

Contents

List of Contributors	xi
Preface	xv
Acknowledgments	xvii
Note on the text	xx
List of Abbreviations	xxi
1 Introduction	1
1.1 How do computers learn?	1
1.2 Historical view on learning machines	3
1.3 Learning machines viewed by a statistical physics	5
1.4 Examples of tasks	6
1.5 Types of learning	8
1.6 How to read this book	10
2 Basics of machine learning	14
2.1 Learning as an optimization problem	14
2.2 Generalization and regularization	19
2.3 Probabilistic view on machine learning	23
2.4 Machine learning models	26
2.4.1 Linear (ridge) regression	28
2.4.2 Logistic regression	31
2.4.3 Support vector machines	32
2.4.4 Neural networks	34
2.4.5 Autoencoders	36
2.4.6 Autoregressive neural networks	38
2.5 Backpropagation	39
3 Phase classification	47
3.1 Prototypical physical systems for the study of phases of matter	47
3.1.1 Ising model	48
3.1.2 Ising gauge theory	49
3.2 Unsupervised phase classification without neural networks	50
3.2.1 Principal component analysis	51
3.2.2 t-Distributed stochastic neighbor embedding	54
3.3 Supervised phase classification with neural networks	56
3.4 Unsupervised phase classification with neural networks	57
3.4.1 Learning with autoencoders	57
3.4.2 Learning by confusion	59
3.4.3 Prediction-based method	61

CONTENTS

3.5	Interpretability of machine learning models	63
3.5.1	Difficulty of interpreting parameters of a model	64
3.5.2	Interpretability via bottlenecks	65
3.5.3	Hessian-based interpretability	70
3.5.4	A probabilistic view on phase classification	73
3.6	Outlook and open problems	75
4	Gaussian processes and other kernel methods	76
4.1	The kernel trick	76
4.1.1	Intuition behind the kernel trick	77
4.1.2	The function space as a Hilbert space	79
4.1.3	Reproducing kernel Hilbert spaces	80
4.1.4	The representer theorem	82
4.1.5	Consequences of the kernel trick	83
4.2	Kernel methods	84
4.2.1	Kernel ridge regression	85
4.2.2	Support vector machines	86
4.2.3	Gaussian processes	90
4.2.4	Training a Gaussian process	94
4.3	Bayesian optimization	95
4.4	Choosing the right model	99
4.4.1	Bayesian information criterion	100
4.4.2	Kernel search	101
4.5	Applications in quantum sciences	102
4.5.1	Inverse problems	102
4.5.2	Improving quantum dynamics, physical models, and experiments	105
4.5.3	Extrapolation problems	106
4.5.4	Bayesian optimization of variational quantum algorithms	108
4.6	Outlook and open problems	108
5	Neural-network quantum states	111
5.1	Variational methods	112
5.1.1	Variational states with exact expectation values	112
5.1.2	Variational states with approximate expectation values	114
5.2	Representing the wave function	117
5.2.1	Restricted Boltzmann machines	118
5.2.2	Autoregressive and recurrent neural networks	120
5.2.3	Capacity and entanglement	122
5.2.4	Implementing symmetries	124
5.2.5	Limitations	125
5.3	Applications	126
5.3.1	Finding the ground state	126
5.3.2	Real-time evolution	128
5.3.3	Imaginary-time evolution	130
5.3.4	Fermionic systems	132
5.3.5	Classical simulation of quantum circuits	134

CONTENTS

5.3.6	Open quantum systems	137
5.3.7	Quantum-state tomography	138
5.4	Outlook and open problems	142
6	Reinforcement learning	145
6.1	Foundations of reinforcement learning	146
6.1.1	Delayed rewards	148
6.1.2	Exploration and exploitation	149
6.1.3	Markov decision processes	150
6.1.4	Model-free versus model-based reinforcement learning	151
6.1.5	Value functions and Bellman equations	152
6.2	Value-based methods	154
6.2.1	Q-learning	155
6.2.2	Double Q-learning	157
6.2.3	Implementing Q-learning with a neural network	157
6.3	Policy gradient methods	159
6.3.1	REINFORCE	160
6.3.2	Implementing REINFORCE with a neural network	163
6.4	Actor-critic methods	165
6.5	Projective simulation	167
6.6	Examples and applications	170
6.6.1	Toy examples	170
6.6.2	Go and Atari games	172
6.6.3	Quantum feedback control	174
6.6.4	Quantum circuit optimization	175
6.6.5	Quantum error correction	175
6.6.6	Quantum experiment design	179
6.6.7	Building optimal relaxations	180
6.7	Outlook and open problems	182
7	Deep learning for quantum sciences: Selected topics	184
7.1	Differentiable programming	184
7.1.1	Automatic differentiation	187
7.1.2	Application to quantum physics problems	193
7.1.3	Outlook and open problems	194
7.2	Generative models in many-body physics	195
7.2.1	Training with or without data	196
7.2.2	Taxonomy of deep generative models	198
7.2.3	An important example: normalizing flows	200
7.2.4	Applications	204
7.2.5	Outlook and open problems	204
7.3	Machine learning for experiments	205
7.3.1	Automation of experimental setups	206
7.3.2	Machine-learning analysis of time-of-flight images	212
7.3.3	Hamiltonian learning	214
7.3.4	Automated design of experiments	218
7.3.5	Outlook and open problems	219

CONTENTS

8 Physics for deep learning	221
8.1 Statistical physics for machine learning	221
8.1.1 Capacity of the perceptron	222
8.1.2 The teacher–student paradigm: A toy model to study the generalization	228
8.1.3 Models of data structure	235
8.1.4 Dynamics of learning	237
8.2 Quantum machine learning	239
8.2.1 Gate-based quantum computing	240
8.2.2 What is quantum machine learning?	240
8.2.3 Ideal quantum computers	242
8.2.4 Noisy intermediate-scale quantum era	243
8.2.5 Support vector machines with quantum kernels	244
8.2.6 Variational approaches	246
8.2.7 Parametrized quantum circuits for quantum machine learning	250
8.2.8 Current experimental and theoretical limitations	251
9 Conclusion and outlook	254
Appendix A Mathematical details on principal component analysis	258
Appendix B Derivation of the kernel trick	260
Appendix C Choosing the kernel matrix as the covariance matrix for a Gaussian process	263
References	265
Index	306