Reasoning with Probabilistic and Deterministic Graphical Models

Exact Algorithms

Second Edition

Synthesis Lectures on Artificial Intelligence and Machine Learning

Lifelong Machine Learning
Zhiyuan Chen and Bing Liu
2016

Statistical Relational Artificial Intelligence: Logic, Probability, and Computation
Luc De Raedt, Kristian Kersting, Sriraam Natarajan, and David Poole
2016

Representing and Reasoning with Qualitative Preferences: Tools and Applications
Ganesh Ram Santhanam, Samik Basu, and Vasant Honavar
2016

Metric Learning
Aurélien Bellet, Amaury Habrard, and Marc Sebban
2015

Graph-Based Semi-Supervised Learning
Amarnag Subramanya and Partha Pratim Talukdar
2014

Robot Learning from Human Teachers
Sonia Chernova and Andrea L. Thomaz
2014

General Game Playing
Michael Genesereth and Michael Thielscher
2014

Judgment Aggregation: A Primer
Davide Grossi and Gabriella Pigozzi
2014

An Introduction to Constraint-Based Temporal Reasoning
Roman Barták, Robert A. Morris, and K. Brent Venable
2014

Reasoning with Probabilistic and Deterministic Graphical Models: Exact Algorithms
Rina Dechter
2013

Introduction to Intelligent Systems in Traffic and Transportation
Ana L.C. Bazzan and Franziska Klügl
2013

A Concise Introduction to Models and Methods for Automated Planning
Hector Geffner and Blai Bonet
2013

Essential Principles for Autonomous Robotics
Henry Hexmoor
2013

Case-Based Reasoning: A Concise Introduction
Beatriz López
2013

Answer Set Solving in Practice
Martin Gebser, Roland Kaminski, Benjamin Kaufmann, and Torsten Schaub
2012

Planning with Markov Decision Processes: An AI Perspective
Mausam and Andrey Kolobov
2012

Active Learning
Burr Settles
2012

Computational Aspects of Cooperative Game Theory
Georgios Chalkiadakis, Edith Elkind, and Michael Wooldridge
2011

Representations and Techniques for 3D Object Recognition and Scene Interpretation
Derek Hoiem and Silvio Savarese
2011

A Short Introduction to Preferences: Between Artificial Intelligence and Social Choice
Francesca Rossi, Kristen Brent Venable, and Toby Walsh
2011

Human Computation
Edith Law and Luis von Ahn
2011

Trading Agents
Michael P. Wellman
2011

Visual Object Recognition
Kristen Grauman and Bastian Leibe
2011

Learning with Support Vector Machines
Colin Campbell and Yiming Ying
2011

Algorithms for Reinforcement Learning
Csaba Szepesvári
2010

Data Integration: The Relational Logic Approach
Michael Genesereth
2010

Markov Logic: An Interface Layer for Artificial Intelligence
Pedro Domingos and Daniel Lowd
2009

Introduction to Semi-Supervised Learning
XiaojinZhu and Andrew B.Goldberg
2009

Action Programming Languages
Michael Thielscher
2008

Representation Discovery using Harmonic Analysis
Sridhar Mahadevan
2008

Essentials of Game Theory: A Concise Multidisciplinary Introduction
Kevin Leyton-Brown and Yoav Shoham
2008

A Concise Introduction to Multiagent Systems and Distributed Artificial Intelligence
Nikos Vlassis
2007

Intelligent Autonomous Robotics: A Robot Soccer Case Study
Peter Stone
2007

Reasoning with Probabilistic and Deterministic Graphical Models: Exact Algorithms, Second Edition

Rina Dechter

ISBN: 978-3-031-00455-1 paperback
ISBN: 978-3-031-01583-0 ebook
ISBN: 978-3-031-00028-7 hardcover

DOI 10.1007/978-3-031-01583-0

A Publication in the Springer series
SYNTHESIS LECTURES ON ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING

Lecture #41
Series Editors: Ronald J. Brachman, *Jacobs Technion-Cornell Institute at Cornell Tech*
Francesca Rossi, *IBM Research AI*
Peter Stone, *University of Texas at Austin*
Series ISSN
Synthesis Lectures on Artificial Intelligence and Machine Learning
Print 1939-4608 Electronic 1939-4616

Reasoning with Probabilistic and Deterministic Graphical Models

Exact Algorithms

Second Edition

Rina Dechter
University of California, Irvine

SYNTHESIS LECTURES ON ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING #41

ABSTRACT

Graphical models (e.g., Bayesian and constraint networks, influence diagrams, and Markov decision processes) have become a central paradigm for knowledge representation and reasoning in both artificial intelligence and computer science in general. These models are used to perform many reasoning tasks, such as scheduling, planning and learning, diagnosis and prediction, design, hardware and software verification, and bioinformatics. These problems can be stated as the formal tasks of constraint satisfaction and satisfiability, combinatorial optimization, and probabilistic inference. It is well known that the tasks are computationally hard, but research during the past three decades has yielded a variety of principles and techniques that significantly advanced the state of the art.

This book provides comprehensive coverage of the primary exact algorithms for reasoning with such models. The main feature exploited by the algorithms is the model's graph. We present inference-based, message-passing schemes (e.g., variable-elimination) and search-based, conditioning schemes (e.g., cycle-cutset conditioning and AND/OR search). Each class possesses distinguished characteristics and in particular has different time vs. space behavior. We emphasize the dependence of both schemes on few graph parameters such as the treewidth, cycle-cutset, and (the pseudo-tree) height. The new edition includes the notion of influence diagrams, which focus on sequential decision making under uncertainty. We believe the principles outlined in the book would serve well in moving forward to approximation and anytime-based schemes. The target audience of this book is researchers and students in the artificial intelligence and machine learning area, and beyond.

KEYWORDS

graphical models, Bayesian networks, constraint networks, Markov networks, influence diagrams, induced-width, treewidth, cycle-cutset, loop-cutset, pseudo-tree, bucket-elimination, variable-elimination, AND/OR search, conditioning, reasoning, inference, knowledge representation

Contents

Preface

Graphical models, including constraint networks (hard and soft), Bayesian networks, Markov random fields, and influence diagrams, have become a central paradigm for knowledge representation and reasoning, and provide powerful tools for solving problems in a variety of application domains, including scheduling and planning, coding and information theory, signal and image processing, data mining, computational biology, and computer vision.

These models can be acquired from experts or learned from data. Once a model is available, we need to be able to make deductions and to extract various types of information. We refer to this as *reasoning* in analogy with the human process of thinking and reasoning. These reasoning problems can be stated as the formal tasks of constraint satisfaction and satisfiability, combinatorial optimization, and probabilistic inference. It is well known that these tasks are computationally hard, but research during the past three decades has yielded a variety of effective principles and led to impressive scalability of exact techniques.

In this book we provide a comprehensive coverage of the main exact algorithms for reasoning with such models. The primary feature exploited by the algorithms is the model's graph structure and they are therefore uniformly applicable across a broad range of models, where dependencies are expressed as constraints, cost functions or probabilistic relationships. We also provide a glimpse into properties of the dependencies themselves, known as context-specific independencies, when treating deterministic functions such as constraints. Clearly, exact algorithms must be complemented by approximations. Indeed, we see this book as the first phase of a broader book that would cover approximation algorithms as well. We believe, however, that in order to have effective approximations we have to start with the best exact algorithms.

The book is organized into seven chapters and a conclusion. Chapter 1 provides an introduction to the book and its contents. Chapter 2 introduces the reader to the formal definition of the general graphical model and then describes the most common models, including constraint networks and probabilistic networks, which are used throughout the book. We distinguish two classes of algorithms: inference-based, message-passing schemes (Chapters 3, 4, and 5) and search-based, conditioning schemes (Chapters 6 and 7). This division is useful because algorithms in each class possesses common and distinguished characteristics and in particular have different behavior with respect to the tradeoff between time and memory. Chapter 7 focuses on this tradeoff, introducing hybrids of search and inference schemes. We emphasize the dependence of both types on few graph parameters such as the treewidth, cycle-cutset, and (the pseudo-tree) height.

The book is based on research done in my lab over the past two decades. It is largely founded on work with my graduate and postdoctoral students including: Dan Frost, Irina

Rish, Kalev Kask, David Larkin, Robert Mateescu, Radu Marinescu, Bozhena Bidyuk, Vibhav Gogate, Lars Otten, Natasha Flerova and William Lam and my postdoctoral students Javier Larrosa, and Emma Rollon. Most heavily it relies on the work of Kalev Kask (Chapter 5) and Robert Mateescu (Chapters 6 and 7). I wish to also thank my colleagues at UCI for providing a supportive environment in our AI and machine learning labs, and especially to Alex Ihler for our recent collaboration that has been particularly inspiring and fruitful.

I owe a great deal to members of my family that took an active role in some parts of this book. First, to my son Eyal who spent several months reading and providing editing, as well as very useful suggestions regarding the book's content and exposition. Thanks also go to my husband Avi on providing editorial comments on large parts of this book and to Anat Gafni for her useful comments on Chapter 1.

The second edition of the book contains many edits and corrections that were suggested by students who took my classes over the last five years and by other readers, and I would like to thank them all. It also includes two new subsections on influence diagrams in Chapter 2 and Chapter 4.

Rina Dechter
Los Angeles, January 2019