Sample MNIST Report

FastEstimator 1.6.0

November 8, 2023

Contents

1 Training Graphs

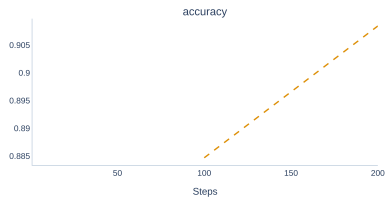

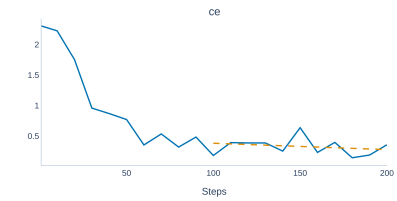

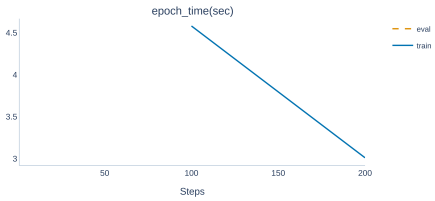

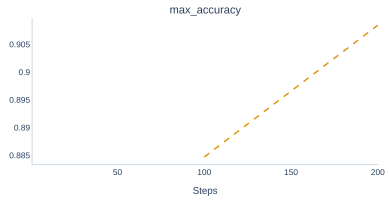

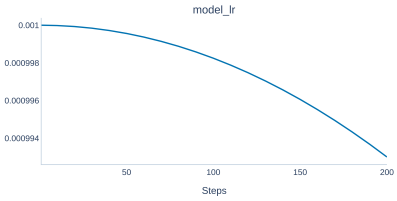

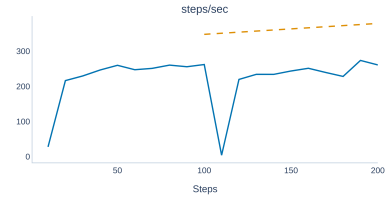

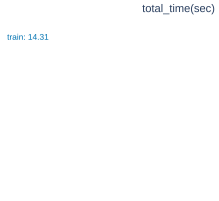

2 FastEstimator Architecture

2.1 Train

2.1.1 Epoch 1

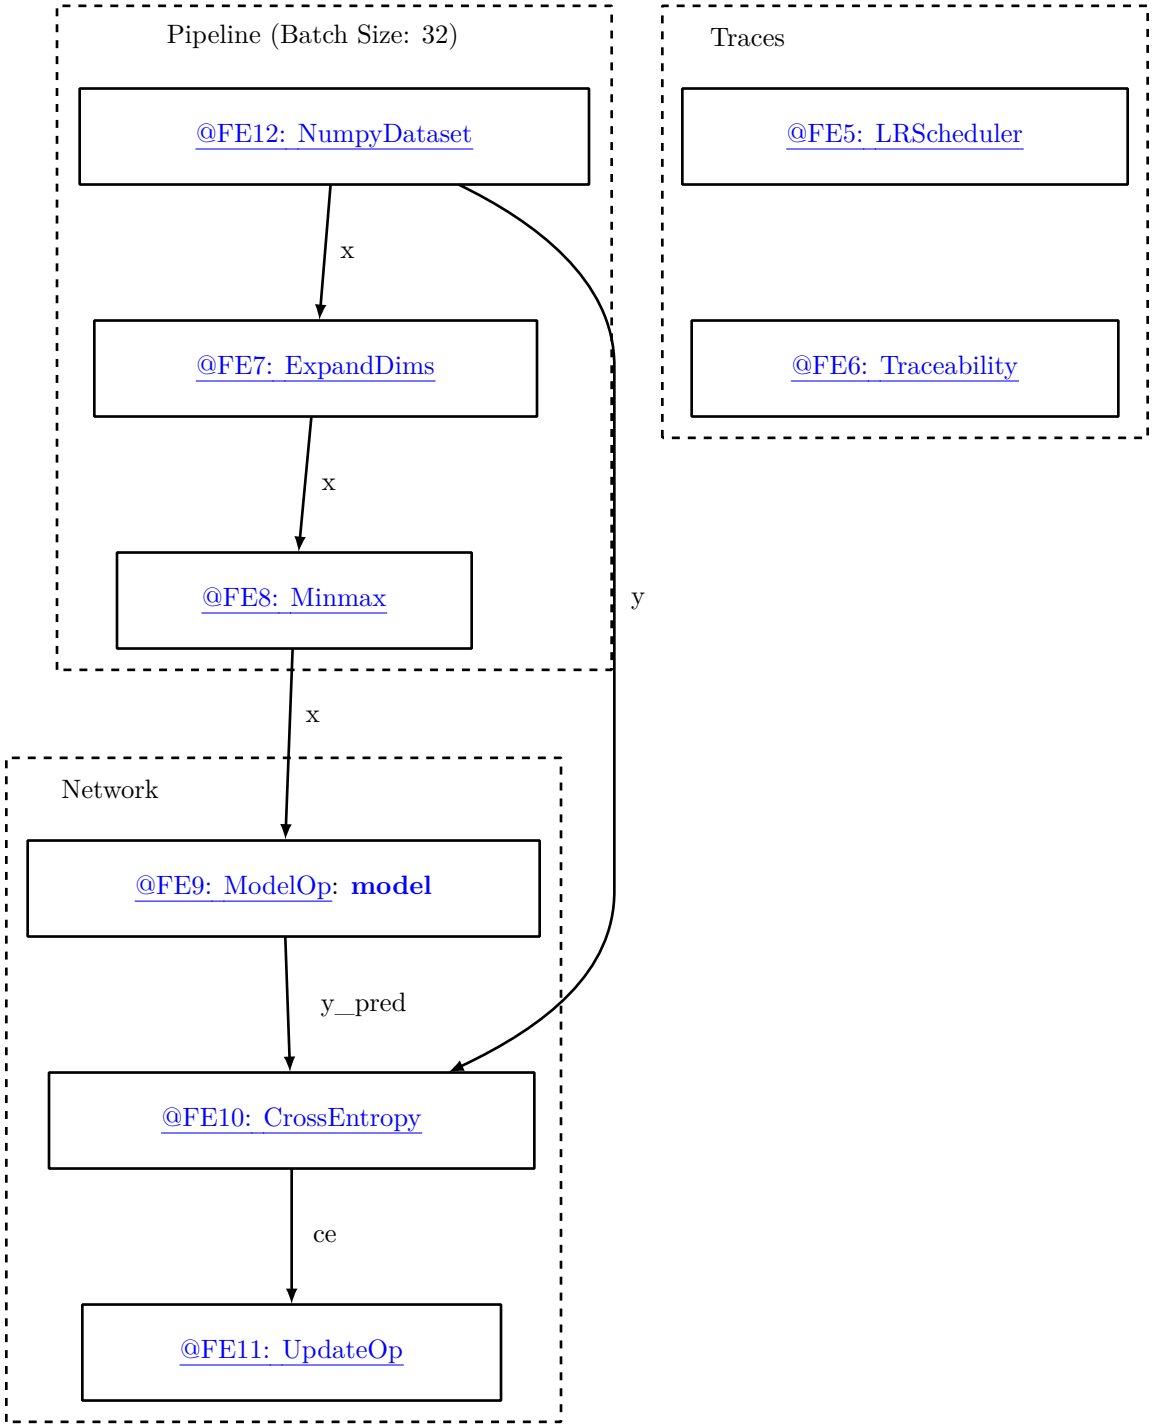

2.2 Eval

2.2.1 Epoch 1

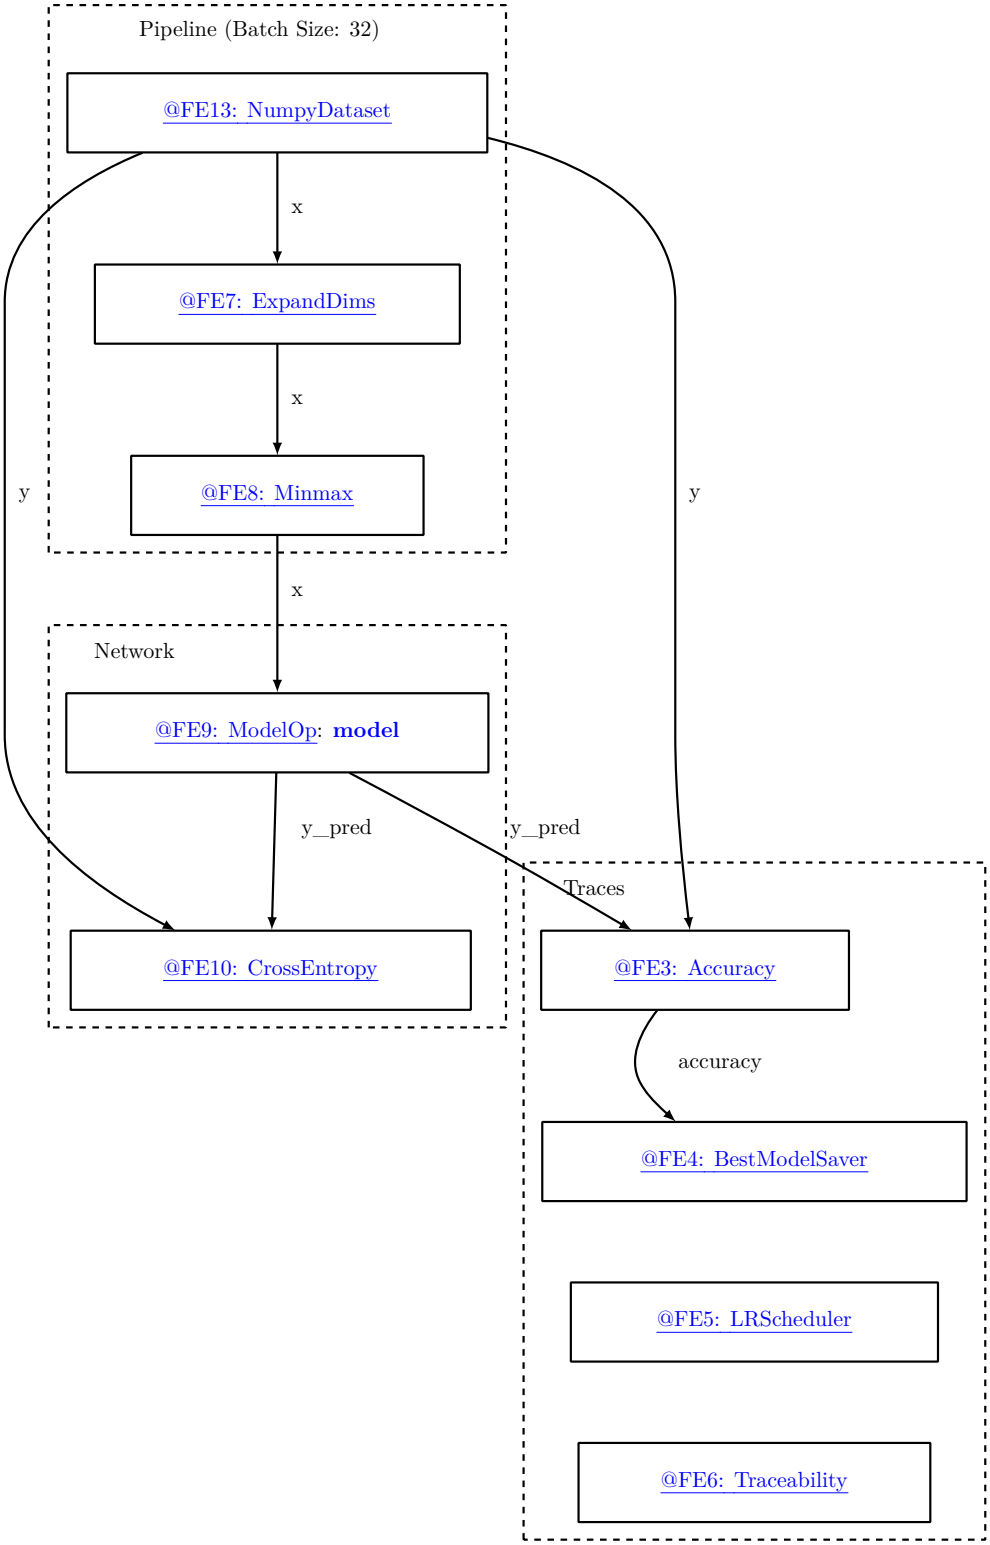

2.3 Test

2.3.1 Epoch 1

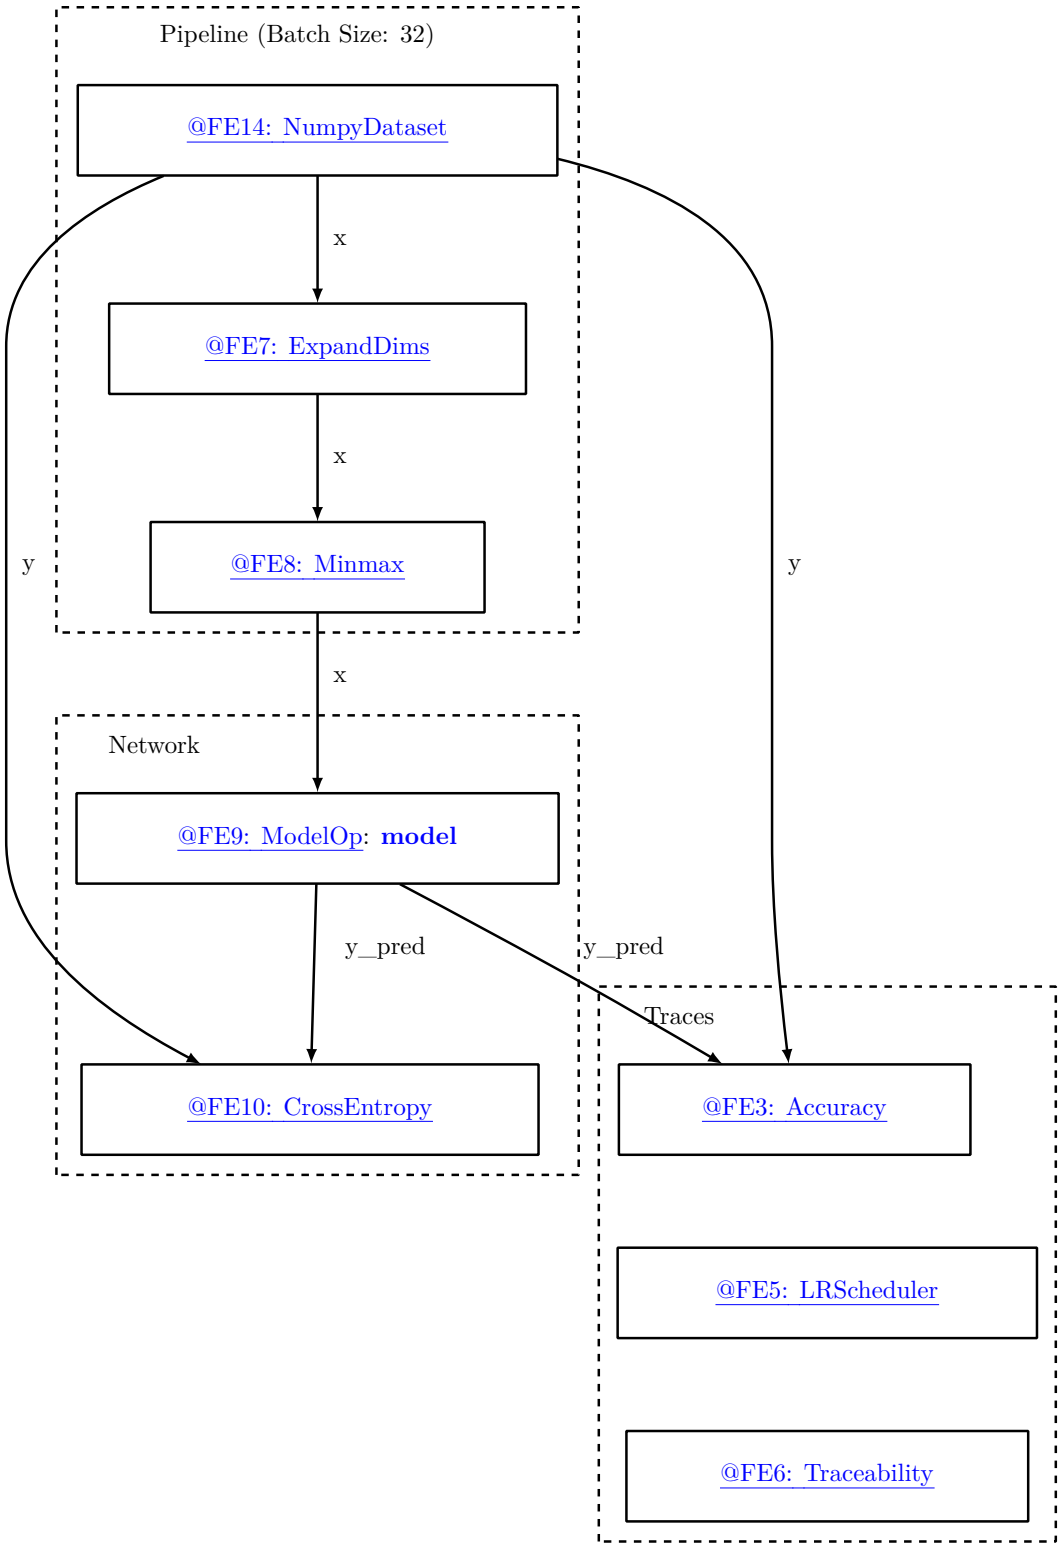

2.4 Infer

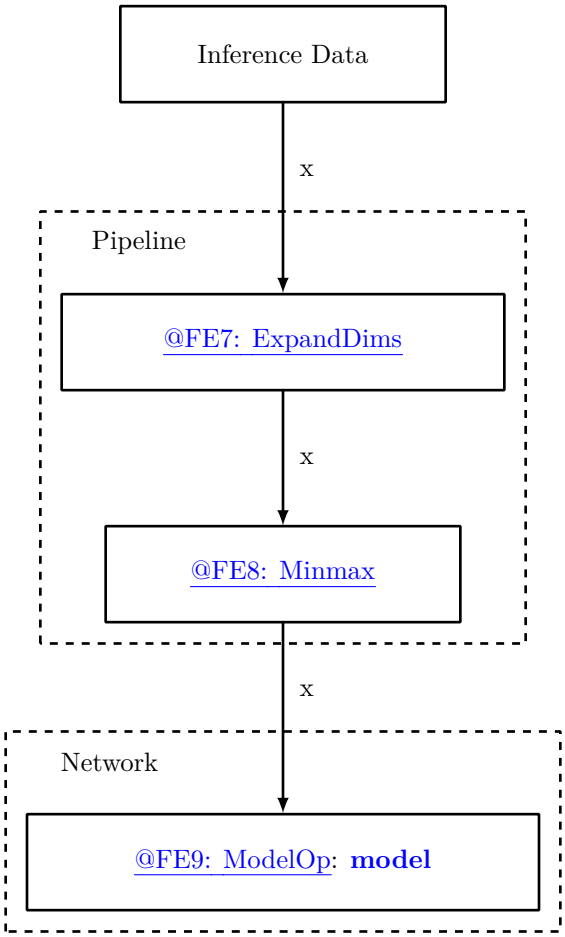
Inference Data
x
Pipeline
@FE7: ExpandDims
x
@FE8: Minmax
x
Network
@FE9: ModelOp: **model**

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
eval_log_steps	()
monitor_names	None

TFNetwork	@FE1
Type:	fastestimator.network.TFNetwork
ops	[@FE9: ModelOp, @FE10: CrossEntropy, @FE11: UpdateOp]
postprocessing	None
slicers	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	'/tmp/tmp4_019_48'
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, warmup=False)
ds_id	None

Traceability	@FE6
Type:	fastestimator.trace.io.traceability.Traceability
save_path	'/tmp/tmp4_019_48/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
new_min	0.0
new_max	1.0

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

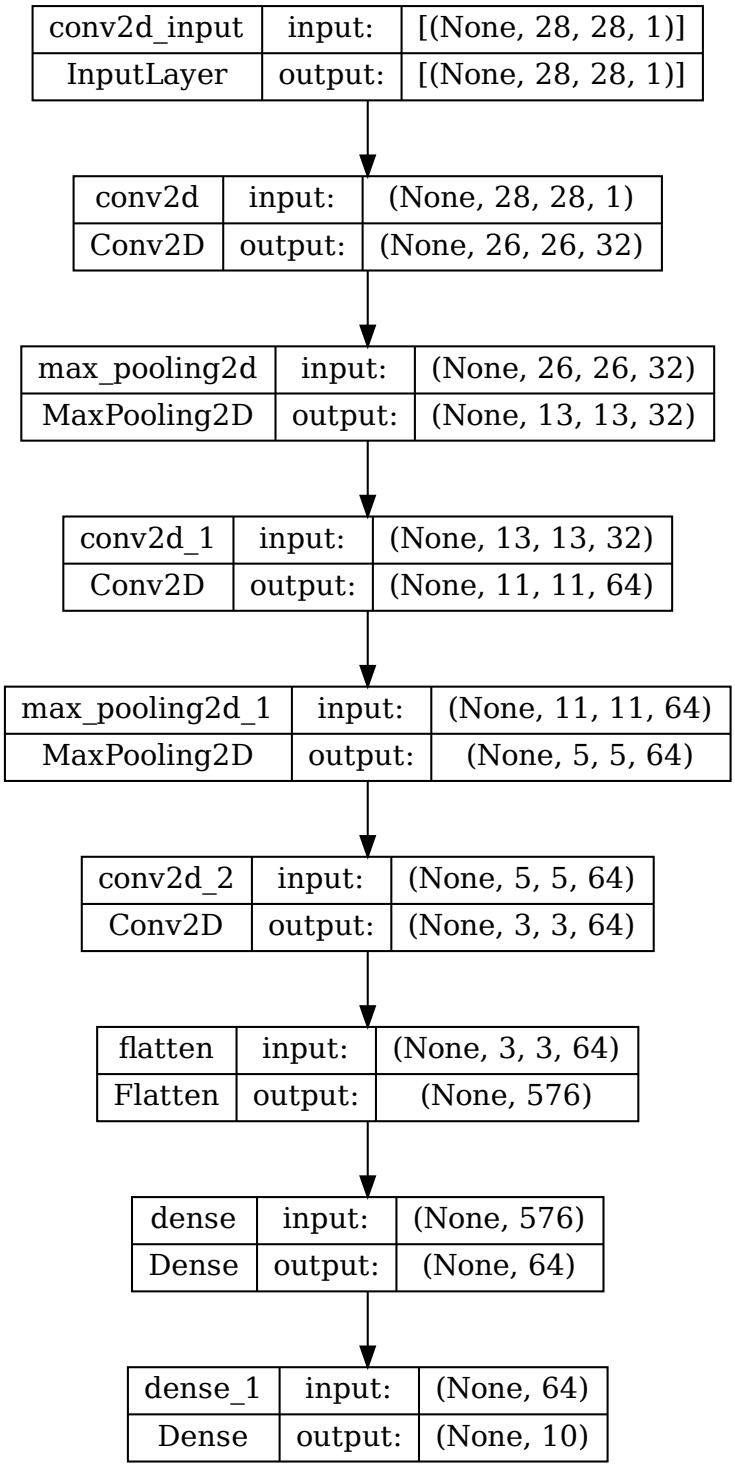

Figure 1: @FE15: model

5 System Configuration

- FastEstimator 1.6.0
- Python 3.8.10
- OS: linux
- CPU Used: 32 Threads
  - AMD EPYC 7742 64-Core Processor (256 Threads)
- GPU(s) Used: 1
  - NVIDIA A100-SXM4-40GB (40536 MiB, Driver=470.161.03)

Module	Version
argparse	1.1
astunparse	1.6.3
backcall	0.2.0
bs4	4.12.2
certifi	2019.11.28
cffi	1.15.1
chardet	3.0.4
charset_normalizer	3.1.0
cloudpickle	2.2.1
csv	1.0
ctypes	1.1.0
cv2	4.7.0
dateutil	2.8.2
debugpy	1.6.7
decimal	1.70
decorator	5.1.1
defusedxml	0.7.1
dill	0.3.6
distutils	3.8.10
dot2tex	2.11.3
entrypoints	0.4
executing	1.2.0
fastestimator	1.6.0
filelock	3.12.2
flatbuffers	23.3.3
fsspec	2023.6.0
gdown	4.6.4
graphviz	0.20.1
h5py	3.8.0
idna	2.8

Continued on Next Page

Module	Version
ipykernel	6.17.1
IPython	8.12.0
jedi	0.18.2
json	2.0.9
jsonpickle	3.0.1
jupyter_client	7.4.9
jupyter_core	5.3.1
kaleido	0.2.1
keras	2.11.0
logging	0.5.1.2
lxml	4.9.3
mpmath	1.3.0
natsort	8.3.1
numpy	1.24.2
opt_einsum	v3.3.0
optparse	1.5.3
ordered_set	4.1.0
orjson	3.8.7
packaging	23.0
pandas	2.0.1
parso	0.8.3
pexpect	4.8.0
pickleshare	0.7.5
platform	1.0.8
platformdirs	3.9.1
plotly	5.13.1
prettytable	3.6.0
prompt_toolkit	3.0.39
psutil	5.9.5
ptyprocess	0.7.0
pure_eval	0.2.2
pydevd	2.9.5
pydot	1.4.2
pyfiglet	0.8.post1
pygments	2.15.1
pylatex	1.4.1
pyparsing	3.1.0
pytz	2023.3
PIL	9.4.0
re	2.2.1
requests	2.28.2
scipy	1.9.1

Continued on Next Page

Module	Version
six	1.14.0
socketserver	0.4
socks	1.7.1
soupsieve	2.4.1
stack_data	0.6.2
sympy	1.12
tensorboard	2.11.2
tensorflow	2.11.1
tensorflow_addons	0.19.0
tensorflow_probability	0.19.0
torch	2.0.1+cu118
torchinfo	1.7.2
torchview	0.2.6
tqdm	4.65.0
traitlets	5.9.0
tree	0.1.8
urllib3	1.25.8
wcwidth	0.2.6
wget	3.2
wrapt	1.15.0
zlib	1.0
zmq	25.1.0