Foundations of Decision-Making Agents

Logic, Probability and Modality

This page intentionally left blank

Foundations of Decision-Making Agents

Logic, Probability and Modality

Subrata Das

World Scientific

Imperial College Press

Published by

World Scientific Publishing Co. Pte. Ltd.

5 Toh Tuck Link, Singapore 596224

USA office: 27 Warren Street, Suite 401-402, Hackensack, NJ 07601

UK office: 57 Shelton Street, Covent Garden, London WC2H 9HE

British Library Cataloguing-in-Publication Data

A catalogue record for this book is available from the British Library.

Cover image by Sebastien Das

FOUNDATIONS OF DECISION-MAKING AGENTS

Logic, Probability and Modality

Copyright © 2008 by World Scientific Publishing Co. Pte. Ltd.

All rights reserved. This book, or parts thereof, may not be reproduced in any form or by any means, electronic or mechanical, including photocopying, recording or any information storage and retrieval system now known or to be invented, without written permission from the Publisher.

For photocopying of material in this volume, please pay a copying fee through the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923, USA. In this case permission to photocopy is not required from the publisher.

ISBN-13 978-981-277-983-0

ISBN-10 981-277-983-3

Editor: Tjan Kwang Wei

Printed in Singapore.

To Janique, Sébastien, and Kabita

This page intentionally left blank

Preface

Artificial Intelligence (AI) systems simulating human behavior are often called *intelligent agents*. By definition, these intelligent agents exhibit some form of human-like intelligence. For example, intelligent agents can learn from the environment, make plans and decisions, react to the environment with appropriate actions, express emotions, or revise beliefs. Intelligent agents typically represent human cognitive states using underlying beliefs and knowledge modeled in a knowledge representation language. We use the term *epistemic state* to refer to an actual or possible cognitive state that drives human behavior at a given point in time; the accurate determination (or estimation) of these epistemic states is crucial to an agent's ability to correctly simulate human behavior.

This book provides three fundamental and generic approaches (logical, probabilistic, and modal) to representing and reasoning with agent epistemic states, specifically in the context of decision making. In addition, the book introduces a formal integration of these three approaches into a single unified approach we call $\mathcal{P3}$ (Propositional, Probabilistic, and Possible World), that combines the advantages of the other approaches. Each of these approaches can be applied to create the foundation for intelligent, decision-making software agents. Each approach is symbolic in nature (unlike the sub-symbolic neural network approach, for example), yielding agent “thought processes” that are sufficiently transparent to provide users with understandable explanations of agent “reasoning.” The classical logic and probability theory thrust of our approach naturally excludes from our consideration other potential symbolic approaches to decision-making such as fuzzy logics. The “symbolic argumentation” approach ($\mathcal{P3}$) to decision making combines logic and probability, and therefore offers several advantages over the traditional approach to decision-making based on simple rule-based expert systems or expected utility theory.

The generic logical, probabilistic, and modal reasoning techniques that we discuss in the book, such as logical deduction and evidence propagation, are applied to various decision-making problems that an agent will encounter in various circumstances. However, the book does not focus on particular applications of these reasoning techniques to such problems as planning,

learning, or belief revision. Neither does the book focus on communication and collaboration within agent societies. Such applications require extensive treatments in themselves and so are beyond the scope of the book; our purpose here is to provide the formal foundation necessary to support decision-making.

This book is divided into four parts by the four reasoning approaches we develop. The first part focuses on logic-based modeling of an epistemic state as a set of propositions expressed by sentences from classical propositional and first-order logics. The second part takes a probabilistic approach to modeling epistemic states. In this approach, states are represented by a probability measure defined over some space of events, which are represented as random variables. The third part focuses on extending logic-based modeling with modalities to yield various systems of modal logics, including epistemic logics. These logics are based on mental constructs that represent information and are suitable for reasoning with possible human cognitive states. Finally, the fourth part describes the combination of these three rather disparate approaches into the single integrated $\mathcal{P3}$ approach within an argumentation framework especially appropriate for simulating human decision making.

The purpose of this book is not to carry out an elaborate philosophical discussion on epistemic state modeling, but rather to provide readers with various practical ways of modeling and reasoning with agent epistemic states used in decision making. Therefore, the resolution-based logic programming paradigm is introduced under logic-based modeling; Bayesian belief networks are introduced as a way of modeling and reasoning with epistemic states where beliefs are represented by probability measures defined over some state space; and modal resolution schemes are introduced under modal logic-based modeling and reasoning with actual and possible worlds.

Frameworks specific to decision making, such as the logic-based production rule or belief-network-based influence diagrams that incorporate the notions of action and utility, are also detailed in the book. The use of logic-based reasoning (especially resolution theorem proving) and probabilistic reasoning (especially Bayesian belief networks) is widespread when implementing agents capable of making decisions, generating plans, revising beliefs, learning, and so on. But the book goes a step further by bridging the gap between these two approaches, augmenting them with informational mental constructs, such as beliefs and knowledge, and then producing a coherent approach that culminates in the penultimate chapter of the book in the form of a symbolic argumentation approach for weighing the pros and cons of decision-making options.

In the book I have tried to balance discussions on theory and practice. As for theory, I have stated results and developed proofs wherever necessary to provide sufficient theoretical foundations for the concepts and procedures introduced. For example, if I have introduced a specific resolution-based theorem-proving procedure, then I also provide the associated soundness and completeness of the procedure with respect to the stated semantics. On the practical side, I have provided detailed algorithms that can be encoded in computer programs in a straightforward manner to incorporate aspects of human behavior into intelligent decision-making agents. For example, an implementation of the evidence propagation algorithms for belief networks and influence diagrams can be incorporated into agents that make decisions within a probabilistic framework.

The decision-making problem is present in every industry; finance, medicine, and defense are a few of the big ones. Computer professionals and researchers in the decision science community within commerce and industry will find the book most useful for building intelligent and practical decision-aiding agents or, more generally, for embedding intelligence in various systems based on sound logical and probabilistic reasoning; the material is especially suited for building safety-critical decision aiding applications based on well-founded theories. This book may also be used as an AI textbook on logic and probability for undergraduate and graduate courses, and as a reference book for researchers in universities.

The chapters in the book and their dependencies are shown in the figure below. Readers are advised not to skip Chapter 1, which includes an informal overview of the rest of the book. A basic understanding of this chapter will prevent the reader from becoming lost later on. The beginner should not be dismayed by the use of unfamiliar terminology in this first chapter, as these terms will be explained in later chapters.

Readers who have been exposed to the concepts of basic mathematics and probability may omit Chapter 2, which provides a background on mathematical preliminaries. Chapters 3 and 4, on classical logics and logic programming, should be read in sequence. Chapter 6, on Bayesian belief networks, can be read independently of Chapters 3 and 4, but Chapter 8, on modal logics, requires an understanding of classical logic as presented in Chapter 3. Chapters 3, 4, 6, and 8 provide the semantic foundations of our three problem modeling approaches, namely, propositional, probabilistic, and possible world. Because production rules are a special form of rules in logic programming with embedded degrees of uncertainty I recommend to reading at least the logic programming section in Chapter 4 and the background material on probability theory before reading

Chapter 5. Chapter 7, on influence diagrams, requires understanding the basics of Bayesian belief networks as it extends this understanding with the concepts of action and utility. The final chapter (Chapter 9) on symbolic argumentation requires an overall understanding of the contents of Chapters 3-8. Ultimately, three different decision-making frameworks are provided, in Chapters 5, 7, and 9.

Figure 1-1: Chapters and their dependencies

There is another suggested way to follow this book for those readers who, with some exposure to mathematical logic and the theory of probability, are looking for some well-founded computational decision support technologies to build deployable systems. They can skip detailed theoretical materials as presented in some chapters (e.g. 3, 4, and 8), and focus more on the computational aspects (e.g. chapter 5, 6, 7, and 9).

My sincere thanks go to my wife Janique, my son Sébastien, and my daughter Kabita, for their love, patience and inspiration throughout the preparation of this book (My children's perception of the book are very different from each other. Sébastien's view is very philosophical, which he manifested in his drawing for the title page of a sage from the Museum of Fine Arts in Boston. However, Kabita is more pragmatic, expressed herself through a dog's decision-making dilemma as depicted below.) Special thanks go to Karen DeSimone for her careful proofreading of the manuscript that enormously enhanced readability and enriched the content of the book. Thanks are due to my colleagues, Dave

Lawless and Josh Introne, with whom I have had many technical discussions on various aspects of this book. Thanks also to all of my colleagues at Charles River Analytics, especially Alex Feinman, Paul Gonsalves, Partha Kanjilal, Rich Reposa, and Greg Zacharias, and Avi Pfeffer at Harvard University, for their part in creating a stimulating intellectual environment which inspired me to write this book. Thanks to the World Scientific/Imperial College Press, especially Tjan Kwang Wei, for their help in producing the book from the beginning. Finally, I thank my parents, brothers, sisters and other family members back in one of several thousands of small villages in India for patiently accepting my absence and showing their encouragement and support through their many phone calls.

This book is based on the work of many researchers, and I express my gratitude to those who have directly or indirectly contributed so much to the fields of deductive logic, probability theory, and artificial intelligence. I have made my best effort to make this book informative, readable, and free from mistakes, and I would welcome any criticism or suggestion for improvement.

Subrata Das
Cambridge, MA
March 2007

This page intentionally left blank

Table of Contents

Preface	vii
Table of Contents	xiii
Chapter 1	1
Modeling Agent Epistemic States: An Informal Overview	
1.1 Models of Agent Epistemic States	1
1.2 Propositional Epistemic Model	3
1.3 Probabilistic Epistemic Model	8
1.4 Possible World Epistemic Model	12
1.5 Comparisons of Models	16
1.6 $\mathcal{P3}$ Model for Decision-Making Agents	17
Chapter 2	23
Mathematical Preliminaries	
2.1 Usage of Symbols	23
2.2 Sets, Relations, and Functions	24
2.3 Graphs and Trees	29
2.4 Probability	34
2.5 Algorithmic Complexity	40
2.6 Further Readings	44
Chapter 3	45
Classical Logics for the Propositional Epistemic Model	
3.1 Propositional Logic	46
3.1.1 Axiomatic Theory for Propositional Logic	54
3.1.2 Soundness and Completeness Theorem	55
3.2 First-Order Logic	57

3.2.1	Axiomatic Theory for First-order Logic.....	62
3.2.2	Soundness and Completeness Theorem	63
3.2.3	Applications.....	70
3.3	Theorem Proving Procedure	71
3.3.1	Clausal Form	71
3.3.2	Herbrand's theorem.....	74
3.3.3	Implementation of Herbrand's theorem	78
3.4	Resolution Theorem Proving	80
3.4.1	Resolution principle and unification.....	80
3.5	Refutation Procedure.....	91
3.6	Complexity Analysis.....	95
3.7	Further Readings.....	96
Chapter 4.....		97
Logic Programming		
4.1	The Concept.....	97
4.2	Program Clauses and Goals	99
4.3	Program Semantics	106
4.4	Definite Programs	108
4.5	Normal Programs	114
4.6	Prolog.....	121
4.6.1	Prolog Syntax	121
4.6.2	Theoretical Background	123
4.6.3	Backtracking.....	126
4.6.4	The Cut.....	127
4.6.5	Special Constructs and Connectives.....	128
4.6.6	Negation	129
4.6.7	Equality	130
4.6.8	List.....	131
4.6.9	Arithmetic.....	132

4.6.10	Input/Output	133
4.6.11	Clause Management	135
4.6.12	Set Evaluation.....	136
4.6.13	Meta Programming.....	139
4.7	Prolog Systems.....	141
4.8	Complexity Analysis.....	141
4.9	Further Readings.....	142
Chapter 5		143
Logical Rules for Making Decisions		
5.1	Evolution of Rules	144
5.2	Bayesian Probability Theory for Handling Uncertainty	146
5.3	Dempster-Shafer Theory for Handling Uncertainty.....	150
5.4	Measuring Consensus	157
5.5	Combining Sources of Varying Confidence	162
5.6	Advantages and Disadvantages of Rule-Based Systems.....	163
5.7	Background and Further Readings.....	164
Chapter 6		165
Bayesian Belief Networks		
6.1	Bayesian Belief Networks.....	166
6.2	Conditional Independence in Belief Networks	171
6.3	Evidence, Belief, and Likelihood.....	179
6.4	Prior Probabilities in Networks without Evidence.....	182
6.5	Belief Revision.....	184
6.6	Evidence Propagation in Polytrees.....	190
6.6.1	Upward Propagation in a Linear Fragment	191
6.6.2	Downward Propagation in a Linear Fragment.....	194
6.6.3	Upward Propagation in a Tree Fragment	198
6.6.4	Downward Propagation in a Tree Fragment.....	200

6.6.5	Upward Propagation in a Polytree Fragment	201
6.6.6	Downward Propagation in a Polytree Fragment	204
6.6.7	Propagation Algorithm	208
6.7	Evidence Propagation in Directed Acyclic Graphs	211
6.7.1	Graphical Transformation	214
6.7.2	Join Tree Initialization	222
6.7.3	Propagation in Join Tree and Marginalization	224
6.7.4	Handling Evidence	227
6.8	Complexity of Inference Algorithms	229
6.9	Acquisition of Probabilities	230
6.10	Advantages and Disadvantages of Belief Networks	234
6.11	Belief Network Tools	235
6.12	Further Readings	235
Chapter 7		237
Influence Diagrams for Making Decisions		
7.1	Expected Utility Theory and Decision Trees	237
7.2	Influence Diagrams	240
7.3	Inferencing in Influence Diagrams	242
7.4	Compilation of Influence Diagrams	248
7.5	Inferencing in Strong Junction Tress	252
7.6	Further Readings	254
Chapter 8		255
Modal Logics for the Possible World Epistemic Model		
8.1	Historical Development of Modal Logics	256
8.2	Systems of Modal Logic	262
8.3	Deductions in Modal Systems	265
8.3.1	Principle of Duality	265
8.3.2	Theorems of $\mathcal{K}$	266

8.3.3	Theorems of $\mathcal{D}$	268
8.3.4	Theorems of $\mathcal{T}$	269
8.3.5	Theorems of $\mathcal{S4}$	269
8.3.6	Theorems of $\mathcal{B}$	270
8.3.7	Theorems of $\mathcal{S5}$	271
8.3.8	Theorems of $\mathcal{S5}'$	272
8.4	Modality.....	272
8.5	Decidability and Matrix Method.....	273
8.6	Relationships among Modal Systems	277
8.7	Possible World Semantics.....	279
8.8	Soundness and Completeness Results.....	286
8.9	Complexity and Decidability of Modal Systems	291
8.10	Modal First-Order Logics	294
8.11	Resolution in Modal First-Order Logics.....	300
8.11.1	Transformation Algorithm.....	302
8.11.2	Unification.....	304
8.12	Modal Epistemic Logics	307
8.13	Logic of Agents Beliefs ($\mathcal{LAB}$)	309
8.13.1	Syntax of $\mathcal{LAB}$	310
8.13.2	Axioms of $\mathcal{LAB}$	312
8.13.3	Possible World Semantics of $\mathcal{LAB}$	313
8.13.4	Soundness and Completeness of $\mathcal{LAB}$	316
8.13.5	Rational Extension of $\mathcal{LAB}$	319
8.13.6	Goals in $\mathcal{LAB}$	320
8.13.7	Dempster-Shafer Interpretation of $\mathcal{LAB}$	320
8.14	Further Readings.....	323
Chapter 9		325
Symbolic Argumentation for Decision-Making		
9.1	Toulmin's Model of Argumentation	327

9.2 Domino Decision-Making Model for $\mathcal{P}3$328

9.3 Knowledge Representation Syntax of $\mathcal{P}3$330

9.4 Formalization of $\mathcal{P}3$ via $\mathcal{LAB}$334

9.5 Aggregation via Dempster-Shafer Theory335

9.6 Aggregation via Bayesian Belief Networks.....339

9.7 Further Readings.....345

References.....347

Index355