

KNOWLEDGE ENGINEERING

Building Cognitive Assistants for Evidence-Based Reasoning

This book presents a significant advancement in the theory and practice of knowledge engineering, the discipline concerned with the development of intelligent agents that use knowledge and reasoning to perform problem-solving and decision-making tasks. It covers the main stages in the development of a knowledge-based agent: understanding the application domain, modeling problem solving in that domain, developing the ontology, learning the reasoning rules, and testing the agent. The book focuses on a special class of agents: cognitive assistants for evidence-based reasoning that learn complex problem-solving expertise directly from human experts, support experts, and nonexperts in problem solving and decision making, and teach their problem-solving expertise to students.

A powerful learning agent shell, Disciple-EBR, is included with the book, enabling students, practitioners, and researchers to develop cognitive assistants rapidly in a wide variety of domains that require evidence-based reasoning, including intelligence analysis, cybersecurity, law, forensics, medicine, and education.

Gheorghe Tecuci (PhD, University of Paris-South, July 1988, and Polytechnic Institute of Bucharest, December 1988) is Professor of Computer Science and Director of the Learning Agents Center in the Volgenau School of Engineering of George Mason University, Member of the Romanian Academy, and former Chair of Artificial Intelligence in the Center for Strategic Leadership of the U.S. Army War College.

Dorin Marcu (PhD, George Mason University, 2009) is Research Assistant Professor, as well as Senior Software and Knowledge Engineer, in the Learning Agents Center, Volgenau School of Engineering, George Mason University.

Mihai Boicu (PhD, George Mason University, 2003) is Associate Professor of Information Sciences and Technology, and Associate Director of the Learning Agents Center, Volgenau School of Engineering, George Mason University.

David A. Schum (PhD, Ohio State University, 1964) is Emeritus Professor of Systems Engineering, Operations Research, and Law, as well as Chief Scientist of the Learning Agents Center at George Mason University. He is also Honorary Professor of Evidence Science at University College London.

Knowledge Engineering

Building Cognitive Assistants for Evidence-Based Reasoning

GHEORGHE TECUCI

George Mason University

DORIN MARCU

George Mason University

MIHAI BOICU

George Mason University

DAVID A. SCHUM

George Mason University



CAMBRIDGE
UNIVERSITY PRESS

CAMBRIDGE UNIVERSITY PRESS

One Liberty Plaza, New York, NY 10006

Cambridge University Press is part of the University of Cambridge.

It furthers the University's mission by disseminating knowledge in the pursuit of education, learning, and research at the highest international levels of excellence.

www.cambridge.org

Information on this title: www.cambridge.org/9781107122567

© Gheorghe Tecuci, Dorin Marcu, Mihai Boicu, and David A. Schum 2016

This publication is in copyright. Subject to statutory exception and to the provisions of relevant collective licensing agreements, no reproduction of any part may take place without the written permission of Cambridge University Press.

First published 2016

Printed in the United States of America by Sheridan Books, Inc.

A catalog record for this publication is available from the British Library.

Library of Congress Cataloging-in-Publication Data

Names: Tecuci, Gheorghe, author. | Marcu, Dorin, author. | Boicu, Mihai, author. | Schum, David A., author.

Title: Knowledge engineering: building cognitive assistants for evidence-based reasoning / Gheorghe Tecuci, George Mason University, Dorin Marcu, George Mason University, Mihai Boicu, George Mason University, David A. Schum, George Mason University.

Description: New York NY : Cambridge University Press, 2016. | Includes bibliographical references and index.

Identifiers: LCCN 2015042941 | ISBN 9781107122567 (Hardback : alk. paper)

Subjects: LCSH: Expert systems (Computer science) | Intelligent agents (Computer software) | Machine learning | Artificial intelligence | Knowledge, Theory of—Data processing.

Classification: LCC QA76.76.E95 T435 2016 | DDC 006.3/3—dc23 LC record available at <http://lcn.loc.gov/2015042941>

ISBN 978-1-107-12256-7 Hardback

Cambridge University Press has no responsibility for the persistence or accuracy of URLs for external or third-party Internet Web sites referred to in this publication and does not guarantee that any content on such Web sites is, or will remain, accurate or appropriate.

Contents

<i>Preface</i>	<i>page xv</i>
<i>Acknowledgments</i>	<i>xxi</i>
<i>About the Authors</i>	<i>xxiii</i>
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

Preface

BOOK PURPOSE

This is a book on knowledge engineering, the discipline concerned with the development of intelligent agents that use knowledge and reasoning to perform problem-solving and decision-making tasks. The book covers the theory and practice of the main stages in the development of a knowledge-based agent: understanding the application domain, modeling problem solving in that domain, developing the ontology, learning the reasoning rules, and testing the agent. However, it does this by focusing on a special class of agents: cognitive assistants that learn complex problem-solving expertise directly from human experts, support experts, and nonexperts in problem solving and decision making and teach their problem-solving expertise to students. These are learning agents that are taught by their users in ways that are similar to how a student, an apprentice, or a new collaborator is taught, through problem-solving examples and explanations and by supervising and correcting their behavior. Because such agents learn to replicate the problem-solving behavior of their users, we have called them Disciple agents.

This book presents a significant advancement in the theory and practice of knowledge engineering, where many tasks are performed by a typical computer user and a learning agent, with only limited support from a knowledge engineer. To simplify further the development of the cognitive assistants by typical users, we have focused on the development of cognitive assistants for evidence-based reasoning. Evidence-based reasoning is at the core of many problem-solving and decision-making tasks in a wide variety of domains, including intelligence analysis, cybersecurity, law, forensics, medicine, physics, chemistry, history, archaeology, education, and many others. Nevertheless, the last part of the book presents Disciple agents for applications that did not involve evidence-based reasoning.

Because knowledge engineering is a practical activity, it is best learned by doing. Therefore, this book presents the theory and methodology of developing cognitive assistants in conjunction with a practical tool, Disciple-EBR, a learning agent shell for evidence-based reasoning (EBR). Consequently, each chapter typically contains a theoretical part presenting general concepts and methods, a methodological part with guidelines on the application of the methods, and a practical part on the actual employment of these methods with Disciple-EBR. It also includes project assignments and review questions.

This book addresses issues of interest to a large spectrum of readers from academia, research, and industry. We have used drafts of this book in our computer science courses on knowledge engineering, expert systems, and knowledge-based agents, at both undergraduate and graduate levels, because it covers the theory and practice of the main stages in the development of knowledge-based agents. We have also used some parts of the book in introductory courses on artificial intelligence, and other parts in the courses on knowledge acquisition and machine learning. These are all examples of courses where this book could be used.

Researchers in knowledge engineering will find in this book an integrated approach that advances the theory and practice in the field. We believe that further research and development of this approach will enable typical computer users to develop their own cognitive assistants without any knowledge engineering assistance. Thus, non-computer scientists will no longer be only users of generic programs developed by others (such as word processors or Internet browsers), as they are today, but also agent developers themselves. They will be able to train their personal assistants to help them with their increasingly complex tasks in the knowledge society, which should have a significant beneficial impact on their work and life.

Practitioners that develop various types of knowledge-based systems will find in this book a detailed, yet intuitive, presentation of an agent development methodology and tool, as well as several case studies of developing intelligent agents that illustrate different types of agents that are relevant to a wide variety of application domains. In fact, a complementary book, *Intelligence Analysis as Discovery of Evidence, Hypotheses, and Arguments: Connecting the Dots* (Tecuci et al., 2016), presents the practical application of Disciple-CD, an agent developed with Disciple-EBR for intelligence analysis problems.

BOOK CONTENTS

Here is a route or map we will follow in the learning venture you will have with the assistance of Disciple-EBR. Chapter 1 is a general introduction to the topics that form the basis of this book. It starts with the problem of understanding the world through evidence-based reasoning. It then presents abductive reasoning, five different conceptions of probability (enumerative, subjective Bayesian, Belief Functions, Baconian, and Fuzzy), and how deductive, abductive, and inductive (probabilistic) reasoning are used in evidence-based reasoning. After that, it introduces artificial intelligence and intelligent agents, and the challenges of developing such agents through conventional knowledge engineering. Afterward, it introduces the development of agents through teaching and learning, which is the approach presented in this book.

Chapter 2 is an overview of evidence-based reasoning, which is a focus of this book. It starts with a discussion of the elements that make evidence-based reasoning an astonishingly complex task. It then introduces a systematic approach that integrates abduction, deduction, and induction to solve a typical evidence-based reasoning task, using intelligence analysis as an example. Finally, it shows the application of the same approach to other

evidence-based reasoning tasks in cybersecurity, geospatial intelligence, and critical thinking education.

Chapter 3 is an overview of the main methodologies and tools for the design and development of knowledge-based agents. It first presents the conventional approach of developing such agents by a knowledge engineer working with a subject matter expert. It then introduces different types of agent shells, which are the typical tools for building agents, and discusses the use of foundational and utility ontologies. The more advanced of these tools, the learning agent shells, are at the basis of a new and more powerful approach to agent design and development, an overview of which is presented in the second part of Chapter 3. This learning-based approach is illustrated with the development of a cognitive assistant for assessing a potential PhD advisor for a student, an example that is used in the following chapters to present in detail each of the main agent development stages.

Chapter 4 presents the modeling of the problem-solving process through analysis and synthesis, which is the most challenging task in the development of a knowledge-based agent. This chapter starts with a more general version of this process for any type of problems, which is used in the Disciple agents presented in Chapter 12. After that, the chapter presents an easier, customized version for evidence-based hypothesis analysis, which is used in the chapters that follow. Chapter 4 also introduces an ontology of evidence and presents the modeling of the believability assessment for different types of evidence.

Chapters 5 and 6 present the representation of knowledge through ontologies, as well as their design and development. They discuss the representation of the objects from the agent's application domain and their relationships. These chapters also discuss the basic reasoning operations of transitivity, inheritance, and matching. They cover several approaches to concept elicitation, as well as a systematic approach to modeling-based ontology specification. The chapters also address the complexity of ontology maintenance.

Chapter 7 presents the agent's reasoning based on ontologies and rules, in the context of a production system architecture. It starts with defining the representation of complex ontology-based concepts and the use of these concepts in the reasoning rules. It defines the reduction and synthesizes rules in general, and the special case of these rules for evidence-based reasoning. It then presents the process of problem solving through analysis and synthesis, which is accomplished through rule and ontology matching. This chapter also introduces the representation of and reasoning with partially learned concepts, features, hypotheses, and rules.

Chapter 8 starts with a general introduction to machine learning and to several learning strategies that are most relevant for knowledge engineering. It continues with the definition of the generalization and specializations of concepts through the use of inductive rules. After that, the chapter defines the basic operations of minimal and maximal generalizations and specialization of concepts, which are at the basis of rule learning and refinement discussed in Chapters 9 and 10. The chapter ends with the presentation of a formal definition of generalization.

Chapter 9 presents the mixed-initiative rule learning method that enables an agent to learn a general rule from a specific example of a

reasoning step. It starts with an overview of the integration of modeling, learning, and problem solving and an illustration of this process. The chapter then introduces the rule learning problem and method. After that, it presents in detail the phases of rule learning, such as the mixed-initiative understanding of the example, and its analogy-based generalizations, which result in a minimal and a maximal generalization forming the plausible version space of the rule. This chapter also discusses the learning of rules involving functions and relational operators.

The refinement of a partially learned rule is presented in Chapter 10. After introducing the incremental rule refinement problem, this chapter presents the refinement of the rule based on positive examples, and then the refinement based on negative examples. This may result in a very complex rule characterized by a main applicability condition and several “except-when” conditions, which capture the reasoning of a subject matter expert. There are various refinement strategies, depending on the position of the examples with respect to the bounds of the rule’s conditions and also on the possible explanations of these examples, but in all cases, they involve simple interactions with the subject matter expert. Changes in the ontology require the learned rules to be updated, the corresponding method being presented and illustrated in the second part of this chapter. The chapter ends with a characterization of rule learning and refinement, which enable a non-computer scientist to teach an agent.

The chapters dedicated to the individual phases of agent development end with Chapter 11, which discusses the abstractions of individual hypotheses and of the reasoning tree to facilitate its browsing, understanding, and further development by the end-user.

As previously indicated, this book focuses on the development of agents for evidence-based reasoning tasks. However, the presented theory and methodology are also applicable to other types of agents. These agents have been developed with learning agent shells representing previous implementations of the Disciple theory and methodology. Four of these agents are presented in Chapter 12. Disciple-WA is an agent that uses expert knowledge from military engineering manuals to develop alternative plans of actions that a military unit can use to work around (WA) damages to a transportation infrastructure, such as a damaged, destroyed, or mined bridge, road, or tunnel. The goal is to find the quickest way for the military unit to bypass the encountered obstacle. There were several cases where the Disciple-WA agent generated better solutions than those of the human expert who evaluated the developed systems, as well as cases where the agent generated new solutions that this expert did not consider.

Disciple-COA is an agent trained by a military expert to critique courses of action (COA) with respect to the principles of war and the tenets of Army operations. It receives as input the description of a course of action and returns a list of strengths and weaknesses of various levels, together with their justifications. A remarkable feature of this agent is that it was judged to exceed the performance of the subject matter experts that defined the problems for its evaluation.

Disciple-COG is an agent that was trained to identify and assess the center of gravity (COG) candidates of the opposing forces in a military scenario. Correctly identifying the centers of gravity of the opposing forces

is of highest importance in any conflict, and Disciple-COG has been used for many years in courses at the U.S. Army War College, as well as at other military institutions, to teach students how to perform a strategic center-of-gravity analysis of a scenario of interest.

The last agent described in Chapter 12 is Disciple-VPT (Virtual Planning Team). Disciple-VPT consists of virtual planning experts that collaborate to develop plans of actions requiring expertise from multiple domains. They are assembled from an extensible library of such agents. The basic component of Disciple-VPT is the Disciple-VE learning agent shell that can be taught how to plan directly by a subject matter expert. Copies of the Disciple-VE shells can be used by experts in different domains to rapidly populate the library of virtual experts (VEs) of Disciple-VPT.

The learning-based approach to knowledge engineering presented in this book illustrates the application of several design principles that are useful in the development of cognitive assistants in general. Therefore, as a conclusion, Chapter 13 summarizes these principles, which are illustrated throughout this book.

The book also includes several appendices that summarize important aspects from the chapters: the list of the knowledge-engineering guidelines for each of the main stages of agent development, the list of operations of Disciple-EBR, and the list of the hands-on exercises. Answers to selected questions from each chapter are made available to the instructors.

BACKGROUND

The learning-based knowledge-engineering theory, methodology, and tools presented in this book are the result of many years of research and experimentation that produced increasingly more general and powerful versions. During this evolution, one may identify four main stages. The first stage corresponds to the PhD work of Gheorghe Tecuci presented in his thesis “Disciple: A Theory, Methodology and System for Learning Expert Knowledge” (Thèse de Docteur en Science, Université de Paris-Sud, July 1988). The main emphasis of that work was on rule learning. The developed methods are among the first multistrategy approaches to learning that later grew into the subfield of multistrategy machine learning (Tecuci, 1993; Michalski and Tecuci, 1994). They also are among the first attempts to integrate machine learning and knowledge acquisition (Tecuci et al., 1994; Tecuci and Kodratoff, 1995). This work benefited from the collaboration of Yves Kodratoff and the support of Mihai Drăgănescu. It was done with the support of the Romanian Research Institute for Informatics, the Romanian Academy, and the French National Center for Scientific Research.

The second stage in the evolution of the Disciple approach is presented in the book *Building Intelligent Agents: An Apprenticeship Multistrategy Learning Theory, Methodology, Tool and Case Studies*, published by Academic Press in 1998. It includes an improvement of Disciple’s multistrategy learning methods and their extension with guided knowledge-elicitation methods. It also includes ontology development tools and illustrations of the application of the Disciple rule-learning approach to a variety of domains, such as teaching of higher-order thinking skills in history and

statistics, engineering design, and military simulation. This work benefited from the collaboration of several of Tecuci's PhD students, particularly Thomas Dybala, Michael Hieb, Harry Keeling, and Kathryn Wright, and it was partially supported by George Mason University, the National Science Foundation, and the Defense Advanced Research Projects Agency.

The third stage in the evolution of the Disciple approach is represented by the Disciple agents described in Chapter 12 of this book. This represents a significant extension of the Disciple approach, with methods for problem solving through analysis and synthesis, modeling of the problem solving process, ontology development, and scenario elicitation. At this stage, Disciple has become a complete, end-to-end agent development methodology that has been applied to develop powerful agents, such as Disciple-COG, used for a period of ten years in courses at the U.S. Army War College and at other institutions. This work also benefited from the collaboration of Tecuci's students, first of all Mihai Boicu and Dorin Marcu, and also Michael Bowman, Vu Le, Cristina Boicu, Bogdan Stanescu, and Marcel Barbulescu. This work was partially supported by George Mason University, the Air Force Office of Scientific Research, the Air Force Research Laboratory, the Defense Advanced Research Projects Agency, the National Science Foundation, and others.

Finally, the latest stage in the evolution of the Disciple approach is represented by the rest of this book. A main conceptual advancement over the previous stage consists in its extension with a theory of evidence-based reasoning, greatly facilitated by the collaboration of David A. Schum. This resulted in a very powerful theory, methodology, and tool for the development of agents for complex evidence-based reasoning applications, such as intelligence analysis. Two of such agents are TIACRITIS (Teaching Intelligence Analysts Critical Thinking Skills) and Disciple-CD (Disciple cognitive assistant for Connecting the Dots). This work was partially supported by George Mason University and by several agencies of the U.S. government, including the Department of Defense.

Acknowledgments

We are very grateful to the many individuals who, in various ways, supported our research, including Kelcy Allwein, Cindy Ayers, Murray Burke, Douglas Campbell, William Cleckner, Jerry Comello, John Donelan, Jim Donlon, Susan Durham, Keri Eustis, Michael Fletcher, Erin Gibbens, Lloyd Griffiths, David Gunning, Ben Hamilton, Sharon Hamilton, Robert Herklotz, Phillip Hwang, Eric Jones, Alex Kilpatrick, David Luginbuhl, Joan McIntyre, Jean-Michel Pomrade, Michelle Quirk, Peter Rocci, William Rzepka, Kimberly Urban, Joan Vallancewhitacre, and Ben Wible.

We also want to thank Heather Bergman, who invited us to write this book for the Cambridge University Press, as well as to senior editor Lauren Cowles, who is a great professional to work with.

About the Authors

Gheorghe Tecuci (PhD, Université de Paris-Sud, July 1988, and Polytechnic Institute of Bucharest, December 1988) is Professor of Computer Science and Director of the Learning Agents Center in the Volgenau School of Engineering of George Mason University, Member of the Romanian Academy, and former Chair of Artificial Intelligence in the Center for Strategic Leadership of the U.S. Army War College. He has published around two hundred papers, including eleven books, with contributions to artificial intelligence, knowledge engineering, cognitive assistants, machine learning, evidence-based reasoning, and intelligence analysis. He has received the U.S. Army Outstanding Civilian Service Medal (for “groundbreaking contributions to the application of artificial intelligence to center of gravity determination”) and the Innovative Application Award from the American Association for Artificial Intelligence.

Dorin Marcu (PhD, George Mason University, 2009) is Research Assistant Professor, as well as Senior Software and Knowledge Engineer, in the Learning Agents Center, Volgenau School of Engineering, George Mason University. He has published more than forty papers, including five books, with contributions to adaptive user interfaces, mixed-initiative human-computer interaction, and cognitive assistants. He has received the Innovative Application Award from the American Association for Artificial Intelligence.

Mihai Boicu (PhD, George Mason University, 2002) is Associate Professor of Information Sciences and Technology and Associate Director of the Learning Agents Center, in the Volgenau School of Engineering of George Mason University. He has published over ninety papers, including five books, with contributions to problem solving and multistrategy learning in dynamic and evolving representation spaces, mixed-initiative interaction, multi-agent systems architecture, collaboration and coordination, abstraction-based reasoning, knowledge representation, and knowledge acquisition. He has received the Innovative Application Award from the American Association for Artificial Intelligence.

David A. Schum (PhD, Ohio State University, 1964) is Emeritus Professor of Systems Engineering, Operations Research, and Law, as well as Chief Scientist of the Learning Agents Center at George Mason University, and Honorary Professor of Evidence Science at University College London.

He has followed a career-long interest in the study of the properties, uses, discovery, and marshaling of evidence in probabilistic reasoning. Dr. Schum has published more than one hundred papers in a variety of journals and eight books, including *The Evidential Foundations of Probabilistic Reasoning*, *Analysis of Evidence*, *Evidence and Inference for the Intelligence Analyst*, and *Probabilistic Analysis of the Sacco and Vanzetti Evidence*, being recognized as one of the founding fathers of the Science of Evidence.