pytorch-igniter-demo

Release 0.1.0

Oct 12, 2020

Contents:

CHAPTER 1

pytorch-igniter-demo

Demo for pytorch-igniter

- Write functions for creating a model, evaluation steps and training steps
- `pytorch-igniter` generates a CLI for training, evaluation and inference
- Logging, checkpointing, iterating, and other features handled automatically
- Training generates a self-contained model package that can run inference

1.1 Installation

```
pip install git+https://github.com/bstriner/pytorch-igniter-demo.git
```

1.2 Documentation

View latest documentation at ReadTheDocs

1.3 GitHub

View source code on GitHub

1.4 Training and Inference

Training generate self-contained models:

- Local training generates a directory

- SageMaker training automatically gzips that directory and uploads to S3

A training package can be run locally using sagemaker libraries or uploaded to SageMaker to run remotely.

- Model data like the neural network are stored in the gz
- Directory `code` within the gz
  - Placed on PYTHONPATH
  - If it contains `requirements.txt` it is installed
  - Module `inference` defines inference
  - Depenencies are automatically uploaded
  - Script copies its own source to `code` when saving a model

The inference module defines how inference happens

- `model_fn` loads your model when the endpoint starts
- `input_fn` reads WAV, JPG, etc content from an HTTP request based on `Content-Type` headers
- `predict_fn` runs your model
- `output_fn` maps your predictions to HTTP responses based on `Accept` headers

The model directory and inference module are created automatically based on an `InferenceSpec` that you define

Locally or remotely trained models can be deployed locally or remotely.

Use models by posting a file in an HTTP request with appropriate `Content-Type` and `Accept` headers.

1.4.1 Local usage

```
# Local training
pytorch-igniter-demo dataprep --output output/data
pytorch-igniter-demo train-and-eval --input output/data

# Invoke local model
# Contents of directory is same as model.gz used for remote invocation
aws-sagemaker-remote endpoint invoke --model-dir output/model --input test/test-image.
↪png --output-type application/json --output output/invoke-local.json

# Upload local model directory as a gzip
# If already gzipped, skip -gz flag
aws-sagemaker-remote upload output/model pytorch-igniter-demo/model.tar.gz --gz
# Register SageMaker model from artifact
aws-sagemaker-remote model create --name pytorch-igniter-demo-local --model-artifact␣
↪pytorch-igniter-demo/model.tar.gz --force
# Create endpoint configuration
aws-sagemaker-remote endpoint-config create --model pytorch-igniter-demo-local --force
# Create endpoint
# This launches servers, takes a while, and begins ongoing costs
aws-sagemaker-remote endpoint create --config pytorch-igniter-demo-local --force
# Wait for the launch
aws sagemaker wait endpoint-in-service --endpoint-name pytorch-igniter-demo-local

# Invoke remote model
# Prefer using boto3 or other libraries to invoke endpoints directly in your own code
aws-sagemaker-remote endpoint invoke --name pytorch-igniter-demo-local --input test/
↪test-image.png --output output/invoke-upload.json --output-type application/json
```

(continues on next page)

(continued from previous page)

```
# Clean up resources
# Only the endpoint itself incurs any significant charges
aws-sagemaker-remote endpoint delete pytorch-igniter-demo-local
aws-sagemaker-remote endpoint-config delete pytorch-igniter-demo-local
aws-sagemaker-remote model delete pytorch-igniter-demo-local
```

1.4.2 Remote usage

```
# Dataprep
pytorch-igniter-demo dataprep --sagemaker-run yes

# Training
pytorch-igniter-demo train-and-eval --sagemaker-run yes --input s3://sagemaker-us-
↪east-1-683880991063/pytorch-igniter-demo-dataprep-2020-10-09-01-20-47-571/output/
↪output

# Deploy model
aws-sagemaker-remote model create --name pytorch-igniter-demo-remote --job training-
↪job-2020-10-12-08-06-48-401
aws-sagemaker-remote endpoint-config create --model pytorch-igniter-demo-remote --
↪force
aws-sagemaker-remote endpoint create --config pytorch-igniter-demo-remote --force
aws sagemaker wait endpoint-in-service --endpoint-name pytorch-igniter-demo-remote

# Invoke remote model
aws-sagemaker-remote endpoint invoke --name pytorch-igniter-demo-remote --input test/
↪test-image.png --output output/invoke-upload.json --output-type application/json

# Clean up resources
aws-sagemaker-remote endpoint delete pytorch-igniter-demo-remote
aws-sagemaker-remote endpoint-config delete pytorch-igniter-demo-remote
aws-sagemaker-remote model delete pytorch-igniter-demo-remote
```

CHAPTER 2

CLI

Installing the package will install the command-line script `pytorch-igniter-demo`

2.1 Features

- Run training, dataprep, evaluation and inference locally or on AWS servers
- Model package is self-contained with all dependencies and configuration
- Integrate with pytorch-igniter for creating an experiment CLI and managing training.
- Integrate with MLflow for tracking training runs, including hyperparameters and metrics.
- Integrate with AWS SageMaker using aws-sagemaker-remote for tracking training runs and executing training remotely on managed containers.

2.2 Command-Line Arguments

Set of arguments and defaults is configured through code. See `pytorch-igniter` documentation.

pytorch-igniter demo script

```
usage: pytorch-igniter-demo [-h] {train,eval,train-and-eval,dataprep} ...
```

2.2.1 command

command Possible choices: train, eval, train-and-eval, dataprep

Command to execute

2.2.2 Sub-commands:

train

Train a model

```
pytorch-igniter-demo train [-h] [--sagemaker-profile SAGEMAKER_PROFILE]
                           [--sagemaker-run [SAGEMAKER_RUN]]
                           [--sagemaker-wait [SAGEMAKER_WAIT]]
                           [--sagemaker-spot-instances [SAGEMAKER_SPOT_INSTANCES]]
                           [--sagemaker-script SAGEMAKER_SCRIPT]
                           [--sagemaker-source SAGEMAKER_SOURCE]
                           [--sagemaker-training-instance SAGEMAKER_TRAINING_INSTANCE]
                           [--sagemaker-training-image SAGEMAKER_TRAINING_IMAGE]
                           [--sagemaker-training-role SAGEMAKER_TRAINING_ROLE]
                           [--sagemaker-base-job-name SAGEMAKER_BASE_JOB_NAME]
                           [--sagemaker-job-name SAGEMAKER_JOB_NAME]
                           [--sagemaker-experiment-name SAGEMAKER_EXPERIMENT_NAME]
                           [--sagemaker-trial-name SAGEMAKER_TRIAL_NAME]
                           [--sagemaker-volume-size SAGEMAKER_VOLUME_SIZE]
                           [--sagemaker-max-run SAGEMAKER_MAX_RUN]
                           [--sagemaker-max-wait SAGEMAKER_MAX_WAIT]
                           [--pytorch-igniter-demo PYTORCH_IGNITER_DEMO]
                           [--model-dir MODEL_DIR] [--output-dir OUTPUT_DIR]
                           [--checkpoint-dir CHECKPOINT_DIR]
                           [--sagemaker-checkpoint-s3 SAGEMAKER_CHECKPOINT_S3]
                           [--sagemaker-checkpoint-container SAGEMAKER_CHECKPOINT_
↪CONTAINER]
                           [--input INPUT] [--device DEVICE]
                           [--classes CLASSES] [--max-epochs N]
                           [--n-saved N_SAVED] [--learning-rate LEARNING_RATE]
                           [--weight-decay WEIGHT_DECAY]
                           [--train-batch-size TRAIN_BATCH_SIZE]
                           [--mlflow-enable [MLFLOW_ENABLE]]
                           [--mlflow-experiment-name MLFLOW_EXPERIMENT_NAME]
                           [--mlflow-run-name MLFLOW_RUN_NAME]
                           [--mlflow-tracking-uri MLFLOW_TRACKING_URI]
                           [--mlflow-tracking-username MLFLOW_TRACKING_USERNAME]
                           [--mlflow-tracking-password MLFLOW_TRACKING_PASSWORD]
                           [--mlflow-tracking-secret-name MLFLOW_TRACKING_SECRET_NAME]
                           [--mlflow-tracking-secret-profile MLFLOW_TRACKING_SECRET_
↪PROFILE]
                           [--mlflow-tracking-secret-region MLFLOW_TRACKING_SECRET_
↪REGION]
```

Named Arguments

--model-dir Directory to save final model (default: output/model)

Default: “output/model”

--output-dir Directory for logs, images, or other output files (default: “output/output”)

Default: “output/output”

SageMaker

SageMaker options

--sagemaker-profile AWS profile for SageMaker session (default: [default])

Default: "default"

--sagemaker-run Run training on SageMaker (yes/no default=False)

Default: False

--sagemaker-wait Wait for SageMaker training to complete and tail logs files (yes/no default=True)

Default: True

--sagemaker-spot-instances Use spot instances for training (yes/no default=False)

Default: False

--sagemaker-script Script to run on SageMaker. (default: [/home/docs/checkouts/readthedocs.org/user_builds/pytorch-igniter-demo/checkouts/latest/pytorch_igniter_demo/main.py])

Default: "/home/docs/checkouts/readthedocs.org/user_builds/pytorch-igniter-demo/checkouts/latest/pytorch_igniter_demo/main.py"

--sagemaker-source Source to upload to SageMaker. Must contain script. If blank, default to directory containing script. (default: [])

Default: ""

--sagemaker-training-instance Instance type for training

Default: "ml.m5.large"

--sagemaker-training-image Docker image for training

Default: "683880991063.dkr.ecr.us-east-1.amazonaws.com/columbo-sagemaker-training:latest"

--sagemaker-training-role Docker image for training

Default: "aws-sagemaker-remote-training-role"

--sagemaker-base-job-name Base job name for tracking and organization on S3. A job name will be generated from the base job name unless a job name is specified.

Default: "training-job"

--sagemaker-job-name Job name for tracking. Use –base-job-name instead and a job name will be automatically generated with a timestamp.

Default: ""

--sagemaker-experiment-name Name of experiment in SageMaker tracking.

--sagemaker-trial-name Name of experiment trial in SageMaker tracking.

--sagemaker-volume-size Volume size in GB.

Default: 30

--sagemaker-max-run Maximum runtime in seconds.

Default: 43200

--sagemaker-max-wait Maximum time to wait for spot instances in seconds.

Default: 86400

Dependencies

Dependencies to upload to SageMaker

--pytorch-igniter-demo Directory for dependency [pytorch_igniter_demo] (default: "/home/docs/checkouts/readthedocs.org/user_builds/pytorch-igniter-demo/checkouts/latest/pytorch_igniter_demo")

Default: "pytorch_igniter_demo"

Checkpoints

Checkpointing options

--checkpoint-dir Local directory to store checkpoints for resuming training (default: "output/checkpoint")

Default: "output/checkpoint"

--sagemaker-checkpoint-s3 Location to store checkpoints on S3 or "default" (default: "default")

Default: "default"

--sagemaker-checkpoint-container Location to store checkpoints on container (default: "/opt/ml/checkpoints")

Default: "/opt/ml/checkpoints"

Inputs

Inputs (local or S3)

--input Input channel [input]. Set to local path and it will be uploaded to S3 and downloaded to SageMaker. Set to S3 path and it will be downloaded to SageMaker. (default: [output/data])

Default: "output/data"

Model

Model arguments

--device device to use (default: None)

--classes Default: 10

Training

Training arguments

--max-epochs number of epochs to train (default: 10)

Default: 10

--n-saved Number of checkpoints to keep (default: `10`)

Default: 10

--learning-rate Default: 0.001

--weight-decay Default: 1e-05

--train-batch-size Default: 32

MLflow

MLflow arguments

--mlflow-enable Enable logging to MLflow (default: True)

Default: True

--mlflow-experiment-name Experiment name in MLflow (default: default)

Default: "default"

--mlflow-run-name Run name in MLflow (default: None)

--mlflow-tracking-uri URI of MLflow tracking server (default: `None`)

--mlflow-tracking-username Username for MLflow tracking server (default: `None`)

--mlflow-tracking-password Password for MLflow tracking server (default: `None`)

--mlflow-tracking-secret-name Secret for accessing MLflow (default: `None`)

--mlflow-tracking-secret-profile Profile for accessing secret for accessing MLflow (default: `None`)

--mlflow-tracking-secret-region Region for accessing secret for accessing MLflow (default: `None`)

eval

Evaluate a model

```
pytorch-igniter-demo eval [-h] [--sagemaker-profile SAGEMAKER_PROFILE]
                          [--sagemaker-run [SAGEMAKER_RUN]]
                          [--sagemaker-wait [SAGEMAKER_WAIT]]
                          [--sagemaker-spot-instances [SAGEMAKER_SPOT_INSTANCES]]
                          [--sagemaker-script SAGEMAKER_SCRIPT]
                          [--sagemaker-source SAGEMAKER_SOURCE]
                          [--sagemaker-training-instance SAGEMAKER_TRAINING_INSTANCE]
                          [--sagemaker-training-image SAGEMAKER_TRAINING_IMAGE]
                          [--sagemaker-training-role SAGEMAKER_TRAINING_ROLE]
                          [--sagemaker-base-job-name SAGEMAKER_BASE_JOB_NAME]
                          [--sagemaker-job-name SAGEMAKER_JOB_NAME]
                          [--sagemaker-experiment-name SAGEMAKER_EXPERIMENT_NAME]
                          [--sagemaker-trial-name SAGEMAKER_TRIAL_NAME]
                          [--sagemaker-volume-size SAGEMAKER_VOLUME_SIZE]
                          [--sagemaker-max-run SAGEMAKER_MAX_RUN]
                          [--sagemaker-max-wait SAGEMAKER_MAX_WAIT]
                          [--pytorch-igniter-demo PYTORCH_IGNITER_DEMO]
                          [--model-dir MODEL_DIR] [--output-dir OUTPUT_DIR]
                          [--checkpoint-dir CHECKPOINT_DIR]
                          [--sagemaker-checkpoint-s3 SAGEMAKER_CHECKPOINT_S3]
                          [--sagemaker-checkpoint-container SAGEMAKER_CHECKPOINT_
↪CONTAINER]
```

(continues on next page)

(continued from previous page)

```
                          [--input INPUT] [--device DEVICE]
                          [--classes CLASSES]
                          [--eval-batch-size EVAL_BATCH_SIZE]
                          [--mlflow-enable [MLFLOW_ENABLE]]
                          [--mlflow-experiment-name MLFLOW_EXPERIMENT_NAME]
                          [--mlflow-run-name MLFLOW_RUN_NAME]
                          [--mlflow-tracking-uri MLFLOW_TRACKING_URI]
                          [--mlflow-tracking-username MLFLOW_TRACKING_USERNAME]
                          [--mlflow-tracking-password MLFLOW_TRACKING_PASSWORD]
                          [--mlflow-tracking-secret-name MLFLOW_TRACKING_SECRET_NAME]
                          [--mlflow-tracking-secret-profile MLFLOW_TRACKING_SECRET_
↪PROFILE]
                          [--mlflow-tracking-secret-region MLFLOW_TRACKING_SECRET_
↪REGION]
```

Named Arguments

--model-dir Directory to save final model (default: output/model)

Default: "output/model"

--output-dir Directory for logs, images, or other output files (default: "output/output")

Default: "output/output"

SageMaker

SageMaker options

--sagemaker-profile AWS profile for SageMaker session (default: [default])

Default: "default"

--sagemaker-run Run training on SageMaker (yes/no default=False)

Default: False

--sagemaker-wait Wait for SageMaker training to complete and tail logs files (yes/no default=True)

Default: True

--sagemaker-spot-instances Use spot instances for training (yes/no default=False)

Default: False

--sagemaker-script Script to run on SageMaker. (default: [/home/docs/checkouts/readthedocs.org/user_builds/pytorch-igniter-demo/checkouts/latest/pytorch_igniter_demo/main.py])

Default: "/home/docs/checkouts/readthedocs.org/user_builds/pytorch-igniter-demo/checkouts/latest/pytorch_igniter_demo/main.py"

--sagemaker-source Source to upload to SageMaker. Must contain script. If blank, default to directory containing script. (default: [])

Default: ""

--sagemaker-training-instance Instance type for training

Default: "ml.m5.large"

--sagemaker-training-image Docker image for training

Default: "683880991063.dkr.ecr.us-east-1.amazonaws.com/columbo-sagemaker-training:latest"

--sagemaker-training-role Docker image for training

Default: "aws-sagemaker-remote-training-role"

--sagemaker-base-job-name Base job name for tracking and organization on S3. A job name will be generated from the base job name unless a job name is specified.

Default: "training-job"

--sagemaker-job-name Job name for tracking. Use –base-job-name instead and a job name will be automatically generated with a timestamp.

Default: ""

--sagemaker-experiment-name Name of experiment in SageMaker tracking.

--sagemaker-trial-name Name of experiment trial in SageMaker tracking.

--sagemaker-volume-size Volume size in GB.

Default: 30

--sagemaker-max-run Maximum runtime in seconds.

Default: 43200

--sagemaker-max-wait Maximum time to wait for spot instances in seconds.

Default: 86400

Dependencies

Dependencies to upload to SageMaker

--pytorch-igniter-demo Directory for dependency [pytorch_igniter_demo] (default: "/home/docs/checkouts/readthedocs.org/user_builds/pytorch-igniter-demo/checkouts/latest/pytorch_igniter_demo")

Default: "pytorch_igniter_demo"

Checkpoints

Checkpointing options

--checkpoint-dir Local directory to store checkpoints for resuming training (default: "output/checkpoint")

Default: "output/checkpoint"

--sagemaker-checkpoint-s3 Location to store checkpoints on S3 or "default" (default: "default")

Default: "default"

--sagemaker-checkpoint-container Location to store checkpoints on container (default: "/opt/ml/checkpoints")

Default: "/opt/ml/checkpoints"

Inputs

Inputs (local or S3)

--input Input channel [input]. Set to local path and it will be uploaded to S3 and downloaded to SageMaker. Set to S3 path and it will be downloaded to SageMaker. (default: [output/data])

Default: "output/data"

Model

Model arguments

--device device to use (default: None)

--classes Default: 10

Evaluation

Evaluation arguments

--eval-batch-size Default: 32

MLflow

MLflow arguments

--mlflow-enable Enable logging to MLflow (default: True)

Default: True

--mlflow-experiment-name Experiment name in MLflow (default: default)

Default: "default"

--mlflow-run-name Run name in MLflow (default: None)

--mlflow-tracking-uri URI of MLflow tracking server (default: `None`)

--mlflow-tracking-username Username for MLflow tracking server (default: `None`)

--mlflow-tracking-password Password for MLflow tracking server (default: `None`)

--mlflow-tracking-secret-name Secret for accessing MLflow (default: `None`)

--mlflow-tracking-secret-profile Profile for accessing secret for accessing MLflow (default: `None`)

--mlflow-tracking-secret-region Region for accessing secret for accessing MLflow (default: `None`)

train-and-eval

Train and evaluate a model

```
pytorch-igniter-demo train-and-eval [-h]
                                    [--sagemaker-profile SAGEMAKER_PROFILE]
                                    [--sagemaker-run [SAGEMAKER_RUN]]
                                    [--sagemaker-wait [SAGEMAKER_WAIT]]
                                    [--sagemaker-spot-instances [SAGEMAKER_SPOT_
↪INSTANCES]]
                                    [--sagemaker-script SAGEMAKER_SCRIPT]
                                    [--sagemaker-source SAGEMAKER_SOURCE]
                                    [--sagemaker-training-instance SAGEMAKER_TRAINING_
↪INSTANCE]
                                    [--sagemaker-training-image SAGEMAKER_TRAINING_
↪IMAGE]
                                    [--sagemaker-training-role SAGEMAKER_TRAINING_
↪ROLE]
                                    [--sagemaker-base-job-name SAGEMAKER_BASE_JOB_
↪NAME]
                                    [--sagemaker-job-name SAGEMAKER_JOB_NAME]
                                    [--sagemaker-experiment-name SAGEMAKER_EXPERIMENT_
↪NAME]
                                    [--sagemaker-trial-name SAGEMAKER_TRIAL_NAME]
                                    [--sagemaker-volume-size SAGEMAKER_VOLUME_SIZE]
                                    [--sagemaker-max-run SAGEMAKER_MAX_RUN]
                                    [--sagemaker-max-wait SAGEMAKER_MAX_WAIT]
                                    [--pytorch-igniter-demo PYTORCH_IGNITER_DEMO]
                                    [--model-dir MODEL_DIR]
                                    [--output-dir OUTPUT_DIR]
                                    [--checkpoint-dir CHECKPOINT_DIR]
                                    [--sagemaker-checkpoint-s3 SAGEMAKER_CHECKPOINT_
↪S3]
                                    [--sagemaker-checkpoint-container SAGEMAKER_
↪CHECKPOINT_CONTAINER]
                                    [--input INPUT] [--device DEVICE]
                                    [--classes CLASSES] [--max-epochs N]
                                    [--n-saved N_SAVED]
                                    [--learning-rate LEARNING_RATE]
                                    [--weight-decay WEIGHT_DECAY]
                                    [--train-batch-size TRAIN_BATCH_SIZE]
                                    [--eval-batch-size EVAL_BATCH_SIZE]
                                    [--mlflow-enable [MLFLOW_ENABLE]]
                                    [--mlflow-experiment-name MLFLOW_EXPERIMENT_NAME]
                                    [--mlflow-run-name MLFLOW_RUN_NAME]
                                    [--mlflow-tracking-uri MLFLOW_TRACKING_URI]
                                    [--mlflow-tracking-username MLFLOW_TRACKING_
↪USERNAME]
                                    [--mlflow-tracking-password MLFLOW_TRACKING_
↪PASSWORD]
                                    [--mlflow-tracking-secret-name MLFLOW_TRACKING_
↪SECRET_NAME]
                                    [--mlflow-tracking-secret-profile MLFLOW_TRACKING_
↪SECRET_PROFILE]
                                    [--mlflow-tracking-secret-region MLFLOW_TRACKING_
↪SECRET_REGION]
```

Named Arguments

--model-dir Directory to save final model (default: output/model)

Default: "output/model"

--output-dir Directory for logs, images, or other output files (default: "output/output")

Default: "output/output"

SageMaker

SageMaker options

--sagemaker-profile AWS profile for SageMaker session (default: [default])

Default: "default"

--sagemaker-run Run training on SageMaker (yes/no default=False)

Default: False

--sagemaker-wait Wait for SageMaker training to complete and tail logs files (yes/no default=True)

Default: True

--sagemaker-spot-instances Use spot instances for training (yes/no default=False)

Default: False

--sagemaker-script Script to run on SageMaker. (default: [/home/docs/checkouts/readthedocs.org/user_builds/pytorch-igniter-demo/checkouts/latest/pytorch_igniter_demo/main.py])

Default: "/home/docs/checkouts/readthedocs.org/user_builds/pytorch-igniter-demo/checkouts/latest/pytorch_igniter_demo/main.py"

--sagemaker-source Source to upload to SageMaker. Must contain script. If blank, default to directory containing script. (default: [])

Default: ""

--sagemaker-training-instance Instance type for training

Default: "ml.m5.large"

--sagemaker-training-image Docker image for training

Default: "683880991063.dkr.ecr.us-east-1.amazonaws.com/columbo-sagemaker-training:latest"

--sagemaker-training-role Docker image for training

Default: "aws-sagemaker-remote-training-role"

--sagemaker-base-job-name Base job name for tracking and organization on S3. A job name will be generated from the base job name unless a job name is specified.

Default: "training-job"

--sagemaker-job-name Job name for tracking. Use –base-job-name instead and a job name will be automatically generated with a timestamp.

Default: ""

--sagemaker-experiment-name Name of experiment in SageMaker tracking.

--sagemaker-trial-name Name of experiment trial in SageMaker tracking.

--sagemaker-volume-size Volume size in GB.

Default: 30

--sagemaker-max-run Maximum runtime in seconds.

Default: 43200

--sagemaker-max-wait Maximum time to wait for spot instances in seconds.

Default: 86400

Dependencies

Dependencies to upload to SageMaker

--pytorch-igniter-demo Directory for dependency [pytorch_igniter_demo] (default: "/home/docs/checkouts/readthedocs.org/user_builds/pytorch-igniter-demo/checkouts/latest/pytorch_igniter_demo")

Default: "pytorch_igniter_demo"

Checkpoints

Checkpointing options

--checkpoint-dir Local directory to store checkpoints for resuming training (default: "output/checkpoint")

Default: "output/checkpoint"

--sagemaker-checkpoint-s3 Location to store checkpoints on S3 or "default" (default: "default")

Default: "default"

--sagemaker-checkpoint-container Location to store checkpoints on container (default: "/opt/ml/checkpoints")

Default: "/opt/ml/checkpoints"

Inputs

Inputs (local or S3)

--input Input channel [input]. Set to local path and it will be uploaded to S3 and downloaded to SageMaker. Set to S3 path and it will be downloaded to SageMaker. (default: [output/data])

Default: "output/data"

Model

Model arguments

--device device to use (default: None)

--classes Default: 10

Training

Training arguments

--max-epochs number of epochs to train (default: 10)

Default: 10

--n-saved Number of checkpoints to keep (default: `10`)

Default: 10

--learning-rate Default: 0.001

--weight-decay Default: 1e-05

--train-batch-size Default: 32

Evaluation

Evaluation arguments

--eval-batch-size Default: 32

MLflow

MLflow arguments

--mlflow-enable Enable logging to MLflow (default: True)

Default: True

--mlflow-experiment-name Experiment name in MLflow (default: default)

Default: "default"

--mlflow-run-name Run name in MLflow (default: None)

--mlflow-tracking-uri URI of MLflow tracking server (default: `None`)

--mlflow-tracking-username Username for MLflow tracking server (default: `None`)

--mlflow-tracking-password Password for MLflow tracking server (default: `None`)

--mlflow-tracking-secret-name Secret for accessing MLflow (default: `None`)

--mlflow-tracking-secret-profile Profile for accessing secret for accessing MLflow (default: `None`)

--mlflow-tracking-secret-region Region for accessing secret for accessing MLflow (default: `None`)

dataprep

Prepare dataset

```
pytorch-igniter-demo dataprep [-h] [--sagemaker-profile SAGEMAKER_PROFILE]
                              [--sagemaker-run [SAGEMAKER_RUN]]
                              [--sagemaker-wait [SAGEMAKER_WAIT]]
                              [--sagemaker-script SAGEMAKER_SCRIPT]
                              [--sagemaker-python SAGEMAKER_PYTHON]
                              [--sagemaker-job-name SAGEMAKER_JOB_NAME]
```

(continues on next page)

(continued from previous page)

```
                                [--sagemaker-base-job-name SAGEMAKER_BASE_JOB_NAME]
                                [--sagemaker-runtime-seconds SAGEMAKER_RUNTIME_SECONDS]
                                [--sagemaker-role SAGEMAKER_ROLE]
                                [--sagemaker-requirements SAGEMAKER_REQUIREMENTS]
                                [--sagemaker-configuration-script SAGEMAKER_
↪CONFIGURATION_SCRIPT]
                                [--sagemaker-configuration-command SAGEMAKER_
↪CONFIGURATION_COMMAND]
                                [--sagemaker-image SAGEMAKER_IMAGE]
                                [--sagemaker-instance SAGEMAKER_INSTANCE]
                                [--sagemaker-volume-size SAGEMAKER_VOLUME_SIZE]
                                [--sagemaker-output-json SAGEMAKER_OUTPUT_JSON]
                                [--sagemaker-output-mount SAGEMAKER_OUTPUT_MOUNT]
                                [--output OUTPUT] [--output-s3 OUTPUT_S3]
                                [--output-mode OUTPUT_MODE]
                                [--sagemaker-module-mount SAGEMAKER_MODULE_MOUNT]
                                [--aws-sagemaker-remote AWS_SAGEMAKER_REMOTE]
```

SageMaker

SageMaker options

--sagemaker-profile AWS profile for SageMaker session (default: [default])

Default: "default"

--sagemaker-run Run processing on SageMaker (yes/no default=False)

Default: False

--sagemaker-wait Wait for SageMaker processing to complete and tail logs (yes/no default=True)

Default: True

--sagemaker-script Python script to execute (default: [/home/docs/checkouts/readthedocs.org/user_builds/pytorch-igniter-demo/checkouts/latest/pytorch_igniter_demo/dataprep.py])

Default: "/home/docs/checkouts/readthedocs.org/user_builds/pytorch-igniter-demo/checkouts/latest/pytorch_igniter_demo/dataprep.py"

--sagemaker-python Python executable to use in container (default: [python3])

Default: "python3"

--sagemaker-job-name Job name for SageMaker processing. If not provided, will be generated from base job name. Leave blank for most use-cases. (default: [])

Default: ""

--sagemaker-base-job-name Base job name for SageMaker processing .Job name will be generated from the base name and a timestamp (default: [pytorch-igniter-demo-dataprep])

Default: "pytorch-igniter-demo-dataprep"

--sagemaker-runtime-seconds SageMaker maximum runtime in seconds (default: [3600])

Default: 3600

--sagemaker-role AWS role for SageMaker execution (default: [aws-sagemaker-remote-processing-role])

Default: "aws-sagemaker-remote-processing-role"

--sagemaker-requirements Requirements file to install on SageMaker (default: [None])

--sagemaker-configuration-script Bash configuration script to source on SageMaker (default: [None])

--sagemaker-configuration-command Bash command to run on SageMaker for configuration (e.g., `pip install aws_sagemaker_remote && export MYVAR=MYVALUE`) (default: [pip3 install –upgrade sagemaker sagemaker-experiments])

Default: "pip3 install –upgrade sagemaker sagemaker-experiments"

--sagemaker-image AWS ECR image URI of Docker image to run SageMaker processing (default: [683880991063.dkr.ecr.us-east-1.amazonaws.com/columbo-compute:latest])

Default: "683880991063.dkr.ecr.us-east-1.amazonaws.com/columbo-compute:latest"

--sagemaker-instance AWS SageMaker instance to run processing (default: [ml.t3.medium])

Default: "ml.t3.medium"

--sagemaker-volume-size AWS SageMaker volume size in GB (default: [30])

Default: 30

--sagemaker-output-json Write SageMaker training details to JSON file (default: [None])

Output

Output options

--sagemaker-output-mount Mount point for outputs. If running on SageMaker, outputs written here are uploaded to S3. If running locally, S3 outputs written here are uploaded to S3. No effect on local outputs when running locally. (default: [/opt/ml/processing/output])

Default: "/opt/ml/processing/output"

--output Output [output] local path. If running locally, set to a local path. (default: [output/data])

Default: "output/data"

--output-s3 Output [output] S3 URI. Upload results to this URI. Empty string automatically generates a URI. (default: [default])

Default: "default"

--output-mode Output [output] mode. Set to Continuous or EndOfJob. (default: [EndOfJob])

Default: "EndOfJob"

Modules

Module options

--sagemaker-module-mount Mount point for modules. If running on SageMaker, modules are mounted here and this directory is added to PYTHONPATH (default: [/opt/ml/processing/modules])

Default: "/opt/ml/processing/modules"

--aws-sagemaker-remote Directory of [aws_sagemaker_remote] module. If running on SageMaker, modules are uploaded and placed on PYTHONPATH. (default: [/home/docs/checkouts/readthedocs.org/user_builds/pytorch-igniter-demo/envs/latest/lib/python3.7/site-packages/aws_sagemaker_remote])

Default: "/home/docs/checkouts/readthedocs.org/user_builds/pytorch-igniter-demo/envs/latest/lib/python3.7/site-packages/aws_sagemaker_remote"

See `pytorch-igniter` documentation for detailed option documentation. See `aws-sagemaker-remote` documentation for detailed SageMaker option documentation.

CHAPTER 3

pytorch_igniter_demo

3.1 pytorch_igniter_demo package

3.1.1 Submodules

3.1.2 pytorch_igniter_demo.config module

class pytorch_igniter_demo.config.**Evaluator**(*args*, *model*)
Bases: object

step(*engine*, *batch*)

class pytorch_igniter_demo.config.**Inferencer**(*args*)
Bases: object

inference(*data*)

class pytorch_igniter_demo.config.**Net**(*args*)
Bases: torch.nn.modules.module.Module

forward(*x*)

class pytorch_igniter_demo.config.**Trainer**(*args*, *model*)
Bases: torch.nn.modules.module.Module

step(*engine*, *batch*)

pytorch_igniter_demo.config.**eval_args**(*parser*)

pytorch_igniter_demo.config.**make_config**()

pytorch_igniter_demo.config.**model_args**(*parser*)

pytorch_igniter_demo.config.**train_args**(*parser*)

3.1.3 pytorch_igniter_demo.dataprep module

class pytorch_igniter_demo.dataprep.**DataprepCommand**(*env=None*)
Bases: aws_sagemaker_remote.processing.main.ProcessingCommand

pytorch_igniter_demo.dataprep.**dataprep**(*args*)

pytorch_igniter_demo.dataprep.**get_dataset**(*path*, *train=True*)

pytorch_igniter_demo.dataprep.**get_loader**(*path*, *batch_size*, *train=True*, *shuffle=True*)

3.1.4 pytorch_igniter_demo.main module

pytorch_igniter_demo.main.**main**(*dry_run=False*)
Generate experiment CLI. Running locally, this function is run by a wrapper created by setuptools

pytorch_igniter_demo.main.**parser_for_docs**()
Create a parser for generating documentation

3.1.5 Module contents

CHAPTER 4

Indices and tables

- genindex
- modindex
- search

Python Module Index

p

Index