularClassifier, TabularData
from pytorch_lightning.metrics.classification import Accuracy, Precision, Recall

# 1. Download the data
download_data("https://pl-flash-data.s3.amazonaws.com/titanic.zip", 'data/')

# 2. Load the data
datamodule = TabularData.from_csv(
    "./data/titanic/titanic.csv",
    test_csv="./data/titanic/test.csv",
    categorical_input=["Sex", "Age", "SibSp", "Parch", "Ticket", "Cabin", "Embarked"],
    numerical_input=["Fare"],
    target="Survived",
    val_size=0.25,
)

# 3. Build the model
model = TabularClassifier.from_data(datamodule, metrics=[Accuracy(), Precision(),
↪ Recall()])

# 4. Create the trainer. Run 10 times on data
trainer = flash.Trainer(max_epochs=10)

# 5. Train the model
trainer.fit(model, datamodule=datamodule)

# 6. Test model
trainer.test()

# 7. Save it!
trainer.save_checkpoint("tabular_classification_model.pt")

# 8. Predict!
predictions = model.predict("data/titanic/titanic.csv")
print(predictions)

```

## 10.3 Inference

You can make predictions on a pretrained model, that has already been trained for the titanic task:

```

from flash.core.data import download_data
from flash.tabular import TabularClassifier

# 1. Download the data
download_data("https://pl-flash-data.s3.amazonaws.com/titanic.zip", 'data/')

# 2. Load the model from a checkpoint
model = TabularClassifier.load_from_checkpoint(
    "https://flash-weights.s3.amazonaws.com/tabular_classification_model.pt"
)

```

(continues on next page)

(continued from previous page)

```
# 3. Generate predictions from a sheet file! Who would survive?
predictions = model.predict("data/titanic/titanic.csv")
print(predictions)
```

Or you can finetune your own model and use that for prediction:

```
import flash
from flash import download_data
from flash.tabular import TabularClassifier, TabularData

# 1. Load the data
datamodule = TabularData.from_csv(
    "my_data_file.csv",
    test_csv="./data/titanic/test.csv",
    categorical_input=["Sex", "Age", "SibSp", "Parch", "Ticket", "Cabin", "Embarked"],
    numerical_input=["Fare"],
    target="Survived",
    val_size=0.25,
)

# 3. Build the model
model = TabularClassifier.from_data(datamodule, metrics=[Accuracy(), Precision(),
↳ Recall()])

# 4. Create the trainer
trainer = flash.Trainer()

# 5. Train the model
trainer.fit(model, datamodule=datamodule)

# 6. Test model
trainer.test()

predictions = model.predict("data/titanic/titanic.csv")
print(predictions)
```

## 10.4 API reference

### 10.4.1 TabularClassifier

```
class flash.tabular.TabularClassifier(num_features, num_classes, embedding_sizes=None, hidden=512,
loss_fn=torch.nn.functional.cross_entropy, optimizer=torch.optim.Adam, metrics=None, learning_rate=0.001)
```

Task that classifies table rows.

#### Parameters

- **num\_features** `(int)` – Number of columns in table (not including target column).
- **num\_classes** `(int)` – Number of classes to classify.

- **embedding\_sizes** (Optional[List[Tuple]]) – List of (num\_classes, emb\_dim) to form categorical embeddings.
- **hidden** (Union[List[int], int]) – Hidden dimension sizes.
- **loss\_fn** (Callable) – Loss function for training, defaults to cross entropy.
- **optimizer** (Type[Optimizer]) – Optimizer to use for training, defaults to *torch.optim.Adam*.
- **metrics** (Optional[List[Metric]]) – Metrics to compute for training and evaluation.
- **learning\_rate** (float) – Learning rate to use for training, defaults to *1e-3*

**static default\_pipeline()**

Pipeline to use when there is no datamodule or it has not defined its pipeline

**Return type** DataPipeline

## 10.4.2 TabularData

```
class flash.tabular.TabularData(train_df, categorical_input, numerical_input, target,
                                valid_df=None, test_df=None, batch_size=2,
                                num_workers=None)
```

Data module for tabular tasks

```
classmethod TabularData.from_csv(train_csv, target, categorical_input, numerical_input,
                                  valid_csv=None, test_csv=None, batch_size=8,
                                  num_workers=None, val_size=None, test_size=None,
                                  **pandas_kwargs)
```

Creates a TextClassificationData object from pandas DataFrames.

### Parameters

- **train\_csv** – train data csv file.
- **target** (str) – The column containing the class id.
- **categorical\_input** (List) – The list of categorical columns.
- **numerical\_input** (List) – The list of numerical columns.
- **valid\_csv** (Optional[str]) – validation data csv file.
- **test\_csv** (Optional[str]) – test data csv file.
- **batch\_size** (int) – the batchsize to use for parallel loading. Defaults to 64.
- **num\_workers** (Optional[int]) – The number of workers to use for parallelized loading. Defaults to None which equals the number of available CPU threads.
- **val\_size** (Optional[float]) – float between 0 and 1 to create a validation dataset from train dataset
- **test\_size** (Optional[float]) – float between 0 and 1 to create a test dataset from train validation

**Returns** The constructed data module.

**Return type** *TabularData*

Examples:

```
text_data = TabularData.from_files("train.csv", label_field="class", text_field=
↪ "sentence")
```

**classmethod** `TabularData.from_df`(*train\_df*, *target*, *categorical\_input*, *numerical\_input*,  
*valid\_df=None*, *test\_df=None*, *batch\_size=8*,  
*num\_workers=None*, *val\_size=None*, *test\_size=None*)

Creates a TabularData object from pandas DataFrames.

#### Parameters

- **train\_df** (DataFrame) – train data DataFrame
- **target** (str) – The column containing the class id.
- **categorical\_input** (List) – The list of categorical columns.
- **numerical\_input** (List) – The list of numerical columns.
- **valid\_df** (Optional[DataFrame]) – validation data DataFrame
- **test\_df** (Optional[DataFrame]) – test data DataFrame
- **batch\_size** (int) – the batchsize to use for parallel loading. Defaults to 64.
- **num\_workers** (Optional[int]) – The number of workers to use for parallelized loading. Defaults to None which equals the number of available CPU threads.
- **val\_size** (Optional[float]) – float between 0 and 1 to create a validation dataset from train dataset
- **test\_size** (Optional[float]) – float between 0 and 1 to create a test dataset from train validation

**Returns** The constructed data module.

**Return type** *TabularData*

Examples:

```
text_data = TextClassificationData.from_files("train.csv", label_field="class", ↪
↪ text_field="sentence")
```





## TRANSLATION

### 11.1 The task

Translation is the task of translating text from a source language to another, such as English to Romanian. This task is a subset of Sequence to Sequence tasks, which requires the model to generate a variable length sequence given an input sequence. In our case the English text would be our input sequence, and the Romanian sentence would be the output sequence from the model.

### 11.2 Inference

The TranslationTask is already pre-trained on [WMT16 English/Romanian](<https://www.statmt.org/wmt16/translation-task.html>), a dataset of English to Romanian samples, based on the Europarl corpora.

Use the TranslationTask pretrained model for inference on any string sequence using `predict()`:

```
# import our libraries
from flash.text import TranslationTask

# 2. Load the model from a checkpoint
model = TranslationTask.load_from_checkpoint("https://flash-weights.s3.amazonaws.com/
↳translation_model_en_ro.pt")

# 2. Perform inference from list of sequences
predictions = model.predict([
    "BBC News went to meet one of the project's first graduates.",
    "A recession has come as quickly as 11 months after the first rate hike and as_
↳long as 86 months.",
])
print(predictions)
```

Or on a given dataset:

```
# import our libraries
from pytorch_lightning import Trainer
from flash import download_data
from flash.text import TranslationData, TranslationTask

# 2. Load the model from a checkpoint
model = TranslationTask.load_from_checkpoint("https://flash-weights.s3.amazonaws.com/
↳translation_model_en_ro.pt")
```

(continues on next page)

(continued from previous page)

```
# 3. Create dataset from file
datamodule = TranslationData.from_file(
    predict_file="data/wmt_en_ro/predict.csv",
    input="input",
)

# 4. generate translations
predictions = Trainer().predict(model, datamodule=datamodule)
print(predictions)
```

For more advanced inference options, see *Predictions (inference)*.

## 11.3 Finetuning

Say you want to finetune to your own translation data. We use the English/Romanian WMT16 dataset as an example which contains a `train.csv` and `valid.csv`, structured like so:

```
input,target
"Written statements and oral questions (tabling): see Minutes", "Declarații scrise și
↳ întrebări orale (depunere): consultați procesul-verbal"
"Closure of sitting", "Ridicarea ședinței"
...
```

In the above the input/target columns represent the English and Romanian translation respectively.

All we need is three lines of code to train our model!

```
# import our libraries
import flash
from flash import download_data
from flash.text import TranslationData, TranslationTask

# 1. Download data
download_data("https://pl-flash-data.s3.amazonaws.com/wmt_en_ro.zip", 'data/')

# Organize the data
datamodule = TranslationData.from_files(
    train_file="data/wmt_en_ro/train.csv",
    valid_file="data/wmt_en_ro/valid.csv",
    test_file="data/wmt_en_ro/test.csv",
    input="input",
    target="target",
)

# 2. Build the task
model = TranslationTask()

# 4. Create trainer
trainer = flash.Trainer(max_epochs=5, gpus=1, precision=16)

# 5. Finetune the task
trainer.finetune(model, datamodule=datamodule)
```

(continues on next page)

(continued from previous page)

```
# 6. Save trainer task
trainer.save_checkpoint("translation_model_en_ro.pt")
```

To run the example:

```
python flash_examples/finetuning/translation.py
```

## 11.4 Changing the backbone

By default, we use [mBART](#) model for translation. You can change the model run by passing in the backbone parameter.

**Note:** When changing the backbone, make sure you pass in the same backbone to the Task and the Data object! Since this is a Seq2Seq task, make sure you use a Seq2Seq model.

```
datamodule = TranslationData.from_files(
    train_file="data/wmt_en_ro/train.csv",
    valid_file="data/wmt_en_ro/valid.csv",
    test_file="data/wmt_en_ro/test.csv",
    input="input",
    target="target",
    backbone="t5-small",
)

model = TranslationTask(backbone="t5-small")
```

## 11.5 API reference

### 11.5.1 TranslationTask

```
class flash.text.seq2seq.translation.model.TranslationTask (backbone='facebook/mbart-large-en-ro',
    loss_fn=None, optimizer=torch.optim.Adam, metrics=None, learning_rate=0.0003, val_target_max_length=128, num_beams=4, n_gram=4, smooth=False)
```

Task for Sequence2Sequence Translation. :type\_sphinx\_paramlinks\_flash.text.seq2seq.translation.model.TranslationTask.backbone: str :param\_sphinx\_paramlinks\_flash.text.seq2seq.translation.model.TranslationTask.backbone: backbone model to use for the task. :type\_sphinx\_paramlinks\_flash.text.seq2seq.translation.model.TranslationTask.loss\_fn: Union[Callable, Mapping, Sequence, None] :param\_sphinx\_paramlinks\_flash.text.seq2seq.translation.model.TranslationTask.loss\_fn:

Loss function for training. :type `_sphinx_paramlinks_flash.text.seq2seq.translation.model.TranslationTask.optimizer: Type[Optimizer]` :param `_sphinx_paramlinks_flash.text.seq2seq.translation.model.TranslationTask.optimizer:` Optimizer to use for training, defaults to `torch.optim.Adam`. :type `_sphinx_paramlinks_flash.text.seq2seq.translation.model.TranslationTask.metrics:` `Union[Metric, Mapping, Sequence, None]` :param `_sphinx_paramlinks_flash.text.seq2seq.translation.model.TranslationTask.metrics:` Metrics to compute for training and evaluation. :type `_sphinx_paramlinks_flash.text.seq2seq.translation.model.TranslationTask.learning_rate:` `float` :param `_sphinx_paramlinks_flash.text.seq2seq.translation.model.TranslationTask.learning_rate:` Learning rate to use for training, defaults to `3e-4` :type `_sphinx_paramlinks_flash.text.seq2seq.translation.model.TranslationTask.val_target_max_length:` `Optional[int]` :param `_sphinx_paramlinks_flash.text.seq2seq.translation.model.TranslationTask.val_target_max_length:` Maximum length of targets in validation. Defaults to `128` :type `_sphinx_paramlinks_flash.text.seq2seq.translation.model.TranslationTask.num_beams:` `Optional[int]` :param `_sphinx_paramlinks_flash.text.seq2seq.translation.model.TranslationTask.num_beams:` Number of beams to use in validation when generating predictions. Defaults to `4` :type `_sphinx_paramlinks_flash.text.seq2seq.translation.model.TranslationTask.n_gram:` `bool` :param `_sphinx_paramlinks_flash.text.seq2seq.translation.model.TranslationTask.n_gram:` Maximum n\_grams to use in metric calculation. Defaults to `4` :type `_sphinx_paramlinks_flash.text.seq2seq.translation.model.TranslationTask.smooth:` `bool` :param `_sphinx_paramlinks_flash.text.seq2seq.translation.model.TranslationTask.smooth:` Apply smoothing in BLEU calculation. Defaults to `True`

#### property task

Override to define AutoConfig task specific parameters stored within the model.

**Return type** `str`

## 11.5.2 TranslationData

```
class flash.text.seq2seq.translation.data.TranslationData (train_ds=None,
                                                         valid_ds=None,
                                                         test_ds=None,
                                                         batch_size=1,
                                                         num_workers=None)
```

Data module for Translation tasks.

```
classmethod TranslationData.from_files (train_file, input='input', target=None, file-
                                         type='csv', backbone='facebook/mbart-large-
                                         en-ro', valid_file=None, test_file=None,
                                         max_source_length=128, max_target_length=128,
                                         padding='max_length', batch_size=8,
                                         num_workers=None)
```

Creates a TranslationData object from files.

#### Parameters

- **train\_file** – Path to training data.
- **input** (`str`) – The field storing the source translation text.
- **target** (`Optional[str]`) – The field storing the target translation text.
- **filetype** – .csv or .json
- **backbone** – tokenizer to use, can use any HuggingFace tokenizer.
- **valid\_file** – Path to validation data.
- **test\_file** – Path to test data.
- **max\_source\_length** (`int`) – Maximum length of the source text. Any text longer will be truncated.

- **max\_target\_length** (int) – Maximum length of the target text. Any text longer will be truncated.
- **padding** (Union[str, bool]) – Padding strategy for batches. Default is pad to maximum length.
- **batch\_size** (int) – the batchsize to use for parallel loading. Defaults to 8.
- **num\_workers** (Optional[int]) – The number of workers to use for parallelized loading. Defaults to None which equals the number of available CPU threads.

**Returns** The constructed data module.

**Return type** TranslateData

Examples:

```
train_df = pd.read_csv("train_data.csv")
tab_data = TabularData.from_df(train_df, target="fraud",
                                numerical_input=["account_value"],
                                categorical_input=["account_type"])
```



## MODEL

```
class flash.core.model.Task(model=None, loss_fn=None, optimizer=torch.optim.Adam, metrics=None, learning_rate=5e-05)
```

A general Task.

### Parameters

- **model** `Optional[Module]` – Model to use for the task.
- **loss\_fn** `Union[Callable, Mapping, Sequence, None]` – Loss function for training
- **optimizer** `Type[Optimizer]` – Optimizer to use for training, defaults to `torch.optim.SGD`.
- **metrics** `Union[Metric, Mapping, Sequence, None]` – Metrics to compute for training and evaluation.
- **learning\_rate** `float` – Learning rate to use for training, defaults to `5e-5`

```
static default_pipeline()
```

Pipeline to use when there is no datamodule or it has not defined its pipeline

**Return type** `DataPipeline`

```
step(batch, batch_idx)
```

The training/validation/test step. Override for custom behavior.

**Return type** `Any`





## 13.1 DataPipeline

To make tasks work for inference, one must create a `DataPipeline`. The `flash.core.data.DataPipeline` exposes 6 hooks to override:

```
class DataPipeline:
    """
        This class purpose is to facilitate the conversion of raw data to processed or_
        ↪ batched data and back.
        Several hooks are provided for maximum flexibility.

        collate_fn:
            - before_collate
            - collate
            - after_collate

        uncollate_fn:
            - before_uncollate
            - uncollate
            - after_uncollate
    """

    def before_collate(self, samples: Any) -> Any:
        """Override to apply transformations to samples"""
        return samples

    def collate(self, samples: Any) -> Any:
        """Override to convert a set of samples to a batch"""
        if not isinstance(samples, Tensor):
            return default_collate(samples)
        return samples

    def after_collate(self, batch: Any) -> Any:
        """Override to apply transformations to the batch"""
        return batch

    def before_uncollate(self, batch: Any) -> Any:
        """Override to apply transformations to the batch"""
        return batch

    def uncollate(self, batch: Any) -> Any:
        """Override to convert a batch to a set of samples"""
        samples = batch
```

(continues on next page)

(continued from previous page)

```
return samples

def after_uncollate(self, samples: Any) -> Any:
    """Override to apply transformations to samples"""
    return samplesA
```

Use these utilities to download data.

---

## 13.2 download\_data

`flash.core.data.utils.download_data(url, path='data/')`

Downloads data automatically from the given url to the path. Defaults to data/ for the path. Automatically handles .csv, .zip

Example:

```
from flash import download_data
```

### Parameters

- `url` *(str)* – path
- `path` *(str)* – local

**Return type** `None`

## TRAINING FROM SCRATCH

Some Flash tasks have been pretrained on large data sets. To accelerate your training, calling the `finetune()` method using a pretrained backbone will fine-tune the backbone to generate a model customized to your data set and desired task. If you want to train the task from scratch instead, pass `pretrained=False` parameter when creating your task. Then, use the `fit()` method to train your model.

```
import flash
from flash import download_data
from flash.vision import ImageClassificationData, ImageClassifier

# 1. download and organize the data
download_data("https://download.pytorch.org/tutorial/hymenoptera_data.zip", 'data/')

data = ImageClassificationData.from_folders(
    train_folder="data/hymenoptera_data/train/",
    valid_folder="data/hymenoptera_data/val/"
)

# 2. build the task, and turn off pre-training
task = ImageClassifier(num_classes=2, pretrained=False)

# 3. train!
trainer = flash.Trainer()
trainer.fit(model, data)
```

### 14.1 Training options

Flash tasks supports many advanced training functionalities out-of-the-box, such as:

- limit number of epochs

```
# train for 10 epochs
trainer.fit(max_epochs=10)
```

- Training on GPUs

```
# train on 1 GPU
trainer.fit(gpus=1)
```

- Training on multiple GPUs

```
# train on multiple GPUs
trainer.fit(gpus=4)
```

```
# train on gpu 1, 3, 5 (3 gpus total)
trainer.fit(gpus=[1, 3, 5])
```

- Using mixed precision training

```
# Multi GPU with mixed precision
trainer.fit(gpus=2, precision=16)
```

- Training on TPUs

```
# Train on TPUs
trainer.fit(tpu_cores=8)
```

You can add to the flash Trainer any argument from the Lightning trainer! Learn more about the Lightning Trainer [here](#).

### 14.1.1 Trainer API

```
class flash.core.trainer.Trainer(*args, **kwargs)
```

```
finetune(model, train_dataloader=None, val_dataloaders=None, datamodule=None, strategy=None)
```

Runs the full optimization routine. Same as `pytorch_lightning.Trainer().fit()`, but unfreezes layers of the backbone throughout training layers of the backbone throughout training.

#### Parameters

- **datamodule** *⚡* (`Optional[LightningDataModule]`) – A instance of `LightningDataModule`.
- **model** *⚡* (`LightningModule`) – Model to fit.
- **train\_dataloader** *⚡* (`Optional[DataLoader]`) – A Pytorch `DataLoader` with training samples. If the model has a predefined `train_dataloader` method this will be skipped.
- **val\_dataloaders** *⚡* (`Union[DataLoader, List[DataLoader], None]`) – Either a single Pytorch `Dataloader` or a list of them, specifying validation samples. If the model has a predefined `val_dataloaders` method this will be skipped
- **strategy** *⚡* (`Union[str, BaseFinetuning, None]`) – Should either be a string or a finetuning callback subclassing `pytorch_lightning.callbacks.BaseFinetuning`.

Currently, default strategies can be enabled with these strings:

- `no_freeze`,
- `freeze`,
- `freeze_unfreeze`,
- `unfreeze_milestones`

```
fit(model, train_dataloader=None, val_dataloaders=None, datamodule=None)
```

Runs the full optimization routine. Same as `pytorch_lightning.Trainer().fit()`

#### Parameters

- **datamodule** (Optional[LightningDataModule]) – A instance of LightningDataModule.
- **model** (LightningModule) – Model to fit.
- **train\_dataloader** (Optional[DataLoader]) – A Pytorch DataLoader with training samples. If the model has a predefined train\_dataloader method this will be skipped.
- **val\_dataloaders** (Union[DataLoader, List[DataLoader], None]) – Either a single Pytorch Dataloader or a list of them, specifying validation samples. If the model has a predefined val\_dataloaders method this will be skipped



## FINETUNING

Finetuning (or transfer-learning) is the process of tweaking a model trained on a large dataset, to your particular (likely much smaller) dataset. All Flash tasks have a pre-trained backbone that was already trained on large datasets such as ImageNet. Finetuning on already pretrained models decrease training time significantly.

You can finetune any Flash task on your own data in just a 3 simple steps:

1. Load your data and organize it using Flash Data Modules. Note that different tasks have different data modules (The *ImageClassificationData* for image classification, *TextClassificationData* for text classification, etc.).
2. Pick a model to run from a variety of Flash tasks: *ImageClassifier*, *TextClassifier*, *TabularClassifier*, all optimized with the latest best practices.
3. Finetune your model using *finetune()* method. You will need to choose a finetune strategy.

Once training is completed, you can use the model for inference to make predictions using the *predict* method.

```
import flash
from flash import download_data
from flash.vision import ImageClassificationData, ImageClassifier

# 1. download and organize the data
download_data("https://download.pytorch.org/tutorial/hymenoptera_data.zip", 'data/')

data = ImageClassificationData.from_folders(
    train_folder="data/hymenoptera_data/train/",
    valid_folder="data/hymenoptera_data/val/"
)

# 2. build the task
task = ImageClassifier(num_classes=2)

# 3. Build the trainer and finetune! In this case, using the no_freeze strategy
trainer = flash.Trainer()
trainer.finetune(model, data, strategy="no_freeze")
```

---

**Tip:** If you have a large dataset and prefer to train from scratch, see the *Training from scratch* guide.

---

## 15.1 Finetune strategies

The flash tasks contain pre-trained models trained on large datasets such as [ImageNet](#), which contains millions of images. These models are called **backbones**. This will be used as the starting point for finetuning.

The model needs to be adapted or refined for the new data available for the task. Usually, the last layers of the backbone need to be modified, to match the backbone output to the number of target classes of the new data. These layers are commonly referred to as the **head**. For example, our backbone might be trained to classify 10 types of animals, but maybe our new dataset only contains images of bees and ants, so we would have to modify our final layer to fit just 2 classes. The head is randomly initialized whereas the backbone conserves its pre-trained weights.

The `finetune()` method trains the new modified model using the new dataset. As the head (new layers) is untrained, the first results (gradients) will be random when training starts and could decrease the backbone performance (by changing its pre-trained parameters). Therefore, it is a good practice to “freeze” the backbone, meaning the parameters of the backbone won’t be updated until they are “unfrozen” a few epochs later.

You can choose a finetuning strategy using `finetune()` *strategy* parameter. Flash finetune *strategy* argument can either a string or an instance of `FlashBaseFinetuning`.

Flash supports 2 built-in Finetuning strategies, that can be passed as strings:

- *no\_freeze*: Don’t freeze anything, the backbone parameters can be modified during finetuning.
- *freeze*: The parameters of the backbone won’t be modified during finetuning.

```
# using the freeze strategy
trainer.finetune(model, data, strategy="freeze")

# using the no_freeze strategy
trainer.finetune(model, data, strategy="no_freeze")
```

For more options, you can pass in an instance of `BaseFinetuning` to the *strategy* parameter.

### 15.1.1 Custom callback finetuning

For more advanced finetuning, you can use flash built-in finetuning callbacks.

- *FreezeUnfreeze*: The backbone parameters will be frozen for a given number of epochs (by default the *unfreeze\_epoch* is set to 10).

```
# import FreezeUnfreeze
from flash.core.finetuning import FreezeUnfreeze

# finetune for 10 epochs. Backbone will be frozen for 5 epochs.
trainer = flash.Trainer(max_epochs=10)
trainer.finetune(model, data, strategy=FreezeUnfreeze(unfreeze_epoch=5))
```

- *UnfreezeMilestones*: This strategy define 2 milestones, one milestone (epoch number) to unfreeze the last layers of the backbone, and a second milestone to unfreeze the remaining layers. For example, by default the first milestone is 5 and the second is 10. So for the first 4 epochs, the backbone parameters will be frozen. In epochs 5-9, only the last layers (5 by default) can be trained. After the 10th epoch, all parameters in all layers can be trained.

```
# import UnfreezeMilestones
from flash.core.finetuning import UnfreezeMilestones

# finetune for 10 epochs. Backbone will be frozen for 3 epochs. The last 2 layers_
will be unfrozen for the first 4 epochs,
```

(continues on next page)



(continued from previous page)

```
# and then the rest will be unfrozen on the 8th epoch
trainer = flash.Trainer(max_epochs=10)
trainer.finetune(model, data, strategy=UnfreezeMilestones(unfreeze_milestones=(5, 8),
↳ num_layers=2))
```

## 15.2 Custom callback finetuning

For even more customization, create your own finetuning callback. Learn more about callbacks [here](#).

```
from flash.core.finetuning import FlashBaseFinetuning

# Create a finetuning callback
class FeatureExtractorFreezeUnfreeze(FlashBaseFinetuning):

    def __init__(self, unfreeze_at_epoch: int = 5, train_bn: bool = True)
        # this will set self.attr_names as ["feature_extractor"]
        super().__init__("feature_extractor", train_bn)
        self._unfreeze_at_epoch = unfreeze_at_epoch

    def finetune_function(self, pl_module, current_epoch, optimizer, opt_idx):
        # unfreeze any module you want by overriding this function

        # When ``current_epoch`` is 5, feature_extractor will start to be trained.
        if current_epoch == self._unfreeze_at_epoch:
            self.unfreeze_and_add_param_group(
                module=pl_module.feature_extractor,
                optimizer=optimizer,
                train_bn=True,
            )

# Init the trainer
trainer = flash.Trainer(max_epochs=10)

# pass the callback to trainer.finetune
trainer.finetune(model, data, strategy=FeatureExtractorFreezeUnfreeze(unfreeze_
↳ epoch=5))
```



## PREDICTIONS (INFERENCE)

You can use Flash to get predictions on pretrained or finetuned models.

### 16.1 Predict on a single sample of data

You can pass in a sample of data (image file path, a string of text, etc) to the `predict()` method.

```
from flash import Trainer
from flash.core.data import download_data
from flash.vision import ImageClassificationData, ImageClassifier

# 1. Download the data set
download_data("https://pl-flash-data.s3.amazonaws.com/hymenoptera_data.zip", 'data/')

# 2. Load the model from a checkpoint
model = ImageClassifier.load_from_checkpoint("https://flash-weights.s3.amazonaws.com/
↪image_classification_model.pt")

# 3. Predict whether the image contains an ant or a bee
predictions = model.predict("data/hymenoptera_data/val/bees/65038344_52a45d090d.jpg")
print(predictions)
```

### 16.2 Predict on a csv file

```
from flash.core.data import download_data
from flash.tabular import TabularClassifier

# 1. Download the data
download_data("https://pl-flash-data.s3.amazonaws.com/titanic.zip", 'data/')

# 2. Load the model from a checkpoint
model = TabularClassifier.load_from_checkpoint(
    "https://flash-weights.s3.amazonaws.com/tabular_classification_model.pt"
)

# 3. Generate predictions from a csv file! Who would survive?
predictions = model.predict("data/titanic/titanic.csv")
print(predictions)
```



## INDICES AND TABLES

- `genindex`
- `modindex`
- `search`



## INDEX

### D

`default_pipeline()` (*flash.core.model.Task* static method), 51  
`default_pipeline()` (*flash.tabular.TabularClassifier* static method), 42  
`default_pipeline()` (*flash.text.classification.model.TextClassifier* static method), 38  
`default_pipeline()` (*flash.vision.ImageClassifier* static method), 22  
`default_pipeline()` (*flash.vision.ImageEmbedder* static method), 27  
`download_data()` (in module *flash.core.data.utils*), 54

### F

`finetune()` (*flash.core.trainer.Trainer* method), 56  
`fit()` (*flash.core.trainer.Trainer* method), 56  
`from_csv()` (*flash.tabular.TabularData* class method), 42  
`from_df()` (*flash.tabular.TabularData* class method), 43  
`from_filepaths()` (*flash.vision.ImageClassificationData* class method), 22  
`from_files()` (*flash.text.classification.data.TextClassificationData* class method), 38  
`from_files()` (*flash.text.seq2seq.summarization.data.SummarizationData* class method), 32  
`from_files()` (*flash.text.seq2seq.translation.data.TranslationData* class method), 48  
`from_folders()` (*flash.vision.ImageClassificationData* class method), 23

### I

*ImageClassificationData* (class in *flash.vision*), 22  
*ImageClassifier* (class in *flash.vision*), 21  
*ImageEmbedder* (class in *flash.vision*), 27

### S

`step()` (*flash.core.model.Task* method), 51

`step()` (*flash.text.classification.model.TextClassifier* method), 38  
*SummarizationData* (class in *flash.text.seq2seq.summarization.data*), 32  
*SummarizationTask* (class in *flash.text.seq2seq.summarization.model*), 32

### T

*TabularClassifier* (class in *flash.tabular*), 41  
*TabularData* (class in *flash.tabular*), 42  
*Task* (class in *flash.core.model*), 51  
`task()` (*flash.text.seq2seq.summarization.model.SummarizationTask* property), 32  
`task()` (*flash.text.seq2seq.translation.model.TranslationTask* property), 48  
*TextClassificationData* (class in *flash.text.classification.data*), 38  
*TextClassifier* (class in *flash.text.classification.model*), 37  
*Trainer* (class in *flash.core.trainer*), 56  
*TranslationData* (class in *flash.text.seq2seq.translation.data*), 48  
*TranslationTask* (class in *flash.text.seq2seq.translation.model*), 47