


The book was found

Probabilistic Graphical Models: Principles And Techniques (Adaptive Computation And Machine Learning Series)



Synopsis

Most tasks require a person or an automated system to reason -- to reach conclusions based on available information. The framework of probabilistic graphical models, presented in this book, provides a general approach for this task. The approach is model-based, allowing interpretable models to be constructed and then manipulated by reasoning algorithms. These models can also be learned automatically from data, allowing the approach to be used in cases where manually constructing a model is difficult or even impossible. Because uncertainty is an inescapable aspect of most real-world applications, the book focuses on probabilistic models, which make the uncertainty explicit and provide models that are more faithful to reality. Probabilistic Graphical Models discusses a variety of models, spanning Bayesian networks, undirected Markov networks, discrete and continuous models, and extensions to deal with dynamical systems and relational data. For each class of models, the text describes the three fundamental cornerstones: representation, inference, and learning, presenting both basic concepts and advanced techniques. Finally, the book considers the use of the proposed framework for causal reasoning and decision making under uncertainty. The main text in each chapter provides the detailed technical development of the key ideas. Most chapters also include boxes with additional material: skill boxes, which describe techniques; case study boxes, which discuss empirical cases related to the approach described in the text, including applications in computer vision, robotics, natural language understanding, and computational biology; and concept boxes, which present significant concepts drawn from the material in the chapter. Instructors (and readers) can group chapters in various combinations, from core topics to more technically advanced material, to suit their particular needs.

Book Information

Series: Adaptive Computation and Machine Learning series

Hardcover: 1280 pages

Publisher: The MIT Press; 1 edition (July 31, 2009)

Language: English

ISBN-10: 0262013193

ISBN-13: 978-0262013192

Product Dimensions: 8 x 1.7 x 9 inches

Shipping Weight: 4.7 pounds (View shipping rates and policies)

Average Customer Review: 4.0 out of 5 stars [See all reviews](#) (35 customer reviews)

Best Sellers Rank: #43,999 in Books (See Top 100 in Books) #4 in [Books > Computers &](#)

Technology > Computer Science > AI & Machine Learning > Natural Language Processing #7
inÂ Books > Computers & Technology > Computer Science > AI & Machine Learning > Computer
Vision & Pattern Recognition #8 inÂ Books > Computers & Technology > Computer Science > AI &
Machine Learning > Machine Theory

Customer Reviews

If you're trying to learn probabilistic graphical models on your own, this is the best book you can buy. The introduction to fundamental probabilistic concepts is better than most probability books out there and the rest of the book has the same quality and in-depth approach. References, discussions and examples are all chosen so that you can take this book as the centre of your learning and make a jump to more detailed treatment of any topic using other resources. Another huge plus is Professor Daphne Koller's online course material. Her course for probabilistic models is available online, and watching the videos alongside the book really helps sometimes. If you have a strong mathematical background, you may find the book a little bit too pedagogic for your taste, but if you're looking for a single resource to learn the topic on your own, then this book is what you need. The only problem with it is that it is a big book to carry around, and if you buy the Kindle edition for the iPad, you'll have to zoom into pages to read comfortably (or maybe I have bad eye sight), and Kindle app on iPad does not keep the zoom level across pages. So my experience is, zoom, pan, read, change page, zoom, pan, go back to previous page to see something, zoom, pan... You get the idea. I'd gladly pay more for a pdf version which I could read with other software on the iPad. Even though my reading experience has been a bit unpleasant due to Kindle app, the book deserves five stars, since it is the content that matters.

Stanford professor, Daphne Koller, and her co-author, Professor Nir Friedman, employed graphical models to motivate thoroughgoing explorations of representation, inference and learning in both Bayesian networks and Markov networks. They do their own bidding at the book's web page, [...], by giving readers a panoramic view of the book in an introductory chapter and a Table of Contents. On the same page, there is a link to an extensive Errata file which lists all the known errors and corrections made in subsequent printings of the book - all the corrections had been incorporated into the copy I have. The authors painstakingly provided necessary background materials from both probability theory and graph theory in the second chapter. Furthermore, in an Appendix, more tutorials are offered on information theory, algorithms and combinatorial optimization. This book is an authoritative extension of Professor Judea Pearl's seminal work on developing the Bayesian

Networks framework for causal reasoning and decision making under uncertainty. Before this book was published, I sent an e-mail to Professor Koller requesting some clarification of her paper on object-oriented Bayesian networks; she was most generous in writing an elaborate reply with deliberate speed.

I bought this book to use for the Coursera course on PGM taught by the author. It was essential to being able to follow the course. I would not say that it is an easy book to pick up and learn from. It was a good reference to use to get more details on the topics covered in the lectures.

This was the book that really got me into AI research. Clearly written and detailed. I especially like that variational inference is taught using discrete variables so you don't need to learn both variational inference and calculus of variations at the same time.

This is a great book on the topic, regardless of whether you are new to probabilistic graphical models or have some familiarity with them but would like a deeper exploration of theory and/or implementation. I have read a number of books and papers on this topic (including Barber's and Bishop's) and I much prefer this one. Dr. Koller's style of writing is to start with simple theory and examples and walk the reader up to the full theory, while adding reminders of relevant topics covered elsewhere. She accomplishes this without condescending to or belittling the reader, or being overly verbose; each of the 1200 pages is concise and well edited. There is an OpenClassroom course that accompanies the book (CS 228), which I highly recommend viewing, as it contains that same style of teaching but in a different format and often with a somewhat different approach.

This popular book makes a noble attempt at unifying the many different types of probabilistic models used in artificial intelligence. It seems like a good reference manual for people who are already familiar with the fundamental concepts of commonly used probabilistic graphical models. However, it contains a lot of rambling and jumping between concepts that will quickly confuse a reader who is not already familiar with the subject. While the book appears to be systematic in introducing the subject with mathematical rigor (definitions and theorems), it actually skips a lot of fundamental concepts and leaves a lot of important proofs as exercises. I would recommend that a beginner in the subject start with another book like that by Jordan and Bishop, while keeping this book around as a reference manual or bank of practice problems for further study. The Coursera class on this

subject is much easier to follow than this book is.

[Download to continue reading...](#)

Probabilistic Graphical Models: Principles and Techniques (Adaptive Computation and Machine Learning series) Machine Learning: A Probabilistic Perspective (Adaptive Computation and Machine Learning series) Foundations of Machine Learning (Adaptive Computation and Machine Learning series) Introduction to Machine Learning (Adaptive Computation and Machine Learning series) Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning series) Bioinformatics: The Machine Learning Approach, Second Edition (Adaptive Computation and Machine Learning) Introduction to Statistical Relational Learning (Adaptive Computation and Machine Learning series) Reinforcement Learning: An Introduction (Adaptive Computation and Machine Learning series) Graphical Models: Foundations of Neural Computation (Computational Neuroscience) Boosting: Foundations and Algorithms (Adaptive Computation and Machine Learning series) Unsupervised Machine Learning in Python: Master Data Science and Machine Learning with Cluster Analysis, Gaussian Mixture Models, and Principal Components Analysis Deep Learning: Recurrent Neural Networks in Python: LSTM, GRU, and more RNN machine learning architectures in Python and Theano (Machine Learning in Python) Unsupervised Deep Learning in Python: Master Data Science and Machine Learning with Modern Neural Networks written in Python and Theano (Machine Learning in Python) Deep Learning in Python Prerequisites: Master Data Science and Machine Learning with Linear Regression and Logistic Regression in Python (Machine Learning in Python) Convolutional Neural Networks in Python: Master Data Science and Machine Learning with Modern Deep Learning in Python, Theano, and TensorFlow (Machine Learning in Python) Deep Learning in Python: Master Data Science and Machine Learning with Modern Neural Networks written in Python, Theano, and TensorFlow (Machine Learning in Python) Biological Sequence Analysis: Probabilistic Models of Proteins and Nucleic Acids IntAR, Interventions Adaptive Reuse, Volume 03; Adaptive Reuse in Emerging Economies Machine Learning with Spark - Tackle Big Data with Powerful Spark Machine Learning Algorithms Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference (Morgan Kaufmann Series in Representation and Reasoning)

[Dmca](#)