




The book was found

Machine Learning: A Probabilistic Perspective (Adaptive Computation And Machine Learning Series)



Synopsis

Today's Web-enabled deluge of electronic data calls for automated methods of data analysis. Machine learning provides these, developing methods that can automatically detect patterns in data and then use the uncovered patterns to predict future data. This textbook offers a comprehensive and self-contained introduction to the field of machine learning, based on a unified, probabilistic approach. The coverage combines breadth and depth, offering necessary background material on such topics as probability, optimization, and linear algebra as well as discussion of recent developments in the field, including conditional random fields, L1 regularization, and deep learning. The book is written in an informal, accessible style, complete with pseudo-code for the most important algorithms. All topics are copiously illustrated with color images and worked examples drawn from such application domains as biology, text processing, computer vision, and robotics. Rather than providing a cookbook of different heuristic methods, the book stresses a principled model-based approach, often using the language of graphical models to specify models in a concise and intuitive way. Almost all the models described have been implemented in a MATLAB software package -- PMTK (probabilistic modeling toolkit) -- that is freely available online. The book is suitable for upper-level undergraduates with an introductory-level college math background and beginning graduate students.

Book Information

Series: Adaptive Computation and Machine Learning series

Hardcover: 1104 pages

Publisher: The MIT Press; 1 edition (August 24, 2012)

Language: English

ISBN-10: 0262018020

ISBN-13: 978-0262018029

Product Dimensions: 8 x 1.4 x 9 inches

Shipping Weight: 4.3 pounds (View shipping rates and policies)

Average Customer Review: 4.0 out of 5 stars 78 customer reviews

Best Sellers Rank: #14,129 in Books (See Top 100 in Books) #5 in [Books > Computers & Technology > Computer Science > AI & Machine Learning > Machine Theory](#) #9 in [Books > Textbooks > Computer Science > Artificial Intelligence](#) #15 in [Books > Computers & Technology > Computer Science > AI & Machine Learning > Intelligence & Semantics](#)

Customer Reviews

An astonishing machine learning book: intuitive, full of examples, fun to read but still comprehensive, strong and deep! A great starting point for any university student -- and a must have for anybody in the field. (Jan Peters, Darmstadt University of Technology; Max-Planck Institute for Intelligent Systems) Kevin Murphy excels at unraveling the complexities of machine learning methods while motivating the reader with a stream of illustrated examples and real world case studies. The accompanying software package includes source code for many of the figures, making it both easy and very tempting to dive in and explore these methods for yourself. A must-buy for anyone interested in machine learning or curious about how to extract useful knowledge from big data. (John Winn, Microsoft Research, Cambridge) This is a wonderful book that starts with basic topics in statistical modeling, culminating in the most advanced topics. It provides both the theoretical foundations of probabilistic machine learning as well as practical tools, in the form of Matlab code. The book should be on the shelf of any student interested in the topic, and any practitioner working in the field. (Yoram Singer, Google Inc.) This book will be an essential reference for practitioners of modern machine learning. It covers the basic concepts needed to understand the field as whole, and the powerful modern methods that build on those concepts. In Machine Learning, the language of probability and statistics reveals important connections between seemingly disparate algorithms and strategies. Thus, its readers will become articulate in a holistic view of the state-of-the-art and poised to build the next generation of machine learning algorithms. (David Blei, Princeton University) This comprehensive book should be of great interest to learners and practitioners in the field of machine learning. (British Computer Society)

Kevin P. Murphy is a Research Scientist at Google. Previously, he was Associate Professor of Computer Science and Statistics at the University of British Columbia.

(Disclaimer: I have worked with a draft of the book and been allowed to use the instructor's review copy for this review. I have bought the book from .co.uk, but apparently this .com review can't be tagged "verified purchase". I don't receive any compensation whatsoever for writing this review. I hope it will help you chose a machine learning textbook.) Similar textbooks on statistical/probabilistic machine learning (links to book websites, not pages):- Barber's Bayesian Reasoning and Machine Learning ("BRML", Cambridge University Press 2012)- Koller and Friedman's Probabilistic Graphical Models ("PGM", MIT Press 2009)- Bishop's Pattern Recognition and Machine Learning ("PRML", Springer 2006)- MacKay's Information Theory, Inference and Learning Algorithms ("ITILA", CUP 2003)- Hastie, Tibshirani and Friedman's Elements of Statistical Learning ("ESL",

Springer 2009)* Perspective: My perspective is that of a machine learning researcher and student, who has used these books for reference and study, but not as classroom textbooks.*

Audience/prerequisites: they are comparable among all the textbooks mentioned. BRML has lower expected commitment and specialization, PGM requires more scrupulous reading. The books differ in their topics and disciplinary approach, some more statistical (ESL), some more Bayesian (PRML, ITILA), some focused on graphical models (PGM, BRML). K Murphy compares MLAPP to others here. For detailed coverage comparison, read the table of contents on the book websites.* Main strength: MLAPP stands out for covering more advanced and current research topics: there is a full chapter on Latent Dirichlet Allocation, learning to rank, L1 regularization, deep networks; in the basics, the decision theory part is quite thorough (e.g. will mention Jeffrey's/uninformative priors). The book is "open" and vivid, doesn't shy away from current research and advanced concepts. This seems to be purposeful, as it shows in many aspects:- quotes liberally from web sources, something usually not done in academic publications- borrows "the best" from other authors (always with permission and acknowledgment, of course): most importantly the best pictures and diagrams, but also tables, recaps, insightful diagrams. Whereas other books will produce their own pictures and diagrams themselves (eg, PRML has a distinctive clarity and style in its illustrations), MLAPP takes many of its colour illustrations from other people's publications; therefore it can select the most pithy and relevant pictures to make a point. You could think that reproductions may be illegible and require extra effort to interpret because they come from a variety of sources; I have found that the bonus coming from having precisely the right image prevails.- frequent references to the literature, mentions of extensions and open questions, as well as computational complexity considerations: for instance, the section on HMMs will mention duration modeling and variable-duration Markov models, and a comparison of the expressive power of hierarchical HMMs versus stochastic context-free grammars, complete with relevant citations, and a brief mention of the computational complexity results from the publications. All this connects the material with research and new ideas in a fine way -- which other textbooks don't achieve, I find. For instance, PGM defers references to a literature section at the end of each chapter, resulting in a more self-contained, but more poorly "linked" text.* Didactic aids: Another distinctive feature is that the author clearly has tried to include didactic aids gathered over the years, such as recaps, comparative tables, diagrams, much in the spirit of the "generative model of generative models" (Roweis and Ghahramani): e.g. table comparing all models discussed, pros and cons of generative vs. discriminative models, recap of operations on HMMs (smoothing, filtering etc), list of parameter estimation methods for CRFs.* Editorial features: Other editorial features worth mentioning are- compared to others, helpful

mentions of terminology, e.g. jargon, nomenclature, concept names, in bold throughout the text ("you could also devise a variant thus; this is called so-and-so")- mathematical notation relatively clear and consistent, occasional obscurities. PGM stands out as excruciatingly precise on this aspect.- boxes/layout: no "skill boxes" or "case study boxes" (PGM), not many roadmap/difficulty indications like ITILA or PGM, examples are present but woven into the text (not separated like PGM or BRML). Layout rather plain and homogeneous, much like PRML.- sadly lacks list of figures and tables, but has index of code* Complete accompanying material:- interesting exercises (yet fewer than PRML, BRML, PGM); solutions, however, are only accessible to instructors (same with BRML, PGM), which in my experience makes them only half as useful for the self-learner. PRML and ITILA have some solutions online resp. in the book.- accompanying Matlab/Octave source code, which I found more readily usable than BRML's. PGM and PRML have no accompanying source code, even though the toolkit distributed with Koller's online PGM class might qualify as one. I find accompanying code a truly useful tool for learning; there's nothing like trying to implement an algorithm, checking your implementation against a reference, having boilerplate/utility code for the parts of the algorithm you're not interested in re-implementing. Also, code may clarify an algorithm, even when presented in pseudo-code. By the way, MLAPP has rather few pseudo-code boxes (like BRML or PRML, while PGM is very good here).- MLAPP is not freely available as a PDF (unlike BRML, closest topic-wise, ESL, or ITILA). This will no doubt reduce its diffusion. My own take on the underlying controversy is in favor of distributing the PDF: makes successful books widely popular and cited (think ITILA or Rasmussen and Williams' Gaussian Processes), increases the book's overall value, equips readers with a weightless copy to annotate with e-ink, or consult on the go. I believe PDF versions positively impact sales, too: impact neutral-to-positive to course textbook/university library sales, indifferent to sales in countries with widely different purchase power, positive to all other segments due to enormous diffusion/popularity.* Conclusion:The closest contender to this book I believe is BRML. Both are excellent textbooks and have accompanying source code.BRML is more accessible, has a free PDF version, and a stronger focus on graphical models.MLAPP has all the qualities of an excellent graduate textbook (unified presentation, valuable learning aids), and yet is unafraid of discussing detail points (e.g. omnipresent results on complexity), as well as advanced and research topics (LDA, L1 regularization).

I'm sure if you are a Google Research Scientist and are not learning the material for the first time, this book is amazing. For everyone else, I would not recommend it. I bought this book for my Fall 2013 COMPSCI 571 class, and I regret it. Before buying this book, consider the following:1. Take a

look at the online Errata. This book is already in its 3rd printing and it just came out. The list of corrections for this (the 3rd edition) is already mind-numbingly long. The 4th printing coming out this month will surely fix some errors, but there are just too many.² Our class has an online forum (for a 100 person class) where we discuss topics, and most questions are either (a) basic topics from the book that no one understood or (b) talking about how one figure in the book has multiple errors associated with it. At first I was really excited to find mistakes and submit them to the Errata - it was like I was part of the book! Now I just get frustrated and have already given up on submitting corrections.³ Our instructor regrets using this book and modifies the examples before giving them to us in class. Our out of class readings now consist mostly of MetaAcademy.com.⁴ There are hardly any worked-through examples, and many of those that are worked through have errors.⁵ Many important concepts are skimmed over way too quickly. For example, there is a whole chapter on Logistic regression. However, Logistic regression is covered for exactly 2 pages. Then a weird 3D graph is presented but not explained (a common theme throughout the book is graphs that look absolutely amazing, but which convey little information as to exactly what's going on to a lay-person like me), then the rest of the chapter presents methods for doing the math, which I'm sure are useful in some sense, but I'm still thinking: "why is this MLE not in closed form, what is a Hessian doing here... and wtf is going on?!" Most students just got the PDF for free online, and I would highly suggest doing something other than paying \$55 for this book.

This review is only related to the Kindle edition of the book. The Kindle edition is based on a PDF and is difficult to read on small devices. You have to zoom into the pages in order to be able to read the tiny text. Not a smooth reading experience. Furthermore the kindle edition is not compatible with the Kindle cloud reader. If you are ordering this book, just go with the hardcover edition. I love Kindle and usually get all my books as e-books but with this Kindle book I returned and ordered the hardcover edition.

I'm just wrapping up a machine learning course using this text, and I'm not a fan. The book is in dramatic need of an editor, and the index is bordering on useless. (Neither of the entire chapters on logistic or linear regression are findable via the index.) Frequently I've encountered a reference to an acronym that isn't defined in the current chapter, and isn't in the book's index. And it seems like every time I read a passage from this book I find some kind of spelling, capitalization, grammar or formatting error. I don't find the book to be very readable or good for initially learning material from, but it does seem to have a pretty complete appraisal of the techniques composing each individual

concept in ML, which is useful as a jumping off point for googling.

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

Machine Learning: A Probabilistic Perspective (Adaptive Computation and Machine Learning series)
Probabilistic Graphical Models: Principles and Techniques (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)
Draw in Perspective: Step by Step, Learn Easily How to Draw in Perspective (Drawing in Perspective, Perspective Drawing, How to Draw 3D, Drawing 3D, Learn to Draw 3D, Learn to Draw in Perspective)
Machine Learning: An Algorithmic Perspective, Second Edition (Chapman & Hall/Crc Machine Learning & Pattern Recognition)
IntAR, Interventions Adaptive Reuse, Volume 03; Adaptive Reuse in Emerging Economies
Human Computation (Synthesis Lectures on Artificial Intelligence and Machine Learning)
Machine Learning: For Beginners: Definitive Guide for Neural Networks, Algorithms, Random Forests and Decision Trees Made Simple (Machine Learning, Book 1)
Probability on Trees and Networks (Cambridge Series in Statistical and Probabilistic Mathematics)
Random Graphs and Complex Networks: Volume 1 (Cambridge Series in Statistical and Probabilistic Mathematics)
Data Analysis and Graphics Using R: An Example-Based Approach (Cambridge Series in Statistical and Probabilistic Mathematics)
Probabilistic Methods of Signal and System Analysis (The Oxford Series in Electrical and Computer Engineering)
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)
Probability and Computing: Randomization and Probabilistic Techniques in Algorithms and Data Analysis
Machine Learning: A Bayesian and Optimization Perspective (Net Developers)
Number Talks: Whole Number Computation, Grades K-5: A Multimedia Professional Learning Resource

[Contact Us](#)

[DMCA](#)

[Privacy](#)

[FAQ & Help](#)