


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: 4.6 pounds (View shipping rates and policies)

Average Customer Review: 4.0 out of 5 stars [See all reviews](#) (61 customer reviews)

Best Sellers Rank: #22,103 in Books (See Top 100 in Books) #4 in [Books > Computers & Technology > Computer Science > AI & Machine Learning > Machine Theory](#) #10 in [Books > Textbooks > Computer Science > Artificial Intelligence](#) #22 in [Books > Computers & Technology > Computer Science > AI & Machine Learning > Intelligence & Semantics](#)

Customer Reviews

(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 choose 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.

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.

[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)
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)
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)
Introduction to Statistical Relational Learning (Adaptive Computation and Machine Learning series)
Reinforcement Learning: An Introduction (Adaptive Computation and Machine Learning series)
Boosting: Foundations and Algorithms (Adaptive Computation and Machine Learning series)
IntAR, Interventions Adaptive Reuse, Volume 03; Adaptive Reuse in Emerging Economies
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)
Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference (Morgan Kaufmann Series in Representation and Reasoning)
Numerical Methods of Statistics (Cambridge Series in Statistical and Probabilistic Mathematics)
Probabilistic Reasoning in Expert Systems: Theory and Algorithms
Biological Sequence Analysis: Probabilistic Models of Proteins and Nucleic Acids

[Dmca](#)