


The book was found

Foundations Of Machine Learning (Adaptive Computation And Machine Learning Series)



Synopsis

This graduate-level textbook introduces fundamental concepts and methods in machine learning. It describes several important modern algorithms, provides the theoretical underpinnings of these algorithms, and illustrates key aspects for their application. The authors aim to present novel theoretical tools and concepts while giving concise proofs even for relatively advanced topics. Foundations of Machine Learning fills the need for a general textbook that also offers theoretical details and an emphasis on proofs. Certain topics that are often treated with insufficient attention are discussed in more detail here; for example, entire chapters are devoted to regression, multi-class classification, and ranking. The first three chapters lay the theoretical foundation for what follows, but each remaining chapter is mostly self-contained. The appendix offers a concise probability review, a short introduction to convex optimization, tools for concentration bounds, and several basic properties of matrices and norms used in the book. The book is intended for graduate students and researchers in machine learning, statistics, and related areas; it can be used either as a textbook or as a reference text for a research seminar.

Book Information

Series: Adaptive Computation and Machine Learning series

Hardcover: 432 pages

Publisher: The MIT Press (August 17, 2012)

Language: English

ISBN-10: 026201825X

ISBN-13: 978-0262018258

Product Dimensions: 7 x 0.8 x 9 inches

Shipping Weight: 2.4 pounds (View shipping rates and policies)

Average Customer Review: 4.1 out of 5 starsÂ See all reviewsÂ (9 customer reviews)

Best Sellers Rank: #224,997 in Books (See Top 100 in Books) #37 inÂ Books > Computers & Technology > Computer Science > AI & Machine Learning > Machine Theory #176 inÂ Books > Computers & Technology > Computer Science > AI & Machine Learning > Intelligence & Semantics #58037 inÂ Books > Reference

Customer Reviews

I picked up this book soon after it came out and found it a wonderful read. Consistent with being a new release, it's more modern than the previous classic ML textbook by Bishop and treats newer subjects that got short shrift there, including PAC learning, VC dimension and Rademacher

complexity. It's very well written and does a great job of covering the material that a new student needs to absorb in order to keep up with the current literature in ML. Highly recommended.

Some textbooks such as those of Chris Bishop and Kevin Murphy present machine learning from the Bayesian perspective, which is a particular point of view. In contrast, this book gives an unbiased presentation of machine learning with solid theoretical justifications. It discusses the principles behind the design of learning algorithms by introducing and using the most modern tools and concepts in learning theory. This helps answering many fundamental questions. The presentation is concise and the topics covered very broad. They include the presentation of several of the most well known binary classification algorithms, multi-class classification, regression, ranking, on-line learning, reinforcement learning, structured prediction, learning theory, and many other topics. In particular, there is a nice and concise presentation of SVMs and boosting. The appendix introduces all the main tools needed, including a brief introduction to convex optimization. I strongly recommend this book to students and researchers. It gives a very modern presentation covering all the main topics in learning, which can serve as a reference for everyone. Perhaps more importantly, it helps us analyze and understand machine learning.

The best book on machine learning theory. This book is extremely clear and is a must-have for any serious machine learning or statistical learning scholar. As the title suggests, this book builds the foundations of machine learning, which are omitted in every other machine learning text book that I've read. This book will prepare you for advanced, research level machine learning papers. There is no other book like it - absolutely incredible! This is the book that experts and professors in the field learn from. Even if you have 10+ years of experience in the field, I'm sure that you will learn something new every time you pick up the book. Furthermore, the book is concise enough that even an beginner could learn from it. Although any beginner should be prepared to read more on their own. A basic understanding of probability theory, linear algebra, and optimization is assumed - although the appendix has the clearest survey of linear algebra, basic probability, and basic optimization that I've ever read. Seriously - this book is incredible.

I wish I could give 0 stars. This "kindle book" is completely unreadable. Sadly, the authors decided they could make a PDF version of the book, charge \$40 and still call it a Kindle Book. Kindle books are legible on the mobile kindle apps. This book is not. shouldn't let them sell it as I just wasted \$40 on something I can't even use. Now I must buy the paper version...

Excellent book. Used for my second year undergraduate learning theory course. Very well written. Recommend this for all CS undergraduates who are interested in learning theory.

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

Foundations of Machine Learning (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) Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning series) Boosting: Foundations and Algorithms (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) Probabilistic Graphical Models: Principles and Techniques (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) The Simple Genetic Algorithm: Foundations and Theory (Complex Adaptive Systems) Graphical Models: Foundations of Neural Computation (Computational Neuroscience) Learning Deep Architectures for AI (Foundations and Trends(r) in Machine Learning) 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

[Dmca](#)