# Modeling and Reasoning with Bayesian Networks

This book provides a thorough introduction to the formal foundations and practical applications of Bayesian networks. It provides an extensive discussion of techniques for building Bayesian networks that model real-world situations, including techniques for synthesizing models from design, learning models from data, and debugging models using sensitivity analysis. It also treats exact and approximate inference algorithms at both theoretical and practical levels. The treatment of exact algorithms covers the main inference paradigms based on elimination and conditioning and includes advanced methods for compiling Bayesian networks, time-space tradeoffs, and exploiting local structure of massively connected networks. The treatment of approximate algorithms covers the main inference paradigms based on sampling and optimization and includes influential algorithms such as importance sampling, MCMC, and belief propagation.

The author assumes very little background on the covered subjects, supplying in-depth discussions for theoretically inclined readers and enough practical details to provide an algorithmic cookbook for the system developer.

Adnan Darwiche is a Professor and Chairman of the Computer Science Department at UCLA. He is also the Editor-in-Chief for the *Journal of Artificial Intelligence Research* (JAIR) and a AAAI Fellow.

# Modeling and Reasoning with Bayesian Networks

**Adnan Darwiche**

*University of California, Los Angeles*

CAMBRIDGE
UNIVERSITY PRESS

To Jinan, Layla, Sarah, Sikna, and Youssef.

CAMBRIDGE UNIVERSITY PRESS
Cambridge, New York, Melbourne, Madrid, Cape Town, Singapore, São Paulo, Delhi

Cambridge University Press
32 Avenue of the Americas, New York, NY 10013-2473, USA

www.cambridge.org
Information on this title: www.cambridge.org/9780521884389



First published 2009

Printed in the United States of America

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

*Library of Congress Cataloging in Publication data*

Darwiche, Adnan, 1966–
Modeling and reasoning with Bayesian networks / Adnan Darwiche.
p. cm.
ISBN 978-0-521-88438-9 (hardback)
1. Bayesian statistical decision theory – Graphic methods. 2. Inference. 3. Probabilities. 4. Modeling.
I. Title.
QA279.5.D37 2009
519.5′42–dc22 2008044605

ISBN 978-0-521-88438-9 (hardback)

# Contents

# Preface

Bayesian networks have received a lot of attention over the last few decades from both scientists and engineers, and across a number of fields, including artificial intelligence (AI), statistics, cognitive science, and philosophy.

Perhaps the largest impact that Bayesian networks have had is on the field of AI, where they were first introduced by Judea Pearl in the midst of a crisis that the field was undergoing in the late 1970s and early 1980s. This crisis was triggered by the surprising realization that a theory of plausible reasoning cannot be based solely on classical logic [McCarthy, 1977], as was strongly believed within the field for at least two decades [McCarthy, 1959]. This discovery has triggered a large number of responses by AI researchers, leading, for example, to the development of a new class of symbolic logics known as non-monotonic logics (e.g., [McCarthy, 1980; Reiter, 1980; McDermott and Doyle, 1980]). Pearl's introduction of Bayesian networks, which is best documented in his book [Pearl, 1988], was actually part of his larger response to these challenges, in which he advocated the use of probability theory as a basis for plausible reasoning and developed Bayesian networks as a practical tool for representing and computing probabilistic beliefs.

From a historical perspective, the earliest traces of using graphical representations of probabilistic information can be found in statistical physics [Gibbs, 1902] and genetics [Wright, 1921]. However, the current formulations of these representations are of a more recent origin and have been contributed by scientists from many fields. In statistics, for example, these representations are studied within the broad class of graphical models, which include Bayesian networks in addition to other representations such as Markov networks and chain graphs [Whittaker, 1990; Edwards, 2000; Lauritzen, 1996; Cowell et al., 1999]. However, the semantics of these models are distinct enough to justify independent treatments. This is why we decided to focus this book on Bayesian networks instead of covering them in the broader context of graphical models, as is done by others [Whittaker, 1990; Edwards, 2000; Lauritzen, 1996; Cowell et al., 1999]. Our coverage is therefore more consistent with the treatments in [Jensen and Nielsen, 2007; Neapolitan, 2004], which are also focused on Bayesian networks.

Even though we approach the subject of Bayesian networks from an AI perspective, we do not delve into the customary philosophical debates that have traditionally surrounded many works on AI. The only exception to this is in the introductory chapter, in which we find it necessary to lay out the subject matter of this book in the context of some historical AI developments. However, in the remaining chapters we proceed with the assumption that the questions being treated are already justified and simply focus on developing the representational and computational techniques needed for addressing them. In doing so, we have taken a great comfort in presenting some of the very classical techniques in ways that may seem unorthodox to the expert. We are driven here by a strong desire to provide the most intuitive explanations, even at the expense of breaking away from norms. We have also made a special effort to appease the scientist, by our emphasis on justification, and the engineer, through our attention to practical considerations.

There are a number of fashionable and useful topics that we did not cover in this book, which are mentioned in the introductory chapter. Some of these topics were omitted because their in-depth treatment would have significantly increased the length of the book, whereas others were omitted because we believe they conceptually belong somewhere else. In a sense, this book is not meant to be encyclopedic in its coverage of Bayesian networks; rather it is meant to be a focused, thorough treatment of some of the core concepts on modeling and reasoning within this framework.

## Acknowledgments

In writing this book, I have benefited a great deal form a large number of individuals who provided help at levels that are too numerous to explicate here. I wish to thank first and foremost members of the automated reasoning group at UCLA for producing quite a bit of the material that is covered in this book, and for their engagement in the writing and proofreading of many of its chapters. In particular, I would like to thank David Allen, Keith Cascio, Hei Chan, Mark Chavira, Arthur Choi, Taylor Curtis, Jinbo Huang, James Park, Knot Pipatsrisawat, and Yuliya Zabiyaka. Arthur Choi deserves special credit for writing the appendices and most of Chapter 14, for suggesting a number of interesting exercises, and for his dedicated involvement in the last stages of finishing the book. I am also indebted to members of the cognitive systems laboratory at UCLA – Blai Bonet, Ilya Shipster, and Jin Tian – who have thoroughly read and commented on earlier drafts of the book. A number of the students who took the corresponding graduate class at UCLA have also come to the rescue whenever called. I would like to especially thank Alex Dow for writing parts of Chapter 9. Moreover, Jason Aten, Omer Bar-or, Susan Chebotariov, David Chen, Hicham Elmongui, Matt Hayes, Anand Panangadan, Victor Shih, Jae-il Shin, Sam Talaie, and Mike Zaloznyy have all provided detailed feedback on numerous occasions.

I would also like to thank my colleagues who have contributed immensely to this work through either valuable discussions, comments on earlier drafts, or strongly believing in this project and how it was conducted. In this regard, I am indebted to Russ Almond, Bozhena Bidyuk, Hans Bodlaender, Gregory Cooper, Rina Dechter, Marek Druzdzel, David Heckerman, Eric Horvitz, Linda van der Gaag, Hector Geffner, Vibhav Gogate, Russ Greiner, Omid Madani, Ole Mengshoel, Judea Pearl, David Poole, Wojtek Przytula, Silja Renooij, Stuart Russell, Prakash Shenoy, Hector Palacios Verdes, and Changhe Yuan.

Finally, I wish to thank my wife, Jinan, and my daughters, Sarah and Layla, for providing a warm and stimulating environment in which I could conduct my work. This book would not have seen the light without their constant encouragement and support.