

Contents

<i>Preface</i>	<i>page</i> xiii
<i>Acknowledgments</i>	xix
<i>About the Author</i>	xxi

Part I Basic Concepts

1 Teaching Computers to Write Programs	3
1.1 What Is Machine Learning?	3
1.2 Branches of Machine Learning	9
1.3 What Do You Need for Machine Learning?	11
1.3.1 Skills	11
1.3.2 Tools	13
1.4 Computational Thinking	15
1.5 Issues of Ethics, Bias, and Privacy in Machine Learning	18
1.6 Summary	19
Key Terms	20
Conceptual Questions	20
Hands-on Problems	21
Further Reading and Resources	25
2 Python	26
2.1 Introduction	26
2.2 Getting Access to Python	26
2.2.1 Download and Install Python	26
2.2.2 Running Python through the Console	27
2.2.3 Using Python through an Integrated Development Environment	27
2.3 Basic Examples	29
2.4 Control Structures	32
2.5 Functions	34
2.6 Making Python Interactive	36
2.7 Installing and Using Python Packages	36
2.7.1 Spyder	37
2.7.2 Jupyter Notebook	37
2.7.3 Google Colab	37

2.8 Statistics Essentials	39
2.8.1 Importing Data	42
2.8.2 Plotting the Data	43
2.8.3 Correlation	44
2.9 Summary	45
Key Terms	45
Conceptual Questions	46
Hands-on Problems	46
Further Reading and Resources	48
3 Cloud Computing	49
3.1 Cloud Computing	49
3.2 Google Cloud Platform	50
3.2.1 Hadoop	53
3.3 Microsoft Azure	59
3.4 Amazon Web Services	66
3.5 Moving between Cloud Platforms	75
3.6 Summary	75
Key Terms	76
Conceptual Questions	76
Hands-on Problems	76
Further Reading and Resources	77
Part II Supervised Learning	
4 Regression	81
4.1 Introduction	81
4.2 Linear Regression	82
4.3 Multiple Linear Regression	91
4.4 Ridge and Lasso Regression	96
4.5 Gradient Descent	98
4.6 Considerations for ML Modeling	110
4.6.1 Accuracy	111
4.6.2 Training–Testing	111
4.6.3 Training Time	111
4.6.4 Linearity	111
4.6.5 Number of Hyperparameters	112
4.6.6 Number of Features	112
4.6.7 Choosing the Right Estimator	112
4.7 Summary	113
Key Terms	113
Conceptual Questions	114
Hands-on Problems	114
Further Reading and Resources	117

5 Classification 1	118
5.1 Introduction	118
5.2 Classification with k -Nearest Neighbors	119
5.3 Decision Tree	123
5.3.1 Decision Rule	127
5.3.2 Classification Rule	127
5.3.3 Association Rule	128
5.4 Random Forest	132
5.5 Summary	137
Key Terms	137
Conceptual Questions	138
Hands-on Problems	138
Further Reading and Resources	143
6 Classification 2	144
6.1 Introduction	144
6.2 Logistic Regression	145
6.3 Softmax Regression	152
6.4 Naïve Bayes	154
6.5 Support Vector Machine	159
6.6 Summary	165
Key Terms	166
Conceptual Questions	166
Hands-on Problems	167
Further Reading and Resources	170
Part III Unsupervised Learning	
7 Clustering	175
7.1 Introduction	175
7.2 Divisive Clustering	176
7.2.1 Divisive Clustering with k -Means	179
7.2.2 Divisive Clustering with k -Modes	182
7.3 Agglomerative Clustering	186
7.4 Expectation Maximization	192
7.5 Density Estimation	196
7.6 Summary	199
Key Terms	200
Conceptual Questions	200
Hands-on Problems	200
Further Reading and Resources	204

8 Dimensionality Reduction	206
8.1 Introduction	206
8.2 Feature Selection	207
8.3 Principal Component Analysis	214
8.4 Linear Discriminant Analysis	218
8.5 Summary	221
Key Terms	222
Conceptual Questions	222
Hands-on Problems	222
Further Reading and Resources	223
Part IV Neural Networks	
9 Neural Networks	227
9.1 Mimicking the Human Brain	227
9.1.1 Basics of a Neural Network	228
9.1.2 Human vs. Artificial Neuron	228
9.2 Architectures of Basic Neural Networks	229
9.2.1 Feed-Forward Network	231
9.2.2 Perceptrons	231
9.2.3 Single-Layer Perceptron	232
9.2.4 Multilayer Perceptron	235
9.3 Autoencoders	243
9.4 Convolution Neural Network	247
9.5 Recurrent Neural Network	251
9.6 Long Short-Term Memory	256
9.7 Summary	259
Key Terms	260
Conceptual Questions	260
Hands-on Problems	261
Further Reading and Resources	263
10 Deep Learning	264
10.1 Introduction	264
10.2 When and How to Use a Deep Model	265
10.3 Building Your First Deep Learning Model with Python	266
10.4 Building a Classifier with Deep Learning	272
10.5 Using Deep Models for Embeddings	278
10.6 Encoders and Transformers	283
10.6.1 Input Embeddings	285
10.6.2 Positional Encoding	285
10.6.3 Encoder Layer	286
10.6.4 Multi-headed Attention	286
10.7 Summary	297

Key Terms	298
Conceptual Questions	298
Hands-on Problems	299
Further Reading and Resources	300

Part V Further Explorations

11 Reinforcement Learning	303
11.1 Introduction	303
11.2 Conceptualizing RL	304
11.3 Doing RL by Hand	306
11.4 Q -learning	311
11.5 Bandits	315
11.5.1 Multi-arm Bandits	316
11.5.2 Contextual Bandits	316
11.6 Summary	322
Key Terms	323
Conceptual Questions	323
Hands-on Problems	323
Further Reading and Resources	324
12 Designing and Evaluating ML Systems	326
12.1 Introduction	326
12.2 Thinking through an ML Solution	327
12.3 Evaluation	331
12.3.1 Offline Evaluation	331
12.3.2 User Research	335
12.3.3 A/B Testing	337
12.3.4 Counterfactual Evaluation	340
12.3.5 Adversarial Learning	344
12.4 Summary	344
Key Terms	345
Conceptual Questions	345
Hands-on Problems	346
Further Reading and Resources	347
13 Responsible AI	348
13.1 Introduction	348
13.2 Diversity vs. Misinformation: A Case Study	349
13.3 Responsible AI in Academia, Industry, and Regulations	352
13.4 Bias and Fairness	353
13.4.1 Where Does Bias Come From?	353
13.4.2 Bias Is Everywhere	354
13.4.3 Bias vs. Fairness	355

13.5	Accountability and Auditing	356
13.6	Transparency and Explainability	358
13.7	Summary	361
	Key Terms	363
	Conceptual Questions	363
	Hands-on Problems	363
	Further Reading and Resources	364
	<i>Appendices</i>	367
	<i>Index</i>	405