

---

## Preface

Ever since the first machines were constructed, artists and scientists have shared a vision of a human-like machine: an autonomous self-moving machine that acts and reasons like a human being. Much effort has been put into this dream, but we are still very far from having androids with even the tiniest similarity to humans.

This does not mean that all of these efforts have been wasted. As a spin-off, we have seen a long series of inventions that can take over very specialized sections of human work. These inventions fall into two categories: machines that can make physical changes in the world and thereby substitute human labor, and machines that can perform activities usually thought of as requiring intellectual skills.

In contemporary science and engineering, we still have this split into two categories. The activity of the first category is mainly concentrated on the construction of robots. The aim is to construct autonomous machines performing “sophisticated” actions such as searching for a cup, finding a way from the office to a lavatory, driving a vehicle in a deserted landscape, or walking on two legs. Construction of such robots requires computers to perform certain kinds of artificial intelligence. Basically, it is the kind of intelligence that humans share with most mammals. It involves skills such as visual recognition of items, sound recognition, learning to abstract crucial items from a scene, or control of balance and position in 3-D space. Although they are very challenging research tasks, and they certainly require enormous computing power and very sophisticated algorithms, you would not say that these skills are intellectual, and the basis for the activity is the physical appearance of a device that moves. To put it another way: the success criterion is how the algorithms work when controlling a physical machine in real time.

The activity in the second category is basically concerned with reasoning and human activities that we presumably do not share with other animals. The activity is separated from matter. When performed, no changes in the physical world need to take place. The first real success was the automated calculator: a machine that can perform very large and complicated arithmetic

calculations. Automated calculation skill is nowadays hardly considered artificial intelligence, and we are now acquainted with computers performing tasks that decades ago were considered highly intellectual (e.g. taking derivatives of functions or performing reduction of mathematical expressions). When an activity has been so well understood that it can be formalized, it will soon be performed by computers, and gradually we acknowledge that this activity does not really require intelligence.

A branch of research in the second category has to do with reasoning. The first successes were in logical reasoning. Propositional logic is fully formalized, and although some tasks are NP-complete and therefore in some situations intractable for a computer, we have for propositional logic completed the transition from “intellectual task” to “we have computers to do this for us.”

Unfortunately, logical reasoning is very limited in scope. It deals with how to infer from propositions that you know are true. Very often you do not know the truth of a proposition for certain, but you still need to perform inference from your incomplete and uncertain knowledge. Actually, this is the most common situation for human reasoning. Reasoning under uncertainty is not yet so well understood that it can be formalized entirely for computers. There are several approaches to reasoning under uncertainty. The approach taken in this book is (subjective) probability theory. When the reasoning ends up in a conclusion on a decision, we use *utilities*, and we assume that the decision taken is the one that maximizes the expected utility. In other words, the approach prescribes a certain behavior. We may not always expect this behavior from human beings, and therefore the approach is also termed *normative*. There are alternative approaches to reasoning under uncertainty. Most prominent is *possibility theory*, which in certain contexts is called *fuzzy logic*. The interested reader may consult the wide literature on these approaches.

The aim of normative systems can in short be termed human *wisdom*: to take decisions on the basis of accumulated and processed experience. The tasks are of the following types:

- using observations to interpret a situation;
- focusing a search for more information;
- choosing an appropriate intervening action;
- adapting to changing environments;
- learning from experience.

A damping factor for properly exploiting the advances in artificial intelligence has for a long time been the lack of successes in robotics. An autonomous agent that moves, observes, and changes the world must carry a not easily controllable body.

Therefore, the advances have been exploited mainly in *decision support systems*, computer systems that provide advice for humans on highly specialized tasks. With the Internet, the scope of artificial intelligence has widened considerably. The Internet is an ideal nonphysical world for intelligent agents,

which are pure spirits without bodies. In the years to come, we will experience a flood of intelligent agents on the Internet, and companies as well as private persons will be able to launch their own agents to explore and collect information on the Internet. Also, we will experience the dark sides of human endeavor. Some agents will destroy, intrude, tell lies and so on, and we will have to defend ourselves against them. Agents will meet agents, and they will have to decide how to treat each other, they will have to learn from past experience, and they will have to adapt to changing environments.

During the 1990s, Bayesian networks and decision graphs attracted a great deal of attention as a framework for building normative systems, not only in research institutions but also in industry. Contrary to most other frameworks for handling uncertainty, a good deal of theoretical insight as well as practical experience is required in order to exploit the opportunities provided by Bayesian networks and decision graphs.

On the other hand, many scientists and engineers wish to exploit the possibilities of normative systems without being experts in the field. This book should meet that demand. It is intended for both classroom use and self-study, and it addresses persons who are interested in exploiting the approach for the construction of decision support systems or bodyless agents.

The theoretical exposition in the book is self-contained, and the mathematical prerequisite is some prior exposure to calculus and elementary graph theory. Throughout the book we alternate between theoretical exposition and practical examples for gaining experience with the use of Bayesian networks and decision graphs, and we have assumed that the reader has access to a computer system for handling Bayesian networks and influence diagrams (the exercises marked with an <sup>E</sup> require such a system). There are many systems, academic as well as commercial. A comprehensive list of systems can be found at [www.cs.berkeley.edu/~murphyk/Bayes/bnsoft.html](http://www.cs.berkeley.edu/~murphyk/Bayes/bnsoft.html). Several of the commercial systems have an academic version, which can be downloaded free of charge. In several chapters the presentation is based on examples, and for overview purposes there is a summary section at the end of each chapter.

A hands-on course could cover Sections 1.1–1.3, Chapter 2, Chapter 3, Sections 6.1–6.2, 7.1, 8.1–8.3, 9.1–9.4, and Sections 11.1–11.2. A first-year graduate course could cover Chapters 1–3, Sections 4.1–4.6, 5.2–5.3, 5.5, 5.7, 6.1–6.3, 7.1–7.3, Chapters 8–9, Sections 10.1–10.2, and Chapter 11.

The book is an introduction to Bayesian networks and decision graphs. Many results are not mentioned or just treated superficially. The following textbooks and monographs can be used for further study:

- Judea Pearl, *Probabilistic Reasoning in Intelligent Systems*, Morgan Kaufmann Publishers, 1988.
- Russell Almond, *Graphical Belief Modelling*, Chapman & Hall, 1995.
- Steffen L. Lauritzen, *Graphical Models*, Oxford University Press, 1996.
- Enrique Castillo, José M. Gutiérrez, and Ali S. Hadi, *Expert Systems and Probabilistic Network Models*, Springer-Verlag, 1997.

- Robert G. Cowell, A. Philip Dawid, and Steffen L. Lauritzen, *Probabilistic Networks and Expert Systems*, Springer-Verlag, 1999.
- Kevin B. Korb and Ann E. Nicholson, *Bayesian Artificial Intelligence*, Chapman & Hall 2004.
- Richard E. Neapolitan, *Learning Bayesian Networks*, Pearson Prentice Hall, 2004.

The annual *Conference on Uncertainty in Artificial Intelligence* ([www.auai.org](http://www.auai.org)) is the main forum for researchers working with Bayesian networks and decision graphs, so the best general references for further reading are the proceedings from these conferences.

Another relevant conference is the biannual *European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty* (ECSQARU). The conference deals with various approaches to uncertainty calculus, and the proceedings are published in the Springer-Verlag series *Lecture Notes in Artificial Intelligence*.

The book is supported by a web site, [bndg.cs.aau.dk](http://bndg.cs.aau.dk), which provides readers with solutions and models for selected exercises, a list of errata, special exercises, and other links relevant to the issues in the book.

## Changes from the First Edition

In the second edition, we have added several subjects. Primarily, we have included chapters presenting commonly used methods for learning graphical models, and we have extended the treatment of graphical languages for modeling decision problems. We have also reorganized the material such that Part I is devoted to Bayesian networks and Part II deals with decision graphs.

The mathematical treatment is intended to be at the same level as in the first edition. However, many of the new issues in the book are mathematically rather demanding, particularly learning. Some of the sections are marked with an asterisk to indicate that they are not required for reading any of the unmarked sections.

## Acknowledgments

We wish to express our gratitude to several people for ideas and comments during the preparation of the book. First, we thank present and previous colleagues in the Machine Intelligence group, Olav Bangsø, Søren L. Dittmer, Uffe Kjærulff, Tomáš Kočka, Anders L. Madsen, Dennis Nilsson, Kristian G. Olesen, Jose Peña, Jiří Vomlel, and Marta Vomlelová. We also thank the many academic colleagues around the world with whom we have had the pleasure of exchanging ideas, in particular Poul S. Eriksen, Linda van der Gaag, Helge Langseth, Steffen L. Lauritzen, Serafin Moral, Prakash Shenoy,

Antonio Salmerón, Claus Skaanning, Marco Valtorta, Yang Xiang, and Nevin Zhang. Special thanks to Søren Holbech Nielsen for assistance with figures, bibliography, and exercises.

We also thank several years' worth of undergraduate students who have had to cope with unfinished drafts of notes for parts of their course on decision support systems.

Aalborg, February 2007

Finn V. Jensen and Thomas D. Nielsen



<http://www.springer.com/978-0-387-68281-5>

Bayesian Networks and Decision Graphs

Nielsen, Th.D.; VERNER JENSEN, F.

2007, XVI, 448 p., Hardcover

ISBN: 978-0-387-68281-5