

## Contents

## Contents

## Figures

## I Agents in the World: What are Agents and How Can They be Built?

<b>1 Artificial Intelligence and Agents</b>	
1.1 What is Artificial Intelligence? . . . . .	3
1.1.1 Artificial and Natural Intelligence . . . . .	5
1.2 A Brief History of Artificial Intelligence . . . . .	7
1.2.1 Relationship to Other Disciplines . . . . .	10
1.3 Agents Situated in Environments . . . . .	11
1.4 Designing Agents . . . . .	13
1.4.1 Design Time, Offline and Online Computation . . . . .	13
1.4.2 Tasks . . . . .	15
1.4.3 Defining a Solution . . . . .	17
1.4.4 Representations . . . . .	19
1.5 Agent Design Space . . . . .	21
1.5.1 Modularity . . . . .	21
1.5.2 Planning Horizon . . . . .	23
1.5.3 Representation . . . . .	23
1.5.4 Computational Limits . . . . .	25
1.5.5 Learning . . . . .	27

1.5.6	Uncertainty . . . . .	28
1.5.7	Preference . . . . .	29
1.5.8	Number of Agents . . . . .	30
1.5.9	Interaction . . . . .	30
1.5.10	Interaction of the Dimensions . . . . .	31
1.6	Prototypical Applications . . . . .	33
1.6.1	An Autonomous Delivery Robot . . . . .	34
1.6.2	A Diagnostic Assistant . . . . .	36
1.6.3	An Intelligent Tutoring System . . . . .	39
1.6.4	A Trading Agent . . . . .	41
1.6.5	Smart House . . . . .	43
1.7	Overview of the Book . . . . .	44
1.8	Review . . . . .	45
1.9	References and Further Reading . . . . .	46
1.10	Exercises . . . . .	47
<b>2</b>	<b>Agent Architectures and Hierarchical Control</b>	<b>49</b>
2.1	Agents . . . . .	50
2.2	Agent Systems . . . . .	51
2.2.1	The Agent Function . . . . .	51
2.3	Hierarchical Control . . . . .	56
2.4	Acting with Reasoning . . . . .	65
2.4.1	Agents Modeling the World . . . . .	65
2.4.2	Knowledge and Acting . . . . .	66
2.4.3	Design Time and Offline Computation . . . . .	67
2.4.4	Online Computation . . . . .	69
2.5	Review . . . . .	70
2.6	References and Further Reading . . . . .	71
2.7	Exercises . . . . .	71
<b>II</b>	<b>Reasoning, Planning and Learning with Certainty</b>	<b>75</b>
<b>3</b>	<b>Searching for Solutions</b>	<b>77</b>
3.1	Problem Solving as Search . . . . .	77
3.2	State Spaces . . . . .	79
3.3	Graph Searching . . . . .	81
3.3.1	Formalizing Graph Searching . . . . .	82
3.4	A Generic Searching Algorithm . . . . .	85
3.5	Uninformed Search Strategies . . . . .	87
3.5.1	Breadth-First Search . . . . .	87
3.5.2	Depth-First Search . . . . .	90
3.5.3	Iterative Deepening . . . . .	94

3.5.4 Lowest-Cost-First Search . . . . .	97
3.6 Heuristic Search . . . . .	98
3.6.1 $A^*$ Search . . . . .	100
3.6.2 Designing a Heuristic Function . . . . .	104
3.7 Pruning the Search Space . . . . .	105
3.7.1 Cycle Pruning . . . . .	105
3.7.2 Multiple-Path Pruning . . . . .	106
3.7.3 Summary of Search Strategies . . . . .	109
3.8 More Sophisticated Search . . . . .	110
3.8.1 Branch and Bound . . . . .	110
3.8.2 Direction of Search . . . . .	113
3.8.3 Dynamic Programming . . . . .	115
3.9 Review . . . . .	119
3.10 References and Further Reading . . . . .	119
3.11 Exercises . . . . .	120
<b>4 Reasoning with Constraints</b> . . . . .	<b>125</b>
4.1 Possible Worlds, Variables, and Constraints . . . . .	125
4.1.1 Variables and Worlds . . . . .	125
4.1.2 Constraints . . . . .	129
4.1.3 Constraint Satisfaction Problems . . . . .	131
4.2 Generate-and-Test Algorithms . . . . .	132
4.3 Solving CSPs Using Search . . . . .	133
4.4 Consistency Algorithms . . . . .	134
4.5 Domain Splitting . . . . .	139
4.6 Variable Elimination . . . . .	141
4.7 Local Search . . . . .	144
4.7.1 Iterative Best Improvement . . . . .	146
4.7.2 Randomized Algorithms . . . . .	148
4.7.3 Local Search Variants . . . . .	149
4.7.4 Evaluating Randomized Algorithms . . . . .	153
4.7.5 Random Restart . . . . .	156
4.8 Population-Based Methods . . . . .	156
4.9 Optimization . . . . .	160
4.9.1 Systematic Methods for Optimization . . . . .	162
4.9.2 Local Search for Optimization . . . . .	164
4.10 Review . . . . .	167
4.11 References and Further Reading . . . . .	167
4.12 Exercises . . . . .	168
<b>5 Propositions and Inference</b> . . . . .	<b>173</b>
5.1 Propositions . . . . .	173
5.1.1 Syntax of Propositional Calculus . . . . .	174

5.1.2	Semantics of the Propositional Calculus . . . . .	175
5.2	Propositional Constraints . . . . .	179
5.2.1	Clausal Form for Consistency Algorithms . . . . .	180
5.2.2	Exploiting Propositional Structure in Local Search . . . . .	181
5.3	Propositional Definite Clauses . . . . .	182
5.3.1	Questions and Answers . . . . .	185
5.3.2	Proofs . . . . .	186
5.4	Knowledge Representation Issues . . . . .	194
5.4.1	Background Knowledge and Observations . . . . .	194
5.4.2	Querying the User . . . . .	194
5.4.3	Knowledge-Level Explanation . . . . .	196
5.4.4	Knowledge-Level Debugging . . . . .	199
5.5	Proving by Contradiction . . . . .	204
5.5.1	Horn Clauses . . . . .	205
5.5.2	Assumables and Conflicts . . . . .	206
5.5.3	Consistency-Based Diagnosis . . . . .	207
5.5.4	Reasoning with Assumptions and Horn Clauses . . . . .	209
5.6	Complete Knowledge Assumption . . . . .	212
5.6.1	Non-monotonic Reasoning . . . . .	216
5.6.2	Proof Procedures for Negation as Failure . . . . .	217
5.7	Abduction . . . . .	220
5.8	Causal Models . . . . .	225
5.9	Review . . . . .	226
5.10	References and Further Reading . . . . .	227
5.11	Exercises . . . . .	228
<b>6</b>	<b>Planning with Certainty</b>	<b>239</b>
6.1	Representing States, Actions, and Goals . . . . .	240
6.1.1	Explicit State-Space Representation . . . . .	241
6.1.2	The STRIPS Representation . . . . .	243
6.1.3	Feature-Based Representation of Actions . . . . .	244
6.1.4	Initial States and Goals . . . . .	246
6.2	Forward Planning . . . . .	246
6.3	Regression Planning . . . . .	249
6.4	Planning as a CSP . . . . .	252
6.4.1	Action Features . . . . .	255
6.5	Partial-Order Planning . . . . .	257
6.6	Review . . . . .	260
6.7	References and Further Reading . . . . .	261
6.8	Exercises . . . . .	262
<b>7</b>	<b>Supervised Machine Learning</b>	<b>267</b>
7.1	Learning Issues . . . . .	268

7.2	Supervised Learning . . . . .	271
7.2.1	Evaluating Predictions . . . . .	274
7.2.2	Types of Errors . . . . .	279
7.2.3	Point Estimates with No Input Features . . . . .	283
7.3	Basic Models for Supervised Learning . . . . .	285
7.3.1	Learning Decision Trees . . . . .	285
7.3.2	Linear Regression and Classification . . . . .	291
7.4	Overfitting . . . . .	298
7.4.1	Pseudocounts . . . . .	301
7.4.2	Regularization . . . . .	304
7.4.3	Cross Validation . . . . .	306
7.5	Neural Networks and Deep Learning . . . . .	308
7.6	Composite Models . . . . .	316
7.6.1	Random Forests . . . . .	317
7.6.2	Ensemble Learning . . . . .	318
7.7	Case-Based Reasoning . . . . .	320
7.8	Learning as Refining the Hypothesis Space . . . . .	323
7.8.1	Version-Space Learning . . . . .	325
7.8.2	Probably Approximately Correct Learning . . . . .	328
7.9	Review . . . . .	331
7.10	References and Further Reading . . . . .	331
7.11	Exercises . . . . .	333
<b>III Reasoning, Learning and Acting with Uncertainty</b>		<b>341</b>
<b>8</b>	<b>Reasoning with Uncertainty</b>	<b>343</b>
8.1	Probability . . . . .	343
8.1.1	Semantics of Probability . . . . .	345
8.1.2	Axioms for Probability . . . . .	347
8.1.3	Conditional Probability . . . . .	350
8.1.4	Expected Values . . . . .	355
8.1.5	Information . . . . .	356
8.2	Independence . . . . .	358
8.3	Belief Networks . . . . .	360
8.3.1	Observations and Queries . . . . .	362
8.3.2	Constructing Belief Networks . . . . .	363
8.4	Probabilistic Inference . . . . .	370
8.4.1	Variable Elimination for Belief Networks . . . . .	372
8.4.2	Representing Conditional Probabilities and Factors . . . . .	381
8.5	Sequential Probability Models . . . . .	384
8.5.1	Markov Chains . . . . .	384
8.5.2	Hidden Markov Models . . . . .	387

8.5.3	Algorithms for Monitoring and Smoothing . . . . .	392
8.5.4	Dynamic Belief Networks . . . . .	393
8.5.5	Time Granularity . . . . .	394
8.5.6	Probabilistic Models of Language . . . . .	395
8.6	Stochastic Simulation . . . . .	402
8.6.1	Sampling from a Single Variable . . . . .	403
8.6.2	Forward Sampling in Belief Networks . . . . .	404
8.6.3	Rejection Sampling . . . . .	405
8.6.4	Likelihood Weighting . . . . .	407
8.6.5	Importance Sampling . . . . .	408
8.6.6	Particle Filtering . . . . .	410
8.6.7	Markov Chain Monte Carlo . . . . .	412
8.7	Review . . . . .	414
8.8	References and Further Reading . . . . .	414
8.9	Exercises . . . . .	415
<b>9</b>	<b>Planning with Uncertainty</b>	<b>425</b>
9.1	Preferences and Utility . . . . .	426
9.1.1	Axioms for Rationality . . . . .	426
9.1.2	Factored Utility . . . . .	433
9.1.3	Prospect Theory . . . . .	435
9.2	One-Off Decisions . . . . .	438
9.2.1	Single-Stage Decision Networks . . . . .	442
9.3	Sequential Decisions . . . . .	444
9.3.1	Decision Networks . . . . .	445
9.3.2	Policies . . . . .	449
9.3.3	Variable Elimination for Decision Networks . . . . .	451
9.4	The Value of Information and Control . . . . .	455
9.5	Decision Processes . . . . .	458
9.5.1	Policies . . . . .	462
9.5.2	Value Iteration . . . . .	464
9.5.3	Policy Iteration . . . . .	468
9.5.4	Dynamic Decision Networks . . . . .	470
9.5.5	Partially Observable Decision Processes . . . . .	474
9.6	Review . . . . .	475
9.7	References and Further Reading . . . . .	476
9.8	Exercises . . . . .	476
<b>10</b>	<b>Learning with Uncertainty</b>	<b>487</b>
10.1	Probabilistic Learning . . . . .	487
10.1.1	Learning Probabilities . . . . .	488
10.1.2	Probabilistic Classifiers . . . . .	491
10.1.3	MAP Learning of Decision Trees . . . . .	496

## Contents

10.1.4	Description Length . . . . .	498
10.2	Unsupervised Learning . . . . .	499
10.2.1	$k$ -Means . . . . .	499
10.2.2	Expectation Maximization for Soft Clustering . . . . .	503
10.3	Learning Belief Networks . . . . .	507
10.3.1	Learning the Probabilities . . . . .	508
10.3.2	Hidden Variables . . . . .	509
10.3.3	Missing Data . . . . .	509
10.3.4	Structure Learning . . . . .	510
10.3.5	General Case of Belief Network Learning . . . . .	512
10.4	Bayesian Learning . . . . .	512
10.5	Review . . . . .	517
10.6	References and Further Reading . . . . .	518
10.7	Exercises . . . . .	518
<b>11</b>	<b>Multiagent Systems</b>	<b>521</b>
11.1	Multiagent Framework . . . . .	521
11.2	Representations of Games . . . . .	523
11.2.1	Normal Form Games . . . . .	523
11.2.2	Extensive Form of a Game . . . . .	524
11.2.3	Multiagent Decision Networks . . . . .	527
11.3	Computing Strategies with Perfect Information . . . . .	528
11.4	Reasoning with Imperfect Information . . . . .	532
11.4.1	Computing Nash Equilibria . . . . .	538
11.5	Group Decision Making . . . . .	541
11.6	Mechanism Design . . . . .	542
11.7	Review . . . . .	544
11.8	References and Further Reading . . . . .	545
11.9	Exercises . . . . .	545
<b>12</b>	<b>Learning to Act</b>	<b>549</b>
12.1	Reinforcement Learning Problem . . . . .	549
12.2	Evolutionary Algorithms . . . . .	553
12.3	Temporal Differences . . . . .	554
12.4	$Q$ -learning . . . . .	555
12.5	Exploration and Exploitation . . . . .	557
12.6	Evaluating Reinforcement Learning Algorithms . . . . .	559
12.7	On-Policy Learning . . . . .	560
12.8	Model-Based Reinforcement Learning . . . . .	562
12.9	Reinforcement Learning with Features . . . . .	565
12.9.1	SARSA with Linear Function Approximation . . . . .	565
12.10	Multiagent Reinforcement Learning . . . . .	569
12.10.1	Perfect-Information Games . . . . .	569

12.10.2 Learning to Coordinate . . . . .	569
12.11 Review . . . . .	574
12.12 References and Further Reading . . . . .	574
12.13 Exercises . . . . .	575
<b>IV Reasoning, Learning and Acting with Individuals and Relations</b>	<b>579</b>
<b>13 Individuals and Relations</b>	<b>581</b>
13.1 Exploiting Relational Structure . . . . .	582
13.2 Symbols and Semantics . . . . .	583
13.3 Datalog: A Relational Rule Language . . . . .	584
13.3.1 Semantics of Ground Datalog . . . . .	587
13.3.2 Interpreting Variables . . . . .	589
13.3.3 Queries with Variables . . . . .	595
13.4 Proofs and Substitutions . . . . .	597
13.4.1 Instances and Substitutions . . . . .	597
13.4.2 Bottom-up Procedure with Variables . . . . .	599
13.4.3 Unification . . . . .	601
13.4.4 Definite Resolution with Variables . . . . .	602
13.5 Function Symbols . . . . .	604
13.5.1 Proof Procedures with Function Symbols . . . . .	610
13.6 Applications in Natural Language . . . . .	612
13.6.1 Using Definite Clauses for Context-Free Grammars . . . . .	614
13.6.2 Augmenting the Grammar . . . . .	618
13.6.3 Building Structures for Non-terminals . . . . .	619
13.6.4 Canned Text Output . . . . .	619
13.6.5 Enforcing Constraints . . . . .	620
13.6.6 Building a Natural Language Interface to a Database . . . . .	621
13.6.7 Limitations . . . . .	627
13.7 Equality . . . . .	628
13.7.1 Allowing Equality Assertions . . . . .	629
13.7.2 Unique Names Assumption . . . . .	630
13.8 Complete Knowledge Assumption . . . . .	633
13.8.1 Complete Knowledge Assumption Proof Procedures . . . . .	637
13.9 Review . . . . .	638
13.10 References and Further Reading . . . . .	638
13.11 Exercises . . . . .	639
<b>14 Ontologies and Knowledge-Based Systems</b>	<b>645</b>
14.1 Knowledge Sharing . . . . .	645
14.2 Flexible Representations . . . . .	646

14.2.1 Choosing Individuals and Relations . . . . .	647
14.2.2 Graphical Representations . . . . .	650
14.2.3 Classes . . . . .	652
14.3 Ontologies and Knowledge Sharing . . . . .	655
14.3.1 Uniform Resource Identifiers . . . . .	661
14.3.2 Description Logic . . . . .	662
14.3.3 Top-Level Ontologies . . . . .	670
14.4 Implementing Knowledge-Based Systems . . . . .	673
14.4.1 Base Languages and Metalanguages . . . . .	674
14.4.2 A Vanilla Meta-Interpreter . . . . .	676
14.4.3 Expanding the Base Language . . . . .	678
14.4.4 Depth-Bounded Search . . . . .	680
14.4.5 Meta-Interpreter to Build Proof Trees . . . . .	681
14.4.6 Delaying Goals . . . . .	682
14.5 Review . . . . .	684
14.6 References and Further Reading . . . . .	684
14.7 Exercises . . . . .	685
<b>15 Relational Planning, Learning, and Probabilistic Reasoning</b>	<b>691</b>
15.1 Planning with Individuals and Relations . . . . .	692
15.1.1 Situation Calculus . . . . .	692
15.1.2 Event Calculus . . . . .	699
15.2 Relational Learning . . . . .	701
15.2.1 Structure Learning: Inductive Logic Programming . . . . .	701
15.2.2 Learning Hidden Properties: Collaborative Filtering . . . . .	706
15.3 Statistical Relational Artificial Intelligence . . . . .	711
15.3.1 Relational Probabilistic Models . . . . .	711
15.4 Review . . . . .	724
15.5 References and Further Reading . . . . .	724
15.6 Exercises . . . . .	725
<b>V Retrospect and Prospect</b>	<b>729</b>
<b>16 Retrospect and Prospect</b>	<b>731</b>
16.1 Dimensions of Complexity Revisited . . . . .	731
16.2 Social and Ethical Consequences . . . . .	736
16.3 References and Further Reading . . . . .	742
16.4 Exercises . . . . .	742
<b>A Mathematical Preliminaries and Notation</b>	<b>745</b>
A.1 Discrete Mathematics . . . . .	745
A.2 Functions, Factors and Arrays . . . . .	746

