

Table of Contents

Preface	xii
Chapter 1: Neural Network Foundations with TensorFlow 2.0	1
What is TensorFlow (TF)?	1
What is Keras?	3
What are the most important changes in TensorFlow 2.0?	3
Introduction to neural networks	5
Perceptron	6
A first example of TensorFlow 2.0 code	7
Multi-layer perceptron – our first example of a network	8
Problems in training the perceptron and their solutions	9
Activation function – sigmoid	10
Activation function – tanh	10
Activation function – ReLU	11
Two additional activation functions – ELU and LeakyReLU	12
Activation functions	13
In short – what are neural networks after all?	13
A real example – recognizing handwritten digits	14
One-hot encoding (OHE)	14
Defining a simple neural network in TensorFlow 2.0	15
Running a simple TensorFlow 2.0 net and establishing a baseline	20
Improving the simple net in TensorFlow 2.0 with hidden layers	21
Further improving the simple net in TensorFlow with Dropout	24
Testing different optimizers in TensorFlow 2.0	26
Increasing the number of epochs	32
Controlling the optimizer learning rate	33
Increasing the number of internal hidden neurons	34

Increasing the size of batch computation	35
Summarizing experiments run for recognizing handwritten charts	36
Regularization	36
Adopting regularization to avoid overfitting	36
Understanding BatchNormalization	38
Playing with Google Colab – CPUs, GPUs, and TPUs	39
Sentiment analysis	42
Hyperparameter tuning and AutoML	45
Predicting output	45
A practical overview of backpropagation	46
What have we learned so far?	48
Towards a deep learning approach	48
References	49
Chapter 2: TensorFlow 1.x and 2.x	51
Understanding TensorFlow 1.x	51
TensorFlow 1.x computational graph program structure	51
Computational graphs	52
Working with constants, variables, and placeholders	54
Examples of operations	55
Constants	55
Sequences	56
Random tensors	56
Variables	57
An example of TensorFlow 1.x in TensorFlow 2.x	59
Understanding TensorFlow 2.x	60
Eager execution	60
AutoGraph	61
Keras APIs – three programming models	63
Sequential API	63
Functional API	64
Model subclassing	66
Callbacks	67
Saving a model and weights	68
Training from tf.data.datasets	69
tf.keras or Estimators?	72
Ragged tensors	74
Custom training	74
Distributed training in TensorFlow 2.x	76
Multiple GPUs	76
MultiWorkerMirroredStrategy	78
TPUStrategy	78
ParameterServerStrategy	78
Changes in namespaces	79

Converting from 1.x to 2.x	80
Using TensorFlow 2.x effectively	80
The TensorFlow 2.x ecosystem	81
Language bindings	82
Keras or tf.keras?	83
Summary	84
Chapter 3: Regression	87
What is regression?	87
Prediction using linear regression	88
Simple linear regression	89
Multiple linear regression	93
Multivariate linear regression	93
TensorFlow Estimators	94
Feature columns	94
Input functions	95
MNIST using TensorFlow Estimator API	95
Predicting house price using linear regression	97
Classification tasks and decision boundaries	101
Logistic regression	102
Logistic regression on the MNIST dataset	103
Summary	107
References	108
Chapter 4: Convolutional Neural Networks	109
Deep Convolutional Neural Network (DCNN)	110
Local receptive fields	110
Shared weights and bias	111
A mathematical example	111
Convnets in TensorFlow 2.x	112
Pooling layers	113
Max pooling	113
Average pooling	113
ConvNets summary	113
An example of DCNN – LeNet	114
LeNet code in TensorFlow 2.0	114
Understanding the power of deep learning	121
Recognizing CIFAR-10 images with deep learning	122
Improving the CIFAR-10 performance with a deeper network	125
Improving the CIFAR-10 performance with data augmentation	128
Predicting with CIFAR-10	130
Very deep convolutional networks for large-scale image recognition	132
Recognizing cats with a VGG16 Net	134

Utilizing tf.keras built-in VGG16 Net module	135
Recycling prebuilt deep learning models for extracting features	136
Summary	137
References	138
Chapter 5: Advanced Convolutional Neural Networks	139
Computer vision	139
Composing CNNs for complex tasks	139
Classification and localization	140
Semantic segmentation	141
Object detection	142
Instance segmentation	145
Classifying Fashion-MNIST with a tf.keras - estimator model	147
Run Fashion-MNIST the tf.keras - estimator model on GPUs	150
Deep Inception-v3 Net used for transfer learning	151
Transfer learning for classifying horses and humans	154
Application Zoos with tf.keras and TensorFlow Hub	157
Keras applications	158
TensorFlow Hub	158
Other CNN architectures	159
AlexNet	159
Residual networks	159
HighwayNets and DenseNets	160
Xception	160
Answering questions about images (VQA)	162
Style transfer	165
Content distance	166
Style distance	167
Creating a DeepDream network	168
Inspecting what a network has learned	172
Video	173
Classifying videos with pretrained nets in six different ways	173
Textual documents	174
Using a CNN for sentiment analysis	175
Audio and music	178
Dilated ConvNets, WaveNet, and NSynth	178
A summary of convolution operations	183
Basic convolutional neural networks (CNN or ConvNet)	183
Dilated convolution	184
Transposed convolution	184
Separable convolution	184
Depthwise convolution	185
Depthwise separable convolution	185
Capsule networks	185

So what is the problem with CNNs?	185
So what is new with Capsule networks?	186
Summary	188
References	188
Chapter 6: Generative Adversarial Networks	191
What is a GAN?	191
MNIST using GAN in TensorFlow	193
Deep convolutional GAN (DCGAN)	198
DCGAN for MNIST digits	200
Some interesting GAN architectures	209
SRGAN	209
CycleGAN	210
InfoGAN	212
Cool applications of GANs	214
CycleGAN in TensorFlow 2.0	218
Summary	228
References	228
Chapter 7: Word Embeddings	231
Word embedding – origins and fundamentals	231
Distributed representations	233
Static embeddings	234
Word2Vec	235
GloVe	238
Creating your own embedding using gensim	239
Exploring the embedding space with gensim	240
Using word embeddings for spam detection	243
Getting the data	244
Making the data ready for use	245
Building the embedding matrix	247
Define the spam classifier	248
Train and evaluate the model	250
Running the spam detector	251
Neural embeddings – not just for words	252
Item2Vec	253
node2vec	253
Character and subword embeddings	259
Dynamic embeddings	260
Sentence and paragraph embeddings	262
Language model-based embeddings	264
Using BERT as a feature extractor	267

28 Fine-tuning BERT	269
29 Classifying with BERT – command line	270
30 Using BERT as part of your own network	271
Summary	275
References	275
Chapter 8: Recurrent Neural Networks	279
The basic RNN cell	280
31 Backpropagation through time (BPTT)	283
32 Vanishing and exploding gradients	284
RNN cell variants	285
33 Long short-term memory (LSTM)	285
34 Gated recurrent unit (GRU)	288
35 Peephole LSTM	288
RNN variants	289
36 Bidirectional RNNs	289
37 Stateful RNNs	290
RNN topologies	291
Example – One-to-Many – learning to generate text	292
Example – Many-to-One – Sentiment Analysis	300
Example – Many-to-Many – POS tagging	307
Encoder-Decoder architecture – seq2seq	316
Example – seq2seq without attention for machine translation	318
Attention mechanism	328
Example – seq2seq with attention for machine translation	330
Transformer architecture	336
Summary	340
References	340
Chapter 9: Autoencoders	345
Introduction to autoencoders	345
Vanilla autoencoders	347
TensorFlow Keras layers – defining custom layers	348
Reconstructing handwritten digits using an autoencoder	350
Sparse autoencoder	354
Denoising autoencoders	356
Clearing images using a Denoising autoencoder	357
Stacked autoencoder	360
Convolutional autoencoder for removing noise from images	360
Keras autoencoder example – sentence vectors	365
Summary	373
References	374

Chapter 10: Unsupervised Learning	375
Principal component analysis	375
PCA on the MNIST dataset	376
TensorFlow Embedding API	379
K-means clustering	380
K-means in TensorFlow 2.0	381
Variations in k-means	384
Self-organizing maps	384
Colour mapping using SOM	387
Restricted Boltzmann machines	392
Reconstructing images using RBM	393
Deep belief networks	397
Variational Autoencoders	399
Summary	404
References	405
Chapter 11: Reinforcement Learning	407
Introduction	407
RL lingo	409
Deep reinforcement learning algorithms	411
Reinforcement success in recent years	414
Introduction to OpenAI Gym	415
Random agent playing Breakout	418
Deep Q-Networks	420
DQN for CartPole	422
DQN to play a game of Atari	427
DQN variants	430
Double DQN	430
Dueling DQN	431
Rainbow	434
Deep deterministic policy gradient	434
Summary	436
References	436
Chapter 12: TensorFlow and Cloud	439
Deep learning on cloud	439
Microsoft Azure	440
Amazon Web Services (AWS)	442
Google Cloud Platform (GCP)	444
IBM Cloud	447
Virtual machines on cloud	447
EC2 on Amazon	448
Compute Instance on GCP	450

Virtual machine on Microsoft Azure	451
Jupyter Notebooks on cloud	451
SageMaker – part of your own network	452
Google Colaboratory	452
Microsoft Azure Notebooks	455
TensorFlow Extended for production	455
TFX Pipelines	456
TFX pipeline components	457
TFX libraries	458
TensorFlow Enterprise	459
Summary	459
References	460
Chapter 13: TensorFlow for Mobile and IoT and TensorFlow.js	461
TensorFlow Mobile	461
TensorFlow Lite	462
Quantization	462
FlatBuffers	463
Mobile converter	463
Mobile optimized interpreter	463
Supported platforms	464
Architecture	464
Using TensorFlow Lite	464
A generic example of application	465
Using GPUs and accelerators	466
An example of application	466
Pretrained models in TensorFlow Lite	468
Image classification	470
Object detection	471
Pose estimation	471
Smart reply	471
Segmentation	471
Style transfer	471
Text classification	472
Question and answering	472
A note about using mobile GPUs	473
An overview of federated learning at the edge	474
TensorFlow FL APIs	476
TensorFlow.js	478
Vanilla TensorFlow.js	478
Converting models	485

Pretrained models	485
Node.js	488
Summary	489
References	489
Chapter 14: An introduction to AutoML	491
What is AutoML?	491
Achieving AutoML	492
Automatic data preparation	493
Automatic feature engineering	493
Automatic model generation	494
AutoKeras	497
Google Cloud AutoML	498
Using Cloud AutoML – Tables solution	499
Using Cloud AutoML – Vision solution	513
Using Cloud AutoML – Text Classification solution	524
Using Cloud AutoML – Translation solution	528
Using Cloud AutoML – Video Intelligence Classification solution	533
Cost	539
Bringing Google AutoML to Kaggle	540
Summary	541
References	542
Chapter 15: The Math Behind Deep Learning	543
History	543
Some mathematical tools	544
Derivatives and gradients everywhere	544
Gradient descent	546
Chain rule	547
A few differentiation rules	547
Matrix operations	547
Activation functions	548
Derivative of the sigmoid	548
Derivative of tanh	549
Derivative of ReLU	550
Backpropagation	550
Forward step	553
Backstep	554
Case 1 – From hidden layer to output layer	556
Case 2 – From hidden layer to hidden layer	557
Limit of backpropagation	561
Cross entropy and its derivative	561

Batch gradient descent, stochastic gradient descent, and mini-batch	563
Batch Gradient Descent (BGD)	564
Stochastic Gradient Descent (SGD)	564
Mini-Batch Gradient Descent (MBGD)	564
Thinking about backpropagation and convnets	564
Thinking about backpropagation and RNNs	565
A note on TensorFlow and automatic differentiation	568
Summary	568
References	570
Chapter 16: Tensor Processing Unit	571
C/G/T processing units	571
CPUs and GPUs	571
TPUs	572
Three generations of TPUs and Edge TPU	573
First-generation TPU	574
Second-generation TPU	576
Third-generation TPU	577
Edge TPU	578
TPU performance	578
How to use TPUs with Colab	580
Checking whether TPUs are available	580
Loading data with <code>tf.data</code>	581
Building a model and loading it into the TPU	581
Using pretrained TPU models	584
Using TensorFlow 2.1 and nightly build	587
Summary	588
References	589
Other Books You May Enjoy	591
Index	597

Smart reply	468
Segmentation	474
Style transfer	474
Text classification	475
Question and answering	476
A note about using mobile GPUs	477
An overview of federated learning at the edge	478
TensorFlow FL APIs	479
TensorFlow.js	479
Vanilla TensorFlow.js	480
Converting models	480