Graphics

Chapter 1: Overview of TensorFlow and Machine Learning

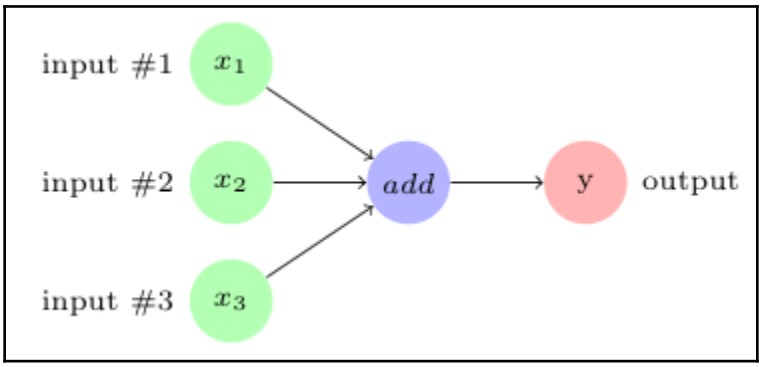

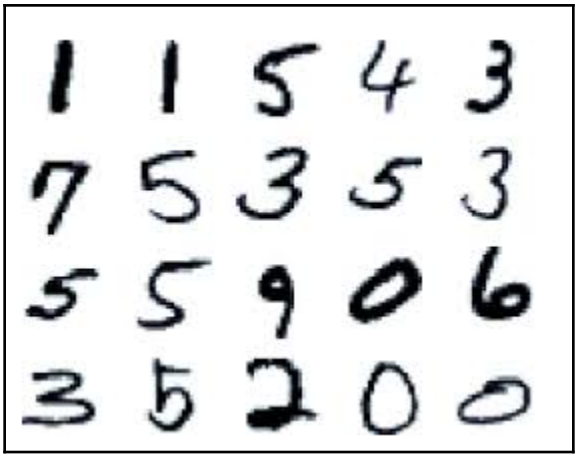

Chapter 2: Using Machine Learning to Detect Exoplanets in Outer Space

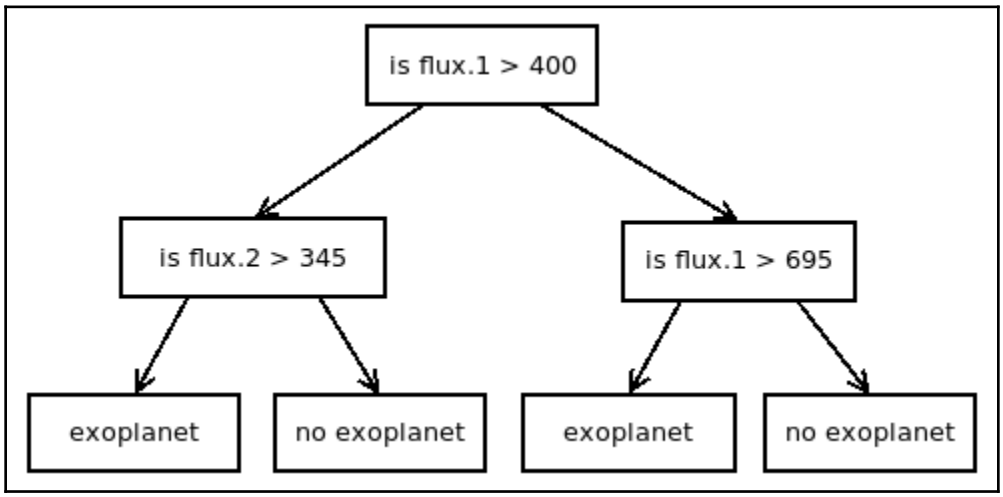

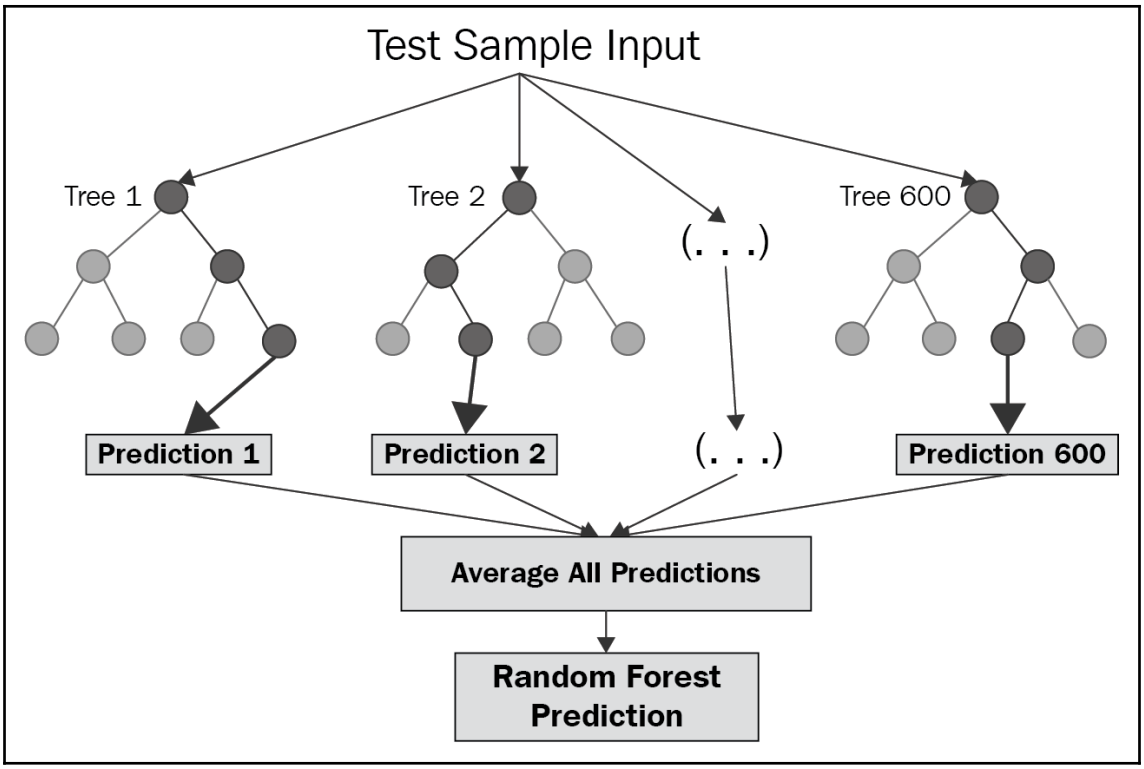
Test Sample Input
Tree 1
Tree 2
(. . .)
Tree 600
Prediction 1
Prediction 2
(. . .)
Prediction 600
Average All Predictions
Random Forest Prediction

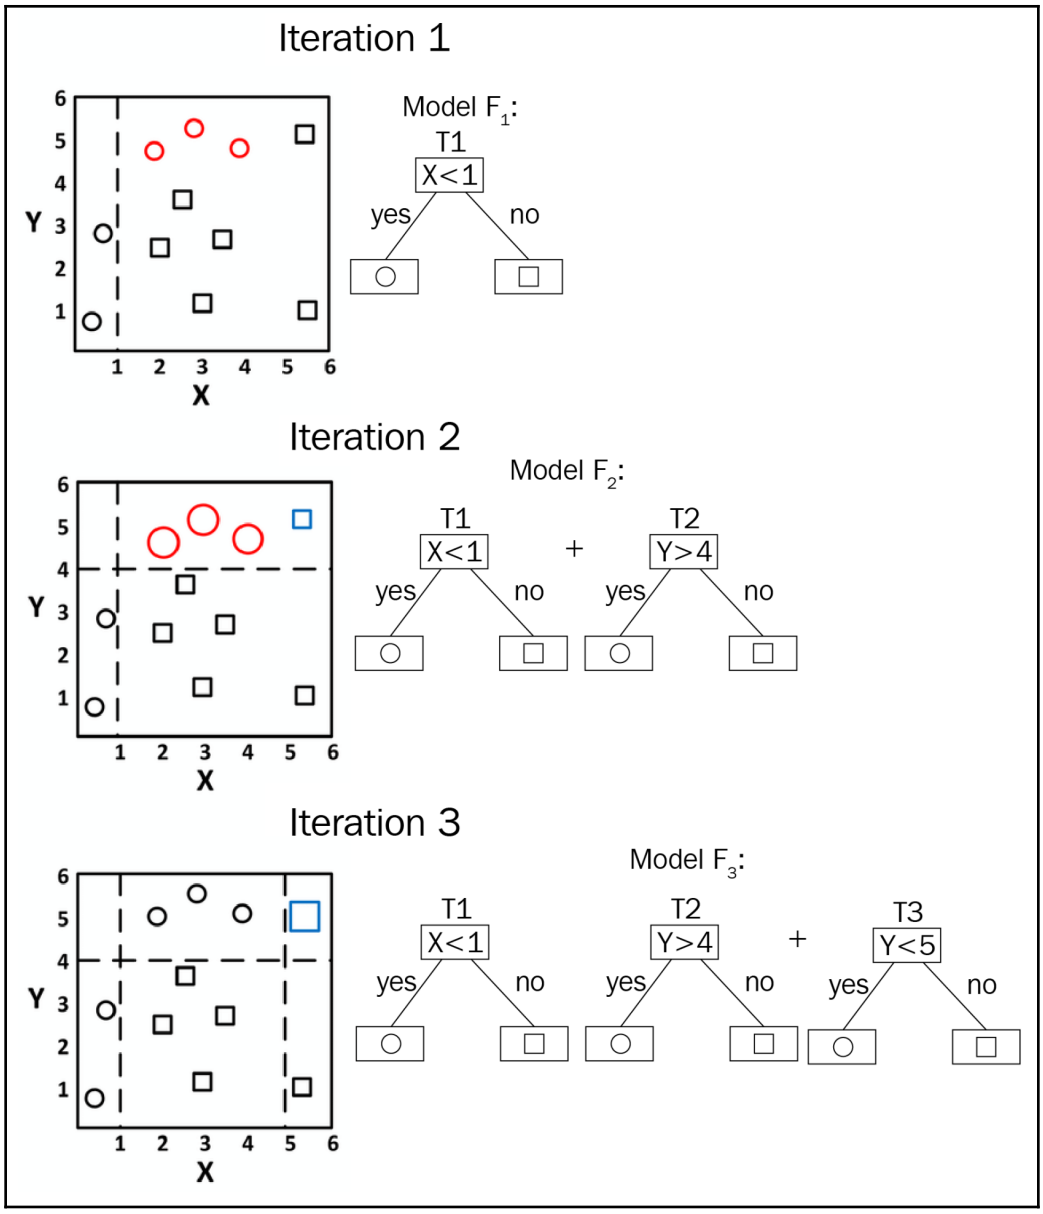
Iteration 1
Model F_1:
T1
X<1
yes
no
Y
X
Iteration 2
Model F_2:
T1
X<1
yes
no
+
T2
Y>4
yes
no
Iteration 3
Model F_3:
T1
X<1
T2
Y>4
T3
Y<5

Lib	D?	Losses	Regularization
scikit-learn	N	*R*: least squares, least absolute dev, huber and quantile. *C*: logistic, Max-Ent and exp	Depth limit, shrinkage, bagging, feature subsampling
GBM	N	*R*: least squares, least absolute dev, t-distribution, quantile, huber. *C*: logistic, Max-Ent, exp, poisson & right censored observations. Supports *ranking*	Shrinkage, bagging, depth limit, min # of examples per node.
MLLib	Y	*R*: least squared and least absolute dev. *C*: logistic.	Shrinkage, early stopping, depth limit, min # of examples per node, min gain, bagging.
Light GBM	Y	*R*: least squares, least absolute dev, huber, fair, poisson. *C*: logistic, Max-Ent. Supports *ranking*.	Dropout, shrinkage, # leafs limit, feature subsampling, bagging, L1 & L2
XGBoost	Y	*R*: least squares, poisson, gamma, tweedie regression. *C*: logistic, Max-Ent. Supports *ranking* and **custom**.	L1 & L2, shrinkage, feature subsampling, dropout, bagging, min child weight and gain, limit on depth and # of nodes, pruning.
TFBT	Y	Any twice differentiable loss from tf.contrib.losses and **custom** losses.	L1 & L2, tree complexity, shrinkage, line search for learning rate, dropout, feature subsampling and bagging, limit on depth and min node weight, pre- post- pruning.

D? is whether a library supports distributed mode. *R* stands for regression, *C* for classification.

Chapter 3: Sentiment Analysis in Your Browser Using TensorFlow.js

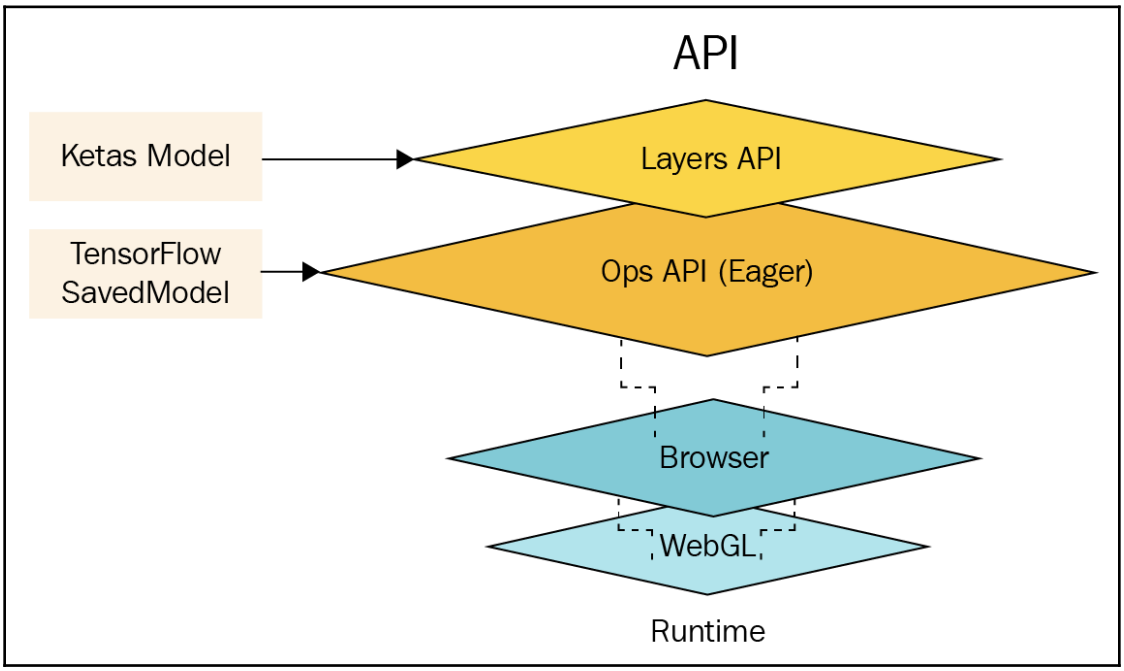

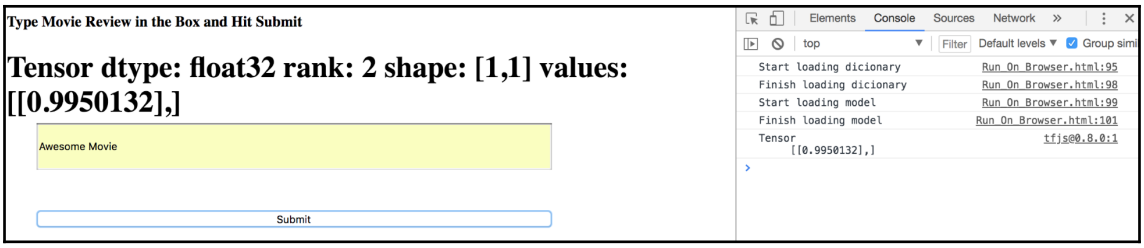
Type Movie Review in the Box and Hit Submit
Tensor dtype: float32 rank: 2 shape: [1,1] values:
[[0.9950132],]
Awesome Movie
Submit
Elements
Console
Sources
Network
top
Filter
Default levels
Group simi
Start loading dicionary
Run_On_Browser.html:95
Finish loading dicionary
Run_On_Browser.html:98
Start loading model
Run_On_Browser.html:99
Finish loading model
Run_On_Browser.html:101
Tensor
[[0.9950132],]
tfjs@0.8.0:1

Chapter 4: Digit Classification Using TensorFlow Lite

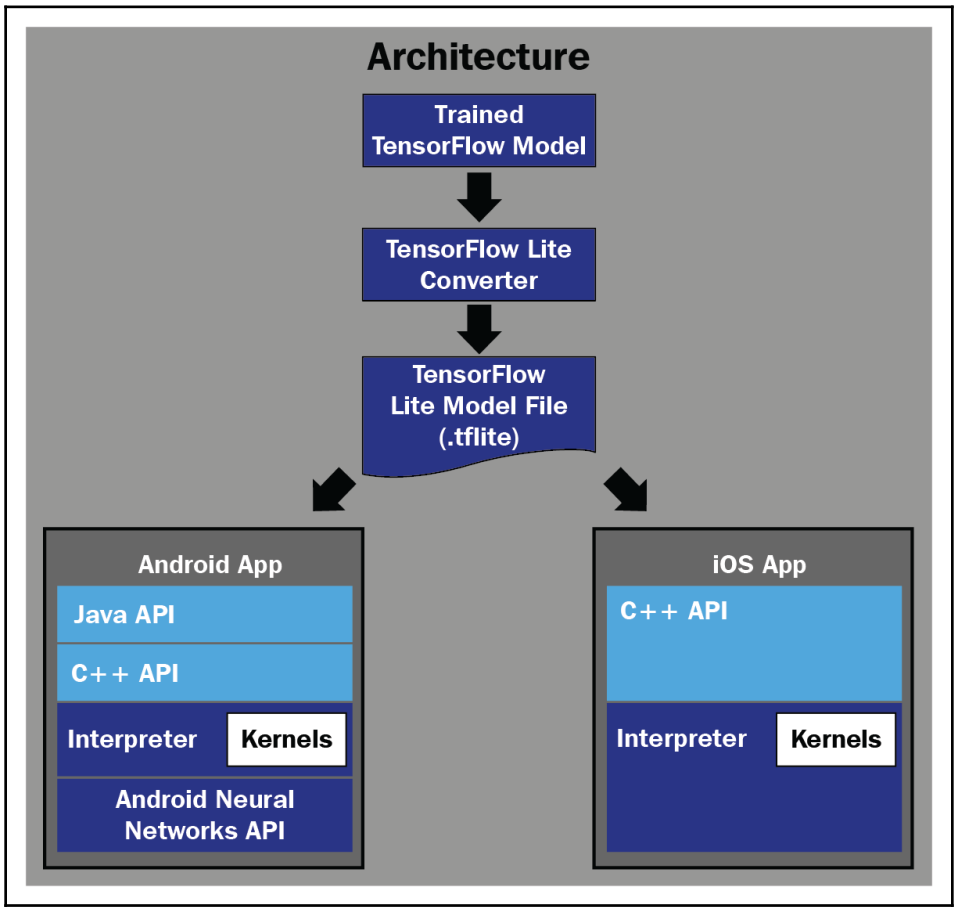

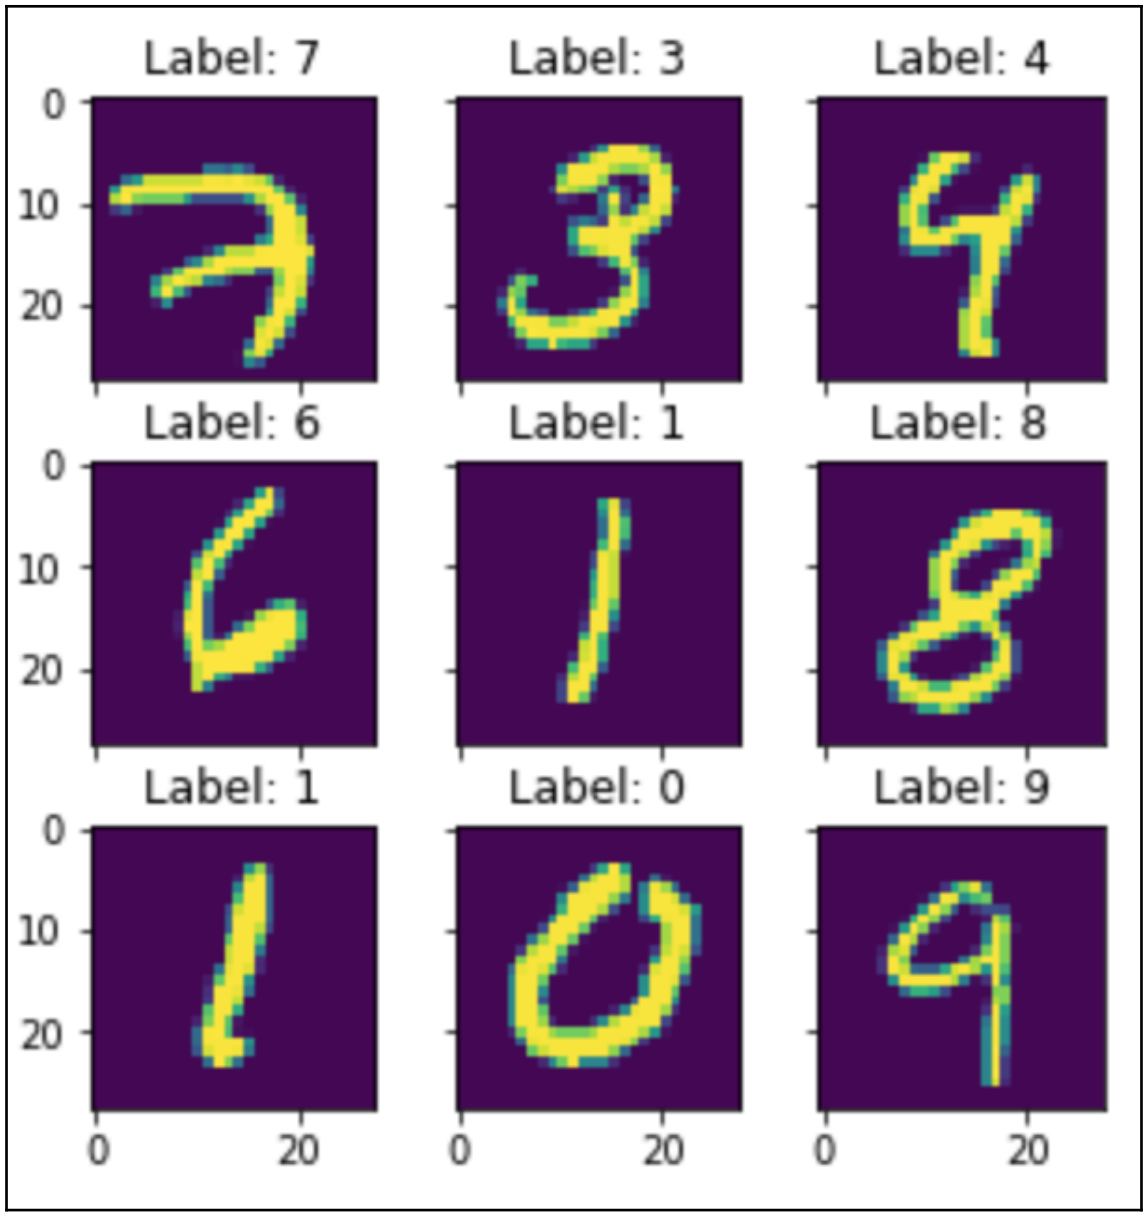
Label: 7
Label: 3
Label: 4
Label: 6
Label: 1
Label: 8
Label: 1
Label: 0
Label: 9
0
10
20
0
20

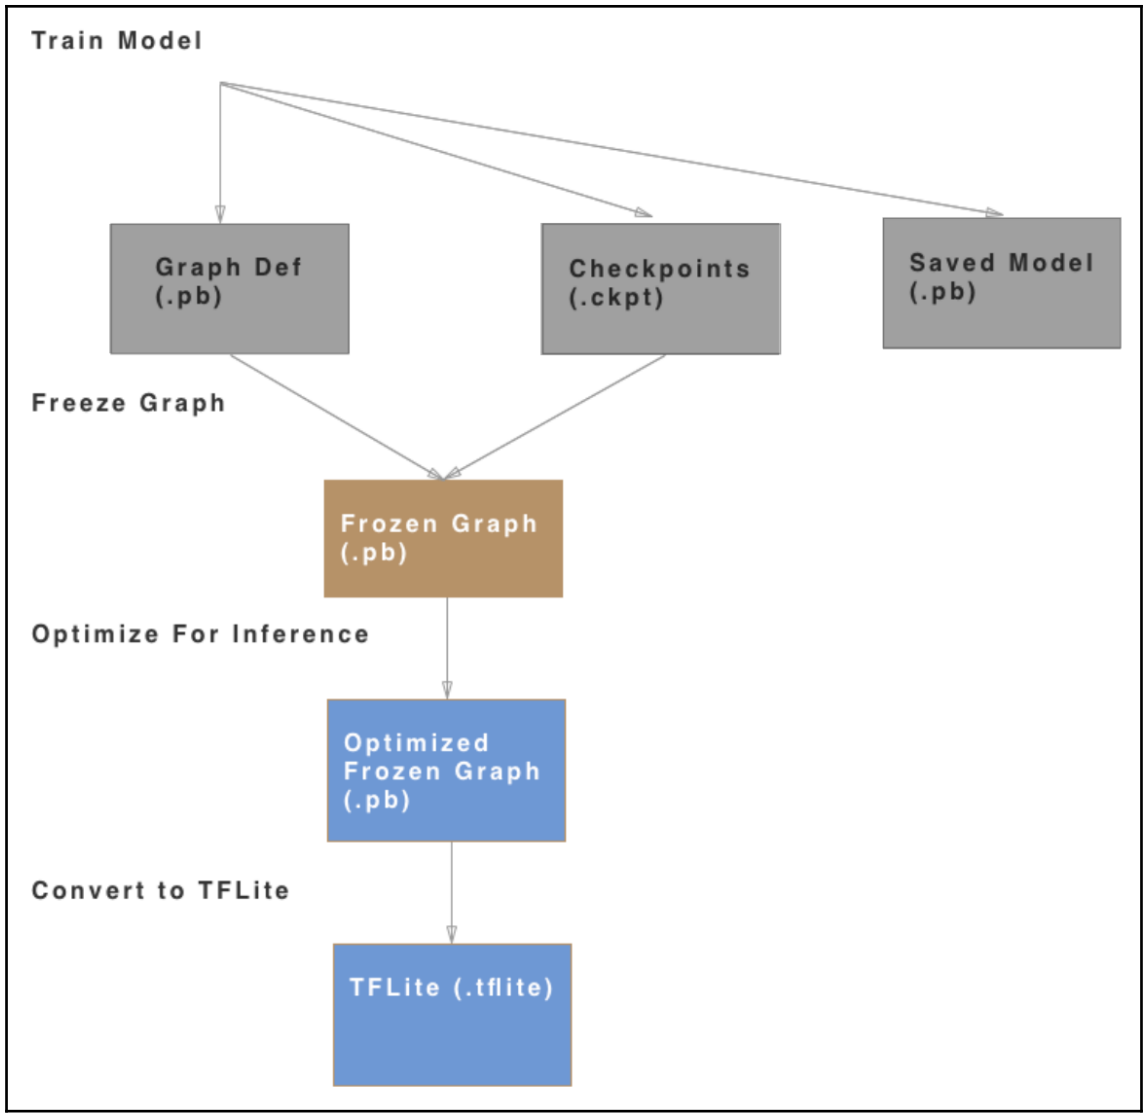
Train Model
Graph Def
(.pb)
Checkpoints
(.ckpt)
Saved Model
(.pb)
Freeze Graph
Frozen Graph
(.pb)
Optimize For Inference
Optimized
Frozen Graph
(.pb)
Convert to TFLite
TFLite (.tflite)

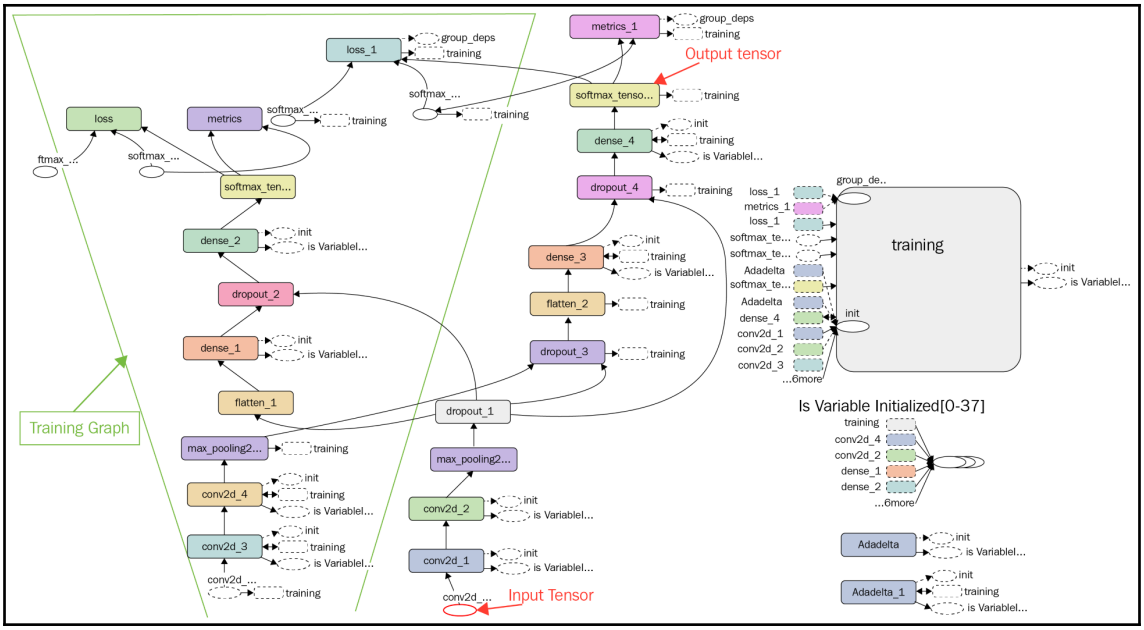
loss_1
group_deps
training
metrics_1
group_deps
training
Output tensor
softmax_...
training
loss
metrics
softmax_...
training
softmax_tenso...
training
ftmax_...
softmax_...
softmax_ten...
dense_4
init
training
is Variable...
dropout_4
training
dense_2
init
is Variable...
dense_3
init
training
is Variable...
dropout_2
flatten_2
training
dense_1
init
is Variable...
dropout_3
training
flatten_1
dropout_1
Training Graph
max_pooling2...
training
max_pooling2...
conv2d_4
init
training
is Variable...
conv2d_2
init
is Variable...
conv2d_3
init
training
is Variable...
conv2d_1
init
is Variable...
conv2d_...
training
conv2d_...
Input Tensor
group_de..
loss_1
metrics_1
loss_1
softmax_te...
softmax_te...
Adadelta
softmax_te...
Adadelta
dense_4
conv2d_1
conv2d_2
conv2d_3
...6more
training
init
init
is Variable...
Is Variable Initialized[0-37]
training
conv2d_4
conv2d_2
dense_1
dense_2
...6more
Adadelta
init
is Variable...
Adadelta_1
init
training
is Variable...

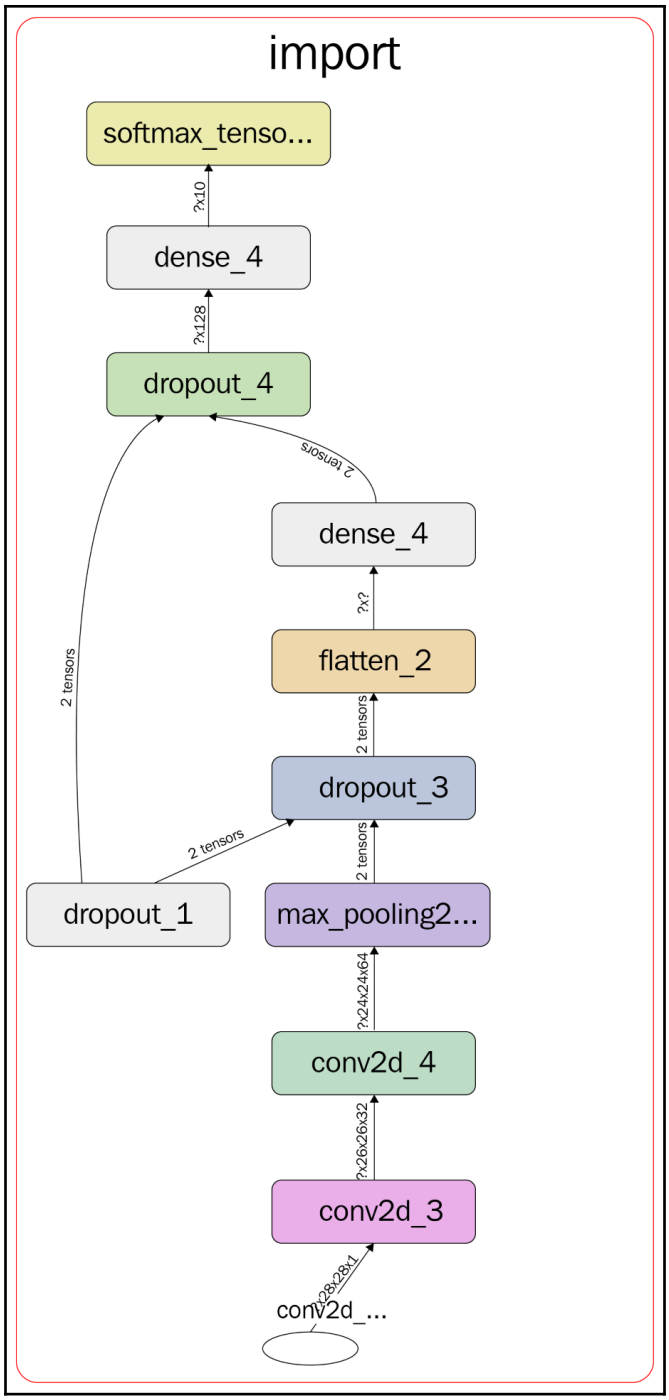
import
softmax_tenso...
?x10
dense_4
?x128
dropout_4
2 tensors
dense_4
?x?
flatten_2
2 tensors
dropout_3
2 tensors
2 tensors
dropout_1
max_pooling2...
?x24x24x64
conv2d_4
?x26x26x32
conv2d_3
?x28x28x1
conv2d_...

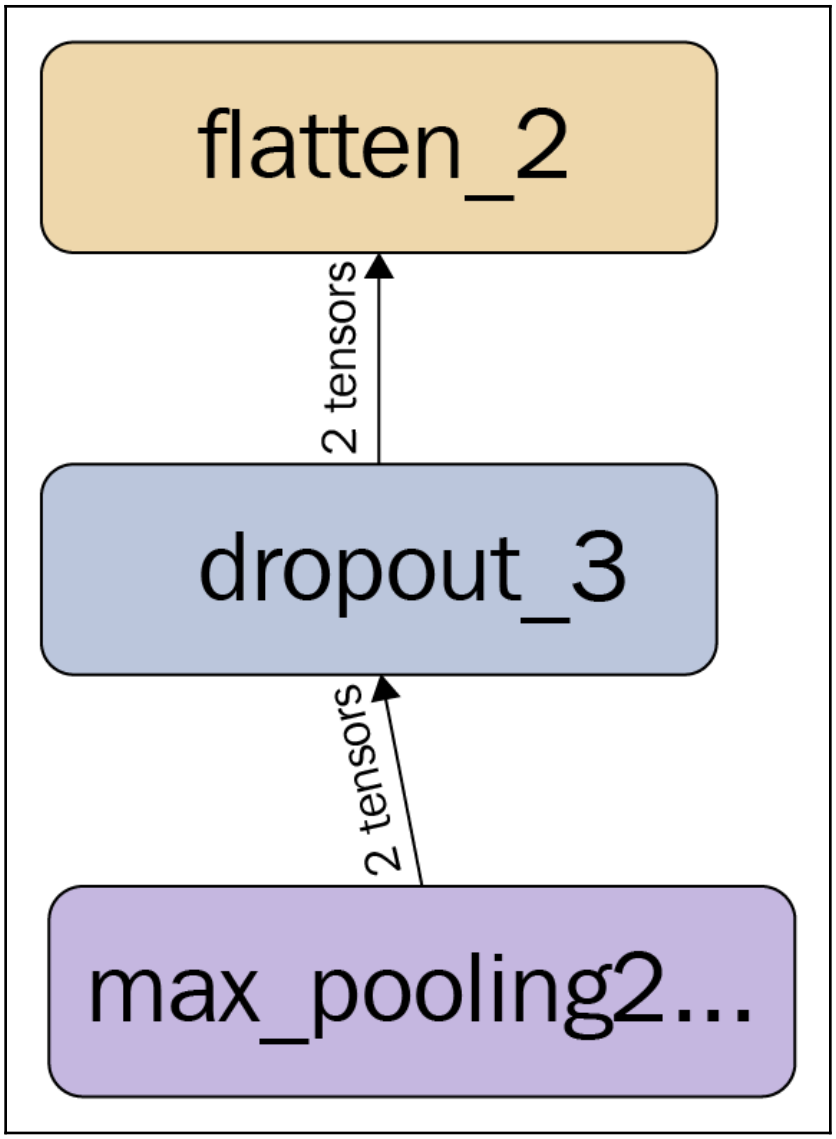
flatten_2
2 tensors
dropout_3
2 tensors
max_pooling2...

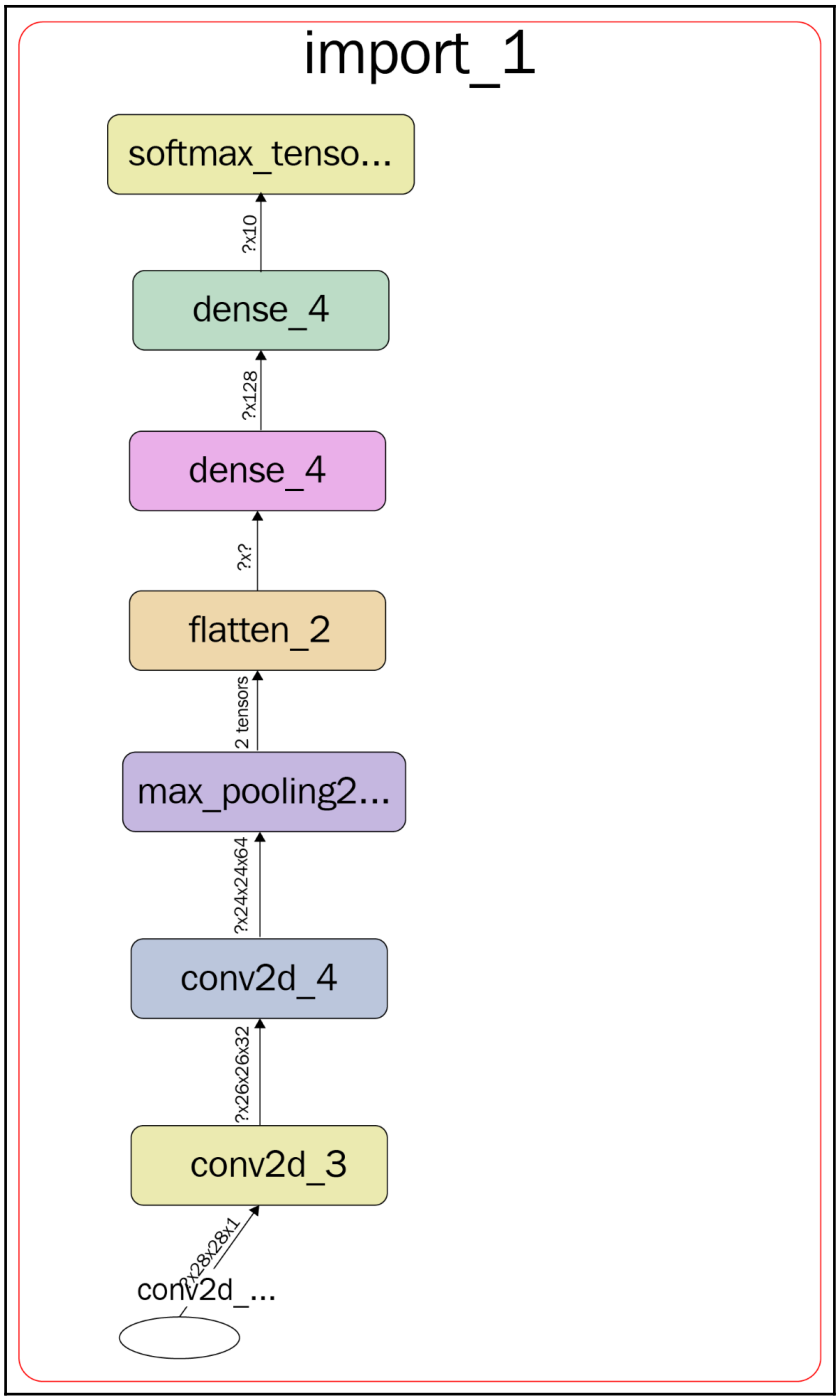
import_1
softmax_tenso...
?x10
dense_4
?x128
dense_4
?x?
flatten_2
2 tensors
max_pooling2...
?x24x24x64
conv2d_4
?x26x26x32
conv2d_3
?x28x28x1
conv2d_...

Chapter 6: Predicting Stock Prices using Gaussian Process Regression

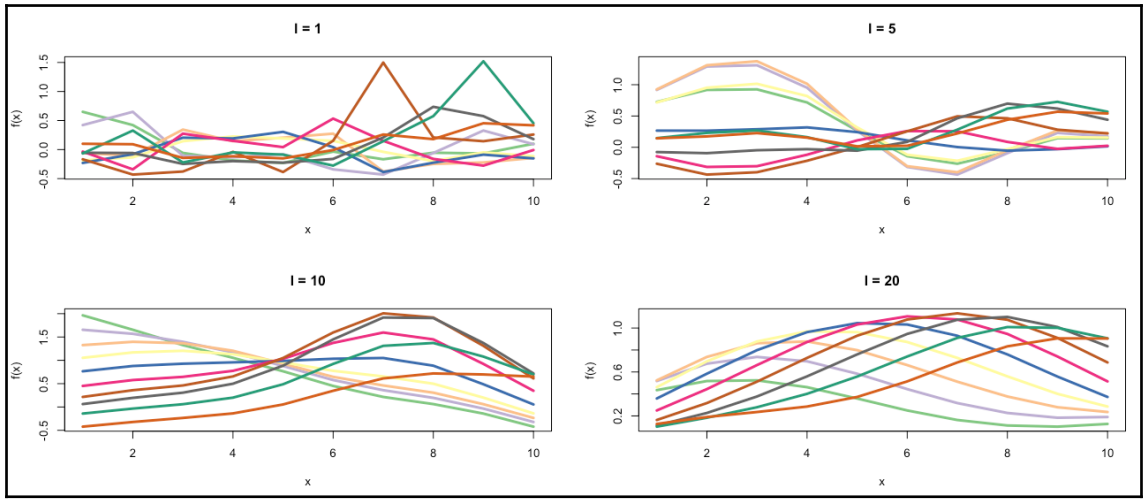

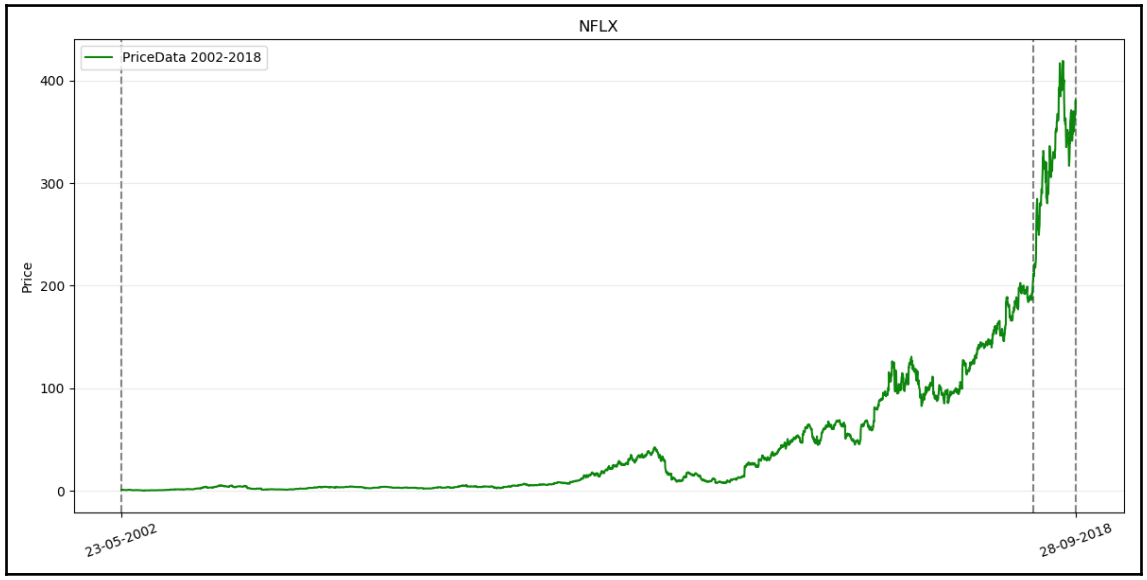

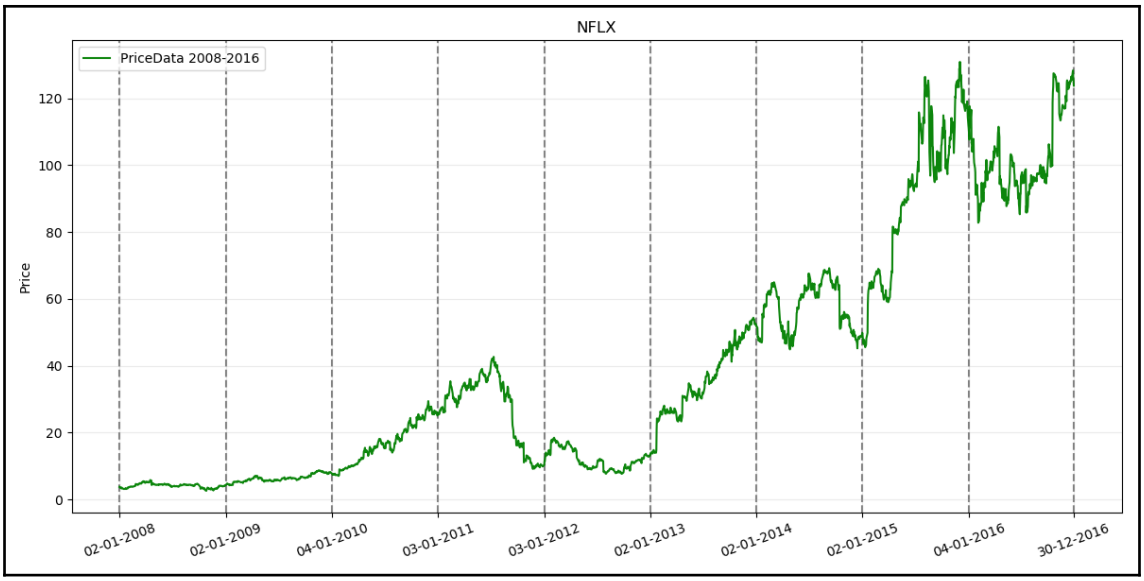
NFLX
PriceData 2008-2016
Price
120
100
80
60
40
20
0
02-01-2008
02-01-2009
04-01-2010
03-01-2011
03-01-2012
02-01-2013
02-01-2014
02-01-2015
04-01-2016
30-12-2016

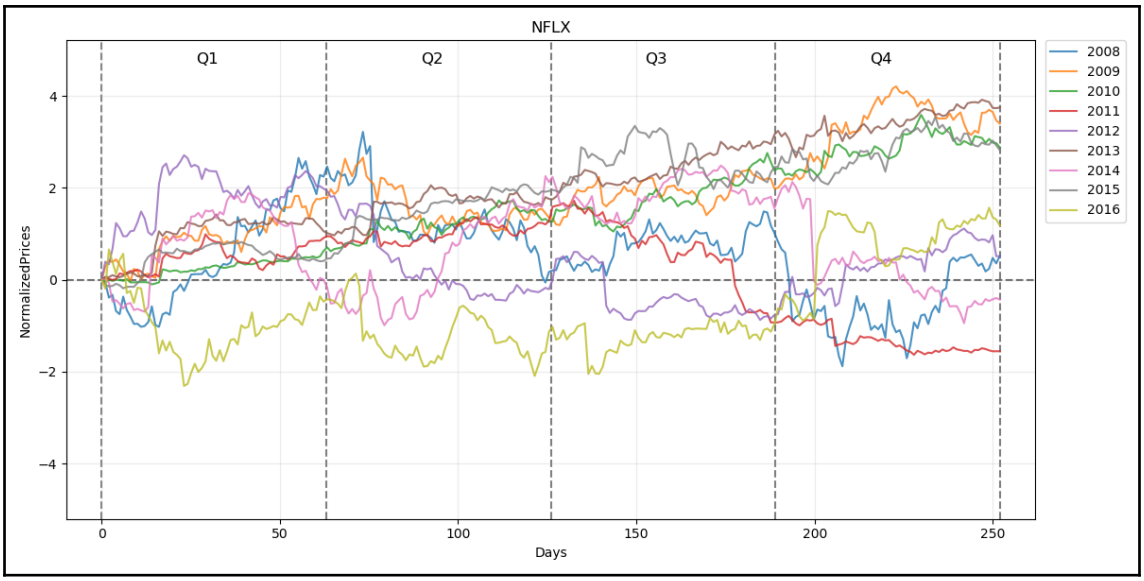
NFLX
Q1
Q2
Q3
Q4
NormalizedPrices
4
2
0
−2
−4
0
50
100
150
200
250
Days
2008
2009
2010
2011
2012
2013
2014
2015
2016

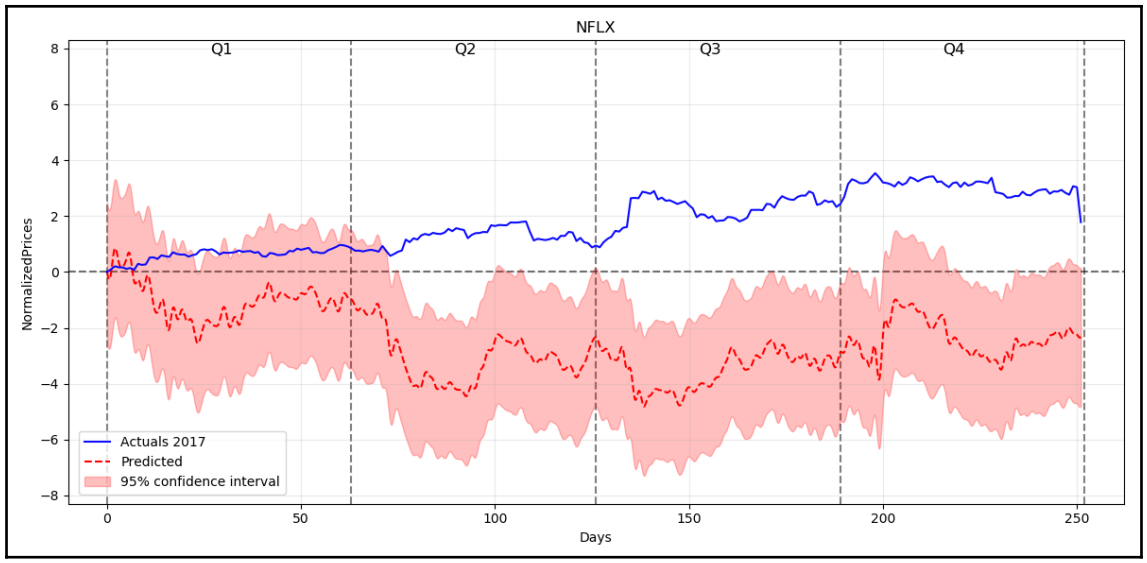
NFLX
Q1
Q2
Q3
Q4
NormalizedPrices
Days
Actuals 2017
Predicted
95% confidence interval

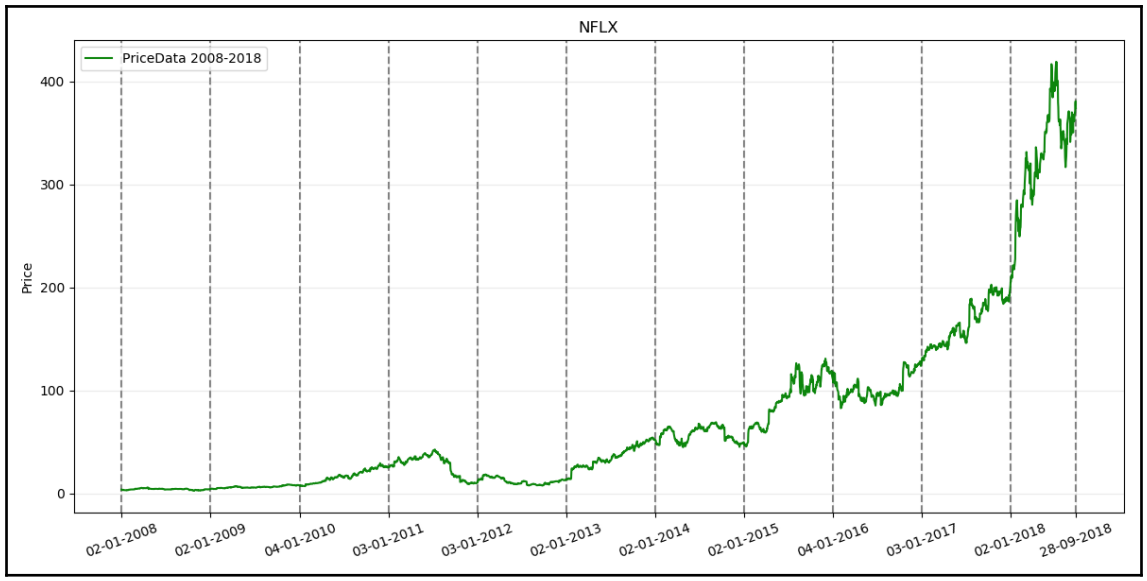
NFLX
PriceData 2008-2018
Price
02-01-2008
02-01-2009
04-01-2010
03-01-2011
03-01-2012
02-01-2013
02-01-2014
02-01-2015
04-01-2016
03-01-2017
02-01-2018
28-09-2018

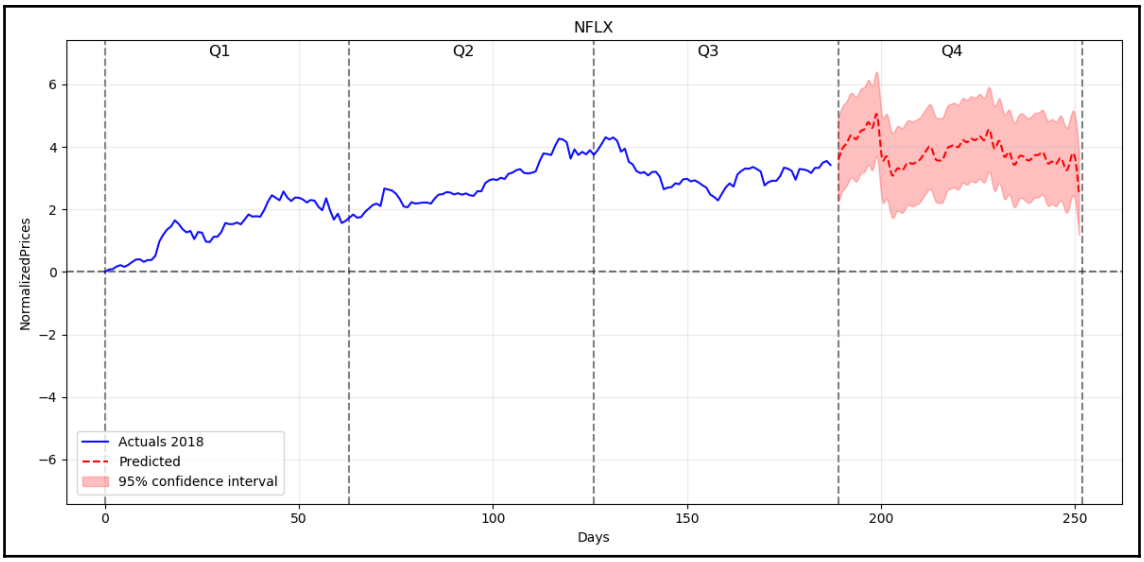
NFLX
Q1
Q2
Q3
Q4
NormalizedPrices
Days
Actuals 2018
Predicted
95% confidence interval


GE
PriceData 1992-2018
Price
30-09-1992
28-09-2018

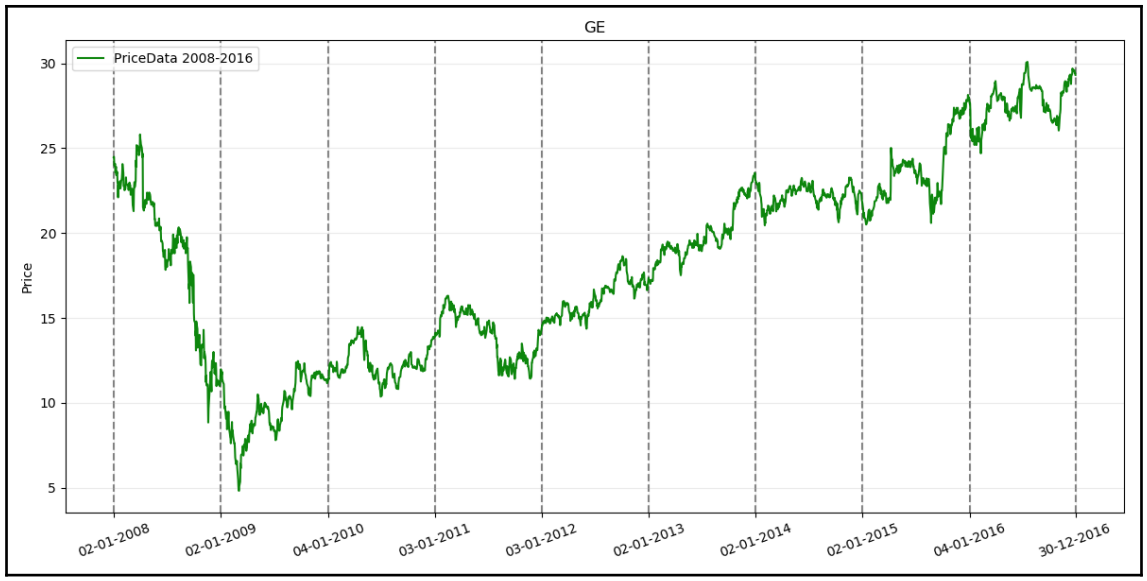
GE
PriceData 2008-2016
Price
30
25
20
15
10
5
02-01-2008
02-01-2009
04-01-2010
03-01-2011
03-01-2012
02-01-2013
02-01-2014
02-01-2015
04-01-2016
30-12-2016

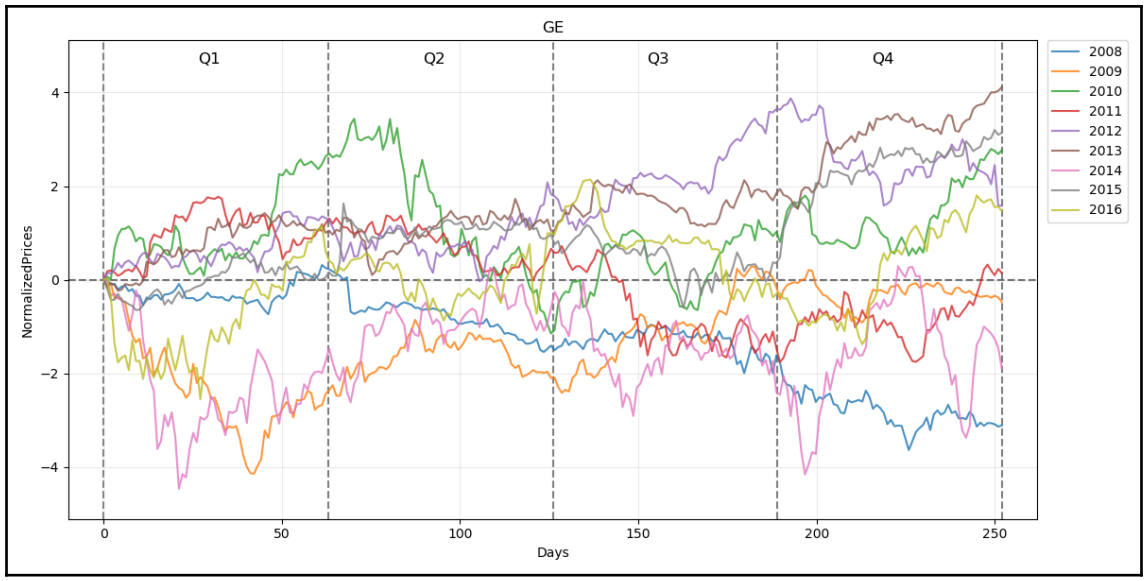
GE
Q1
Q2
Q3
Q4
NormalizedPrices
4
2
0
−2
−4
0
50
100
150
200
250
Days
2008
2009
2010
2011
2012
2013
2014
2015
2016

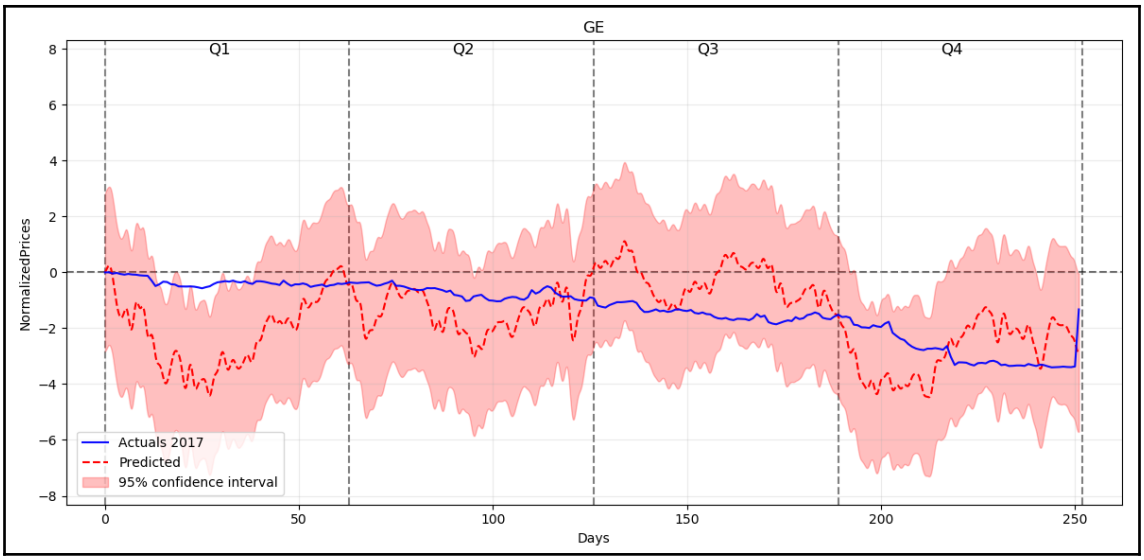
GE
Q1
Q2
Q3
Q4
NormalizedPrices
Days
Actuals 2017
Predicted
95% confidence interval

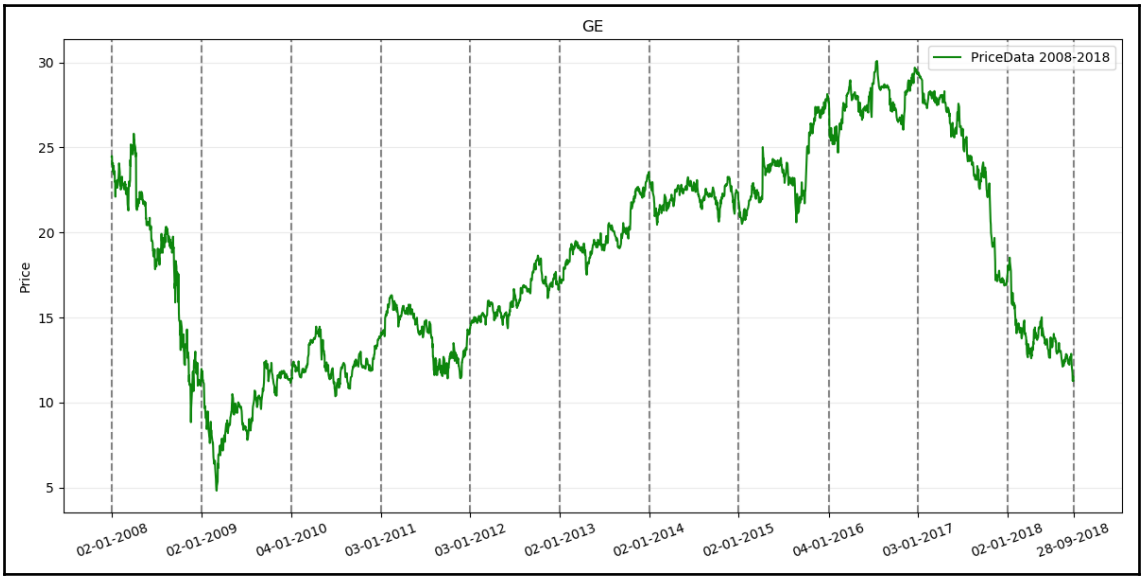
GE
PriceData 2008-2018
Price
02-01-2008
02-01-2009
04-01-2010
03-01-2011
03-01-2012
02-01-2013
02-01-2014
02-01-2015
04-01-2016
03-01-2017
02-01-2018
28-09-2018

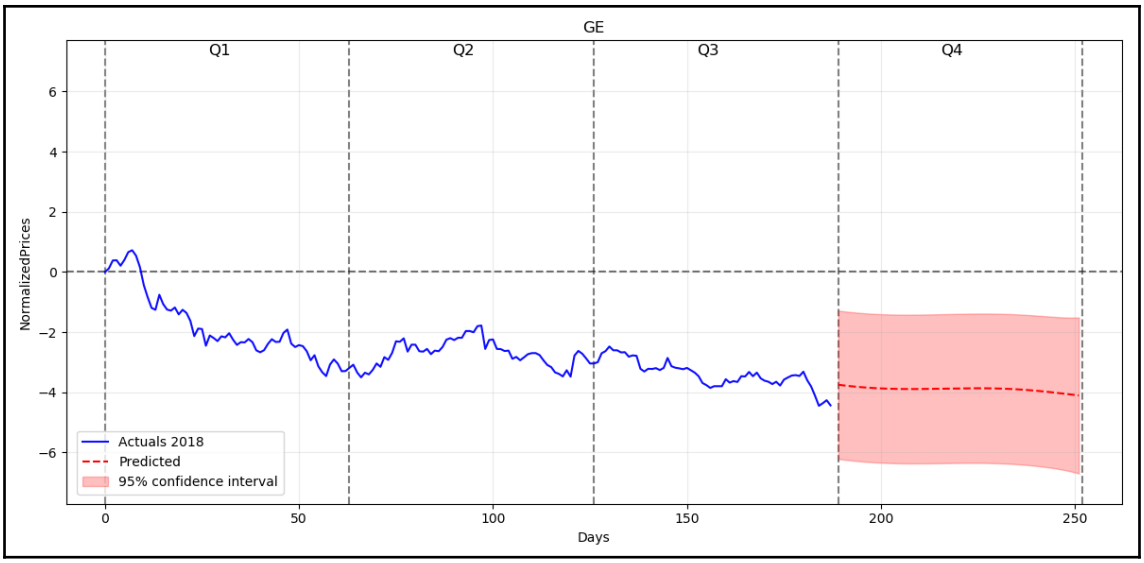
GE
Q1
Q2
Q3
Q4
NormalizedPrices
Days
Actuals 2018
Predicted
95% confidence interval

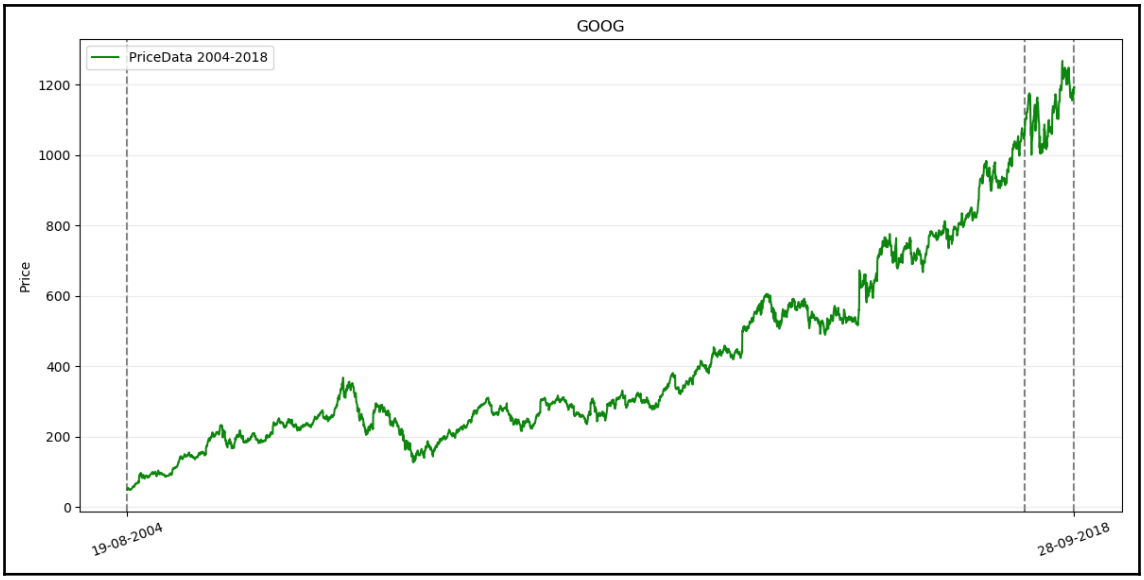
GOOG
PriceData 2004-2018
Price
19-08-2004
28-09-2018


GOOG
PriceData 2008-2016
Price
800
700
600
500
400
300
200
100
02-01-2008
02-01-2009
04-01-2010
03-01-2011
03-01-2012
02-01-2013
02-01-2014
02-01-2015
04-01-2016
30-12-2016

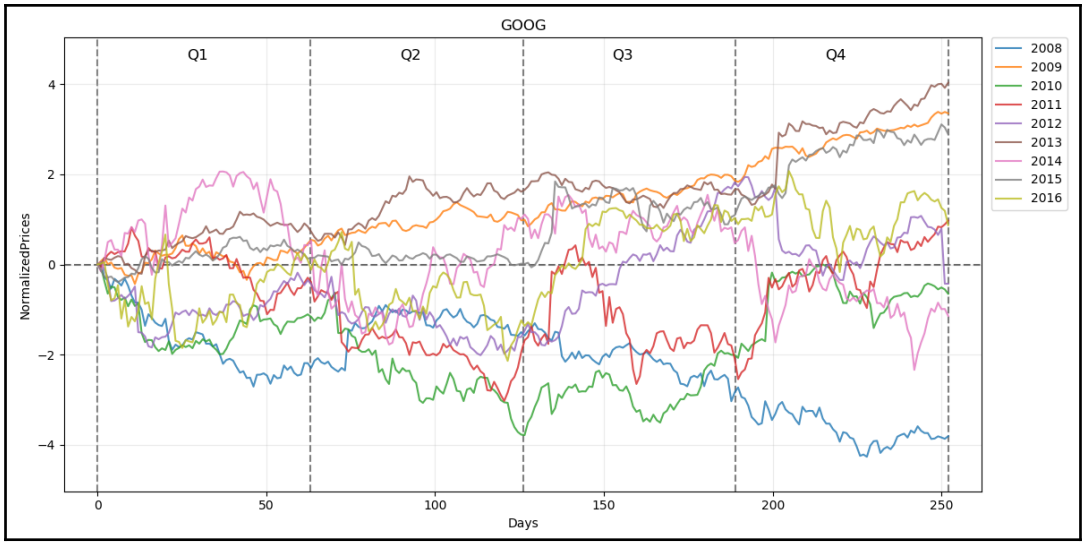
GOOG
Q1
Q2
Q3
Q4
NormalizedPrices
4
2
0
−2
−4
0
50
100
150
200
250
Days
2008
2009
2010
2011
2012
2013
2014
2015
2016

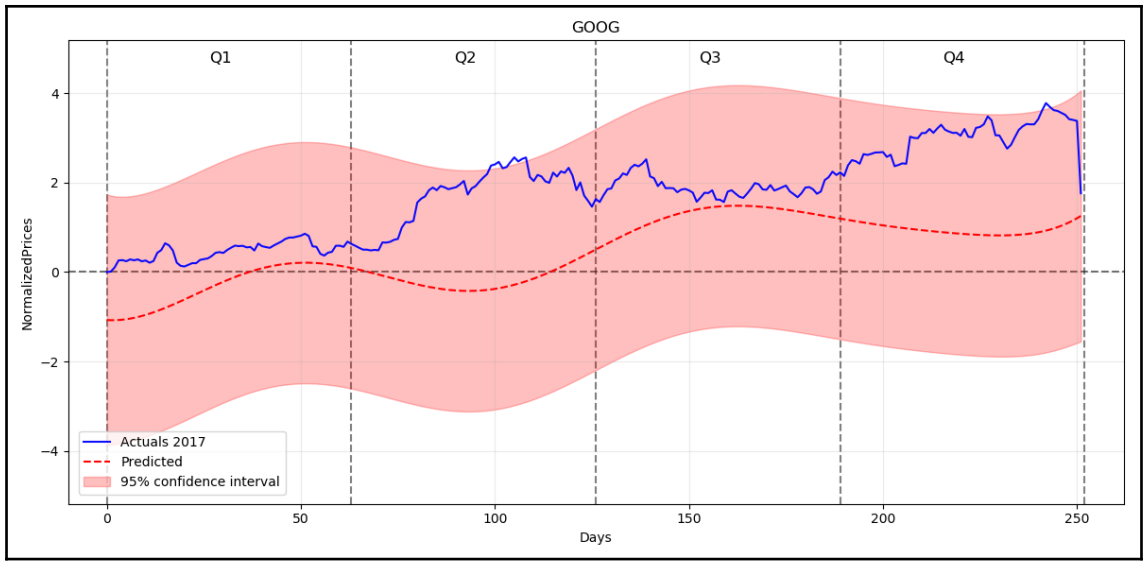
GOOG
Q1
Q2
Q3
Q4
NormalizedPrices
Days
Actuals 2017
Predicted
95% confidence interval

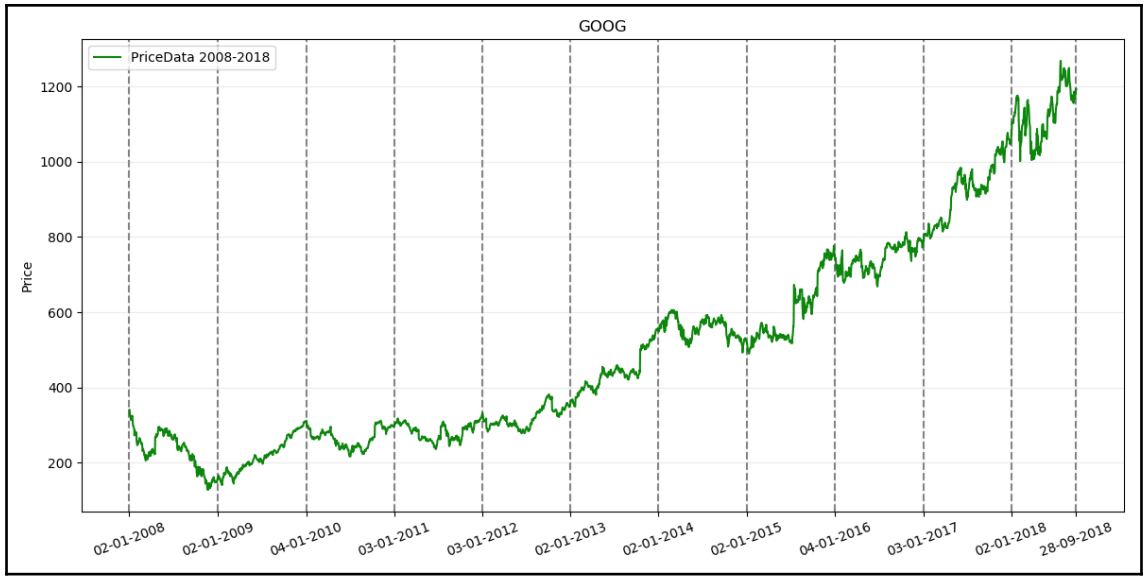
GOOG
PriceData 2008-2018
Price
02-01-2008
02-01-2009
04-01-2010
03-01-2011
03-01-2012
02-01-2013
02-01-2014
02-01-2015
04-01-2016
03-01-2017
02-01-2018
28-09-2018

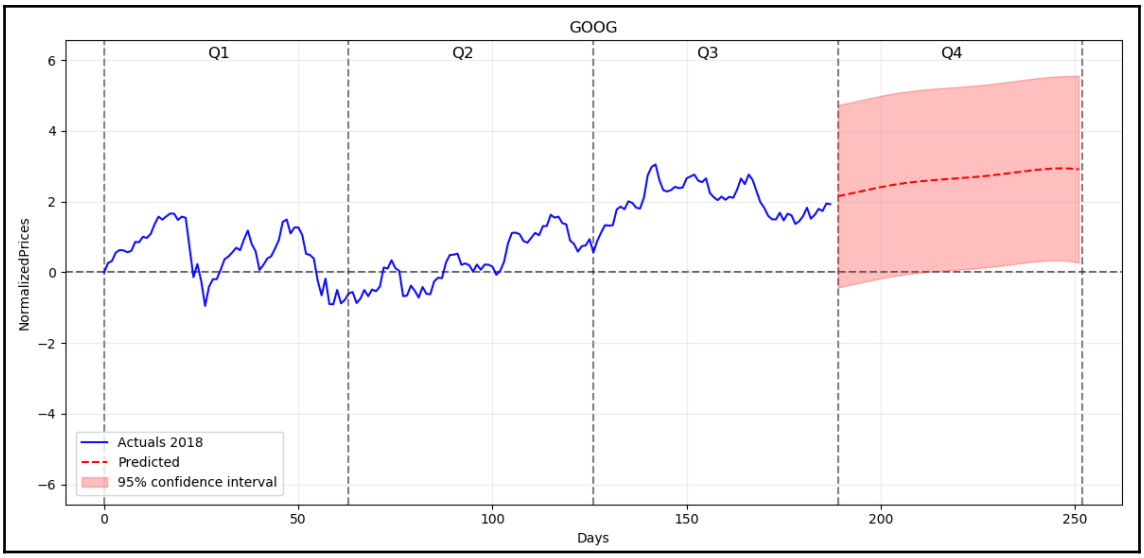
GOOG
Q1
Q2
Q3
Q4
NormalizedPrices
Days
Actuals 2018
Predicted
95% confidence interval

Chapter 7: Credit Card Fraud Detection using Autoencoders

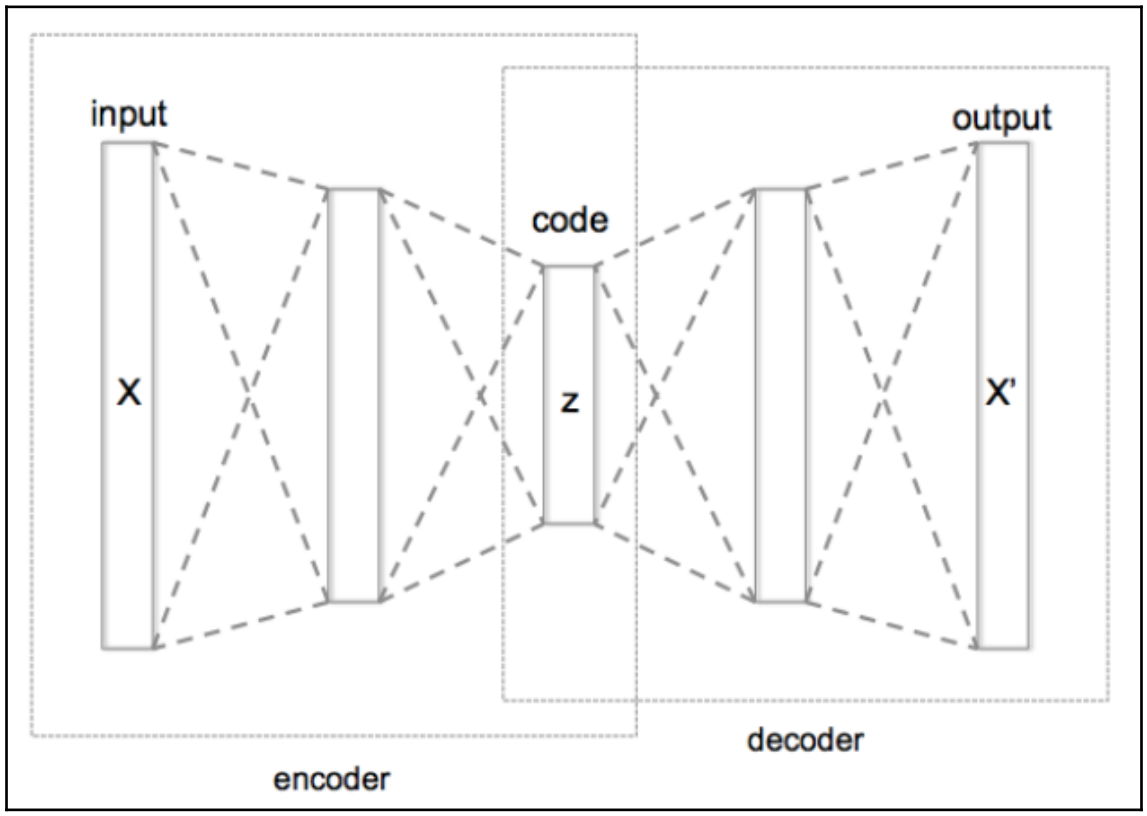

$$\phi : \mathcal{X} \rightarrow \mathcal{F}$$

$$\psi : \mathcal{F} \rightarrow \mathcal{X}$$

$$\phi, \psi = \underset{\phi,\psi}{\arg\min} \, \|X - (\psi \circ \phi)X\|^2$$

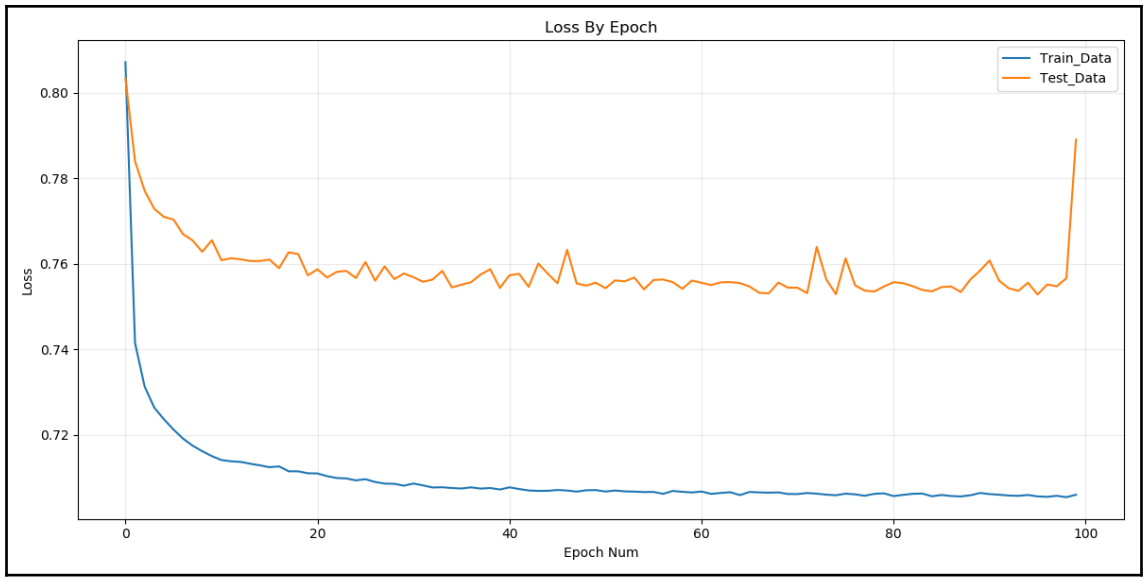
Loss By Epoch
Train_Data
Test_Data
Loss
0.80
0.78
0.76
0.74
0.72
0
20
40
60
80
100
Epoch Num

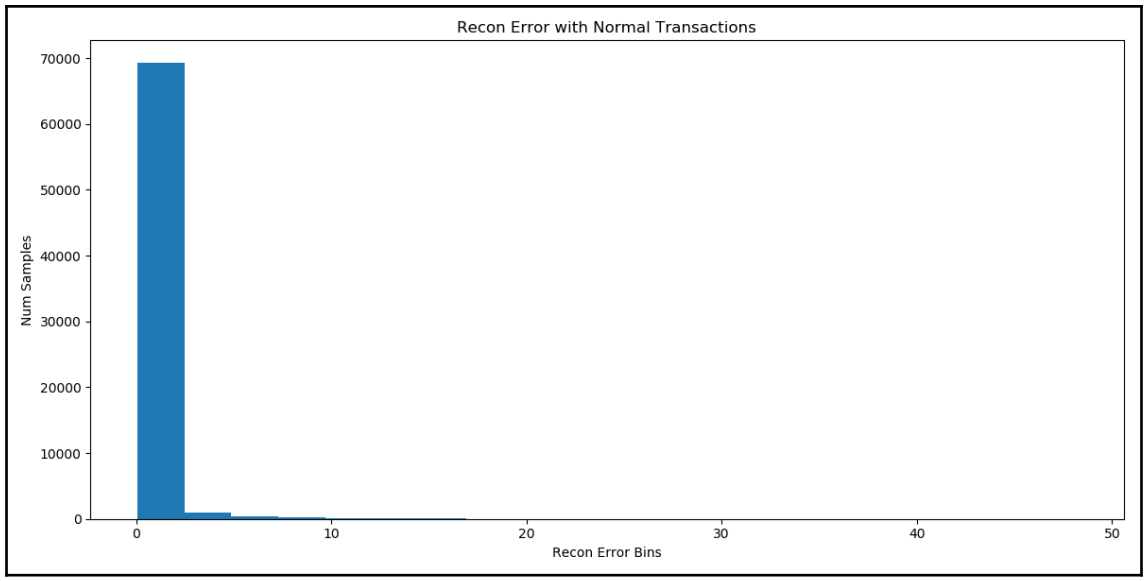
Recon Error with Normal Transactions
Num Samples
70000
60000
50000
40000
30000
20000
10000
0
0
10
20
30
40
50
Recon Error Bins

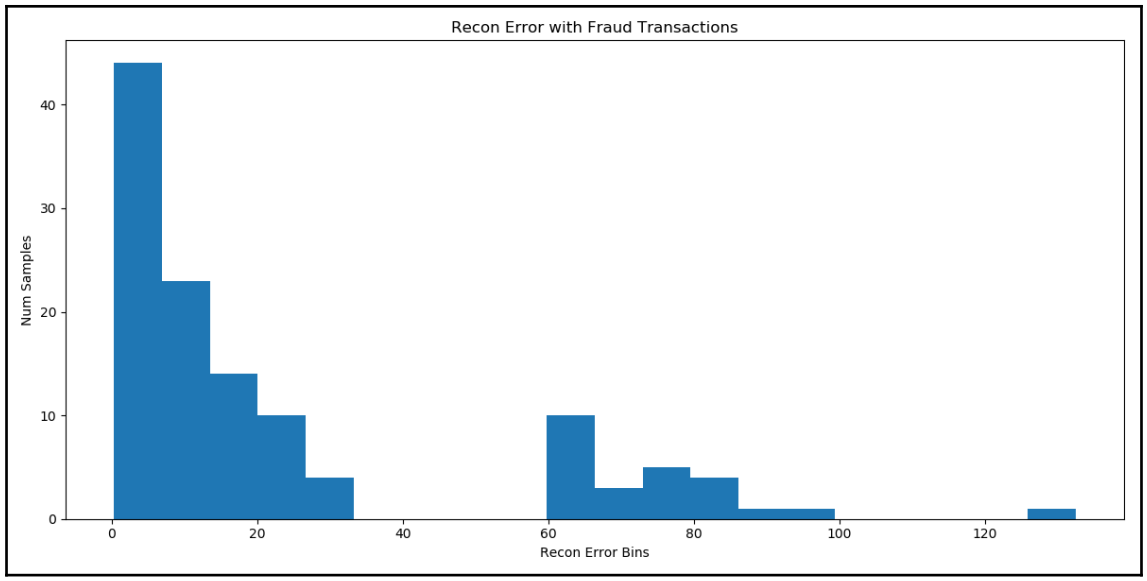
Recon Error with Fraud Transactions
Num Samples
Recon Error Bins
0
10
20
30
40
0
20
40
60
80
100
120

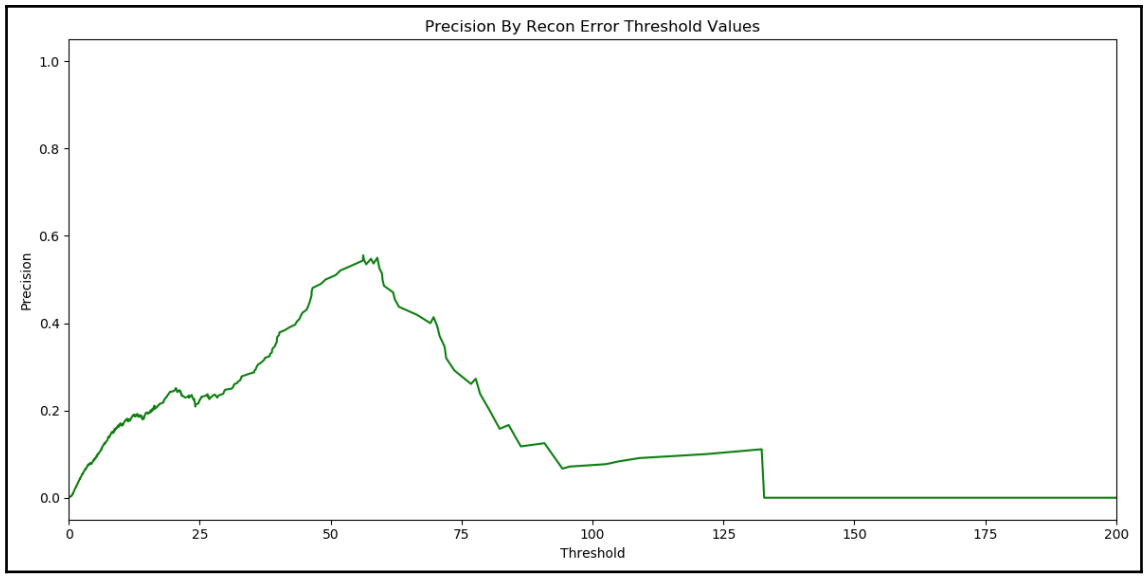
Precision By Recon Error Threshold Values
Precision
Threshold
0.0
0.2
0.4
0.6
0.8
1.0
0
25
50
75
100
125
150
175
200

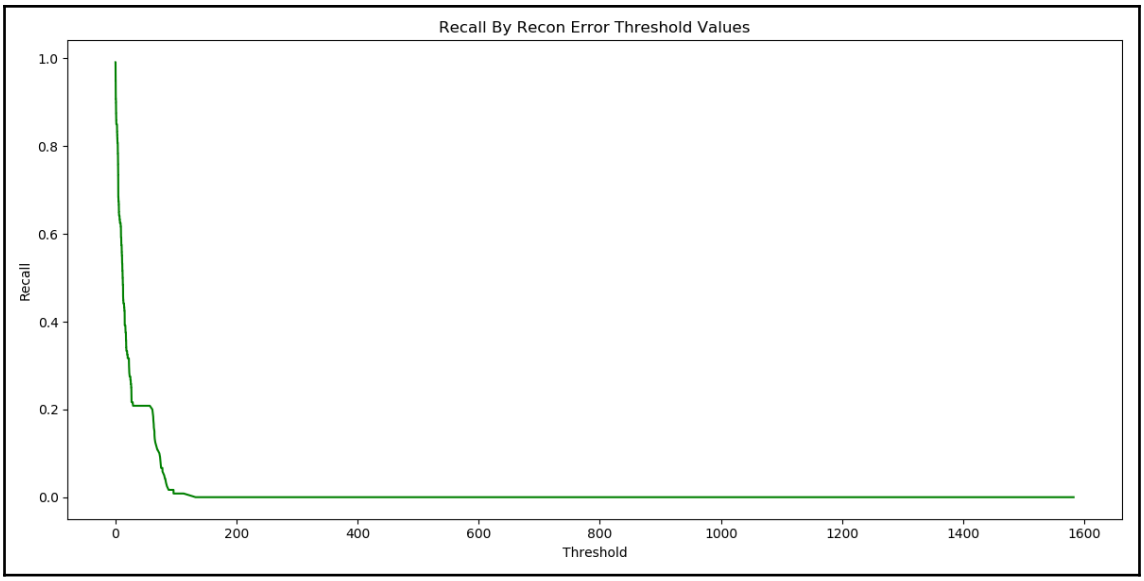
Recall By Recon Error Threshold Values
Recall
1.0
0.8
0.6
0.4
0.2
0.0
0
200
400
600
800
1000
1200
1400
1600
Threshold

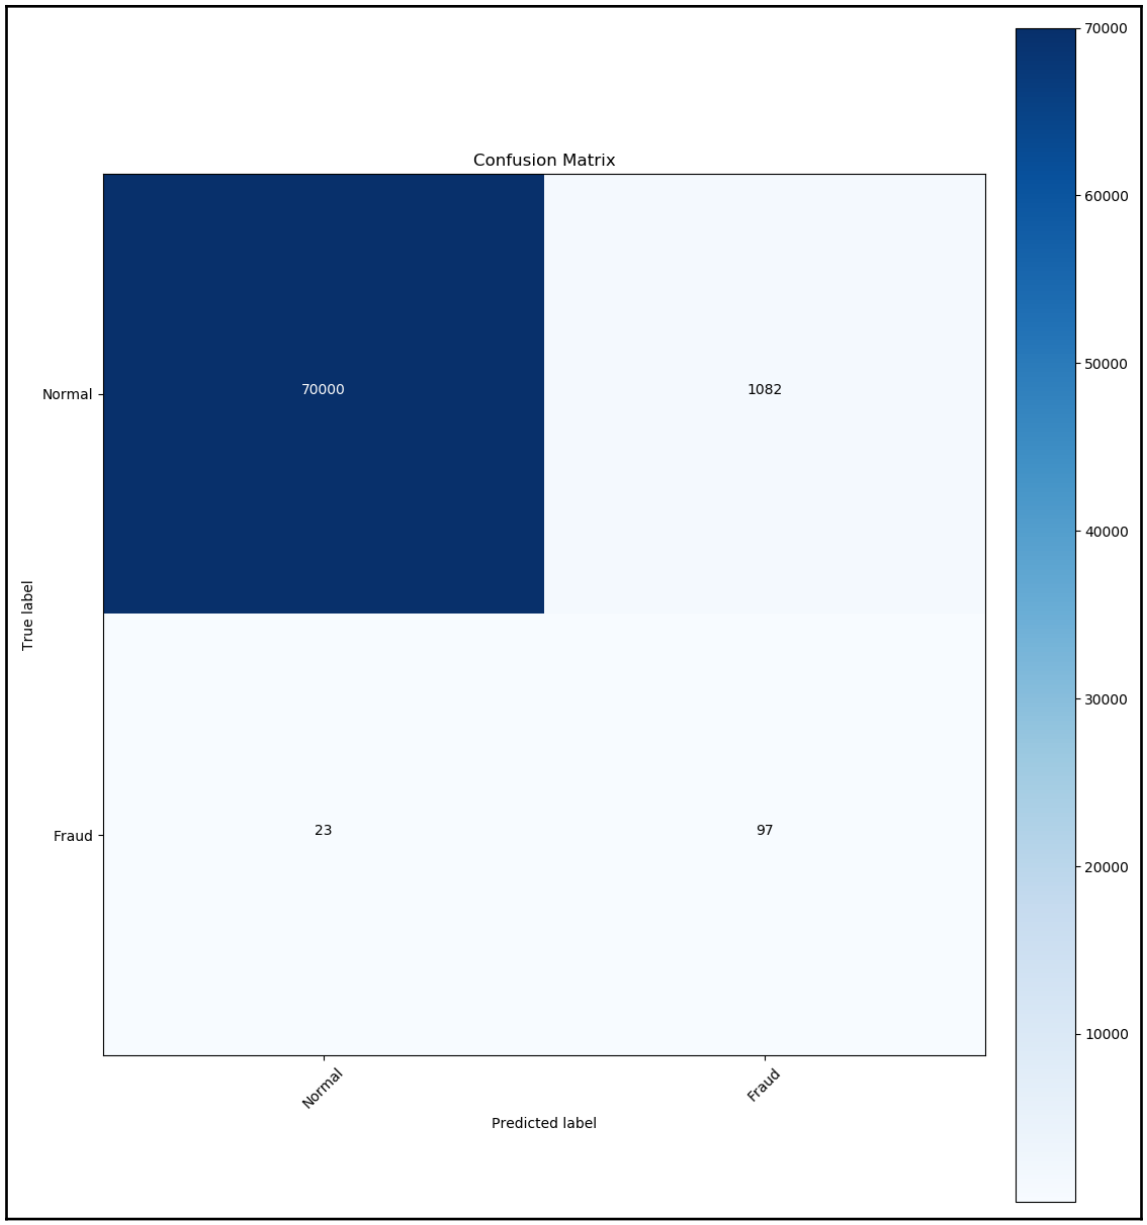
Confusion Matrix
Normal
Fraud
True label
70000
1082
23
97
Normal
Fraud
Predicted label
70000
60000
50000
40000
30000
20000
10000

Chapter 8: Generating Uncertainty in Traffic Signs Classifier Using Bayesian Neural Networks


30
00000_00000.ppm

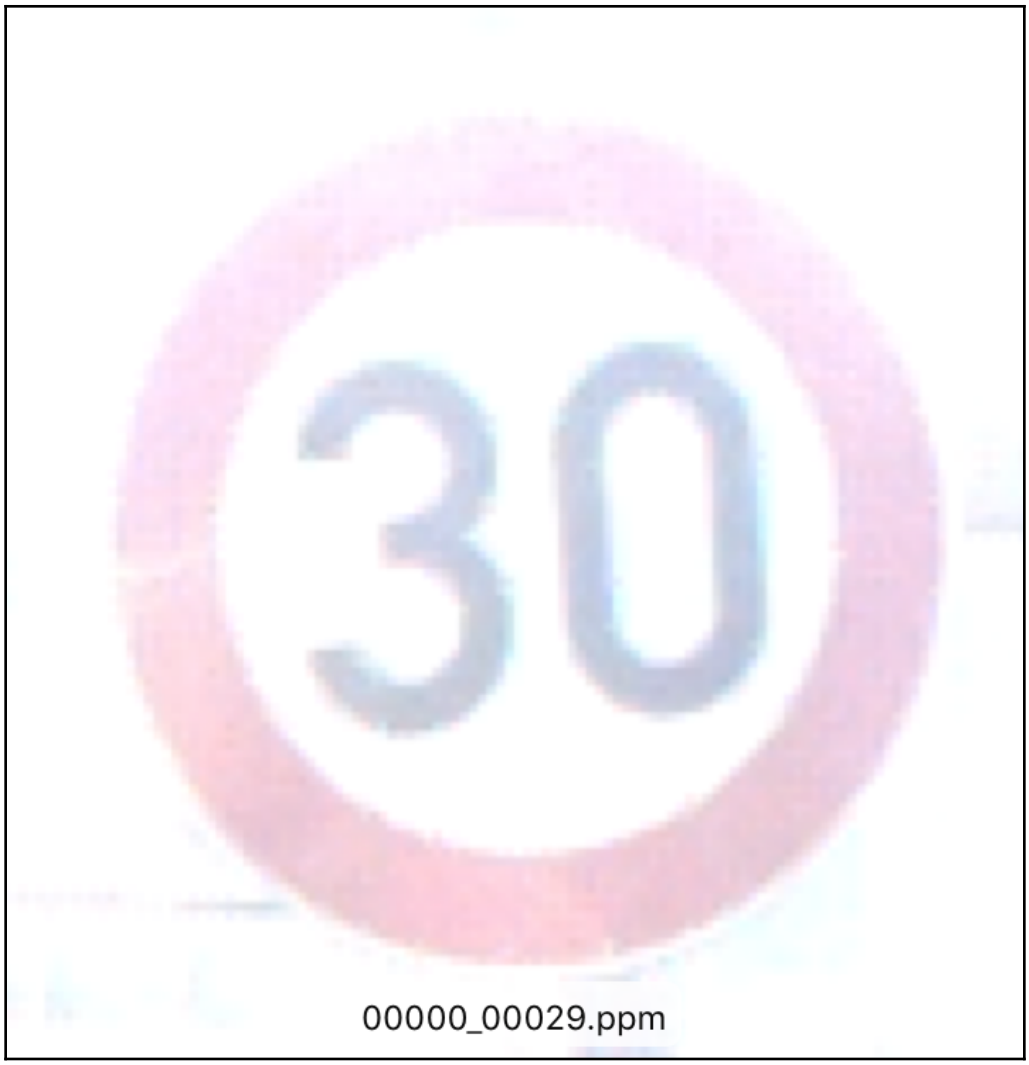
30
00000_00029.ppm

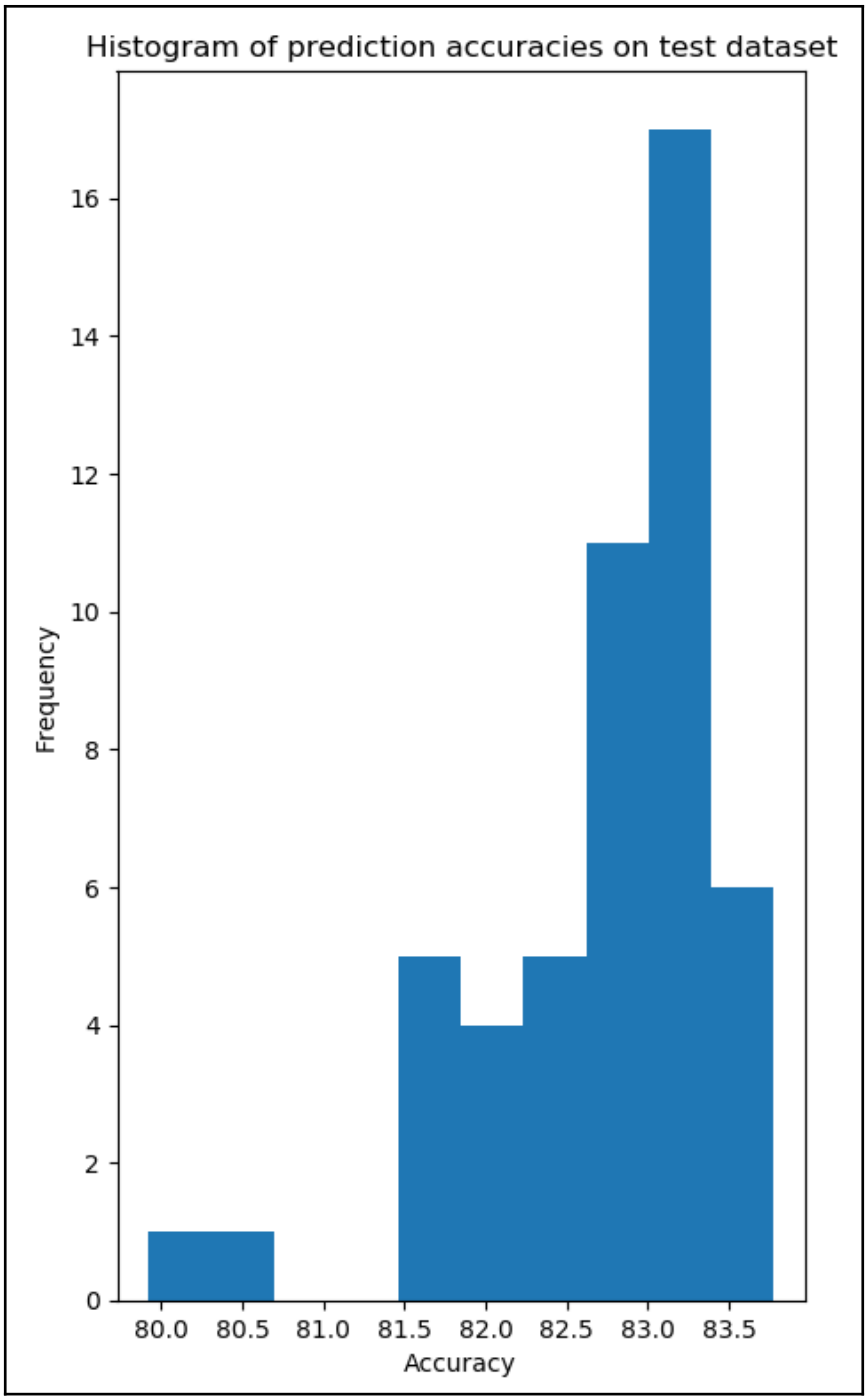
Histogram of prediction accuracies on test dataset
Frequency
0
2
4
6
8
10
12
14
16
80.0
80.5
81.0
81.5
82.0
82.5
83.0
83.5
Accuracy

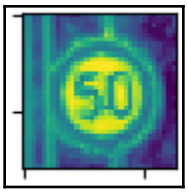

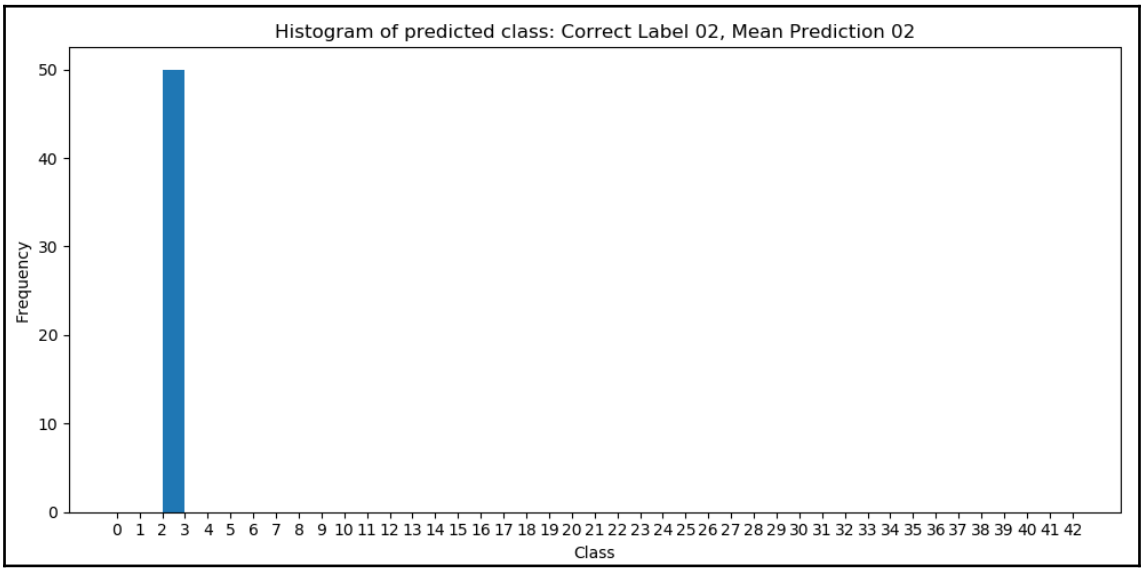
Histogram of predicted class: Correct Label 02, Mean Prediction 02
Frequency
Class
0
10
20
30
40
50
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31 32 33 34 35 36 37 38 39 40 41 42

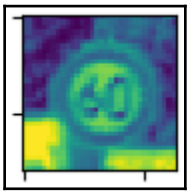

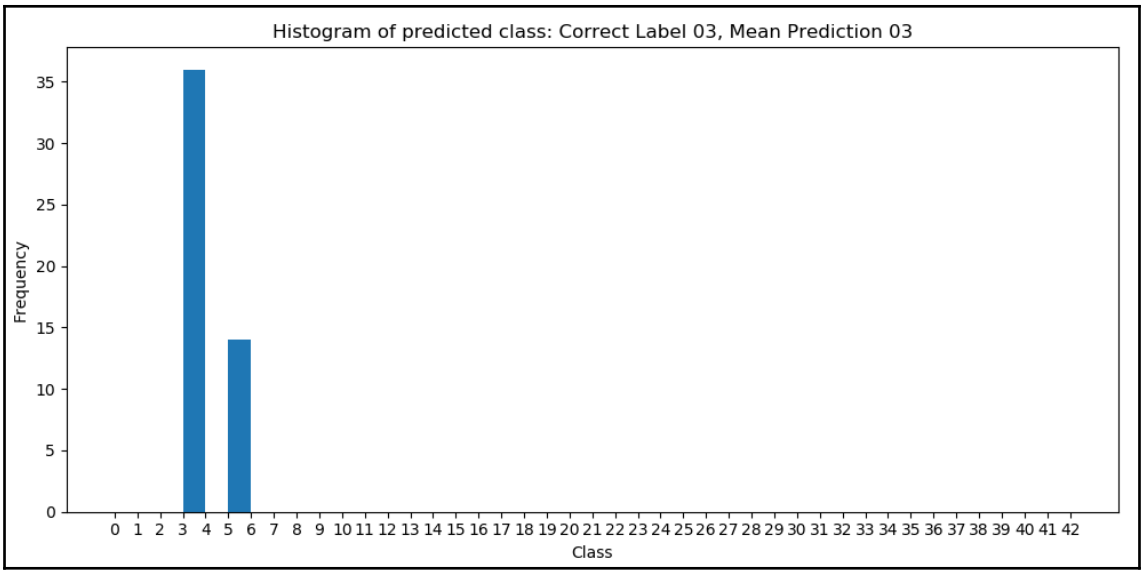
Histogram of predicted class: Correct Label 03, Mean Prediction 03
Frequency
Class

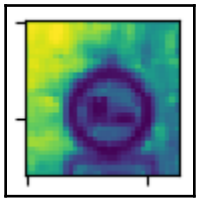

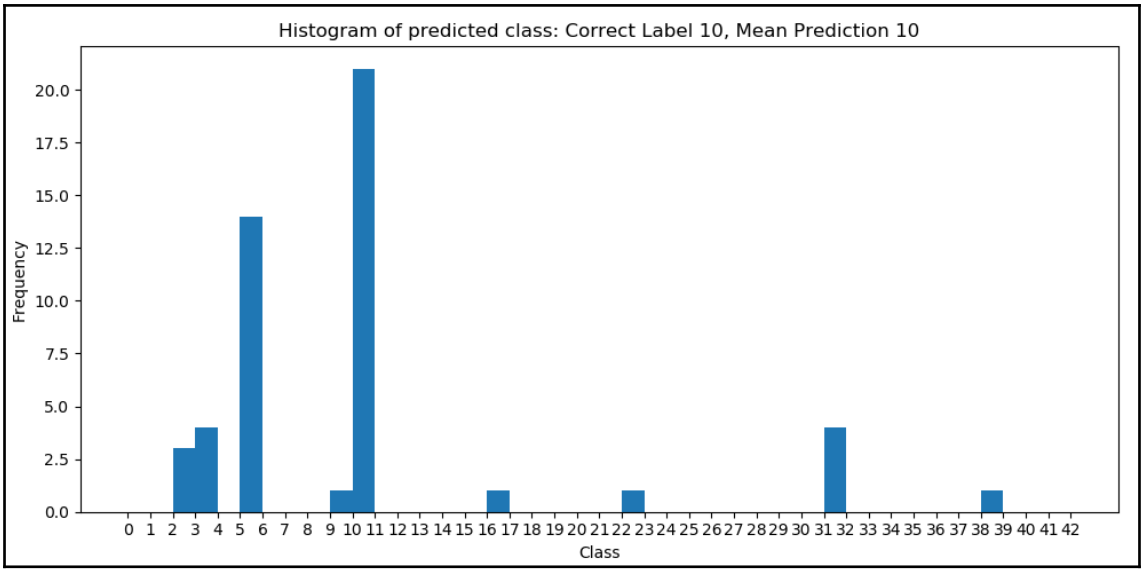
Histogram of predicted class: Correct Label 10, Mean Prediction 10
Frequency
Class

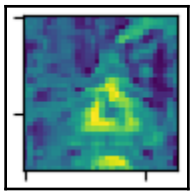

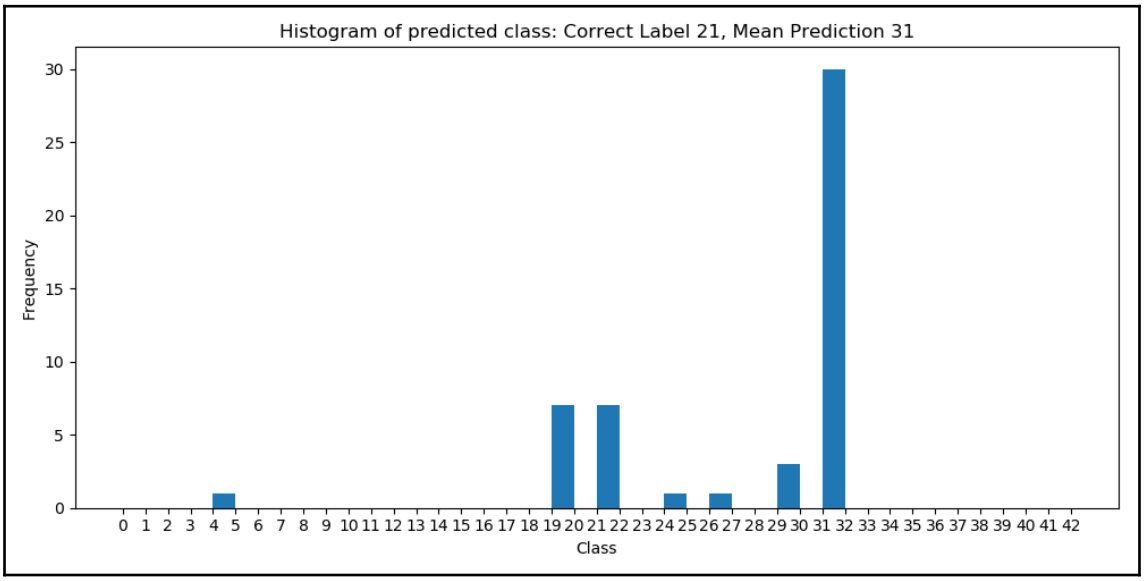
Histogram of predicted class: Correct Label 21, Mean Prediction 31
Frequency
Class
0
5
10
15
20
25
30
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31 32 33 34 35 36 37 38 39 40 41 42

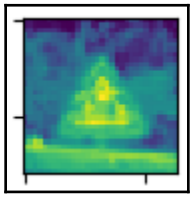

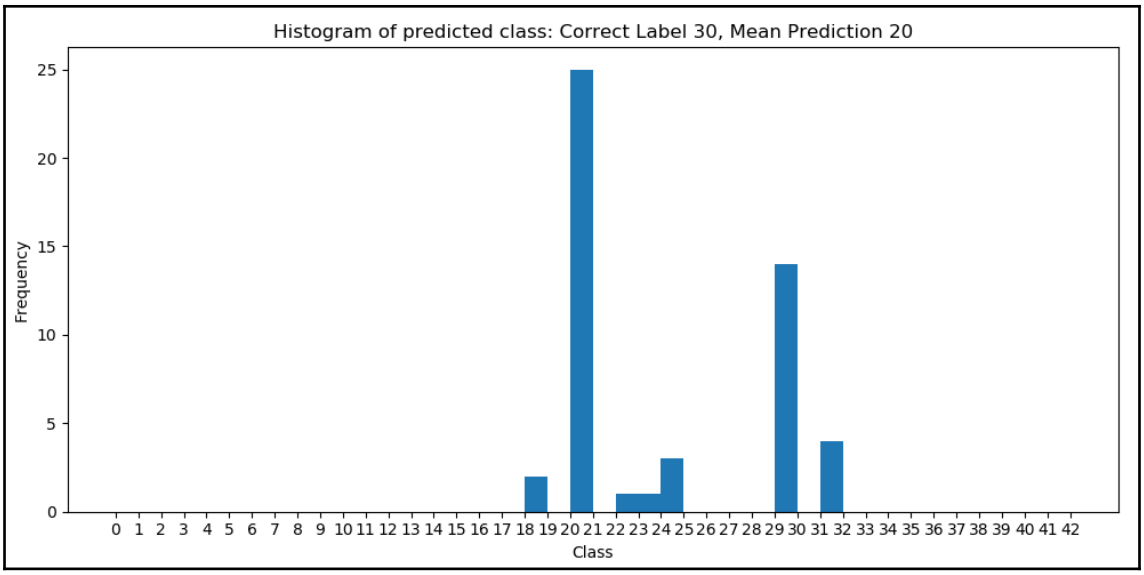
Histogram of predicted class: Correct Label 30, Mean Prediction 20
Frequency
Class
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31 32 33 34 35 36 37 38 39 40 41 42
0 5 10 15 20 25

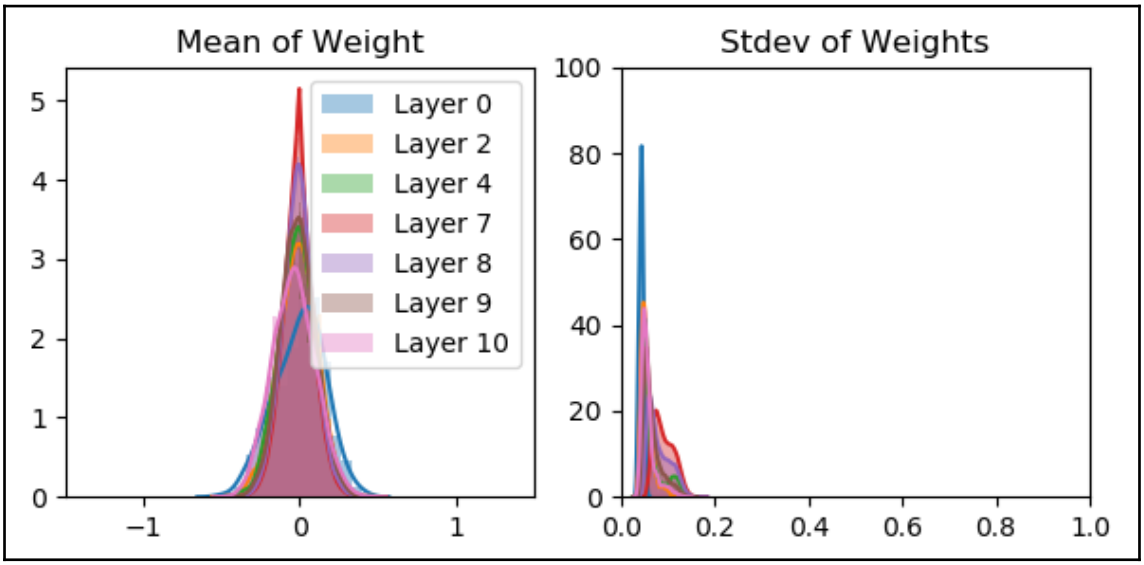
Mean of Weight
Stdev of Weights
Layer 0
Layer 2
Layer 4
Layer 7
Layer 8
Layer 9
Layer 10
0 1 2 3 4 5
−1 0 1
0 20 40 60 80 100
0.0 0.2 0.4 0.6 0.8 1.0

Chapter 9: Generating Matching Shoe Bags from Shoe Images Using DiscoGANs

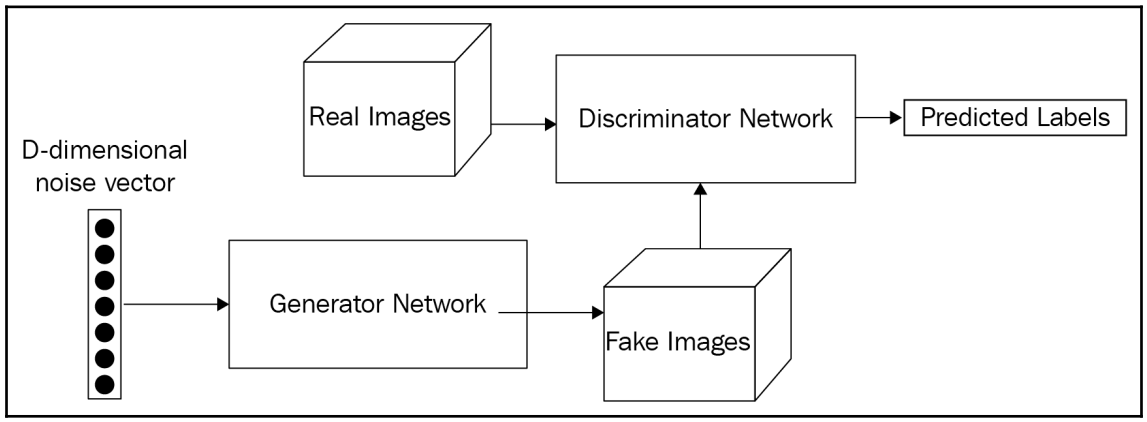

for number of training iterations **do**

for k steps **do**

Minimise D

- Sample minibatch of m noise samples $\{z^{(1)},\ldots,z^{(m)}\}$ from noise prior $p_g(z)$.
- Sample minibatch of m examples $\{x^{(1)},\ldots,x^{(m)}\}$ from data generating distribution $p_{\text{data}}(x)$.
- Update the discriminator by ascending its stochastic gradient:

$$\nabla_{\theta_d}\frac{1}{m}\sum_{i=1}^{m}\left[\log D\left(x^{(i)}\right)+\log\left(1-D\left(G\left(z^{(i)}\right)\right)\right)\right]$$

end for

Maximise G

- Sample minibatch of m noise samples $\{z^{(1)},\ldots,z^{(m)}\}$ from noise prior $p_g(z)$.
- Update the generator by descending its stochastic gradient:

$$\nabla_{\theta_a}\frac{1}{m}\sum_{i=1}^{m}\log\left(1-D\left(G\left(z^{(i)}\right)\right)\right)$$

end for

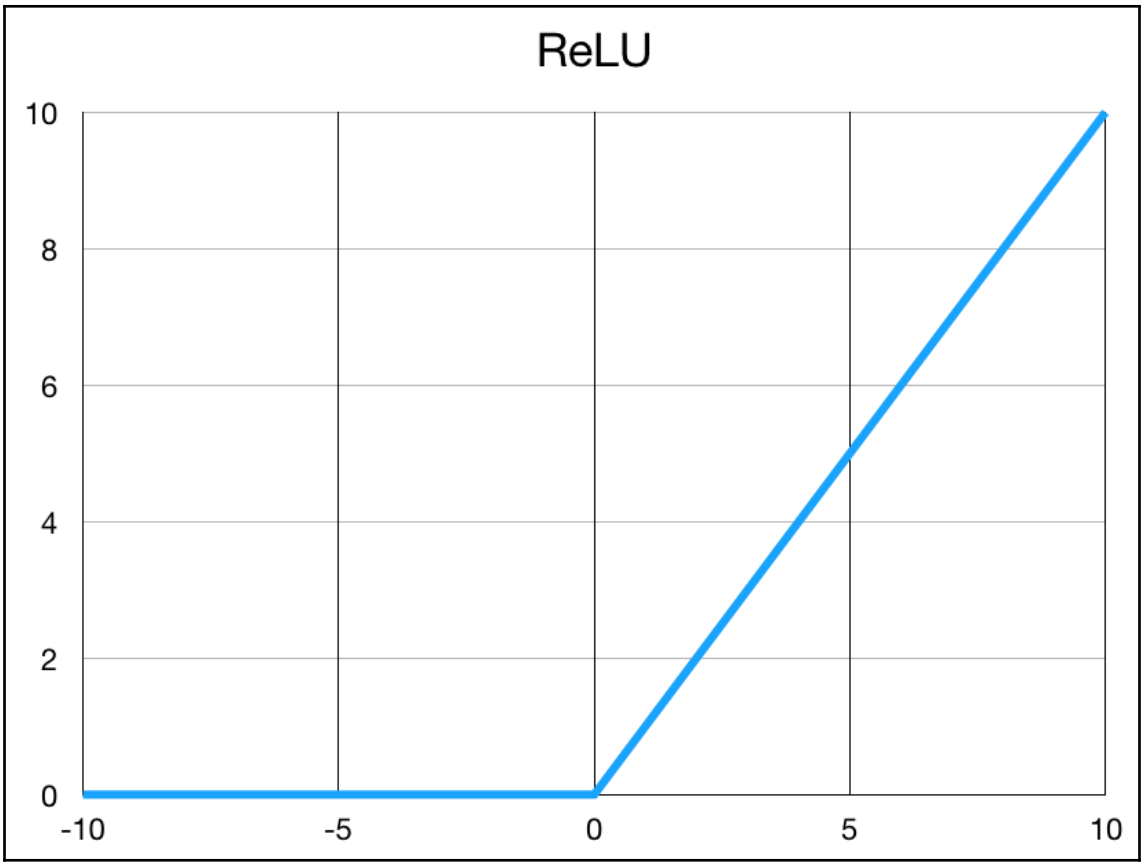
ReLU
10
8
6
4
2
0
-10
-5
0
5
10

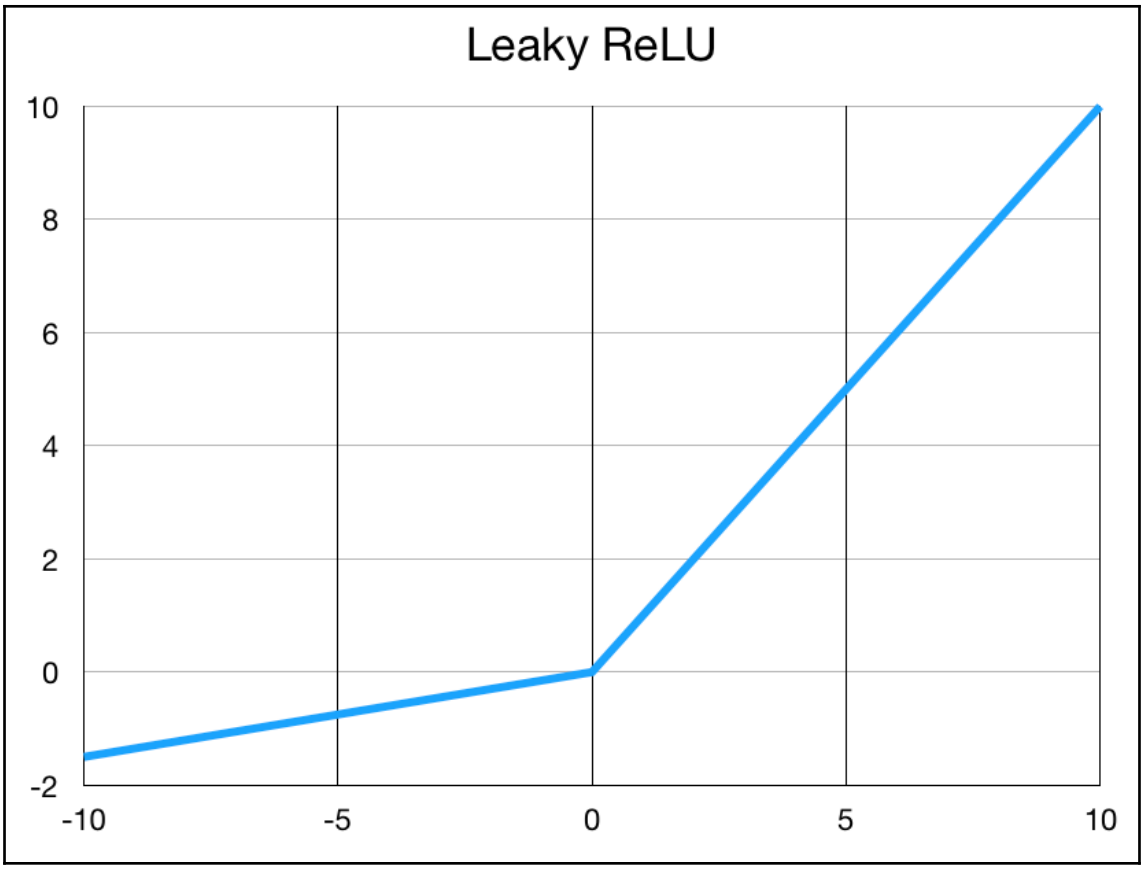
Leaky ReLU
10
8
6
4
2
0
-2
-10
-5
0
5
10


MK

Chapter 10: Classifying Clothing Images using Capsule Networks

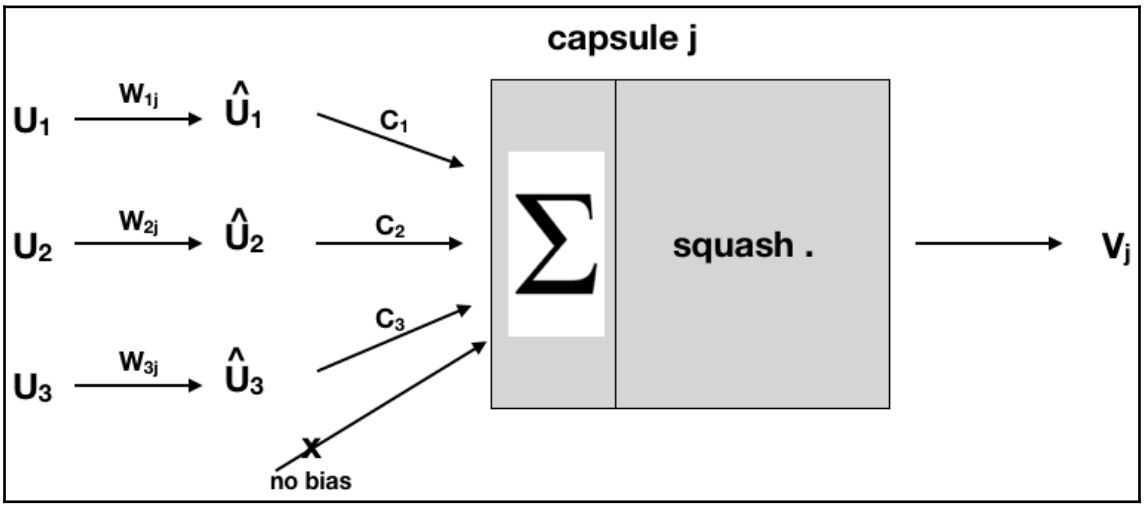

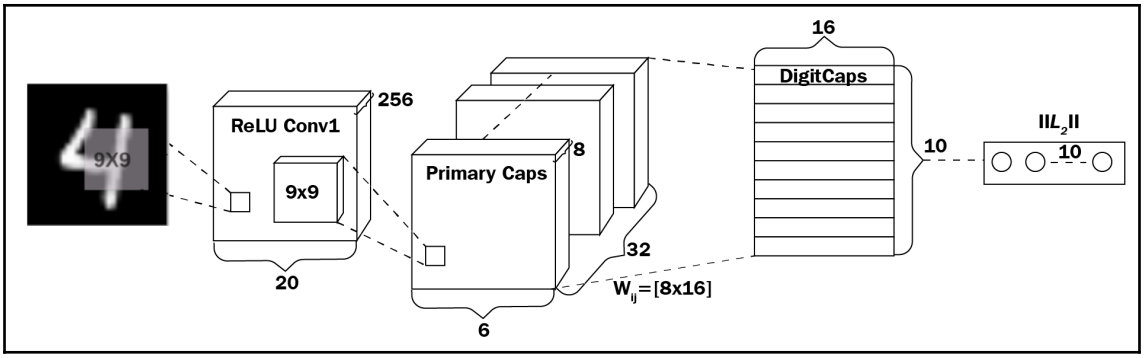
9X9
ReLU Conv1
256
9x9
20
Primary Caps
8
32
6
W_{ij}=[8x16]
16
DigitCaps
10
$||L_2||$
10

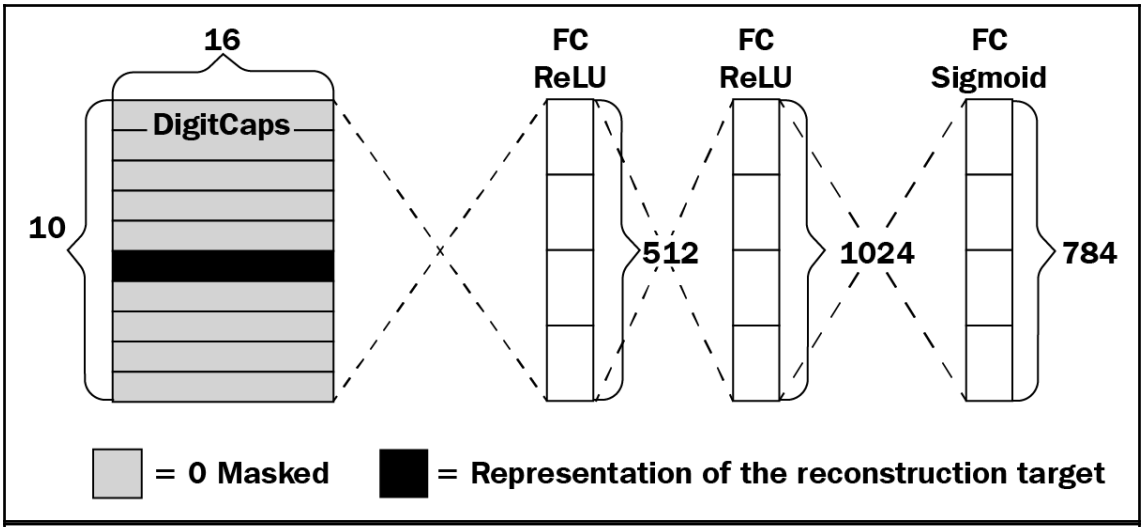
16
DigitCaps
10
FC
ReLU
512
FC
ReLU
1024
FC
Sigmoid
784
= 0 Masked
= Representation of the reconstruction target

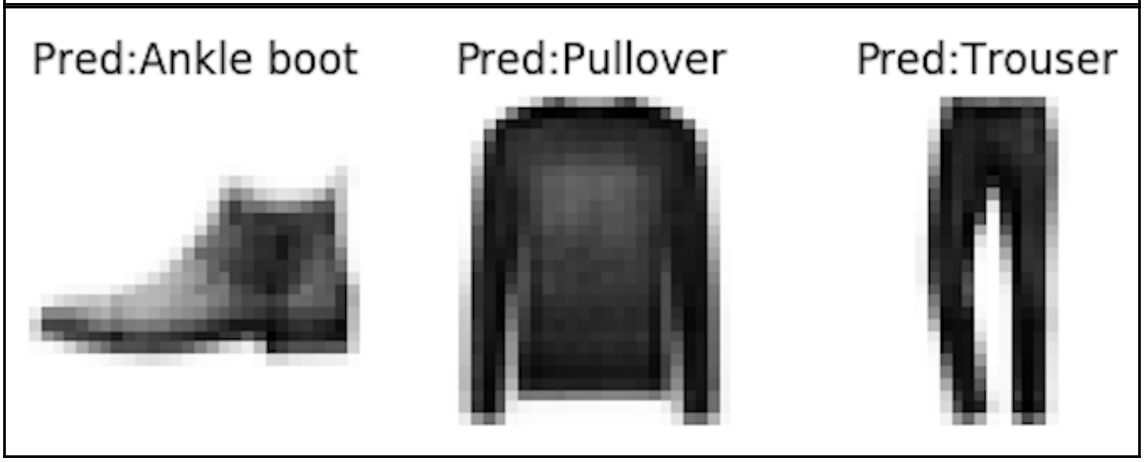
Pred:Ankle boot
Pred:Pullover
Pred:Trouser

Chapter 11: Making Quality Product Recommendations Using TensorFlow

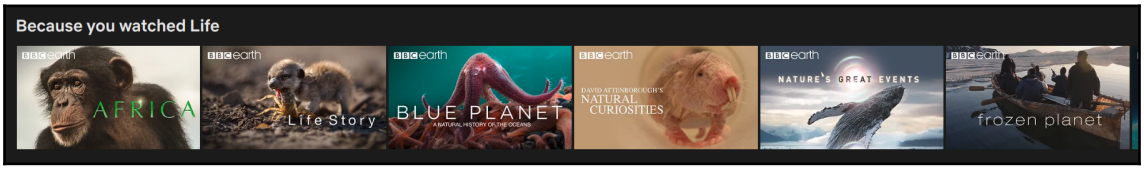

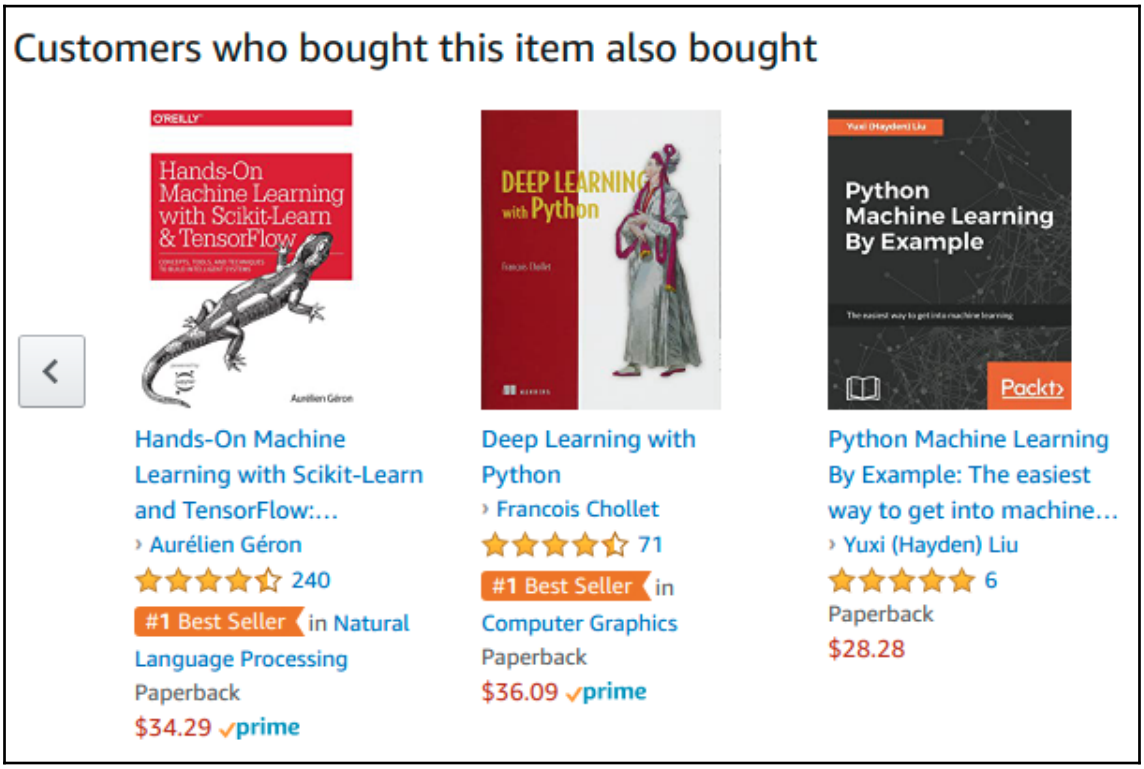

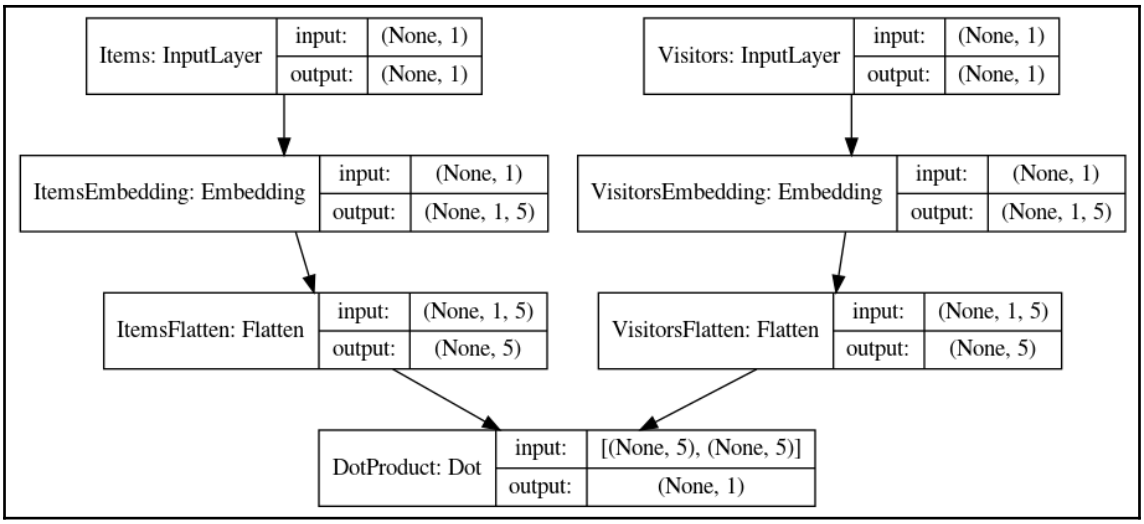
Items: InputLayer
input: (None, 1)
output: (None, 1)
Visitors: InputLayer
input: (None, 1)
output: (None, 1)
ItemsEmbedding: Embedding
input: (None, 1)
output: (None, 1, 5)
VisitorsEmbedding: Embedding
input: (None, 1)
output: (None, 1, 5)
ItemsFlatten: Flatten
input: (None, 1, 5)
output: (None, 5)
VisitorsFlatten: Flatten
input: (None, 1, 5)
output: (None, 5)
DotProduct: Dot
input: [(None, 5), (None, 5)]
output: (None, 1)

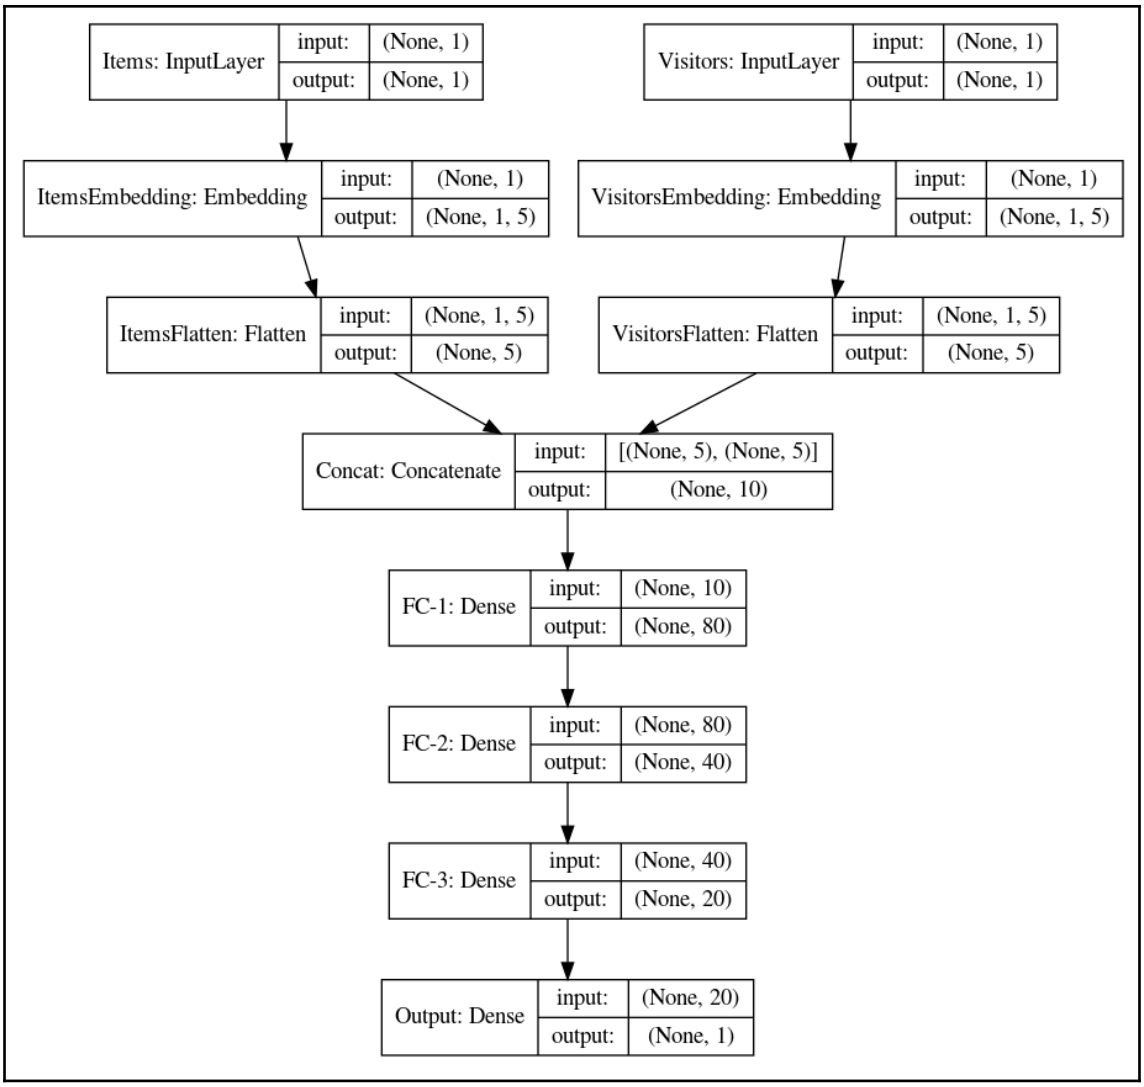
Items: InputLayer
input: (None, 1)
output: (None, 1)
Visitors: InputLayer
input: (None, 1)
output: (None, 1)
ItemsEmbedding: Embedding
input: (None, 1)
output: (None, 1, 5)
VisitorsEmbedding: Embedding
input: (None, 1)
output: (None, 1, 5)
ItemsFlatten: Flatten
input: (None, 1, 5)
output: (None, 5)
VisitorsFlatten: Flatten
input: (None, 1, 5)
output: (None, 5)
Concat: Concatenate
input: [(None, 5), (None, 5)]
output: (None, 10)
FC-1: Dense
input: (None, 10)
output: (None, 80)
FC-2: Dense
input: (None, 80)
output: (None, 40)
FC-3: Dense
input: (None, 40)
output: (None, 20)
Output: Dense
input: (None, 20)
output: (None, 1)

Chapter 12: Object Detection at a Large Scale with TensorFlow

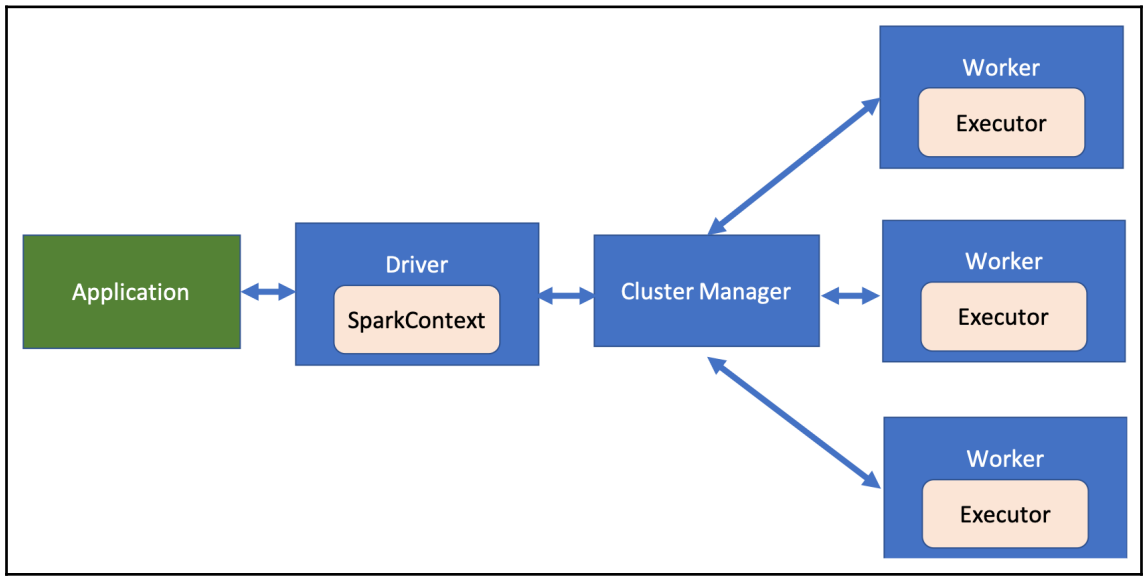


```
[13.0, 13.0]
(tensorflow) AKs-Mac-mini:Desktop am$ python distributed.py
2018-11-11 08:56:11.697963: I tensorflow/core/platform/cpu_feature_guard.cc:141] Your CPU supports instructions that this TensorFlow binary was not compiled to use
: SSE4.1 SSE4.2 AVX
2018-11-11 08:56:11.701184: I tensorflow/core/distributed_runtime/rpc/grpc_channel.cc:215] Initialize GrpcChannelCache for job worker -> {0 -> localhost:2222, 1 ->
 localhost:2223}
2018-11-11 08:56:11.702310: I tensorflow/core/distributed_runtime/rpc/grpc_server_lib.cc:375] Started server with target: grpc://localhost:2222
2018-11-11 08:56:11.705961: I tensorflow/core/distributed_runtime/rpc/grpc_channel.cc:215] Initialize GrpcChannelCache for job worker -> {0 -> localhost:2222, 1 ->
 localhost:2223}
2018-11-11 08:56:11.707046: I tensorflow/core/distributed_runtime/rpc/grpc_server_lib.cc:375] Started server with target: grpc://localhost:2223
2018-11-11 08:56:11.813805: I tensorflow/core/distributed_runtime/master_session.cc:1165] Start master session 3dce8682f06562a5 with config:
[7.0, 12.0]
```

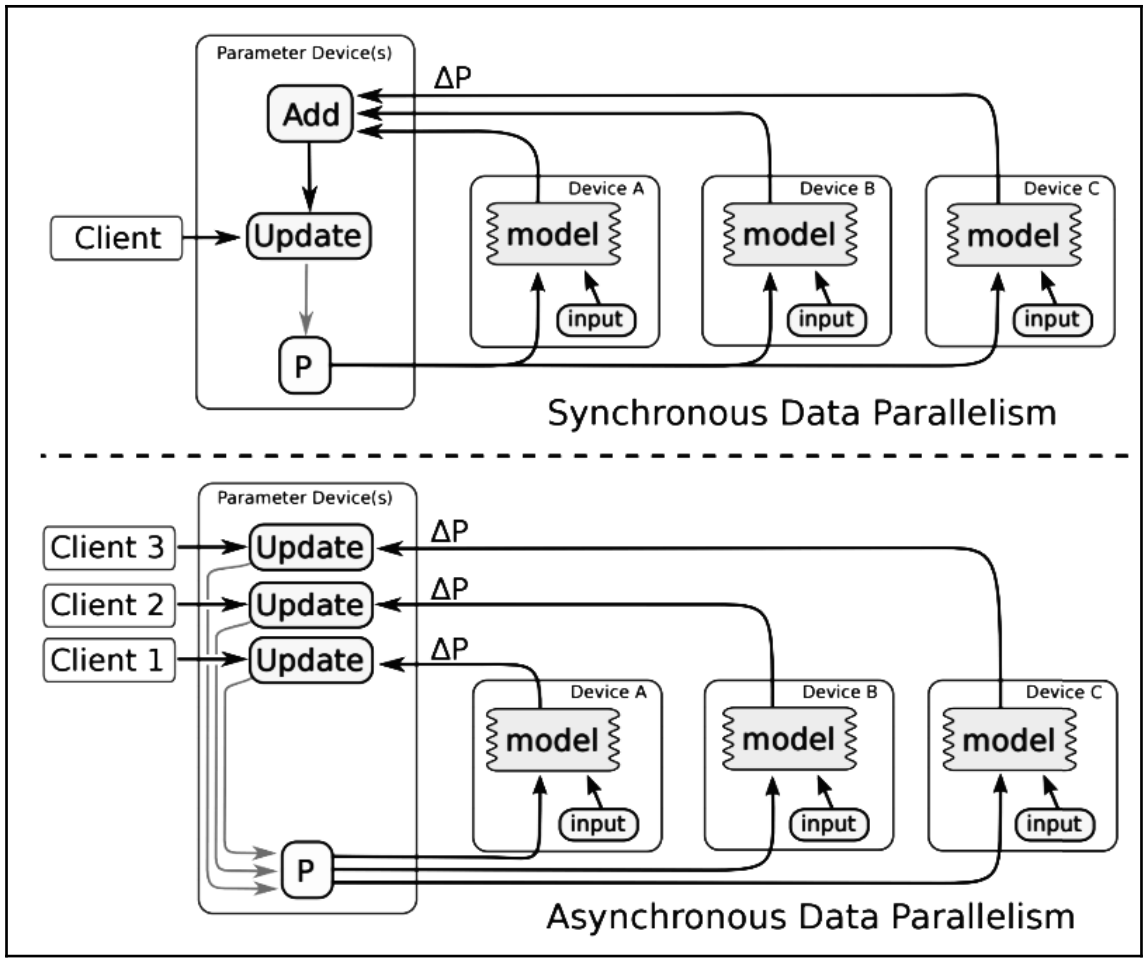
Parameter Device(s)
Add
ΔP
Client
Update
P
Device A
Device B
Device C
model
input
Synchronous Data Parallelism
Parameter Device(s)
Client 3
Client 2
Client 1
Update
ΔP
Asynchronous Data Parallelism

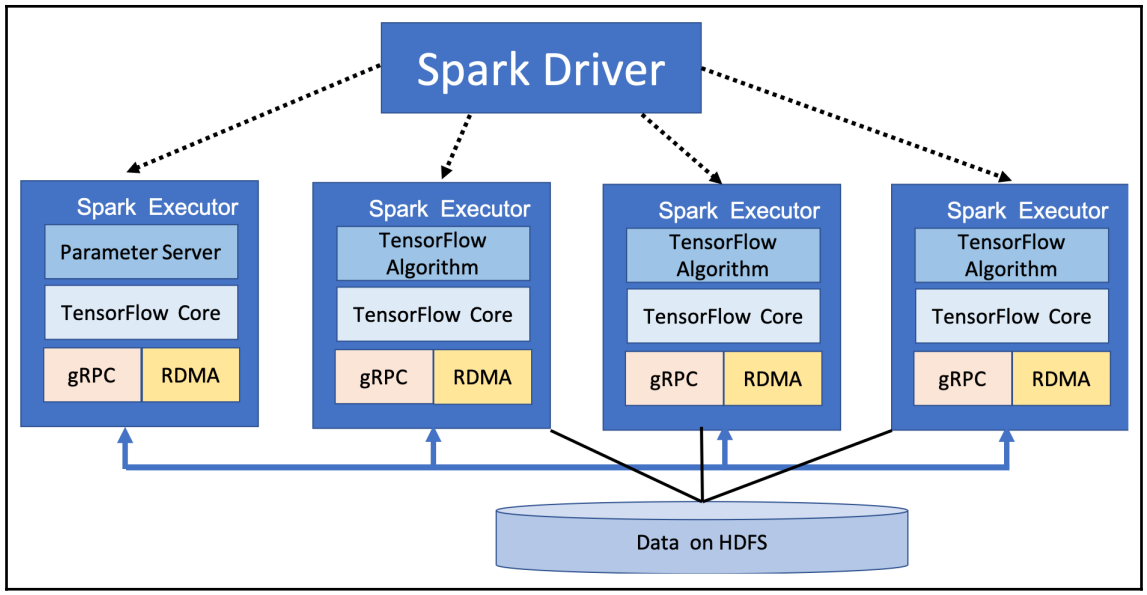

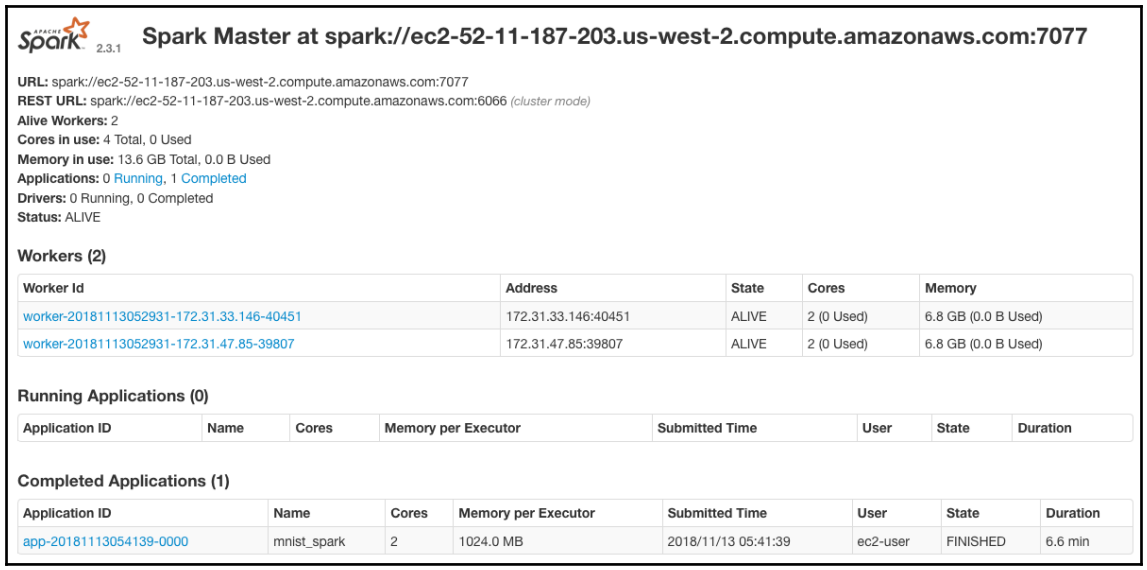

Spark Master at spark://ec2-52-11-187-203.us-west-2.compute.amazonaws.com:7077

URL: spark://ec2-52-11-187-203.us-west-2.compute.amazonaws.com:7077
REST URL: spark://ec2-52-11-187-203.us-west-2.compute.amazonaws.com:6066 *(cluster mode)*
Alive Workers: 2
Cores in use: 4 Total, 0 Used
Memory in use: 13.6 GB Total, 0.0 B Used
Applications: 0 Running, 1 Completed
Drivers: 0 Running, 0 Completed
Status: ALIVE

Workers (2)

Worker Id	Address	State	Cores	Memory
worker-20181113052931-172.31.33.146-40451	172.31.33.146:40451	ALIVE	2 (0 Used)	6.8 GB (0.0 B Used)
worker-20181113052931-172.31.47.85-39807	172.31.47.85:39807	ALIVE	2 (0 Used)	6.8 GB (0.0 B Used)

Running Applications (0)

Application ID	Name	Cores	Memory per Executor	Submitted Time	User	State	Duration

Completed Applications (1)

Application ID	Name	Cores	Memory per Executor	Submitted Time	User	State	Duration
app-20181113054139-0000	mnist_spark	2	1024.0 MB	2018/11/13 05:41:39	ec2-user	FINISHED	6.6 min

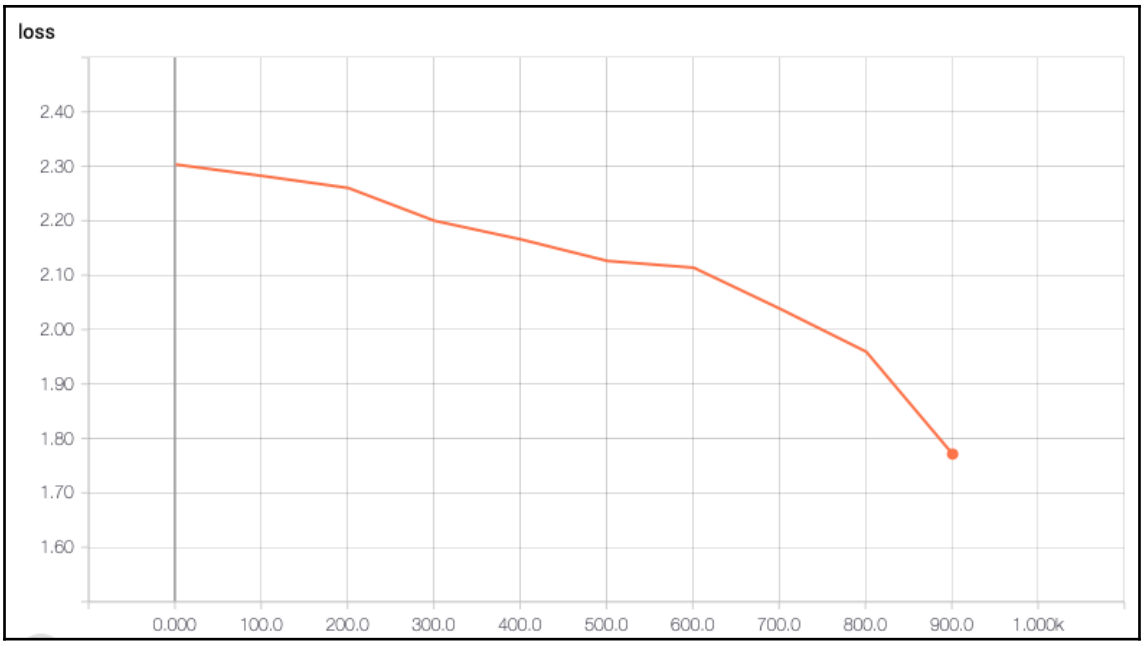
loss
2.40
2.30
2.20
2.10
2.00
1.90
1.80
1.70
1.60
0.000
100.0
200.0
300.0
400.0
500.0
600.0
700.0
800.0
900.0
1.000k

```
+--------------------+-----+
|               image|label|
+--------------------+-----+
|[file:/home/ubunt...|    0|
|[file:/home/ubunt...|    0|
|[file:/home/ubunt...|    0|
|[file:/home/ubunt...|    0|
|[file:/home/ubunt...|    0|
+--------------------+-----+
only showing top 5 rows

+--------------------+-----+
|               image|label|
+--------------------+-----+
|[file:/home/ubunt...|    1|
|[file:/home/ubunt...|    1|
|[file:/home/ubunt...|    1|
|[file:/home/ubunt...|    1|
|[file:/home/ubunt...|    1|
+--------------------+-----+
only showing top 5 rows
```

```
+----------------+---+---+
|prediction_label|  0|  1|
+----------------+---+---+
|             1.0|  0| 14|
|             0.0| 22|  0|
+----------------+---+---+
```

Chapter 13: Generating Book Scripts Using LSTMs

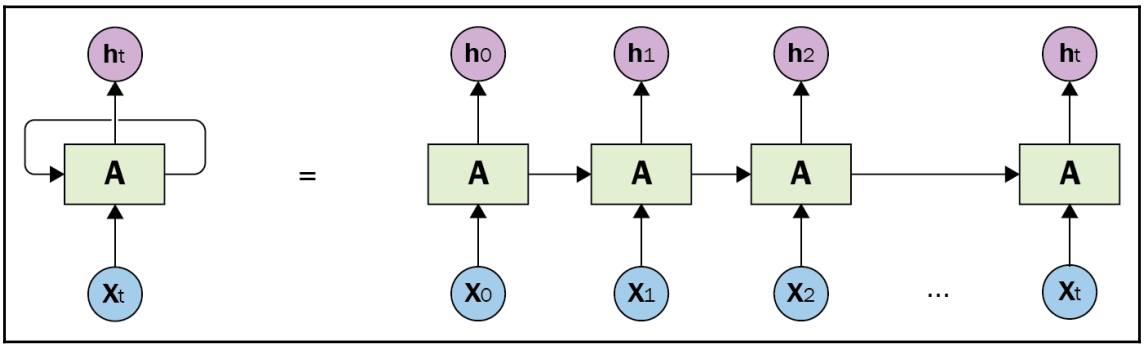

Chapter 14: Playing Pacman Using Deep Reinforcement Learning

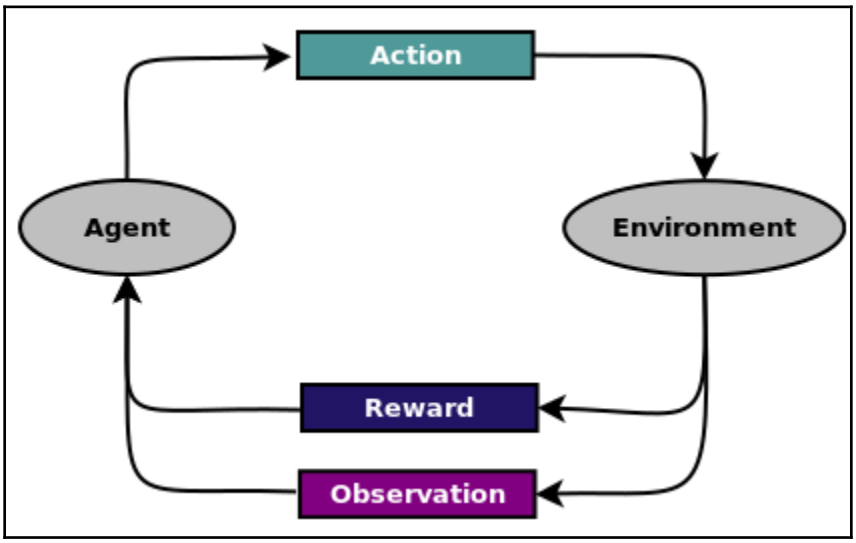

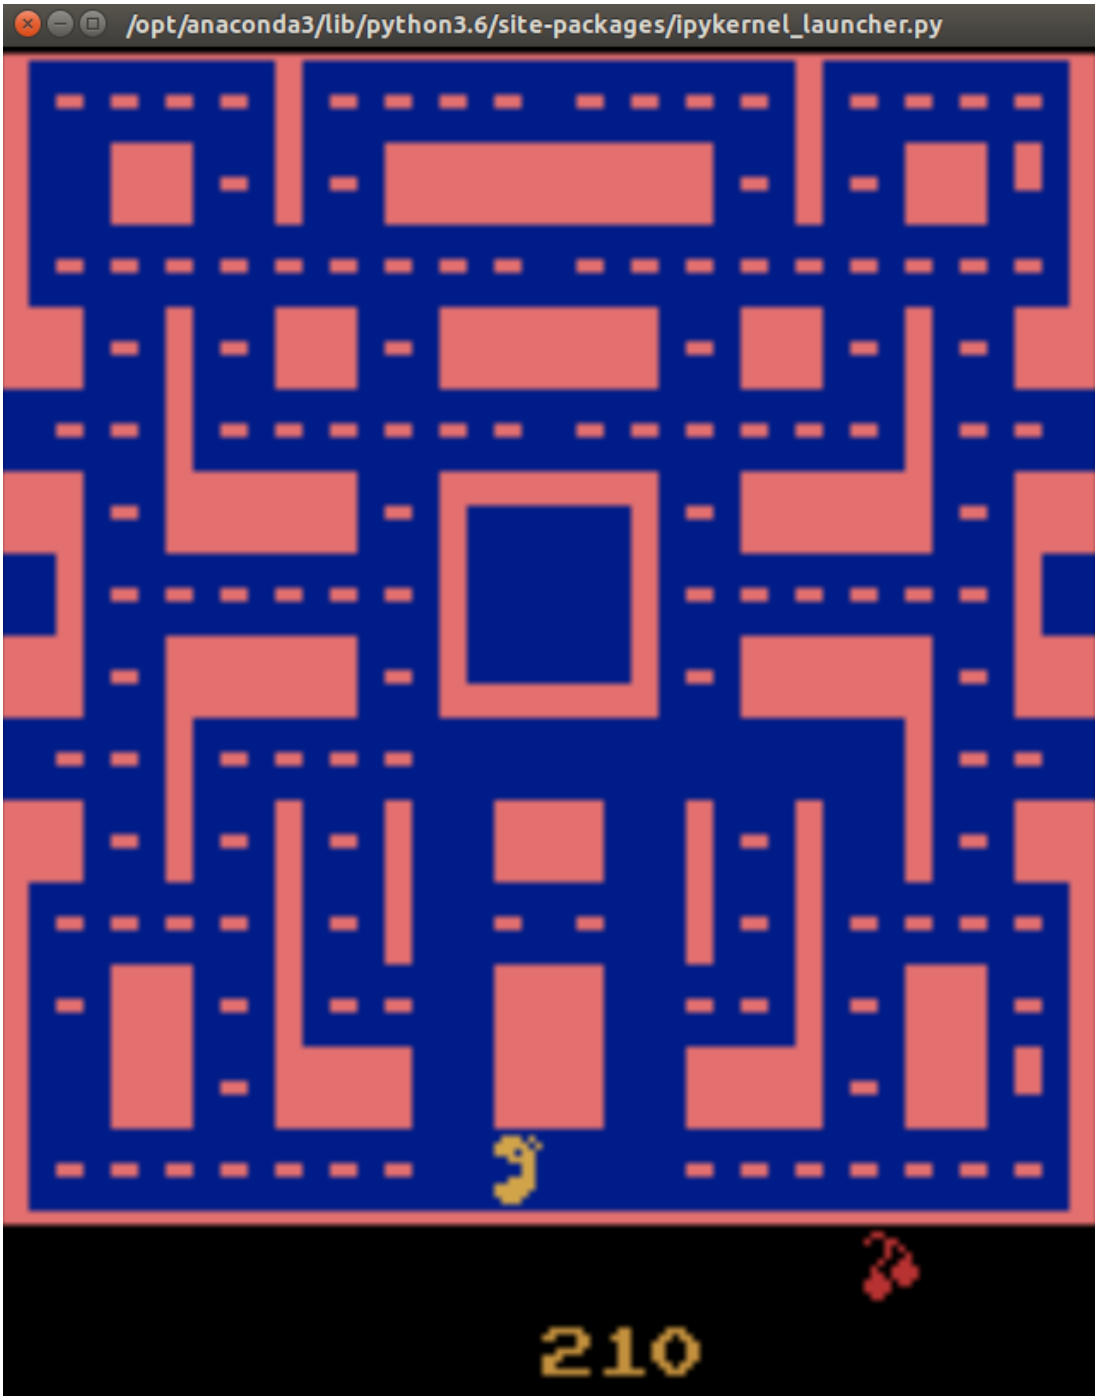
/opt/anaconda3/lib/python3.6/site-packages/ipykernel_launcher.py
210

Chapter 15: What is Next?

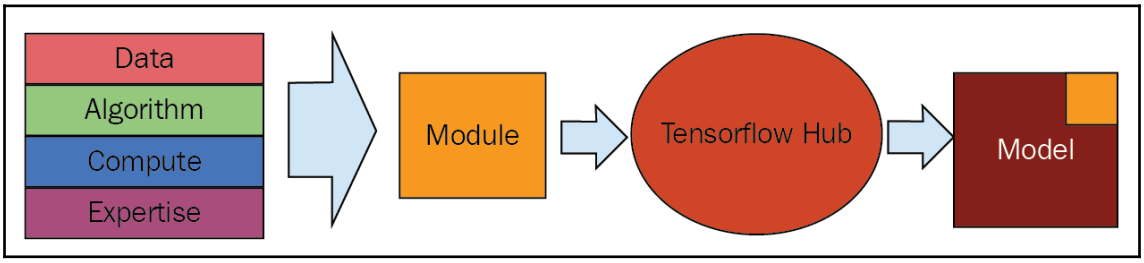

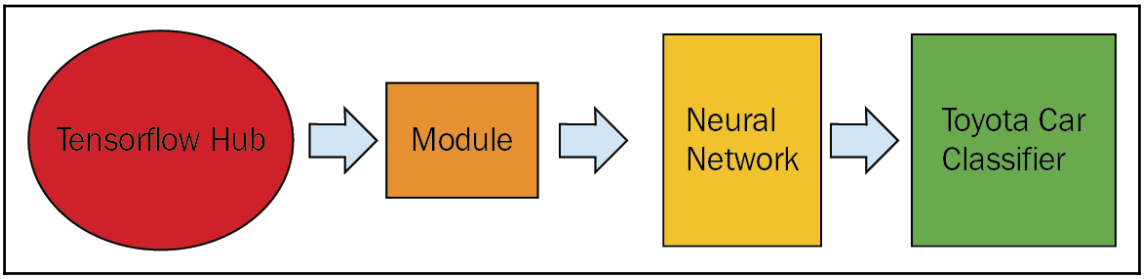

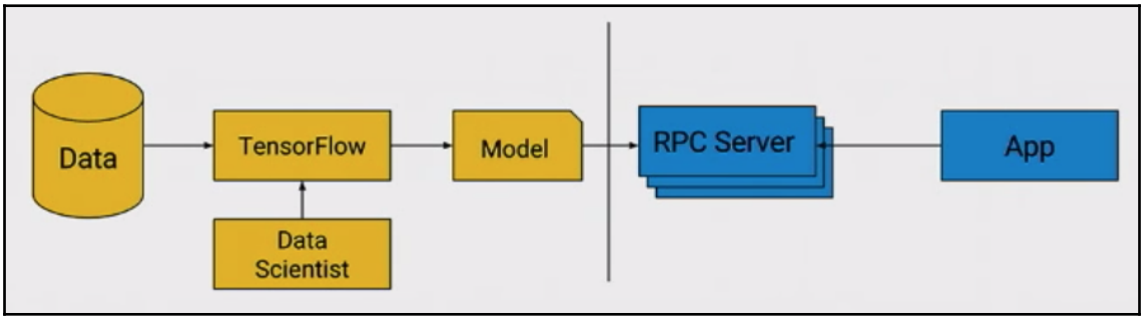

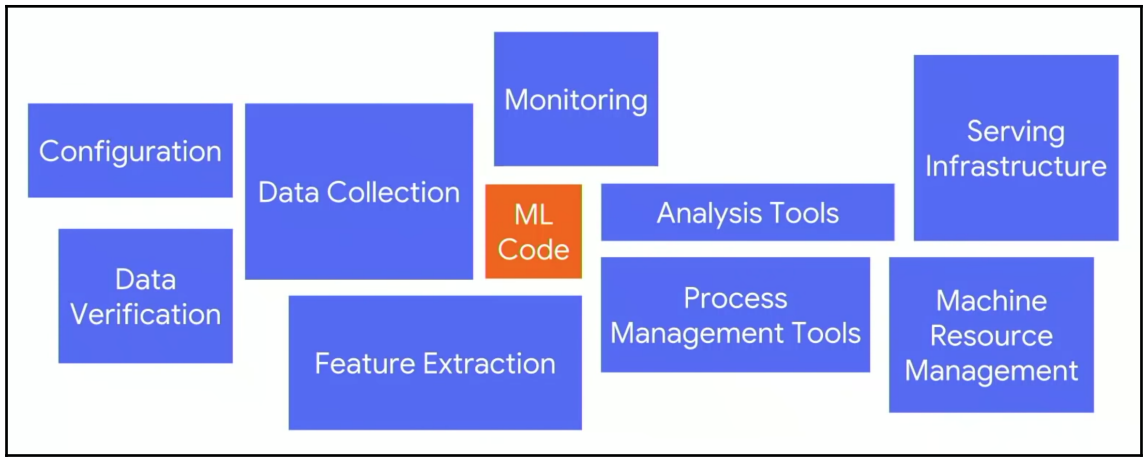
Configuration
Data Collection
Monitoring
Serving Infrastructure
Data Verification
ML Code
Analysis Tools
Process Management Tools
Machine Resource Management
Feature Extraction

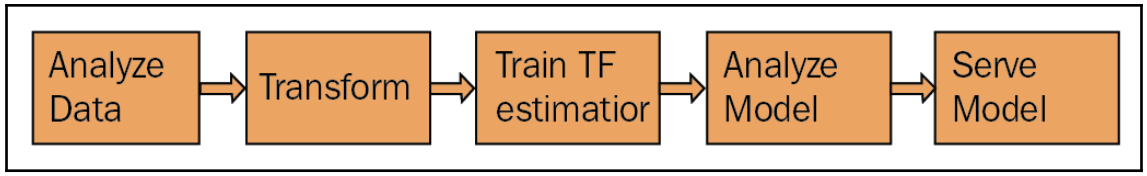
Analyze Data
Transform
Train TF estimatior
Analyze Model
Serve Model