

Contents

<i>Preface</i>	<i>page xv</i>
<i>Acknowledgments</i>	xxi
<i>About the Authors</i>	xxiii
1 Introduction	1
1.1 Understanding the World through Evidence-based Reasoning	1
1.1.1 What Is Evidence?	1
1.1.2 Evidence, Data, and Information	1
1.1.3 Evidence and Fact	2
1.1.4 Evidence and Knowledge	2
1.1.5 Ubiquity of Evidence	5
1.2 Abductive Reasoning	5
1.2.1 From Aristotle to Peirce	5
1.2.2 Peirce and Sherlock Holmes on Abductive Reasoning	6
1.3 Probabilistic Reasoning	9
1.3.1 Enumerative Probabilities: Obtained by Counting	9
1.3.1.1 Aleatory Probability	9
1.3.1.2 Relative Frequency and Statistics	9
1.3.2 Subjective Bayesian View of Probability	11
1.3.3 Belief Functions	13
1.3.4 Baconian Probability	16
1.3.4.1 Variative and Eliminative Inferences	16
1.3.4.2 Importance of Evidential Completeness	17
1.3.4.3 Baconian Probability of Boolean Expressions	20
1.3.5 Fuzzy Probability	20
1.3.5.1 Fuzzy Force of Evidence	20
1.3.5.2 Fuzzy Probability of Boolean Expressions	21
1.3.5.3 On Verbal Assessments of Probabilities	22
1.3.6 A Summary of Uncertainty Methods and What They Best Capture	23
1.4 Evidence-based Reasoning	25
1.4.1 Deduction, Induction, and Abduction	25
1.4.2 The Search for Knowledge	26
1.4.3 Evidence-based Reasoning Everywhere	27

1.5	Artificial Intelligence	29
1.5.1	Intelligent Agents	30
1.5.2	Mixed-Initiative Reasoning	32
1.6	Knowledge Engineering	33
1.6.1	From Expert Systems to Knowledge-based Agents and Cognitive Assistants	33
1.6.2	An Ontology of Problem-Solving Tasks	35
1.6.2.1	Analytic Tasks	36
1.6.2.2	Synthetic Tasks	36
1.6.3	Building Knowledge-based Agents	37
1.6.3.1	How Knowledge-based Agents Are Built and Why It Is Hard	37
1.6.3.2	Teaching as an Alternative to Programming: Disciple Agents	39
1.6.3.3	Disciple-EBR, Disciple-CD, and TIACRITIS	40
1.7	Obtaining Disciple-EBR	41
1.8	Review Questions	42
2	Evidence-based Reasoning: Connecting the Dots	46
2.1	How Easy Is It to Connect the Dots?	46
2.1.1	How Many Kinds of Dots Are There?	47
2.1.2	Which Evidential Dots Can Be Believed?	48
2.1.3	Which Evidential Dots Should Be Considered?	50
2.1.4	Which Evidential Dots Should We Try to Connect?	50
2.1.5	How to Connect Evidential Dots to Hypotheses?	52
2.1.6	What Do Our Dot Connections Mean?	54
2.2	Sample Evidence-based Reasoning Task: Intelligence Analysis	56
2.2.1	Evidence in Search of Hypotheses	56
2.2.2	Hypotheses in Search of Evidence	58
2.2.3	Evidentiary Testing of Hypotheses	60
2.2.4	Completing the Analysis	62
2.3	Other Evidence-based Reasoning Tasks	64
2.3.1	Cyber Insider Threat Discovery and Analysis	64
2.3.2	Analysis of Wide-Area Motion Imagery	68
2.3.3	Inquiry-based Teaching and Learning in a Science Classroom	70
2.3.3.1	Need for Inquiry-based Teaching and Learning	70
2.3.3.2	Illustration of Inquiry-based Teaching and Learning	71
2.3.3.3	Other Examples of Inquiry-based Teaching and Learning	74
2.4	Hands On: Browsing an Argumentation	76
2.5	Project Assignment 1	81
2.6	Review Questions	81

3	Methodologies and Tools for Agent Design and Development	83
3.1	A Conventional Design and Development Scenario	83
3.1.1	Conventional Design and Development Phases	83
3.1.2	Requirements Specification and Domain Understanding	83
3.1.3	Ontology Design and Development	85
3.1.4	Development of the Problem-Solving Rules or Methods	86
3.1.5	Verification, Validation, and Certification	87
3.2	Development Tools and Reusable Ontologies	88
3.2.1	Expert System Shells	88
3.2.2	Foundational and Utility Ontologies and Their Reuse	89
3.2.3	Learning Agent Shells	90
3.2.4	Learning Agent Shell for Evidence-based Reasoning	91
3.3	Agent Design and Development Using Learning Technology	93
3.3.1	Requirements Specification and Domain Understanding	93
3.3.2	Rapid Prototyping	93
3.3.3	Ontology Design and Development	100
3.3.4	Rule Learning and Ontology Refinement	101
3.3.5	Hierarchical Organization of the Knowledge Repository	104
3.3.6	Learning-based Design and Development Phases	105
3.4	Hands On: Loading, Saving, and Closing Knowledge Bases	107
3.5	Knowledge Base Guidelines	111
3.6	Project Assignment 2	111
3.7	Review Questions	112
4	Modeling the Problem-Solving Process	113
4.1	Problem Solving through Analysis and Synthesis	113
4.2	Inquiry-driven Analysis and Synthesis	113
4.3	Inquiry-driven Analysis and Synthesis for Evidence-based Reasoning	119
4.3.1	Hypothesis Reduction and Assessment Synthesis	119
4.3.2	Necessary and Sufficient Conditions	120
4.3.3	Sufficient Conditions and Scenarios	120
4.3.4	Indicators	121
4.4	Evidence-based Assessment	122
4.5	Hands On: Was the Cesium Stolen?	124
4.6	Hands On: Hypothesis Analysis and Evidence Search and Representation	130
4.7	Believability Assessment	133
4.7.1	Tangible Evidence	133
4.7.2	Testimonial Evidence	135

4.7.3	Missing Evidence	137
4.7.4	Authoritative Record	137
4.7.5	Mixed Evidence and Chains of Custody	138
4.8	Hands On: Believability Analysis	140
4.9	Drill-Down Analysis, Assumption-based Reasoning, and What-If Scenarios	143
4.10	Hands On: Modeling, Formalization, and Pattern Learning	144
4.11	Hands On: Analysis Based on Learned Patterns	146
4.12	Modeling Guidelines	147
4.13	Project Assignment 3	151
4.14	Review Questions	152
5	Ontologies	155
5.1	What Is an Ontology?	155
5.2	Concepts and Instances	156
5.3	Generalization Hierarchies	157
5.4	Object Features	158
5.5	Defining Features	158
5.6	Representation of N-ary Features	160
5.7	Transitivity	161
5.8	Inheritance	162
	5.8.1 Default Inheritance	162
	5.8.2 Multiple Inheritance	162
5.9	Concepts as Feature Values	163
5.10	Ontology Matching	164
5.11	Hands On: Browsing an Ontology	165
5.12	Project Assignment 4	168
5.13	Review Questions	168
6	Ontology Design and Development	174
6.1	Design and Development Methodology	174
6.2	Steps in Ontology Development	174
6.3	Domain Understanding and Concept Elicitation	176
	6.3.1 Tutorial Session Delivered by the Expert	177
	6.3.2 Ad-hoc List Created by the Expert	177
	6.3.3 Book Index	177
	6.3.4 Unstructured Interviews with the Expert	177
	6.3.5 Structured Interviews with the Expert	177
	6.3.6 Protocol Analysis (Think-Aloud Technique)	178
	6.3.7 The Card-Sort Method	179
6.4	Modeling-based Ontology Specification	179
6.5	Hands On: Developing a Hierarchy of Concepts and Instances	180
6.6	Guidelines for Developing Generalization Hierarchies	186
	6.6.1 Well-Structured Hierarchies	186
	6.6.2 Instance or Concept?	187
	6.6.3 Specific Instance or Generic Instance?	188
	6.6.4 Naming Conventions	188
	6.6.5 Automatic Support	189

Contents

ix

6.7	Hands On: Developing a Hierarchy of Features	189
6.8	Hands On: Defining Instances and Their Features	192
6.9	Guidelines for Defining Features and Values	195
6.9.1	Concept or Feature?	195
6.9.2	Concept, Instance, or Constant?	196
6.9.3	Naming of Features	196
6.9.4	Automatic Support	197
6.10	Ontology Maintenance	197
6.11	Project Assignment 5	198
6.12	Review Questions	198
7	Reasoning with Ontologies and Rules	202
7.1	Production System Architecture	202
7.2	Complex Ontology-based Concepts	203
7.3	Reduction and Synthesis Rules and the Inference Engine	204
7.4	Reduction and Synthesis Rules for Evidence-based Hypotheses Analysis	206
7.5	Rule and Ontology Matching	207
7.6	Partially Learned Knowledge	212
7.6.1	Partially Learned Concepts	212
7.6.2	Partially Learned Features	213
7.6.3	Partially Learned Hypotheses	214
7.6.4	Partially Learned Rules	214
7.7	Reasoning with Partially Learned Knowledge	215
7.8	Review Questions	216
8	Learning for Knowledge-based Agents	222
8.1	Introduction to Machine Learning	222
8.1.1	What Is Learning?	222
8.1.2	Inductive Learning from Examples	223
8.1.3	Explanation-based Learning	224
8.1.4	Learning by Analogy	225
8.1.5	Multistrategy Learning	226
8.2	Concepts	227
8.2.1	Concepts, Examples, and Exceptions	227
8.2.2	Examples and Exceptions of a Partially Learned Concept	228
8.3	Generalization and Specialization Rules	229
8.3.1	Turning Constants into Variables	230
8.3.2	Turning Occurrences of a Variable into Different Variables	230
8.3.3	Climbing the Generalization Hierarchies	231
8.3.4	Dropping Conditions	231
8.3.5	Extending Intervals	231
8.3.6	Extending Ordered Sets of Intervals	232
8.3.7	Extending Symbolic Probabilities	232
8.3.8	Extending Discrete Sets	232
8.3.9	Using Feature Definitions	233
8.3.10	Using Inference Rules	233

8.4	Types of Generalizations and Specializations	234
8.4.1	Definition of Generalization	234
8.4.2	Minimal Generalization	234
8.4.3	Minimal Specialization	235
8.4.4	Generalization of Two Concepts	236
8.4.5	Minimal Generalization of Two Concepts	236
8.4.6	Specialization of Two Concepts	237
8.4.7	Minimal Specialization of Two Concepts	237
8.5	Inductive Concept Learning from Examples	238
8.6	Learning with an Incomplete Representation Language	242
8.7	Formal Definition of Generalization	243
8.7.1	Formal Representation Language for Concepts	243
8.7.2	Term Generalization	245
8.7.3	Clause Generalization	245
8.7.4	BRU Generalization	246
8.7.5	Generalization of Concepts with Negations	247
8.7.6	Substitutions and the Generalization Rules	247
8.8	Review Questions	247
9	Rule Learning	252
9.1	Modeling, Learning, and Problem Solving	252
9.2	An Illustration of Rule Learning and Refinement	253
9.3	The Rule-Learning Problem	257
9.4	Overview of the Rule-Learning Method	258
9.5	Mixed-Initiative Example Understanding	260
9.5.1	What Is an Explanation of an Example?	260
9.5.2	Explanation Generation	262
9.6	Example Reformulation	264
9.7	Analogy-based Generalization	265
9.7.1	Analogical Problem Solving Based on Explanation Similarity	265
9.7.2	Upper Bound Condition as a Maximally General Analogy Criterion	266
9.7.3	Lower Bound Condition as a Minimally General Analogy Criterion	268
9.8	Rule Generation and Analysis	270
9.9	Generalized Examples	270
9.10	Hypothesis Learning	271
9.11	Hands On: Rule and Hypotheses Learning	275
9.12	Explanation Generation Operations	279
9.12.1	Guiding Explanation Generation	279
9.12.2	Fixing Values	280
9.12.3	Explanations with Functions	280
9.12.4	Explanations with Comparisons	283
9.12.5	Hands On: Explanations with Functions and Comparisons	285
9.13	Guidelines for Rule and Hypothesis Learning	285
9.14	Project Assignment 6	289
9.15	Review Questions	289

Contents

xi

10	Rule Refinement	294
10.1	Incremental Rule Refinement	294
10.1.1	The Rule Refinement Problem	294
10.1.2	Overview of the Rule Refinement Method	295
10.1.3	Rule Refinement with Positive Examples	296
10.1.3.1	Illustration of Rule Refinement with a Positive Example	296
10.1.3.2	The Method of Rule Refinement with a Positive Example	298
10.1.3.3	Summary of Rule Refinement with a Positive Example	300
10.1.4	Rule Refinement with Negative Examples	300
10.1.4.1	Illustration of Rule Refinement with Except-When Conditions	300
10.1.4.2	The Method of Rule Refinement with Except-When Conditions	305
10.1.4.3	Illustration of Rule Refinement through Condition Specialization	305
10.1.4.4	The Method of Rule Refinement through Condition Specialization	307
10.1.4.5	Summary of Rule Refinement with a Negative Example	308
10.2	Learning with an Evolving Ontology	309
10.2.1	The Rule Regeneration Problem	309
10.2.2	On-Demand Rule Regeneration	310
10.2.3	Illustration of the Rule Regeneration Method	312
10.2.4	The Rule Regeneration Method	316
10.3	Hypothesis Refinement	316
10.4	Characterization of Rule Learning and Refinement	317
10.5	Hands On: Rule Refinement	319
10.6	Guidelines for Rule Refinement	321
10.7	Project Assignment 7	322
10.8	Review Questions	322
11	Abstraction of Reasoning	329
11.1	Statement Abstraction	329
11.2	Reasoning Tree Abstraction	331
11.3	Reasoning Tree Browsing	331
11.4	Hands On: Abstraction of Reasoning	331
11.5	Abstraction Guideline	334
11.6	Project Assignment 8	335
11.7	Review Questions	335
12	Disciple Agents	338
12.1	Introduction	338
12.2	Disciple-WA: Military Engineering Planning	338
12.2.1	The Workaround Planning Problem	338
12.2.2	Modeling the Workaround Planning Process	341
12.2.3	Ontology Design and Development	343

12.2.4	Rule Learning	345
12.2.5	Experimental Results	346
12.3	Disciple-COA: Course of Action Critiquing	348
12.3.1	The Course of Action Critiquing Problem	348
12.3.2	Modeling the COA Critiquing Process	351
12.3.3	Ontology Design and Development	352
12.3.4	Training the Disciple-COA Agent	355
12.3.5	Experimental Results	360
12.4	Disciple-COG: Center of Gravity Analysis	364
12.4.1	The Center of Gravity Analysis Problem	364
12.4.2	Overview of the Use of Disciple-COG	367
12.4.3	Ontology Design and Development	376
12.4.4	Script Development for Scenario Elicitation	376
12.4.5	Agent Teaching and Learning	380
12.4.6	Experimental Results	383
12.5	Disciple-VPT: Multi-Agent Collaborative Planning	387
12.5.1	Introduction	387
12.5.2	The Architecture of Disciple-VPT	388
12.5.3	The Emergency Response Planning Problem	389
12.5.4	The Disciple-VE Learning Agent Shell	390
12.5.5	Hierarchical Task Network Planning	394
12.5.6	Guidelines for HTN Planning	396
12.5.7	Integration of Planning and Inference	400
12.5.8	Teaching Disciple-VE to Perform Inference Tasks	403
12.5.9	Teaching Disciple-VE to Perform Planning Tasks	409
12.5.9.1	Why Learning Planning Rules Is Difficult	409
12.5.9.2	Learning a Set of Correlated Planning Rules	409
12.5.9.3	The Learning Problem and Method for a Set of Correlated Planning Rules	413
12.5.9.4	Learning Correlated Planning Task Reduction Rules	413
12.5.9.5	Learning Correlated Planning Task Concretion Rules	414
12.5.9.6	Learning a Correlated Action Concretion Rule	415
12.5.10	The Virtual Experts Library	416
12.5.11	Multidomain Collaborative Planning	420
12.5.12	Basic Virtual Planning Experts	421
12.5.13	Evaluation of Disciple-VPT	422
12.5.14	Final Remarks	422
13	Design Principles for Cognitive Assistants	426
13.1	Learning-based Knowledge Engineering	426
13.2	Problem-Solving Paradigm for User-Agent Collaboration	427
13.3	Multi-Agent and Multidomain Problem Solving	427
13.4	Knowledge Base Structuring for Knowledge Reuse	427
13.5	Integrated Teaching and Learning	428

Contents**xiii**

13.6	Multistrategy Learning	428
13.7	Knowledge Adaptation	429
13.8	Mixed-Initiative Modeling, Learning, and Problem Solving	429
13.9	Plausible Reasoning with Partially Learned Knowledge	430
13.10	User Tutoring in Problem Solving	430
13.11	Agent Architecture for Rapid Agent Development	430
13.12	Design Based on a Complete Agent Life Cycle	431
	<i>References</i>	433
	<i>Appendixes</i>	443
	Summary: Knowledge Engineering Guidelines	443
	Summary: Operations with Disciple-EBR	444
	Summary: Hands-On Exercises	446
	<i>Index</i>	447