

The book was found

Machine Learning



Synopsis

This book covers the field of machine learning, which is the study of algorithms that allow computer programs to automatically improve through experience. The book is intended to support upper level undergraduate and introductory level graduate courses in machine learning.

Book Information

Series: McGraw-Hill Series in Computer Science

Hardcover: 432 pages

Publisher: McGraw-Hill Education; 1 edition (March 1, 1997)

Language: English

ISBN-10: 0070428077

ISBN-13: 978-0070428072

Product Dimensions: 6.3 x 1 x 9.6 inches

Shipping Weight: 2 pounds (View shipping rates and policies)

Average Customer Review: 4.2 out of 5 stars [See all reviews](#) (59 customer reviews)

Best Sellers Rank: #131,508 in Books (See Top 100 in Books) #20 in [Books > Computers & Technology > Computer Science > AI & Machine Learning > Machine Theory](#) #23 in [Books > Computers & Technology > Computer Science > AI & Machine Learning > Neural Networks](#) #31 in [Books > Textbooks > Computer Science > Algorithms](#)

Customer Reviews

I agree with some of the previous reviews which criticize the book for its lack of depth, but I believe this to be an asset rather than a liability given its target audience (seniors and beginning grad. students). The average college senior typically knows very little about subjects like neural networks, genetic algorithms, or Bayesian networks, and this book goes a long way in demystifying these subjects in a very clear, concise, and understandable way. Moreover, the first-year grad. student who is interested in possibly doing research in this field needs more of an overview than to dive deeply into one of the many branches which themselves have had entire books written about them. This is one of the few if only books where one will find diverse areas of learning (e.g. analytical, reinforcement, Bayesian, neural-network, genetic-algorithmic) all within the same cover. But more than just an encyclopedic introduction, the author makes a number of connections between the different paradigms. For example, he explains that associated with each paradigm is the notion of an inductive-learning bias, i.e. the underlying assumptions that lend validity to a given learning approach. These end-of-chapter discussions on bias seem very interesting and unique to this

book. Finally, I used this book for part of the reading material for an intro. AI class, and received much positive feedback from the students, although some did find the presentation a bit too abstract for their undergraduate tastes

It's pretty well done, it covers theory and core areas but - maybe it was more the state of the field when it was written - I found it unsatisfyingly un-synthesized, unconnected, and short of detail (but this is subjective). I found the 2nd edition of Russell and Norvig to be a better introduction where it covers the same topic, which it does for everything I can think of, except VC dimension. The book sorely needs an update, it was written in 1997 and the field has moved fast. A comparison with Mitchell's current course (materials generously available online) shows that about 1/4 of the topics taught have arisen since the book was published; Boosting, Support Vector Machines and Hidden Markov Models to name the best-known. The book also does not cover statistical or data mining methods. Despite the subjective complaint about lack of depth it does give the theoretical roots and many fundamental techniques decently and readably. For many purposes though it may have been superceded by R&N 2nd ed.

I first used this book as the required text for my course in ML in 1997 and got rave reviews from the students. I will be using it again in 1999. I found ALL of the major topics and issues in ML addressed. The book is easily readable with anyone with a computer science background, and the book works quite well in a wide variety of approaches to presentation at the advanced undergraduate and graduate levels.

I am a graduate student at a major research university. I am currently taking my fifth AI/Machine Learning graduate course. This is the one book everyone grabs for when they need a reference. I had to mark the spine of my book with tape so I could find it more easily on my colleagues shelves. Other books are either not as accessible or too niche-specific. This is the only book out there that covers all of the major machine learning techniques (with the possible exception of support vector machines) and covers them in a manner that can be well understood. Every discipline has one book that must be on your shelf. If you are planning on doing serious research in Machine Learning - this is the one book.

I teach AI at the graduate level in a major US research University, and I specialize in the area. The book does cover many different areas of Machine Learning. Unfortunately, the treatment is quite

superficial. A student would find it extremely difficult to grasp important concepts without referring to other material. It may be a good reference, but I would definitely not recommend it as the main textbook. Unfortunately, there seem to be very few books in this area adequate for a senior or graduate level course.

This is extremely intuitive and general point of view on ML. good for quick reading and getting introduced to the topic. I'd recommend this to people starting ML. then move on to more mathematically rigorous and specific books such as "Pattern Classification" / "Pattern Recognition and Machine Learning" / Hastie's "Element of Statistical Learning" i never say this for a book. but it is too pricey for what it is offering. FYI i think they should increase the price of Chris Bishop's book.

This is a great book if you're starting out with machine learning. It's rare to come across a book like this that is very well written and has technical depth. The writing is to the point, maybe even a bit terse, but all that you need to know is in there. It's a bit old so doesn't cover kernel methods or SVM's but is still a great first machine learning book.

I bought this book because it was part of the required reading material for a class I am taking at Georgia Tech, Machine Learning. It is certainly not a beginners level book and requires a lot of math and statistics to make sense of it. I also was shocked at the price for this book considering how old it is and would have thought by now there would have been a revision but there hasn't been one. I may be critical of this book because I ended up dropping the class due to the fact I felt I lacked the background required to succeed and this book didn't help fill the gaps.

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

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) Machine Learning: A Probabilistic Perspective (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 Machine Learning with Spark - Tackle Big Data with Powerful Spark Machine Learning Algorithms 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) First-Time Machine Applique: Learning to Machine Applique in Nine Easy Lessons A collection of Advanced Data Science and Machine Learning Interview Questions Solved in Python and Spark (II): Hands-on Big Data and Machine ... Programming Interview Questions (Volume 7) Introduction to Statistical Relational Learning (Adaptive Computation and Machine Learning series) Reinforcement Learning: An Introduction (Adaptive Computation and Machine Learning series) Learning Deep Architectures for AI (Foundations and Trends(r) in Machine Learning) The Bread Lover's Bread Machine Cookbook: A Master Baker's 300 Favorite Recipes for Perfect-Every-Time Bread-From Every Kind of Machine More Bread Machine Magic : More Than 140 New Recipes From the Authors of Bread Machine Magic for Use in All Types of Sizes of Bread Machines Bread Machine Cookbook: Delicious And Simple Bread Machine Recipes

[Dmca](#)