


The book was found

Understanding Machine Learning: From Theory To Algorithms



Synopsis

Machine learning is one of the fastest growing areas of computer science, with far-reaching applications. The aim of this textbook is to introduce machine learning, and the algorithmic paradigms it offers, in a principled way. The book provides an extensive theoretical account of the fundamental ideas underlying machine learning and the mathematical derivations that transform these principles into practical algorithms. Following a presentation of the basics of the field, the book covers a wide array of central topics that have not been addressed by previous textbooks. These include a discussion of the computational complexity of learning and the concepts of convexity and stability; important algorithmic paradigms including stochastic gradient descent, neural networks, and structured output learning; and emerging theoretical concepts such as the PAC-Bayes approach and compression-based bounds. Designed for an advanced undergraduate or beginning graduate course, the text makes the fundamentals and algorithms of machine learning accessible to students and non-expert readers in statistics, computer science, mathematics, and engineering.

Book Information

Hardcover: 410 pages

Publisher: Cambridge University Press; 1 edition (May 19, 2014)

Language: English

ISBN-10: 1107057132

ISBN-13: 978-1107057135

Product Dimensions: 7 x 1.1 x 10 inches

Shipping Weight: 1.8 pounds (View shipping rates and policies)

Average Customer Review: 4.3 out of 5 starsÂ See all reviewsÂ (11 customer reviews)

Best Sellers Rank: #98,552 in Books (See Top 100 in Books) #16 inÂ Books > Computers & Technology > Computer Science > AI & Machine Learning > Computer Vision & Pattern Recognition #1323 inÂ Books > Textbooks > Computer Science #27312 inÂ Books > Reference

Customer Reviews

This book provides a great story line along with solid proofs of machine learning theories and algorithms. Each chapter is rather short (15-20 pages), yet is well written to convey the topic in detail, making the book comfortable to read. Moreover, the connection among consecutive chapters is strong, giving an excellent coarse-to-fine introduction on sophisticated theories. Over the past few years, I have read several machine learning books, and this is the one solidly based on "statistical learning theory". Compared to other books that give only brief description to this aspect, this book

does a good job not only on providing the basic proofs, but also on extending the theories to well-known practical algorithms, supporting the success of these algorithms and showing how theories can be used to design or analyze practical algorithms. For whom eager to know more about learning theory, this is a must-read book.

First, let me just say I regret purchasing the kindle version, as it is difficult to read the math symbols on the kindle, and even somewhat difficult to read them on the kindle for mac app on a big screen. Zoomed in leaves the symbols the same size (it appears as though they're images), with the surrounding text large. Perhaps this is a problem on most math texts, but I was disappointed. I'm enjoying the book. It reads like a textbook that one might find at a university, and has exercises and notes for the order you'd go through it while teaching a class. I find it well-written and for the most part, easy to digest--a bit heavy on the math for what I was looking for, but you can skim over it for the ideas.

This is an excellent introduction to the theory of Machine Learning (ML). I would like though to stress the word "theory". This is probably not the first introductory book in ML (the readers with strong mathematical background can disregard this reservation, they indeed can use this as !), for the beginners who want to learn the basic concepts of the ML and to understand the motivation behind the mathematical concepts I would recommend something like "Learning from Data" by Yaser Abu-Mostafa et al.

http://www..com/Learning-Data-Yaser-S-Abu-Mostafa/dp/1600490069/ref=sr_1_1?s=books&ie=UTF8&qid=1416599862&sr=1-1&keywords=learning+from+data complemented by the e-chapters in the online forum <http://book.caltech.edu/bookforum/> and, probably, by online course (see <http://edx.org>). But for those who have already got some basic ideas about the concepts of ML and the motivation for the theoretical justification of the algorithms, this is definitely should be the next book to read: it provides the rigorous proofs and presents the concepts and algorithms in clear mathematical language. There is no need to be scared though: the presentation of the stuff is excellent, the chapters are short enough in order to enable the reader to advance in reasonable steps (the book is derived from the lectures presented by both of the authors), there are excellent exercises. The theory is indeed well connected to practical algorithms and real applications as promised by the subtitle :)

This is an excellent introduction to machine learning which fills an important gap in the literature by

introducing students to formal broad conceptual frameworks for understanding, comparing, analyzing, and designing large classes of popular machine learning algorithms. These frameworks are explicitly presented as mathematical theorems but the authors are careful about explaining the underlying assumptions of key theorems and interpreting the conclusions of such theorems. Richard M. Golden.

Very well written! This book emphasizes the essential principles of machine learning, useful both for reading theoretical papers in machine learning and putting the algorithms to work on real data.

This is a very well written book. The chapters are short (bite-sized), but are very lucid. Almost all proofs are given in detail. I am enjoying reading it a lot.

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

Machine Learning with Spark - Tackle Big Data with Powerful Spark Machine Learning Algorithms
Understanding Machine Learning: From Theory to Algorithms 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 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) Boosting: Foundations and Algorithms (Adaptive Computation and Machine Learning series) Java Artificial Intelligence: Made Easy, w/ Java Programming; Learn to Create your * Problem Solving * Algorithms! TODAY! w/ Machine Learning & Data ... engineering, r programming, iOS development) Javascript Artificial

Intelligence: Made Easy, w/ Essential Programming; Create your * Problem Solving * Algorithms! TODAY! w/ Machine Learning & Data ... engineering, r programming, iOS development) Artificial Intelligence: Made Easy w/ Ruby Programming; Learn to Create your * Problem Solving * Algorithms! TODAY! w/ Machine Learning & Data ... engineering, r programming, iOS development) Fundamentals of Machine Learning for Predictive Data Analytics: Algorithms, Worked Examples, and Case Studies (MIT Press) Java: Artificial Intelligence; Made Easy, w/ Java Programming; Learn to Create your * Problem Solving * Algorithms! TODAY! w/ Machine Learning & Data Structures (Artificial Intelligence Series) Machine Learning: The Art and Science of Algorithms that Make Sense of Data

[Dmca](#)