

# **Machine Learning in Business:**

## **An Introduction to the World of Data Science**

John C. Hull

*University Professor  
Joseph L. Rotman School of Management  
University of Toronto*

# Contents

<b>Preface</b>	<b>x</b>
<b>Chapter 1 Introduction</b>	<b>1</b>
1.1 This book and the ancillary material	3
1.2 Types of machine learning models	4
1.3 Validation and testing	7
1.4 Data cleaning	14
1.5 Bayes theorem	16
Summary	19
Short concept questions	20
<b>Exercises</b>	<b>21</b>
<b>Chapter 2 Unsupervised Learning</b>	<b>23</b>
2.1 Feature scaling	24
2.2 The $k$ -means algorithm	25
2.3 Choosing $k$	28
2.4 The curse of dimensionality	31
2.5 Country risk	32
2.6 Alternative clustering algorithms	36
2.7 Principal components analysis	40
Summary	44
Short concept questions	46
Exercises	46
<b>Chapter 3 Supervised Learning: Linear Regression</b>	<b>49</b>
3.1 Linear regression: one feature	50
3.2 Linear regression: multiple features	51
3.3 Categorical features	54
3.4 Regularization	55
3.5 Ridge regression	55
3.6 Lasso regression	60
3.7 Elastic Net regression	63
3.8 Results for house price data	64
3.9 Logistic regression	70

3.10 Accuracy of logistic regression	72
3.11 Application to credit decisions	73
3.12 The $k$ -nearest neighbor algorithm	79
Summary	79
Short concept questions	81
Exercises	81
<b>Chapter 4 Supervised Learning: Decision Trees</b>	<b>83</b>
4.1 Nature of decision trees	84
4.2 Information gain measures	84
4.3 Application to credit decisions	87
4.4 The na'ive Bayes classifier	92
4.5 Continuous target variables	97
4.6 Ensemble methods	100
Summary	102
Short concept questions	104
Exercises	104
<b>Chapter 5 Supervised Learning: SVMs</b>	<b>105</b>
5.1 Linear SVM classification	105
5.2 Modification for soft margin	111
5.3 Non-linear separation	113
5.4 Predicting a continuous variable	115
Summary	118
Short concept questions	119
Exercises	120
<b>Chapter 6 Supervised Learning: Neural Networks</b>	<b>121</b>
6.1 Single layer ANNs	121
6.2 Multi-layer ANNs	125
6.3 Gradient descent algorithm	126
6.4 Variations on gradient descent	131
6.5 The stopping rule	133
6.6 Applications to derivatives	133
6.7 Convolutional neural networks	135
6.8 Recurrent neural networks	136
Summary	137
Short concept questions	138
Exercises	139
Appendix: Backpropagation	141

<b>Chapter 7</b>	<b>Reinforcement Learning</b>	<b>143</b>
	7.1 The multi-armed bandit problem	144
	7.2 Q-learning	148
	7.3 The game of Nim	150
	7.4 Temporal difference learning	154
	7.5 Deep Q-learning	155
	7.6 Applications	156
	Summary	158
	Short concept questions	159
	Exercises	160
<b>Chapter 8</b>	<b>Issues for Society</b>	<b>161</b>
	8.1 Data privacy	162
	8.2 Biases	163
	8.3 Ethics	164
	8.4 Transparency	166
	8.5 Adversarial machine learning	167
	8.6 Legal issues	168
	8.7 Man vs. machine	169
	<b>Answers to End of Chapter Questions</b>	<b>171</b>
	<b>Glossary of Terms</b>	<b>185</b>
	<b>Index</b>	<b>191</b>