


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: 1270 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: 4.8 pounds (View shipping rates and policies)

Average Customer Review: 3.9 out of 5 stars 40 customer reviews

Best Sellers Rank: #79,862 in Books (See Top 100 in Books) #10 in Books > Computers &

Technology > Computer Science > AI & Machine Learning > Computer Vision & Pattern Recognition #14 in Books > Computers & Technology > Computer Science > AI & Machine Learning > Natural Language Processing #18 in Books > Computers & Technology > Computer Science > AI & Machine Learning > Machine Theory

Customer Reviews

This landmark book provides a very extensive coverage of the field, ranging from basic representational issues to the latest techniques for approximate inference and learning. As such, it is likely to become a definitive reference for all those who work in this area. Detailed worked examples and case studies also make the book accessible to students. (Kevin Murphy, Department of Computer Science, University of British Columbia)

Daphne Koller is Professor in the Department of Computer Science at Stanford University. Nir Friedman is Professor in the Department of Computer Science and Engineering at Hebrew University. Nir Friedman is Professor in the Department of Computer Science and Engineering at Hebrew University.

Judging by the first few chapters, the text is cumbersome and not as clear as it could have been under a more disciplined writing style; Sentences and paragraphs are longer than they should be, and the English grammar is most of the time improper or just a little odd. Reads too much like a transcript of a free speech lecture. Hopefully this alleviates later on in the book.

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.

If you want a very close look under the hood of Bayesian Networks, I can highly recommend Probabilistic Graphical Models. It's extremely comprehensive (1,200+ pages), well structured and clearly written. Theory, computation and application - including how to think about causation - are all covered in depth. Not light reading and not suited for those with limited stats background, but all in all one of the best textbooks on analytics topics I've ever read. Very impressive.

Very usefull book, and te best. companion for the course about.

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.

Very detailed and in depth material. Not recommended as an introduction to probabilistic models.

I just started reading this book for a course I want to do the coming semester. It seems like a very interesting book, exploring in depth probabilistic graphical models. Some topics are pretty tough to understand, and require additional background knowledge (such as statistics, algorithms, machine learning and AI). Another issue is that the book doesn't come with answers to the exercises, so you never know if you are on the right track when solving them.

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.

[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) Introduction to Machine Learning (Adaptive Computation and Machine Learning series) Reinforcement Learning with Python: An Introduction (Adaptive Computation and Machine Learning series) Deep Learning (Adaptive Computation and Machine Learning series) Semi-Supervised Learning (Adaptive Computation and Machine Learning series) Reinforcement Learning: An Introduction (Adaptive Computation and Machine Learning) Markov Models: Understanding Data Science, Markov Models, and Unsupervised Machine Learning in Python Human Computation (Synthesis Lectures on Artificial Intelligence and Machine Learning) IntAR,

Interventions Adaptive Reuse, Volume 03; Adaptive Reuse in Emerging Economies Machine Learning: For Beginners: Definitive Guide for Neural Networks, Algorithms, Random Forests and Decision Trees Made Simple (Machine Learning, Book 1) Machine Learning: An Algorithmic Perspective, Second Edition (Chapman & Hall/Crc Machine Learning & Pattern Recognition) Project Economics and Decision Analysis: Probabilistic Models Biological Sequence Analysis: Probabilistic Models of Proteins and Nucleic Acids Basics of R/C Model Aircraft Design: Practical Techniques for Building Better Models: Practical Techniques for Building Better Models BREAD MACHINE COOKBOOK: 120 Most Delicious Bread Machine Recipes (bread, bread bible, bread makers, breakfast, bread machine cookbook, bread baking, bread making, healthy, healthy recipes) Introduction to Deep Learning Using R: A Step-by-Step Guide to Learning and Implementing Deep Learning Models Using R Practical Machine Learning with H2O: Powerful, Scalable Techniques for Deep Learning and AI Complex Adaptive Systems: An Introduction to Computational Models of Social Life (Princeton Studies in Complexity) Probability and Computing: Randomization and Probabilistic Techniques in Algorithms and Data Analysis

[Contact Us](#)

[DMCA](#)

[Privacy](#)

[FAQ & Help](#)