

Contents

I	Introducing Artificial Intelligence	
1	The Pre-History of Artificial Intelligence	3
1.1	Mary Shelley, Frankenstein, and Prometheus	4
1.2	The Age of Rationalism	6
1.3	The Empiricist Tradition	8
1.4	Immanuel Kant: Bridging the Rationalist/Empiricist Viewpoints	9
1.5	The Reverend Thomas Bayes	11
1.6	The Mathematical Foundations for Artificial Intelligence	11
1.7	American Pragmatism	16
1.8	The Turing Test: Can a Machine Be “Intelligent”?	17
1.9	The 1956 Dartmouth Summer Workshop	20
1.10	Summary	24
1.11	Exercises	24
2	Computing, Representations, and Definitions of Artificial Intelligence	27
2.1	Artificial Intelligence: Attempting a Definition	28
2.2	Computer-Based Representations of the World	30
2.3	The General Themes of Current AI Practice	40
2.4	Summary of Part I and an Introduction to Part II	48
2.5	Exercises	49
II	Symbol-Based AI: Foundations	
3	The State Space, Finite State Machines, and Artificial Life	53
3.1	Graph Theory: The Origins of the State Space Representation	54
3.2	The State Space Representation	58
3.3	The Finite State Machine	64
3.4	Artificial Life: The Emergence of Complexity	67
3.5	Summary	71
3.6	Exercises	72
4	Searching the State Space	73
4.1	Data-Driven and Goal-Driven Search	74
4.2	Implementing Graph Search with the Backtrack Algorithm	77
4.3	Breadth-First and Depth-First Search	80
4.4	Extending Search Strategies to and/or Graphs	88
4.5	Summary	94
4.6	Exercises	95
5	Heuristic Search	97
5.1	An Introduction to Heuristics	98
5.2	Hill-Climbing	101
5.3	Best-First Search	104

5.4	Genetic Algorithms: Evaluating Multiple “Best” States	113
5.5	In Summary	121
5.6	Exercises	122
6	Heuristics: 2-Person Games and Theoretical Constraints	125
6.1	The Minimax Procedure on Exhaustively Searchable Graphs.....	126
6.2	Using Minimax to a Fixed Ply Depth	128
6.3	The Alpha-Beta Procedure	133
6.4	Multi-Person Games, Admissibility, Monotonicity, and Informedness.....	134
6.5	Heuristics and Complexity	141
6.6	Summary	145
6.7	Exercises	145
III	The Propositional and Predicate Calculi and Resolution-Based Reasoning	
7	Introduction to the Propositional and Predicate Calculi.....	151
7.1	The Propositional Calculus.....	152
7.2	The Predicate Calculus.....	156
7.3	Summary	169
7.4	Exercises	169
8	The Predicate Calculus and Unification.....	171
8.1	Using Reasoning Rules to Produce New Predicate Expressions.....	172
8.2	The Unification Algorithm.....	176
8.3	Summary	189
8.4	Exercises	189
9	Resolution: Reasoning with the Propositional and Predicate Calculi.....	191
9.1	Introduction to Resolution.....	192
9.2	Resolution Refutation Systems.....	195
9.3	Strategies and Simplification Techniques for Resolution.....	204
9.4	Extracting Answers from Resolution Refutations	209
9.5	Logic Programming and Prolog	212
9.6	Summary	219
9.7	Exercises	219
IV	Advanced Applications of Symbol-Based AI	
10	The Production System Representation and Search Engine.....	223
10.1	The Production System: An Architecture for Organizing Search.....	224
10.2	More Examples of Production System Problem-Solving.....	227
10.3	The Expert System	233
10.4	The Physical Symbol System Hypothesis and the Birth of Cognitive Science.....	238
10.5	Summary	240
10.6	Exercises	241

11	Advanced Applications of Symbol-Based AI: Planning and Learning	243
11.1	Introduction to Planning and Robotics	244
11.2	Using Planning Macros: STRIPS	250
11.3	Model-Based Planning: A NASA Example (Williams and Nayak 1996, 1997)	252
11.4	Symbol-Based Learning	256
11.5	Bacon: Modeling the Celestial Environment	268
11.6	Expertise Wherever It Is Needed	269
11.7	Summary	271
11.8	Exercises	271
12	Uncertain Reasoning: Symbol Based	273
12.1	Logic-Based Reasoning in Uncertain Situations	274
12.2	Uncertain Reasoning: Alternatives to Logic	286
12.3	Summary and Pointers to Parts VI and VII	295
12.4	Exercises	295
V	Symbol-Based Associational Models for AI	
13	Introduction to Association-Based Knowledge Representations	299
13.1	The Behaviorist Tradition and Semantic Networks	300
13.2	Conceptual Dependencies	307
13.3	Scripts	313
13.4	Summary	318
13.5	Exercises	318
14	Association-Based Representations: Frames, Conceptual Graphs, WordNet, and FrameNet	321
14.1	Frames	322
14.2	Conceptual Graphs	326
14.3	Current Uses of Association-Based Representations	336
14.4	Summary	338
14.5	Exercises	338
VI	Neural or Connectionist Networks	
15	An Introduction to Neural Networks	345
15.1	An Artificial Neuron and Applications	346
15.2	Early Research: McCulloch, Pitts, and Hebb	348
15.3	Perceptrons	354
15.4	Summary	360
15.5	Exercises	361
16	The Delta Rule, Backpropagation, and Matrix Representations	363
16.1	The Generalized Delta Rule	364
16.2	The Backpropagation Algorithm	368
16.3	Matrix Representations for Network Processing	374
16.4	Matrix Representations and Neural Network Solutions	377
16.5	Summary	382
16.6	Exercises	382

17	Deep Learning: Introduction and Representations	383
17.1	Toward Deep Learning	384
17.2	Meta-parameters for Very Large Networks	392
17.3	Alternative Architectures: Convolutional and Recurrent Networks	397
17.4	Autoencoders and Transfer Learning	403
17.5	Summary	406
17.6	Exercises	407
18	Building Language Models and Transformers	409
18.1	Latent Semantic Analysis: Distributed Semantic Representations	410
18.2	Building Large Language Models	412
18.3	Toward Transformer-Based Large Language Models	417
18.4	The Transformer in Practice	428
18.5	Summary	437
18.6	Exercises	438
19	Alternative Network Architectures: Prototypes and Classifiers	441
19.1	A Kohonen Network: Winner-Takes-All Classification	442
19.2	A Kohonen Network: Learning Prototypes	444
19.3	Outstar Networks and Counterpropagation	446
19.4	Supervised Hebbian Learning	450
19.5	Associative Memories and the Linear Associator	452
19.6	Summary	457
19.7	Exercises	457
20	Alternative Network Architectures: Attractor Networks and Memories	459
20.1	Introduction to Associative Memories	460
20.2	BAM, the Bidirectional Associative Memory	461
20.3	Autoassociative Memories and Hopfield Networks	467
20.4	Summary	472
20.5	Exercises	472
VII	Probabilistic Artificial Intelligence	
21	Counting, the Foundation for Probabilities	475
21.1	Introduction to Probabilistic Reasoning	476
21.2	The Elements of Counting	478
21.3	Elements of Probability Theory	482
21.4	Summary	488
21.5	Exercises	488
22	Bayes' Theorem	491
22.1	Random Variables	492
22.2	Conditional Probability and an Introduction to Bayesian Reasoning	495
22.3	Bayes' Theorem	499
22.4	Two Examples of Bayesian Reasoning	503
22.5	Summary	505
22.6	Exercises	505

23	Bayesian Belief Networks and Observable Markov Models	507
23.1	Introduction to Stochastic Models	508
23.2	A Directed Graphical Model: The Bayesian Belief Network	509
23.3	Dynamic Bayesian Networks	517
23.4	Observable Markov Models	518
23.5	Summary	522
23.6	Exercises	522
24	Hidden Markov and Alternative Probabilistic Models	525
24.1	Hidden Markov Models	526
24.2	Important Variants of Hidden Markov Models	529
24.3	A Short Survey of Alternative Markov Models	534
24.4	First-Order Alternatives to BBNs and HMMs	536
24.5	Two Stochastic Engineering Examples: An MDP and a DBN	539
24.6	Summary	544
24.7	Exercises	545
 VIII AI: Ethical Issues, Fundamental Limitations, and Future Promise		
25	Artificial Intelligence: User’s Ethical Issues	551
25.1	Artificial and Human Intelligence	552
25.2	Ethical Issues from the Perspective of the AI User	556
25.3	We Are Already Embedded in an AI World	563
25.4	Summary	570
25.5	Exercises	571
26	AI Ethical Issues: From a Social Perspective	573
26.1	Building Society-Oriented AI Projects	574
26.2	If We Do Build These AI Projects, What Levels of Protection Are Needed?	577
26.3	What Is Data?	581
26.4	Data and Algorithm Transparency and Explainability	582
26.5	Design Honesty, Transparency, Explanations, Data Bias, and Privacy	592
26.6	Safeguarding AI Research and Practice	595
26.7	Summary	597
26.8	Exercises	598
27	AI: Philosophical Perspectives, Current Limitations, and Future Promise	601
27.1	Modern AI: A Psychological, Mathematical, and Philosophical Perspective	602
27.2	Several Fundamental Limitations of Current AI Technology	607
27.3	Artificial Intelligence: Where Are We Going?	610
	Supplementary Information	
	Bibliography	616
	Index	633