

# Contents

Preface	xi
Acknowledgments	xv
List of Figures	xvii
List of Symbols	xix
<b>1 Introduction</b>	<b>1</b>
1.1 The Online Convex Optimization Setting	1
1.2 Examples of Problems That Can Be Modeled via Online Convex Optimization	3
1.3 A Gentle Start: Learning from Expert Advice	7
1.4 Bibliographic Remarks	12
1.5 Exercises	13
<b>2 Basic Concepts in Convex Optimization</b>	<b>15</b>
2.1 Basic Definitions and Setup	15
2.2 Gradient Descent	19
2.3 Constrained Gradient/Subgradient Descent	25
2.4 Reductions to Nonsmooth and Nonstrongly Convex Functions	27
2.5 Example: Support Vector Machine Training	31
2.6 Bibliographic Remarks	33
2.7 Exercises	34
<b>3 First-Order Algorithms for Online Convex Optimization</b>	<b>37</b>
3.1 Online Gradient Descent	38
3.2 Lower Bounds	40
3.3 Logarithmic Regret	41
3.4 Application: Stochastic Gradient Descent	43
3.5 Bibliographic Remarks	46
3.6 Exercises	46

<b>4</b>	<b>Second-Order Methods</b>	49
4.1	Motivation: Universal Portfolio Selection	49
4.2	Exp-Concave Functions	52
4.3	Exponentially Weighted Online Convex Optimization	53
4.4	The Online Newton Step Algorithm	55
4.5	Bibliographic Remarks	60
4.6	Exercises	61
<b>5</b>	<b>Regularization</b>	63
5.1	Regularization Functions	64
5.2	The RFTL Algorithm and Its Analysis	65
5.3	Online Mirror Descent	69
5.4	Application and Special Cases	72
5.5	Randomized Regularization	74
5.6	Adaptive Gradient Descent	81
5.7	Bibliographic Remarks	86
5.8	Exercises	87
<b>6</b>	<b>Bandit Convex Optimization</b>	89
6.1	The Bandit Convex Optimization Setting	89
6.2	The Multiarmed Bandit (MAB) Problem	90
6.3	A Reduction from Limited Information to Full Information	94
6.4	Online Gradient Descent without a Gradient	99
6.5	Optimal Regret Algorithms for Bandit Linear Optimization	101
6.6	Bibliographic Remarks	105
6.7	Exercises	106
<b>7</b>	<b>Projection-Free Algorithms</b>	107
7.1	Review: Relevant Concepts from Linear Algebra	107
7.2	Motivation: Recommender Systems	108
7.3	The Conditional Gradient Method	110
7.4	Projections versus Linear Optimization	114
7.5	The Online Conditional Gradient Algorithm	116
7.6	Bibliographic Remarks	119
7.7	Exercises	120
<b>8</b>	<b>Games, Duality, and Regret</b>	123
8.1	Linear Programming and Duality	124
8.2	Zero-Sum Games and Equilibria	125

8.3	Proof of von Neumann Theorem	128
8.4	Approximating Linear Programs	130
8.5	Bibliographic Remarks	131
8.6	Exercises	131
<b>9</b>	<b>Learning Theory, Generalization, and Online Convex Optimization</b>	<b>133</b>
9.1	Statistical Learning Theory	133
9.2	Agnostic Learning Using Online Convex Optimization	138
9.3	Learning and Compression	143
9.4	Bibliographic Remarks	145
9.5	Exercises	145
<b>10</b>	<b>Learning in Changing Environments</b>	<b>147</b>
10.1	A Simple Start: Dynamic Regret	148
10.2	The Notion of Adaptive Regret	149
10.3	Tracking the Best Expert	151
10.4	Efficient Adaptive Regret for Online Convex Optimization	153
10.5	Computationally Efficient Methods	155
10.6	Bibliographic Remarks	159
10.7	Exercises	160
<b>11</b>	<b>Boosting and Regret</b>	<b>163</b>
11.1	The Problem of Boosting	164
11.2	Boosting by Online Convex Optimization	165
11.3	Bibliographic Remarks	169
11.4	Exercises	170
<b>12</b>	<b>Online Boosting</b>	<b>171</b>
12.1	Motivation: Learning from a Huge Set of Experts	171
12.2	The Contextual Learning Model	173
12.3	The Extension Operator	174
12.4	The Online Boosting Method	175
12.5	Bibliographic Remarks	179
12.6	Exercises	180
<b>13</b>	<b>Blackwell Approachability and Online Convex Optimization</b>	<b>181</b>
13.1	Vector-Valued Games and Approachability	181
13.2	From OCO to Approachability	183

13.3 From Approachability to OCO	186
13.4 Bibliographic Remarks	189
13.5 Exercises	190
Notes	191
References	193
Index	207