
Table of Contents

Foreword..... **xiii**

Preface..... **xv**

Part I. Building Models

1. Introduction to TensorFlow.....	1
What Is Machine Learning?	1
Limitations of Traditional Programming	3
From Programming to Learning	5
What Is TensorFlow?	7
Using TensorFlow	8
Installing TensorFlow in Python	9
Using TensorFlow in PyCharm	9
Using TensorFlow in Google Colab	12
Getting Started with Machine Learning	13
Seeing What the Network Learned	18
Summary	19
2. Introduction to Computer Vision.....	21
Recognizing Clothing Items	21
The Data: Fashion MNIST	22
Neurons for Vision	23
Designing the Neural Network	25
The Complete Code	26
Training the Neural Network	29

Exploring the Model Output	29
Training for Longer—Discovering Overfitting	30
Stopping Training	30
Summary	32
3. Going Beyond the Basics: Detecting Features in Images.	33
Convolutions	34
Pooling	35
Implementing Convolutional Neural Networks	37
Exploring the Convolutional Network	39
Building a CNN to Distinguish Between Horses and Humans	42
The Horses or Humans Dataset	42
The Keras ImageDataGenerator	43
CNN Architecture for Horses or Humans	44
Adding Validation to the Horses or Humans Dataset	47
Testing Horse or Human Images	49
Image Augmentation	52
Transfer Learning	55
Multiclass Classification	59
Dropout Regularization	63
Summary	65
4. Using Public Datasets with TensorFlow Datasets.	67
Getting Started with TFDS	68
Using TFDS with Keras Models	70
Loading Specific Versions	73
Using Mapping Functions for Augmentation	73
Using TensorFlow Addons	74
Using Custom Splits	74
Understanding TFRecord	76
The ETL Process for Managing Data in TensorFlow	78
Optimizing the Load Phase	80
Parallelizing ETL to Improve Training Performance	81
Summary	83
5. Introduction to Natural Language Processing.	85
Encoding Language into Numbers	85
Getting Started with Tokenization	86
Turning Sentences into Sequences	87
Removing Stopwords and Cleaning Text	91
Working with Real Data Sources	93

Getting Text from TensorFlow Datasets	93
Getting Text from CSV Files	97
Getting Text from JSON Files	99
Summary	102
6. Making Sentiment Programmable Using Embeddings.....	103
Establishing Meaning from Words	103
A Simple Example: Positives and Negatives	104
Going a Little Deeper: Vectors	105
Embeddings in TensorFlow	106
Building a Sarcasm Detector Using Embeddings	106
Reducing Overfitting in Language Models	109
Using the Model to Classify a Sentence	119
Visualizing the Embeddings	120
Using Pretrained Embeddings from TensorFlow Hub	123
Summary	125
7. Recurrent Neural Networks for Natural Language Processing.....	127
The Basis of Recurrence	127
Extending Recurrence for Language	130
Creating a Text Classifier with RNNs	132
Stacking LSTMs	134
Using Pretrained Embeddings with RNNs	139
Summary	146
8. Using TensorFlow to Create Text.....	147
Turning Sequences into Input Sequences	148
Creating the Model	152
Generating Text	154
Predicting the Next Word	154
Compounding Predictions to Generate Text	155
Extending the Dataset	156
Changing the Model Architecture	157
Improving the Data	158
Character-Based Encoding	162
Summary	163
9. Understanding Sequence and Time Series Data.....	165
Common Attributes of Time Series	167
Trend	167
Seasonality	167

Autocorrelation	168
Noise	168
Techniques for Predicting Time Series	169
Naive Prediction to Create a Baseline	169
Measuring Prediction Accuracy	172
Less Naive: Using Moving Average for Prediction	172
Improving the Moving Average Analysis	173
Summary	174
10. Creating ML Models to Predict Sequences.....	175
Creating a Windowed Dataset	176
Creating a Windowed Version of the Time Series Dataset	178
Creating and Training a DNN to Fit the Sequence Data	180
Evaluating the Results of the DNN	181
Exploring the Overall Prediction	183
Tuning the Learning Rate	185
Exploring Hyperparameter Tuning with Keras Tuner	187
Summary	191
11. Using Convolutional and Recurrent Methods for Sequence Models.....	193
Convolutions for Sequence Data	193
Coding Convolutions	194
Experimenting with the Conv1D Hyperparameters	197
Using NASA Weather Data	200
Reading GISS Data in Python	201
Using RNNs for Sequence Modeling	202
Exploring a Larger Dataset	205
Using Other Recurrent Methods	207
Using Dropout	208
Using Bidirectional RNNs	211
Summary	213

Part II. Using Models

12. An Introduction to TensorFlow Lite.....	217
What Is TensorFlow Lite?	217
Walkthrough: Creating and Converting a Model to TensorFlow Lite	219
Step 1. Save the Model	220
Step 2. Convert and Save the Model	221
Step 3. Load the TFLite Model and Allocate Tensors	221

Step 4. Perform the Prediction	222
Walkthrough: Transfer Learning an Image Classifier and Converting to TensorFlow Lite	224
Step 1. Build and Save the Model	224
Step 2. Convert the Model to TensorFlow Lite	225
Step 3. Optimize the Model	227
Summary	229
13. Using TensorFlow Lite in Android Apps.....	231
What Is Android Studio?	231
Creating Your First TensorFlow Lite Android App	232
Step 1. Create a New Android Project	232
Step 2. Edit Your Layout File	234
Step 3. Add the TensorFlow Lite Dependencies	236
Step 4. Add Your TensorFlow Lite Model	238
Step 5. Write the Activity Code to Use TensorFlow Lite for Inference	238
Moving Beyond “Hello World”—Processing Images	241
TensorFlow Lite Sample Apps	244
Summary	246
14. Using TensorFlow Lite in iOS Apps.....	247
Creating Your First TensorFlow Lite App with Xcode	247
Step 1. Create a Basic iOS App	247
Step 2. Add TensorFlow Lite to Your Project	249
Step 3. Create the User Interface	250
Step 4. Add and Initialize the Model Inference Class	252
Step 5. Perform the Inference	255
Step 6. Add the Model to Your App	256
Step 7. Add the UI Logic	258
Moving Beyond “Hello World”—Processing Images	260
TensorFlow Lite Sample Apps	263
Summary	263
15. An Introduction to TensorFlow.js.....	265
What Is TensorFlow.js?	265
Installing and Using the Brackets IDE	267
Building Your First TensorFlow.js Model	268
Creating an Iris Classifier	272
Summary	276

16. Coding Techniques for Computer Vision in TensorFlow.js.....	277
JavaScript Considerations for TensorFlow Developers	278
Building a CNN in JavaScript	279
Using Callbacks for Visualization	281
Training with the MNIST Dataset	284
Running Inference on Images in TensorFlow.js	290
Summary	290
17. Reusing and Converting Python Models to JavaScript.....	293
Converting Python-Based Models to JavaScript	293
Using the Converted Models	295
Using Preconverted JavaScript Models	297
Using the Toxicity Text Classifier	297
Using MobileNet for Image Classification in the Browser	300
Using PoseNet	303
Summary	306
18. Transfer Learning in JavaScript.....	307
Transfer Learning from MobileNet	308
Step 1. Download MobileNet and Identify the Layers to Use	308
Step 2. Create Your Own Model Architecture with the Outputs from MobileNet as Its Input	310
Step 3. Gather and Format the Data	312
Step 4. Train the Model	318
Step 5. Run Inference with the Model	319
Transfer Learning from TensorFlow Hub	321
Using Models from TensorFlow.org	324
Summary	326
19. Deployment with TensorFlow Serving.....	329
What Is TensorFlow Serving?	329
Installing TensorFlow Serving	332
Installing Using Docker	332
Installing Directly on Linux	333
Building and Serving a Model	334
Exploring Server Configuration	337
Summary	340
20. AI Ethics, Fairness, and Privacy.....	341
Fairness in Programming	342
Fairness in Machine Learning	345

Tools for Fairness	347
The What-If Tool	347
Facets	348
Federated Learning	351
Step 1. Identify Available Devices for Training	351
Step 2. Identify Suitable Available Devices for Training	352
Step 3. Deploy a Trainable Model to Your Training Set	352
Step 4. Return the Results of the Training to the Server	353
Step 5. Deploy the New Master Model to the Clients	353
Secure Aggregation with Federated Learning	354
Federated Learning with TensorFlow Federated	355
Google's AI Principles	356
Summary	357
Index.....	359