LOC Page

Contents

Foreword

Probabilistic models based on directed acyclic graphs have a long and rich tradition, beginning with work by the geneticist Sewall Wright in the 1920s. Variants have appeared in many fields. Within statistics, such models are known as directed graphical models; within cognitive science and artificial intelligence, such models are known as Bayesian networks. The name honors the Reverend Thomas Bayes (1702–1761), whose rule for updating probabilities in the light of new evidence is the foundation of the approach. The initial development of Bayesian networks in the late 1970s was motivated by the need to model the top-down (semantic) and bottom-up (perceptual) combination of evidence in reading. The capability for bidirectional inferences, combined with a rigorous probabilistic foundation, led to the rapid emergence of Bayesian networks as the method of choice for uncertain reasoning in AI and expert systems, replacing earlier ad-hoc rule-based schemes. Perhaps the most important aspect of Bayesian networks is that they are direct representations of the world, not of reasoning processes. The arrows in the diagrams represent real causal connections and not the flow of information during reasoning (as in rule-based systems or neural networks). Reasoning processes can operate on Bayesian networks by propagating information in any direction. For example, if the sprinkler is on, then the pavement is probably wet (prediction); if someone slips on the pavement, that also provides evidence that it is wet (abduction, or reasoning to a probable cause). On the other hand, if we see that the pavement is wet, that makes it more likely that the sprinkler is on or that it is raining (abduction); but if we then observe that the sprinkler is on, that reduces the likelihood that it is raining. It is the ability to perform this last form of reasoning—called explaining away—that makes Bayesian networks so powerful compared to rule-based systems or neural networks. They are especially useful and important for risk assessment and decision-making.

Although Bayesian networks are now used widely in many disciplines, those responsible for developing (as opposed to using) Bayesian network models typically require highly specialist knowledge of mathematics, probability, statistics, and computing. Part of the reason for this is that, although there have been several excellent books dedicated to Bayesian Networks and related methods, these books tend to be aimed at readers who already have a high level of mathematical sophistication—typically they are books that would be used at graduate or advanced undergraduate level in mathematics, statistics, or computer science. As such they are not accessible to readers who are not already proficient in those subjects. This book is an exciting development because it addresses this problem. While I am sure it would be suitable for undergraduate courses on probability and risk, it should be understandable by any numerate reader interested in risk assessment and decision making. The book provides sufficient motivation and examples (as well as the mathematics and probability where needed from scratch) to enable readers to understand the core principles and power of Bayesian networks. However, the focus is on ensuring that readers can build practical Bayesian network models, rather than understand in depth the underlying algorithms and theory. Indeed readers are provided with a tool that performs the propagation, so they will be able to build their own models to solve real-world risk assessment problems.

Judea Pearl
UCLA Computer Science Department
Los Angeles, California

Preface

Businesses and governments must often assess and manage risk in areas where there is little or no direct historical data to draw upon, or where relevant data is difficult to identify. The international financial crisis that started in 2008 was not predicted by the majority of the world's leading economists, because they relied on models based on historical statistical data that could not adapt to new circumstances even when those circumstances (in this case the collapse of the mortgage subprime market) were foreseeable by some experts. In short, analysts are too dependent on models that are good for predicting the past but poor at predicting the future.

The challenges are similarly acute when the source of the risk is novel: terrorist attacks, ecological disasters, major project failures, and more general failures of novel systems, marketplaces, and business models. Even though we may have little or no historical data in these cases, there is often an abundance of expert (but subjective) judgment, as well as diverse information and data on indirectly related risks. These are the types of situations that can be successfully addressed using Bayesian networks (BNs), even when classical, data-driven approaches to risk assessment are not possible.

BNs describe networks of causes and effects, using a graphical framework that provides for the rigorous quantification of risks and the clear communication of results. They can combine historical data and expert judgment, using calculations that are based on a theorem created by the Reverend Thomas Bayes dating back to 1763. The theorem provides the only rational and consistent way to solve the problem of how to update a belief in some uncertain event (such as a decline in a company's share price of more than 10% next year) when we observe new evidence about the event (such as better than expected first quarter figures). The problem of correctly updating beliefs in the light of new evidence is one that is central to many disciplines (law, medicine, and engineering as well as finance). It is also prevalent in our everyday decision making (even those of us who are not gamblers).

BNs are now widely recognized as an exciting and powerful technology for handling risk assessment, uncertainty, and decision making. During the last decade, researchers have incorporated BN techniques into easy-to-use toolsets, which in turn have enabled the development of decision support systems in a diverse set of application domains, including medical diagnosis, safety assessment, the law, forensics, procurement, equipment fault diagnosis and software quality. Further technology and tool advancements mean that end users, rather than just researchers, are now able to develop and deploy their own BN-based solutions. Recent commercial case studies provide evidence of impressive returns on investment from these techniques.

But, although many BN applications are now in everyday use, BNs have not yet achieved the mainstream penetration that the approach deserves. Although there have been several excellent books dedicated to BNs and related methods (these are described in the Further Reading sections), these books are largely aimed at the mathematical and statistical community and focus more on algorithms and theory rather than practical real-world problem solving and model building. Hence, there is a shortage of material that is truly accessible to people who would benefit from BN solutions but who are not already qualified mathematicians or statisticians.

This book aims to fix this problem. Although it is suitable for undergraduate courses on probability and risk, it is written to be understandable by other professional people generally interested in risk assessment and decision making. Unlike other BN books, this one makes no assumptions about previous probability and statistical knowledge. It is driven by real examples that are introduced early on as motivation, with the probability and statistics introduced (from scratch) as and when necessary. The more mathematical topics are separated from the main text by comprehensive use of boxes and appendices. The focus is on applications and practical model building. A free version of the powerful commercial software tool AgenaRisk is available for those who have purchased the book, as we think the only real way to learn about BNs is to build and use them.

Many of the examples in this book are influenced by our academic research but also by our experience in putting the ideas into practice with commercial and government decision and policy makers. Together we have consulted for a wide variety of commercial and government organizations, including Absa Bank, Bosch, Ericsson, Motorola, NATS, Philips, QinetiQ, Dstl (UK Ministry of Defence), Royal Bank of Canada, TNO, and VocaLink. We are firm believers in technology transfer and putting useful decision-making theory into the hands of those at the sharp end.

The book has two parts. The first ten chapters of the book teache all the basics of probability and risk, and about building and using BN models, whereas the last three chapters go into the detailed applications. The underlying BN algorithms appear in appendices rather than the main text since there is no need to understand these to build and use BN models. Although purists have argued that only by understanding the algorithms can you understand the limitations and hence build efficient BN models, we overcome this by providing pragmatic advice about model building to ensure models are built efficiently. Our approach means that the main body of the text is free of the intimidating mathematics that has been a major impediment to the more widespread use of BNs.

The book has a dedicated Web site, www.BayesianRisk.com, which contains executable versions of all of the models described, exercises and worked solutions for all chapters, PowerPoint slides, numerous other resources, and a free downloadable copy of the AgenaRisk software to readers who have purchased a copy of the book.

For our parents, wives and children.

Acknowledgments

Numerous people have helped in various ways to ensure that we achieved the best possible outcome for this book. We would especially like to thank the following:

Colin Aitken, Carol Alexander, Peter Ayton, David Balding, Beth Bateman, George Bearfield, Daniel Berger, Nic Birtles, Robin Bloomfield, Bill Boyce, Rob Calver, Neil Cantle, Patrick Cates, Chris Chapman, Xiaoli Chen, Keith Clarke, Julie Cooper, Robert Cowell, Anthony Constantinou, Paul Curzon, Phil Dawid, Chris Eagles, Shane Cooper, Eugene Dementiev, Itiel Dror, John Elliott, Phil Evans, Ian Evett, Geir Fagerhus, Simon Forey, Duncan Gillies, Jean-Jacques Gras, Gerry Graves, David Hager, George Hanna, David Hand, Roger Harris, Peter Hearty, Joan Hunter, Jose Galan, Steve Gilmour, Shlomo Gluck, James Gralton, Richard Jenkinson, Adrian Joseph, Ian Jupp, Agnes Kaposi, Paul Kaye, Kevin Korb, Paul Krause, Dave Lagnado, Helge Langseth, Steffen Lauritzen, Robert Leese, Peng Lin, Bev Littlewood, Paul Loveless, Peter Lucas, Bob Malcom, Amber Marks, David Marquez, William Marsh, Peter McOwan, Tim Menzies, Phil Mercy, Martin Newby, Richard Nobles, Magda Osman, Max Parmar, Judea Pearl, Elena Perez-Minana, Andrej Peitschker, Ursula Martin, Shoaib Qureshi, Lukasz Radlinksi, Soren Riis, Edmund Robinson, Thomas Roelleke, Angela Saini, Thomas Schulz, Jamie Sherrah, Leila Schneps, David Schiff, Bernard Silverman, Adrian Smith, Ian Smith, Jim Smith, Julia Sonander, David Spiegelhalter, Andrew Stuart, Alistair Sutcliffe, Lorenzo Strigini, Nigel Tai, Manesh Tailor, Franco Taroni, Ed Tranham, Marc Trepanier, Keith van Rijsbergen, Richard Tonkin, Sue White, Robin Whitty, Rosie Wild, Patricia Wiltshire, Rob Wirszycz, David Wright and Barbaros Yet.