# Sample MNIST Report

FastEstimator 1.5.0

April 13, 2022

## Contents

# 1 Training Graphs

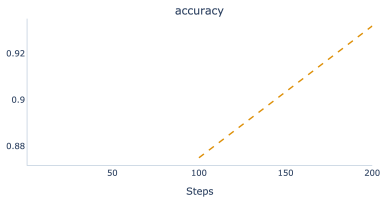


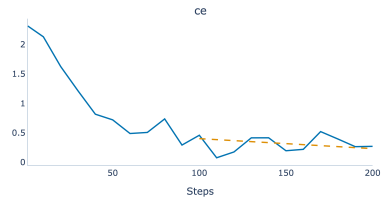


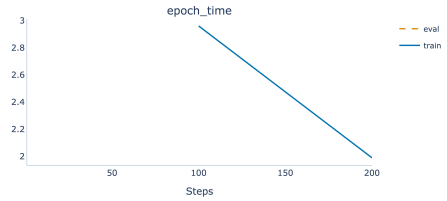


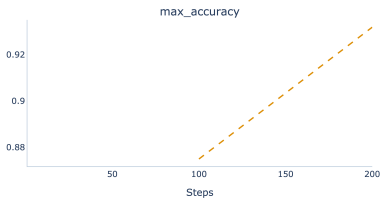


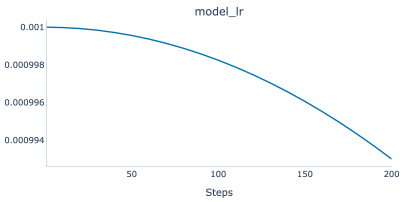



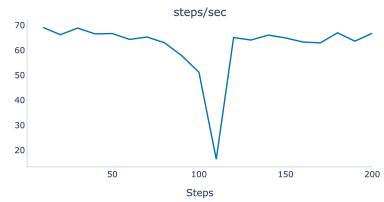

# 2 FastEstimator Architecture

## 2.1 Train

### 2.1.1 Epoch 1

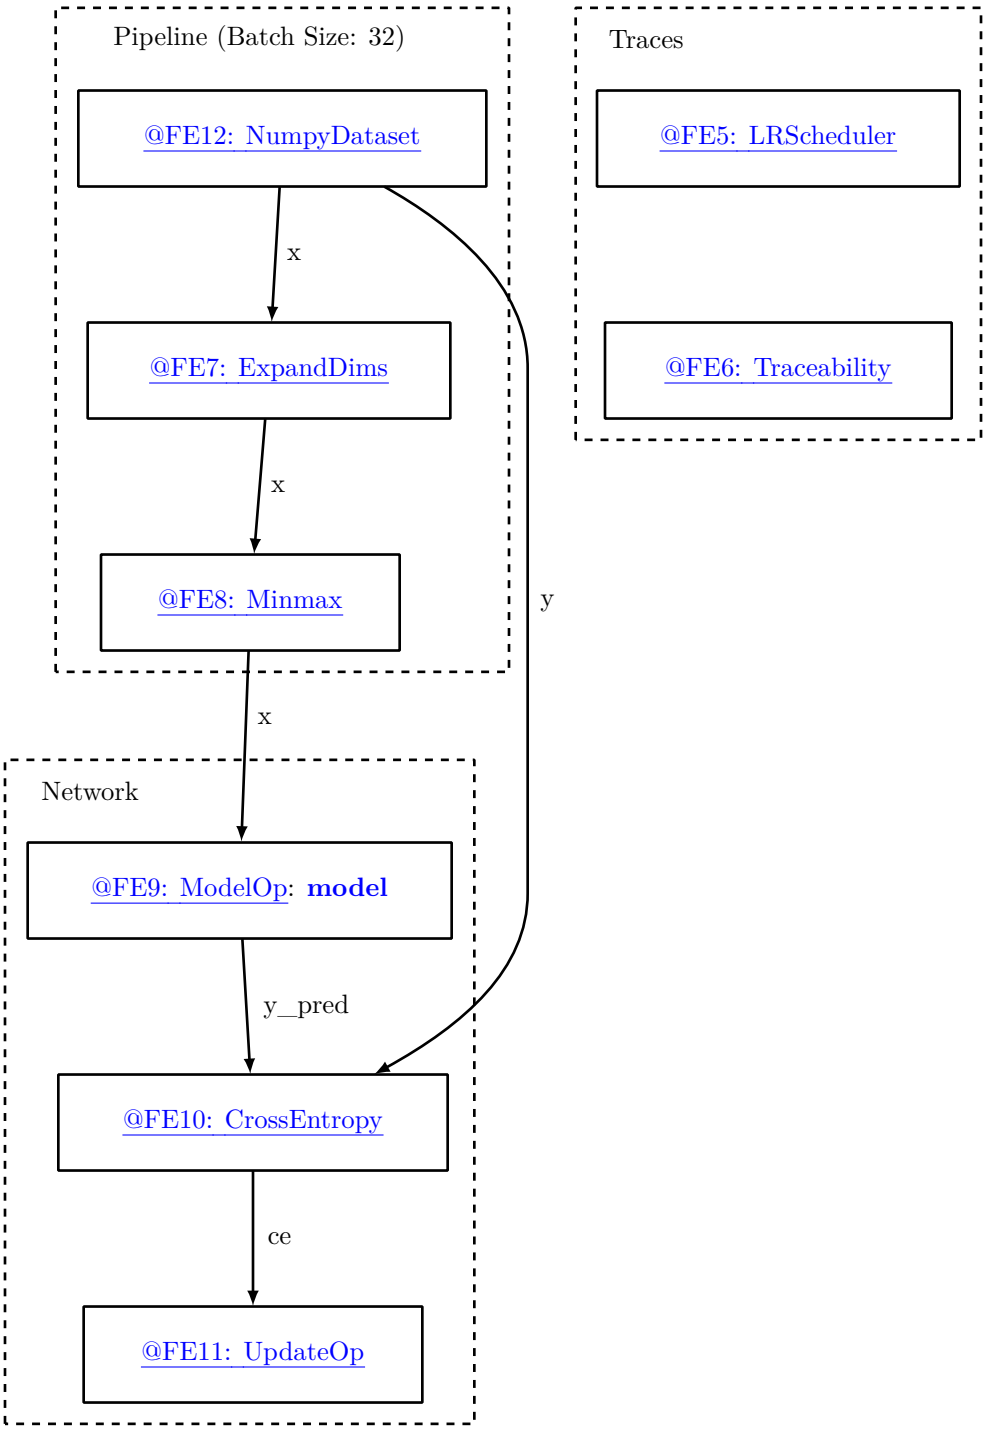

## 2.2 Eval

### 2.2.1 Epoch 1

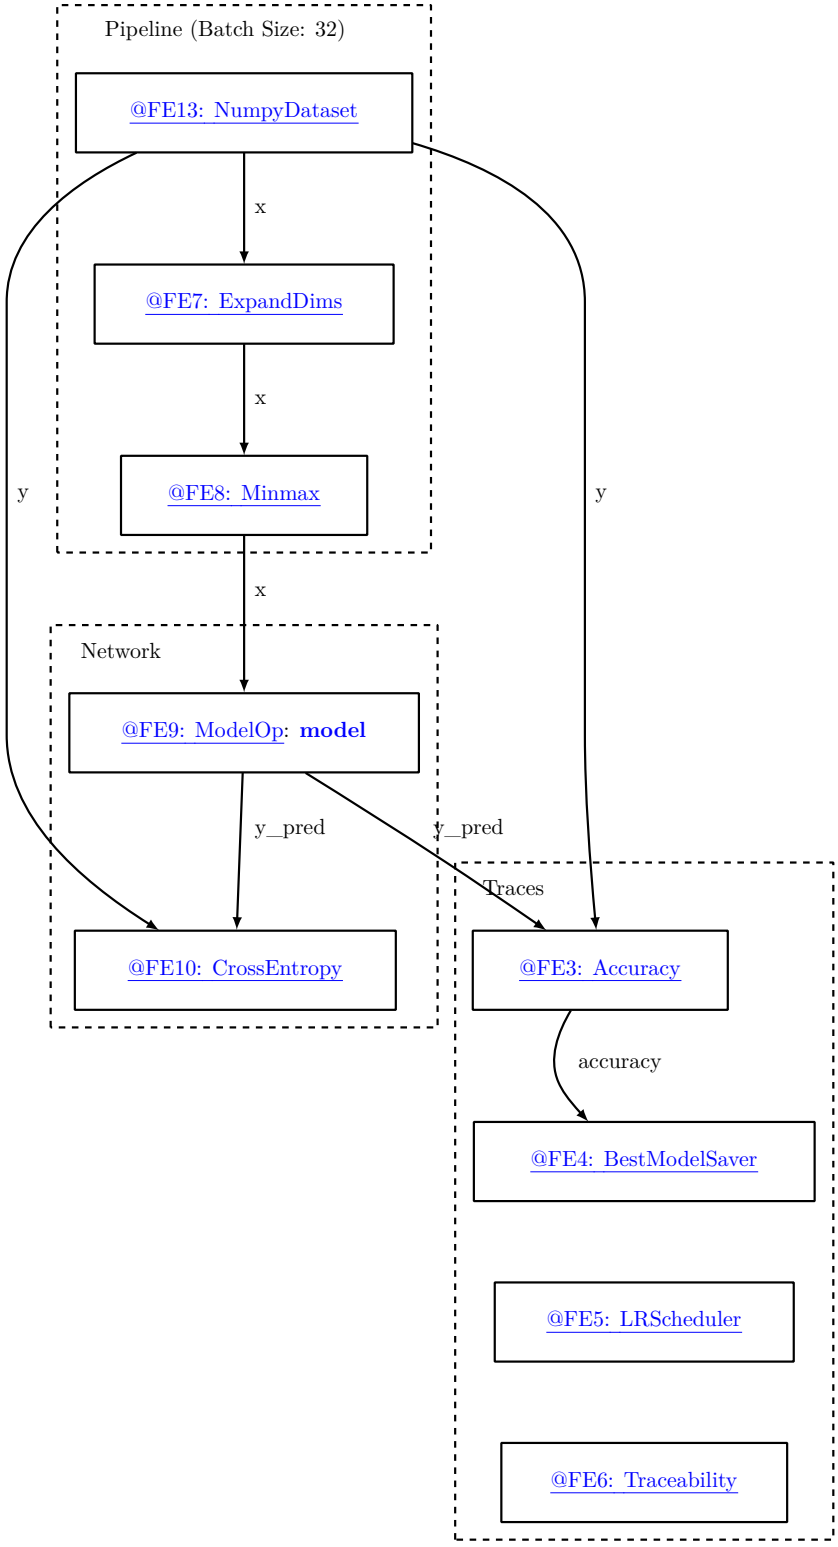

## 2.3 Test

### 2.3.1 Epoch 1

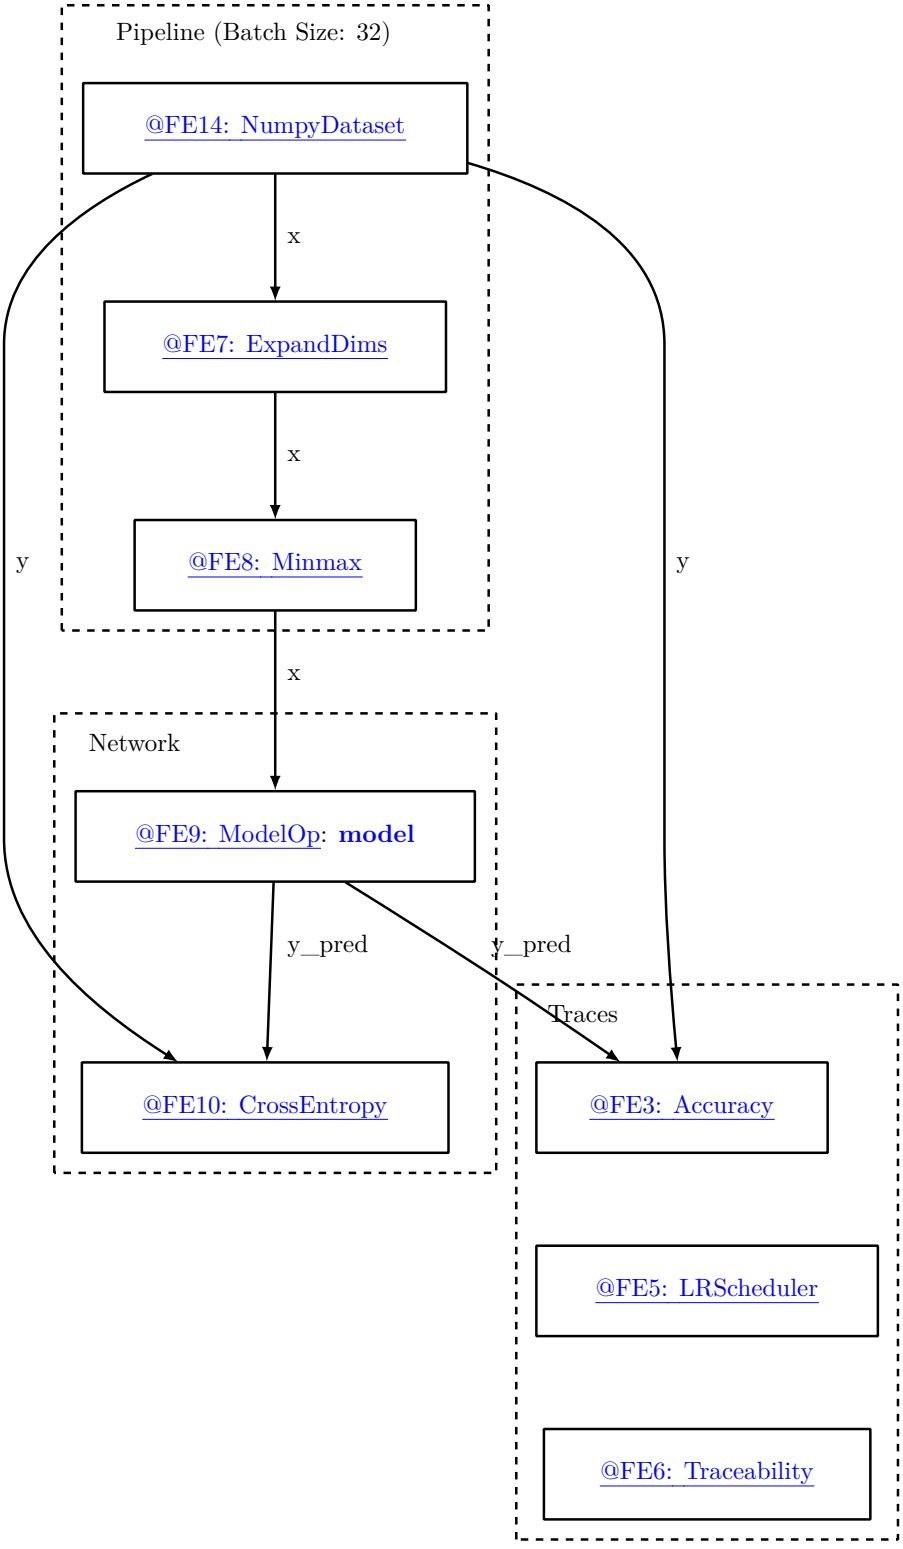

# 3 Parameters

## 3.1 Base Classes

| **Estimator** | @FE0 |
|---|---|
| Type: | fastestimator.estimator.Estimator |
| *pipeline* | @FE2: Pipeline |
| *network* | @FE1: TFNetwork |
| *epochs* | 2 |
| *train_steps_per_epoch* | 100 |
| *eval_steps_per_epoch* | 100 |
| *traces* | [ @FE3: Accuracy, @FE4: BestModelSaver, @FE5: LRScheduler, @FE6: Traceability] |
| *log_steps* | 10 |
| *monitor_names* | None |

| **TFNetwork** | @FE1 |
|---|---|
| Type: | fastestimator.network.TFNetwork |
| *ops* | [ @FE9: ModelOp, @FE10: CrossEntropy, @FE11: UpdateOp] |
| *postprocessing* | None |

| **Pipeline** | @FE2 |
|---|---|
| Type: | fastestimator.pipeline.Pipeline |
| *train_data* | @FE12: NumpyDataset |
| *eval_data* | @FE13: NumpyDataset |
| *test_data* | @FE14: NumpyDataset |
| *batch_size* | 32 |
| *ops* | [ @FE7: ExpandDims, @FE8: Minmax] |
| *num_process* | None |

## 3.2 Traces

| **Accuracy** | @FE3 |
|---|---|
| Type: | fastestimator.trace.metric.accuracy.Accuracy |
| *true_key* | 'y' |
| *pred_key* | 'y_pred' |
| *mode* | ( 'eval', 'test') |
| *ds_id* | None |
| *from_logits* | False |
| *output_name* | 'accuracy' |
| *per_ds* | True |

| **BestModelSaver** | @FE4 |
|---|---|
| Type: | fastestimator.trace.io.best_model_saver.BestModelSaver |
| *model* | @FE15: model |
| *save_dir* | '/var/folders/3r/h9kh47050gv6rbt_pgf8cl540000gn/T/tmpo70xpgjr' |
| *metric* | 'accuracy' |
| *save_best_mode* | 'max' |
| *load_best_final* | False |
| *save_architecture* | False |

| **LRScheduler** | @FE5 |
|---|---|
| Type: | fastestimator.trace.adapt.lr_scheduler.LRScheduler |
| *model* | @FE15: model |
| *lr_fn* | lambda step: cosine_decay(time=step, cycle_length=3750, init_lr=0.001, min_lr=1e-06, start=1, cycle_multiplier=1) |
| *ds_id* | None |

| **Traceability** | @FE6 |
|---|---|
| Type: | fastestimator.trace.io.traceability.Traceability |
| *save_path* | '/var/folders/3r/h9kh47050gv6rbt_pgf8cl540000gn/T/tmpo70xpgjr/report' |
| *extra_objects* | None |

## 3.3 Operators

| **ExpandDims** | @FE7 |
|---|---|
| Type: | fastestimator.op.numpyop.univariate.expand_dims.ExpandDims |
| *inputs* | 'x' |
| *outputs* | 'x' |
| *mode* | None |
| *ds_id* | None |
| *axis* | -1 |

| **Minmax** | @FE8 |
|---|---|
| Type: | fastestimator.op.numpyop.univariate.minmax.Minmax |
| *inputs* | 'x' |
| *outputs* | 'x' |
| *mode* | None |
| *ds_id* | None |
| *epsilon* | 1e-07 |

| **ModelOp** | @FE9 |
|---|---|
| Type: | fastestimator.op.tensorop.model.model.ModelOp |
| *model* | @FE15: model |
| *inputs* | 'x' |
| *outputs* | 'y_pred' |
| *mode* | None |
| *ds_id* | None |
| *trainable* | True |
| *intermediate_layers* | None |

| **CrossEntropy** | @FE10 |
|---|---|
| Type: | fastestimator.op.tensorop.loss.cross_entropy.CrossEntropy |
| *inputs* | ( 'y_pred', 'y') |
| *outputs* | 'ce' |
| *mode* | '!infer' |
| *ds_id* | None |
| *from_logits* | False |
| *average_loss* | True |
| *form* | None |
| *class_weights* | None |

| **UpdateOp** | @FE11 |
|---|---|
| Type: | fastestimator.op.tensorop.model.update.UpdateOp |
| *model* | @FE15: model |
| *loss_name* | 'ce' |
| *gradients* | None |
| *mode* | 'train' |
| *ds_id* | None |
| *merge_grad* | 1 |
| *defer* | False |

## 3.4 Datasets

| **NumpyDataset** | @FE12 |
|---|---|
| Type: | fastestimator.dataset.numpy_dataset.NumpyDataset |
| *data* | { 'x': @FE17: tensor, 'y': @FE18: tensor} |

| **NumpyDataset** | @FE13 |
|---|---|
| Type: | fastestimator.dataset.numpy_dataset.NumpyDataset |
| Split: | self(-100) |
| *data* | { 'x': @FE19: tensor, 'y': @FE20: tensor} |

| **NumpyDataset** | @FE14 |
|---|---|
| Type: | fastestimator.dataset.numpy_dataset.NumpyDataset |
| Split: | @FE13(100) |
| *data* | { 'x': @FE19: tensor, 'y': @FE20: tensor} |

## 3.5 Models

| **model** | @FE15 |
|---|---|
| Type: | keras.engine.sequential.Sequential |
| Definition: | @FE16: LeNet |
| Optimizer: | 'adam' |

## 3.6 Functions

| **LeNet** | @FE16 |
|---|---|
| Type: | function<br>fastestimator.architecture.tensorflow.lenet.LeNet |

## 3.7 Tensors

| **tensor** | @FE17 |
|---|---|
| Type: | numpy.ndarray |
| Shape: | (60000, 28, 28) |

| **tensor** | @FE18 |
|---|---|
| Type: | numpy.ndarray |
| Shape: | (60000,) |

| **tensor** | @FE19 |
|---|---|
| Type: | numpy.ndarray |
| Shape: | (10000, 28, 28) |

| **tensor** | @FE20 |
|---|---|
| Type: | numpy.ndarray |
| Shape: | (10000,) |

# 4 Models

## 4.1 Model

```
Model: "sequential"
_________________________________________________________________
 Layer (type)                Output Shape              Param #
=================================================================
 conv2d (Conv2D)             (None, 26, 26, 32)        320

 max_pooling2d (MaxPooling2D)  (None, 13, 13, 32)      0

 conv2d_1 (Conv2D)           (None, 11, 11, 64)        18496

 max_pooling2d_1 (MaxPooling2D)  (None, 5, 5, 64)      0

 conv2d_2 (Conv2D)           (None, 3, 3, 64)          36928

 flatten (Flatten)           (None, 576)               0

 dense (Dense)               (None, 64)                36928

 dense_1 (Dense)             (None, 10)                650

=================================================================
Total params: 93,322
Trainable params: 93,322
Non-trainable params: 0
_________________________________________________________________
```

@FE15: model

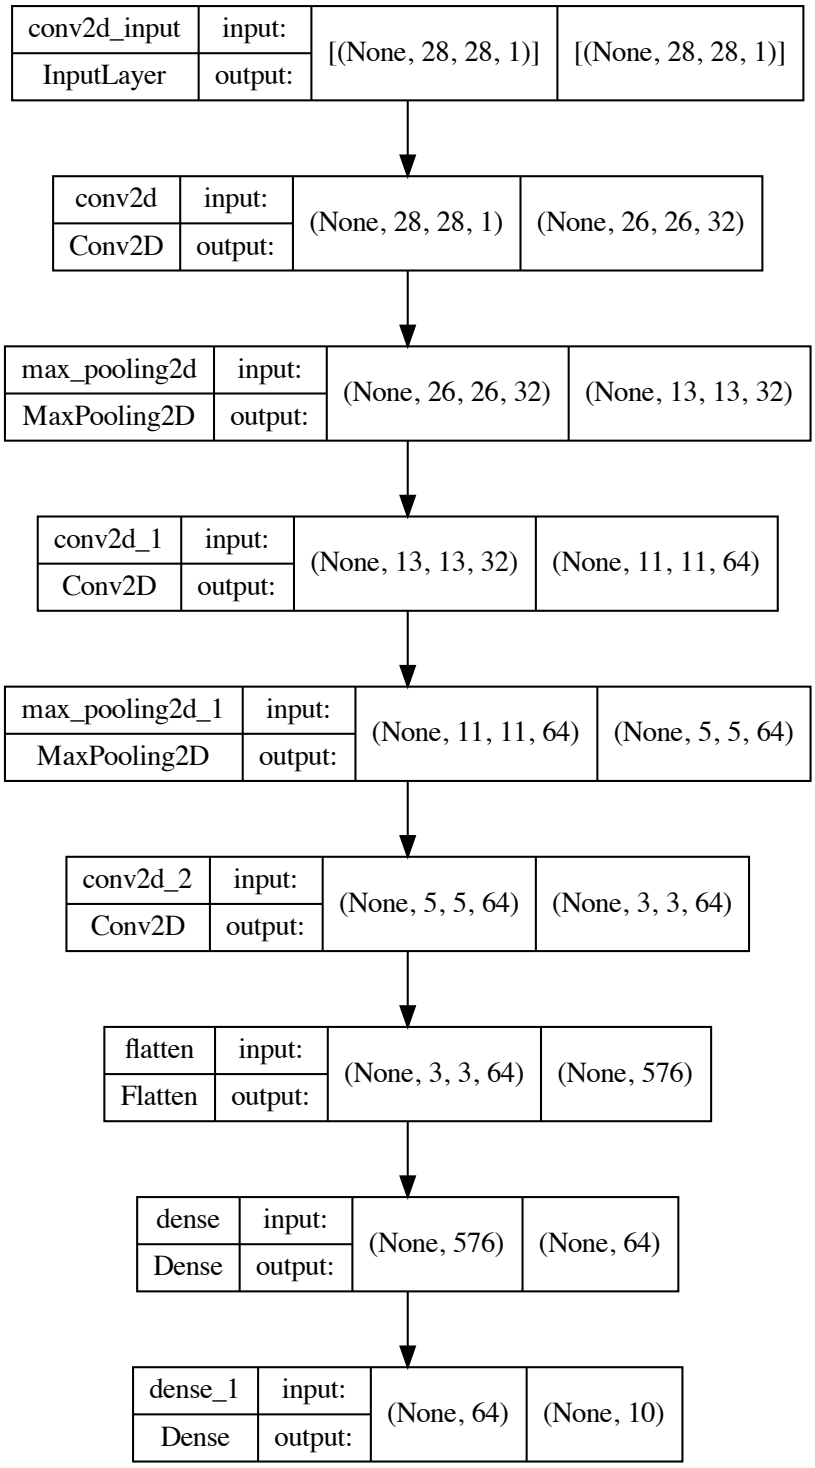


Figure 1: @FE15: model

# 5 System Configuration

- FastEstimator 1.5.0
- Python 3.8.12
- OS: darwin
- Number of GPUs: 0

| Module | Version |
|---|---|
| albumentations | 1.1.0 |
| appnope | 0.1.2 |
| argparse | 1.1 |
| astunparse | 1.6.3 |
| backcall | 0.2.0 |
| brotli | 1.0.9 |
| certifi | 2021.10.08 |
| cffi | 1.15.0 |
| charset_normalizer | 2.0.12 |
| cloudpickle | 2.0.0 |
| csv | 1.0 |
| ctypes | 1.1.0 |
| cv2 | 4.5.5 |
| cycler | 0.10.0 |
| dateutil | 2.8.2 |
| debugpy | 1.5.1 |
| decimal | 1.70 |
| decorator | 5.1.1 |
| defusedxml | 0.7.1 |
| dill | 0.3.4 |
| distutils | 3.8.12 |
| dot2tex | 2.11.3 |
| entrypoints | 0.4 |
| executing | 0.8.2 |
| fastestimator | 1.5.0 |
| filelock | 3.5.0 |
| flatbuffers | 2.0 |
| gast | 0.5.3 |
| gdown | 3.12.0 |
| h5py | 3.6.0 |
| idna | 3.3 |
| imageio | 2.16.0 |
| imgaug | 0.4.0 |
| ipaddress | 1.0 |
| ipykernel | 6.9.1 |

Continued on Next Page

| Module | Version |
|---|---|
| ipython_genutils | 0.2.0 |
| ipywidgets | 7.6.5 |
| IPython | 8.0.1 |
| jedi | 0.18.1 |
| joblib | 1.1.0 |
| json | 2.0.9 |
| jsonpickle | 2.1.0 |
| jupyter_client | 7.1.2 |
| jupyter_core | 4.9.2 |
| kaleido | 0.2.1 |
| keras | 2.8.0 |
| keras_preprocessing | 1.1.2 |
| kiwisolver | 1.3.2 |
| logging | 0.5.1.2 |
| matplotlib | 3.4.3 |
| natsort | 8.1.0 |
| nltk | 3.7 |
| numpy | 1.22.2 |
| opt_einsum | v3.3.0 |
| optparse | 1.5.3 |
| ordered_set | 4.1.0 |
| packaging | 21.3 |
| pandas | 1.4.1 |
| parso | 0.8.3 |
| pexpect | 4.8.0 |
| pickleshare | 0.7.5 |
| platform | 1.0.8 |
| plotly | 5.7.0 |
| prettytable | 3.1.0 |
| prompt_toolkit | 3.0.28 |
| ptyprocess | 0.7.0 |
| pure_eval | 0.2.2 |
| pydevd | 2.6.0 |
| pydot | 1.4.2 |
| pyfiglet | 0.8.post1 |
| pygments | 2.11.2 |
| pylatex | 1.4.1 |
| pyparsing | 3.0.7 |
| pytz | 2021.3 |
| pywt | 1.2.0 |
| PIL | 9.0.1 |
| re | 2.2.1 |

Continued on Next Page

| Module | Version |
|---|---|
| regex | 2.5.110 |
| requests | 2.27.1 |
| scipy | 1.8.0 |
| seaborn | 0.11.2 |
| setuptools | 56.0.0 |
| six | 1.16.0 |
| skimage | 0.19.1 |
| sklearn | 1.0.2 |
| socketserver | 0.4 |
| socks | 1.7.1 |
| stack_data | 0.2.0 |
| tensorboard | 2.8.0 |
| tensorflow | 2.8.0 |
| tensorflow_addons | 0.16.1 |
| tensorflow_probability | 0.16.0 |
| termcolor | (1, 1, 0) |
| threadpoolctl | 3.1.0 |
| torch | 1.10.2 |
| tqdm | 4.62.3 |
| traitlets | 5.1.1 |
| tree | 0.1.6 |
| urllib3 | 1.26.8 |
| wcwidth | 0.2.5 |
| wget | 3.2 |
| wrapt | 1.13.3 |
| yaml | 6.0 |
| zlib | 1.0 |
| zmq | 22.3.0 |