

Cambridge University Press

978-0-521-15390-4 - Probabilistic Reasoning in Multiagent Systems: A Graphical Models Approach

Yang Xiang

Frontmatter

[More information](#)

PROBABILISTIC REASONING IN MULTIAGENT SYSTEMS

This book investigates the opportunities in building intelligent decision support systems offered by multiagent distributed probabilistic reasoning.

Probabilistic reasoning with graphical models, also known as Bayesian networks or belief networks, has become an active field of research and practice in artificial intelligence, operations research, and statistics in the past two decades. The success of this technique in modeling intelligent decision support systems under the centralized and single-agent paradigm has been striking. In this book, the author extends graphical dependence models to the distributed and multiagent paradigm. He identifies the major technical challenges involved in such an endeavor and presents the results from a decade's research.

The framework developed in the book allows distributed representation of uncertain knowledge on a large and complex environment embedded in multiple cooperative agents and effective, exact, and distributed probabilistic inference.

Yang Xiang is Associate Professor of Computing and Information Science at the University of Guelph, Canada, where he directs the Intelligent Decision Support System Laboratory. He received his Ph.D. from the University of British Columbia and developed the Java-based toolkit WebWeavr, which has been distributed to registered users in more than 20 countries. He also serves as Principal Investigator in the Institute of Robotics and Intelligent Systems (IRIS), Canada.

Cambridge University Press

978-0-521-15390-4 - Probabilistic Reasoning in Multiagent Systems: A Graphical Models Approach

Yang Xiang

Frontmatter

[More information](#)

PROBABILISTIC REASONING IN MULTIAGENT SYSTEMS

A Graphical Models Approach

YANG XIANG

University of Guelph



CAMBRIDGE
UNIVERSITY PRESS

Cambridge University Press

978-0-521-15390-4 - Probabilistic Reasoning in Multiagent Systems: A Graphical Models Approach

Yang Xiang

Frontmatter

[More information](#)

CAMBRIDGE UNIVERSITY PRESS

Cambridge, New York, Melbourne, Madrid, Cape Town, Singapore,
São Paulo, Delhi, Dubai, Tokyo, Mexico City

Cambridge University Press
The Edinburgh Building, Cambridge CB2 8RU, UK

Published in the United States of America by Cambridge University Press, New York

www.cambridge.org

Information on this title: www.cambridge.org/9780521153904

© Yang Xiang 2002

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 2002
First paperback printing 2010

A catalogue record for this publication is available from the British Library

Library of Congress Cataloguing in Publication data

Xiang, Yang, 1954–

Probabilistic reasoning in multi-agent systems : a graphical models approach / Yang Xiang.

p. cm.

Includes bibliographical references and index.

ISBN 0-521-81308-5

1. Distributed artificial intelligence. 2. Bayesian statistical decision theory – Data processing.
3. Intelligent agents (Computer software) I. Title.

Q337 .X53 2002

006.3 – dc21

2001052874

ISBN 978-0-521-81308-2 Hardback

ISBN 978-0-521-15390-4 Paperback

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 ix</i>
1 Introduction	1
1.1 Intelligent Agents	1
1.2 Reasoning about the Environment	4
1.3 Why Uncertain Reasoning?	5
1.4 Multiagent Systems	7
1.5 Cooperative Multiagent Probabilistic Reasoning	11
1.6 Application Domains	13
1.7 Bibliographical Notes	14
2 Bayesian Networks	16
2.1 Guide to Chapter 2	16
2.2 Basics on Bayesian Probability Theory	19
2.3 Belief Updating Using JPD	23
2.4 Graphs	24
2.5 Bayesian Networks	27
2.6 Local Computation and Message Passing	30
2.7 Message Passing over Multiple Networks	31
2.8 Approximation with Massive Message Passing	33
2.9 Bibliographical Notes	35
2.10 Exercises	36
3 Belief Updating and Cluster Graphs	37
3.1 Guide to Chapter 3	38
3.2 Cluster Graphs	40
3.3 Conventions for Message Passing in Cluster Graphs	43
3.4 Relation with $\lambda - \pi$ Message Passing	44
3.5 Message Passing in Nondegenerate Cycles	47
3.6 Message Passing in Degenerate Cycles	53

vi	<i>Contents</i>	
3.7	Junction Trees	56
3.8	Bibliographical Notes	59
3.9	Exercises	59
4	Junction Tree Representation	61
4.1	Guide to Chapter 4	62
4.2	Graphical Separation	64
4.3	Sufficient Message and Independence	68
4.4	Encoding Independence in Graphs	69
4.5	Junction Trees and Chordal Graphs	71
4.6	Triangulation by Elimination	76
4.7	Junction Trees as I-maps	78
4.8	Junction Tree Construction	80
4.9	Bibliographical Notes	83
4.10	Exercises	84
5	Belief Updating with Junction Trees	86
5.1	Guide to Chapter 5	86
5.2	Algebraic Properties of Potentials	88
5.3	Potential Assignment in Junction Trees	94
5.4	Passing Belief over Separators	97
5.5	Passing Belief through a Junction Tree	100
5.6	Processing Observations	104
5.7	Bibliographical Notes	105
5.8	Exercises	105
6	Multiply Sectioned Bayesian Networks	107
6.1	Guide to Chapter 6	108
6.2	The Task of Distributed Uncertain Reasoning	112
6.3	Organization of Agents during Communication	117
6.4	Agent Interface	124
6.5	Multiagent Dependence Structure	128
6.6	Multiply Sectioned Bayesian Networks	133
6.7	Bibliographical Notes	137
6.8	Exercises	140
7	Linked Junction Forests	142
7.1	Guide to Chapter 7	143
7.2	Multiagent Distributed Compilation of MSBNs	146
7.3	Multiagent Moralization of MSDAG	147
7.4	Effective Communication Using Linkage Trees	152
7.5	Linkage Trees as I-maps	155
7.6	Multiagent Triangulation	158

<i>Contents</i>	<i>vii</i>
7.7 Constructing Local Junction Trees and Linkage Trees	174
7.8 Bibliographical Notes	181
7.9 Exercises	181
8 Distributed Multiagent Inference	182
8.1 Guide to Chapter 8	183
8.2 Potentials in a Linked Junction Forest	186
8.3 Linkage Trees over the Same d-sepset	190
8.4 Extended Linkage Potential	192
8.5 E-message Passing between Agents	194
8.6 Multiagent Communication	196
8.7 Troubleshooting a Digital System	201
8.8 Complexity of Multiagent Communication	207
8.9 Regional Multiagent Communication	208
8.10 Alternative Methods for Multiagent Inference	209
8.11 Bibliographical Notes	212
8.12 Exercises	213
9 Model Construction and Verification	215
9.1 Guide to Chapter 9	216
9.2 Multiagent MSBN System Integration	217
9.3 Verification of Subdomain Division	219
9.4 Agent Registration	221
9.5 Distributed Verification of Acyclicity	223
9.6 Verification of Agent Interface	237
9.7 Complexity of Cooperative d-sepset Testing	271
9.8 Bibliographical Notes	272
9.9 Exercises	272
10 Looking into the Future	274
10.1 Multiagent Reasoning in Dynamic Domains	274
10.2 Multiagent Decision Making	277
10.3 What If Verification Fails?	279
10.4 Dynamic Formation of MSBNs	279
10.5 Knowledge Adaptation and Learning	280
10.6 Negotiation over Agent Interfaces	281
10.7 Relaxing Hypertree Organization	283
10.8 Model Approximation	284
10.9 Mixed Models	285
10.10 Bibliographical Notes	285
<i>Bibliography</i>	287
<i>Index</i>	293

Preface

This book investigates opportunities for building intelligent decision support systems offered by multiagent, distributed probabilistic reasoning. Probabilistic reasoning with graphical models, known as Bayesian networks or belief networks, has become an active field of research and practice in artificial intelligence, operations research, and statistics in the last two decades. Inspired by the success of Bayesian networks and other graphical dependence models under the centralized and single-agent paradigm, this book extends them to representation formalisms under the distributed and multiagent paradigm. The major technical challenges to such an endeavor are identified and the results from a decade's research are presented. The framework developed allows distributed representation of uncertain knowledge on a large and complex environment embedded in multiple cooperative agents and effective, exact, and distributed probabilistic inference.

Under the single-agent paradigm, many exact or approximate methods have been proposed for probabilistic reasoning using graphical models. Not all of them admit effective extension into the multiagent paradigm. Concise message passing in a compiled, treelike graphical structure has emerged from a decade's research as one class of methods that extends well into the multiagent paradigm. How to structure multiple agents' diverse knowledge on a complex environment as a set of coherent probabilistic graphical models, how to compile these models into graphical structures that support concise message passing, and how to perform concise message passing to accomplish tasks in model verification, model compilation, and distributed inference are the foci of the book. The advantages of concise message passing over alternative methods are also analyzed.

It would be impossible to present multiagent probabilistic reasoning without an exposition of its single-agent counterpart. The results from single-agent inference have been the subject of several books (Pearl [52]; Neapolitan [43]; Lauritzen [36]; Jensen [29]; Shafer [62]; Castillo, Gutierrez, and Hadi [6]; and Cowell et al. [9]). Only a small subset of these results, which were most influential

Cambridge University Press

978-0-521-15390-4 - Probabilistic Reasoning in Multiagent Systems: A Graphical Models Approach

Yang Xiang

Frontmatter

[More information](#)

x

Preface

to the work presented on multiagent probabilistic reasoning, is included in this book. In particular, only the theory and algorithms central to concise message-passing methods are covered in detail. These results are attributed mainly to the work of J. Pearl and his disciples as well as the Hugin researchers in Denmark. In presenting these results, instead of describing them as given solutions, the book is structured to emphasize why essential aspects of these solutions are necessary. Results from the author's own research in this regard are presented.

The book is organized into two parts. The first part includes Chapters 1 through 5 and covers probabilistic inference by concise message passing under the single-agent paradigm. Readers are prepared for comprehension of the second half of the book on multiagent probabilistic inference. The second part comprises Chapters 6 through 10 in which a formal framework is elaborated for distributed representation of probabilistic knowledge in a cooperative multiagent system and for effective, exact, and distributed inference by the agents.

Chapter 1 outlines the roles of intelligent agents and multiagent systems in decision support systems and substantiates the need for uncertain reasoning. Chapter 2 introduces Bayesian networks as a concise representation of probabilistic knowledge and demonstrates the idea of belief updating by concise message passing. Chapter 3 introduces cluster graphs as alternative models for effective belief updating by concise message passing. Through analyses of possible types of cycles in cluster graphs, this chapter formally establishes that belief updating by concise message passing requires cluster trees and, in particular, junction tree models. Chapter 4 defines graphical separation criteria in three types of graphical models and the concept of I-maps. The chapter describes stepwise how to compile a Bayesian network into a junction tree model while preserving the I-mapness as much as possible. Chapter 5 defines common operations on potentials and presents laws governing mixed operations. Algorithms for belief updating by passing potentials as messages in a junction tree are presented. Chapter 6 sets forth five basic assumptions on uncertain reasoning in a cooperative multiagent system. The logic consequences of these assumptions, which imply a particular knowledge representation formalism termed a multiply sectioned Bayesian network (MSBN), are derived. Chapter 7 presents a set of distributed algorithms used to compile an MSBN into a collection of related junction tree models, termed a linked junction forest, for effective multiagent belief updating. Chapter 8 describes a set of algorithms for performing multiagent communication and belief updating by concise message passing. The material presented in this chapter establishes that multiagent probabilistic reasoning using an MSBN is exact, distributed, and efficient (when the MSBN is sparse). Chapter 9 addresses the issues of model construction and verification and presents distributed algorithms for ensuring the integration of independently developed agents into a syntactically and semantically correct MSBN.

Cambridge University Press

978-0-521-15390-4 - Probabilistic Reasoning in Multiagent Systems: A Graphical Models Approach

Yang Xiang

Frontmatter

[More information](#)*Preface*

xi

Chapter 10 puts the state of affairs in cooperative multiagent probabilistic reasoning in perspective and outlines several research issues in extending MSBNs into more powerful frameworks for future intelligent decision support systems.

The book is intended for researchers, educators, practitioners, and graduate students in artificial intelligence, multiagent systems, uncertain reasoning, operations research, and statistics. It can be used for self-study, as a handbook for practitioners, or as a supplemental text for graduate-level courses on multiagent systems or uncertain reasoning with graphical models. A set of exercises is included at the end of most chapters for teaching and learning. Familiarity with algorithms and mathematical exposure from a typical computer science undergraduate curriculum (discrete structure and probability) are sufficient background. Previous exposure to artificial intelligence and distributed systems is beneficial but not required.

The book treats major results formally with the underlying ideas motivated and explained intuitively, and the algorithms as well as other results are demonstrated through many examples. All algorithms are presented at sufficient levels of detail for implementation. They are written in pseudocode and can be implemented with languages of the reader's choice. The executable code of a Java-based toolkit *WebWeavr*, which implements most of the algorithms in the book, can be downloaded from

http://snowwhite.cis.uoguelph.ca/faculty_info/yxiang/

Most of the chapters (Chapters 2 through 9) contain a Guide to Chapter section as a short roadmap to the chapter. Styled differently from the rest of the chapter, this section presents no formal materials. Instead, the main issues, ideas, and results are intuitively described and often illustrated with simple examples. These sections can be used collectively as a quick tour of the more formal content of the book. They can also be used by practitioners to determine the right focus of materials for their needs.

The following convention is followed in numbering theorem-like structures: Within each chapter, all algorithms are numbered with a single sequence, and all other formal structures are numbered with another sequence, including definitions, lemmas, propositions, theorems, and corollaries.

The input, inspiration, and support of many people were critical in making this book a reality, and I am especially grateful to them: David Poole introduced me to the literature on uncertain reasoning with graphical models. Michael Beddoes made the PainULim project, during which the framework of single-agent oriented MSBNs was born, possible. Andrew Eisen and Bhanu Pant provided domain expertise in the PainULim project, and their intuition inspired the ideas behind the formal MSBN framework. Judea Pearl acted as the external examiner for my Ph.D. dissertation in which the theory of MSBNs was first documented. I owe a great deal to Bill

Cambridge University Press

978-0-521-15390-4 - Probabilistic Reasoning in Multiagent Systems: A Graphical Models Approach

Yang Xiang

Frontmatter

[More information](#)

Havens for supporting my postdoctoral research. Nick Cercone has been a long-time colleague and has given me much support and encouragement over the years. Finn Jensen invited me for a research trip to Aalborg University during which many interesting interactions and exchanges of ideas took place. Victor Lesser was the host of my one-year sabbatical at the University of Massachusetts, and for years he has encouraged and supported the work toward a multiagent inference framework based on MSBNs. Michael Wong taught me much when I was a junior faculty member.

The work reported has benefited from my interaction with many academic colleagues, mostly in the fields of multiagent systems and uncertain reasoning: Craig Boutilier, Brahim Chaib-draa, Bruce D'Ambrosio, Keith Decker, Abhijit Deshmukh, Robert Dodier, Edmund Durfee, Piotr Gmytrasiewicz, Randy Goebel, Carl Hewitt, Michael Huhns, Stephen Joseph, Uffe Kjaerulff, Burton Lee, Alan Mackworth, Eric Neufeld, Kristian Olesen, Simon Parsons, Gregory Provan, Tuomas Sandholm, Eugene Santos, Jr., Paul Snow, Michael Wellman, Nevin Lianwen Zhang, and Shlomo Zilberstein. Students in the Intelligent Decision Support Systems Laboratory have been very helpful. Xiaoyun Chen read and commented on the drafts. I thank the users of the *WebWeavr* toolkit throughout the world for their interest and encouragement, and I hope to make an enhanced version of the toolkit available soon. My thanks also go to anonymous reviewers whose comments on the early draft proved valuable.

The Natural Sciences and Engineering Research Council (NSERC) deserves acknowledgment for sponsoring the research that has led to these results. Additional funding was provided by the Institute of Robotics and Intelligent Systems (IRIS) in the Networks of Centres of Excellence program. A significant portion of the research presented was conducted while I was a faculty member at the University of Regina. I am grateful to my many colleagues there, chaired by Brien Maguire at the time of my departure, for years of cooperation and friendship. Some of the manuscript was completed while I was visiting the University of Massachusetts at Amherst, and it was partially funded by the National Science Foundation (NSF) and the Defense Advanced Research Projects Agency (DARPA).

I would like to thank the editorial and production staffs, Lauren Cowles and Caitlin Daggart at Cambridge University Press, and Eleanor Umali and John Joswick at TechBooks for their hard work.

I am greatly indebted to my parents for their caring and patience during my extended absence. I especially would like to thank my wife, Zoe, for her love, encouragement, and support. I am also grateful to my children, Jing and Jasen, for learning to live out of a cardboard box as we moved across the country.