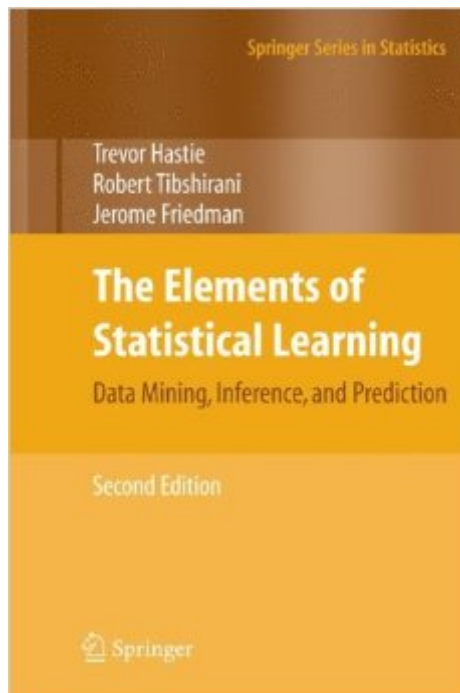


The book was found

The Elements Of Statistical Learning: Data Mining, Inference, And Prediction, Second Edition (Springer Series In Statistics)



Synopsis

During the past decade there has been an explosion in computation and information technology. With it have come vast amounts of data in a variety of fields such as medicine, biology, finance, and marketing. The challenge of understanding these data has led to the development of new tools in the field of statistics, and spawned new areas such as data mining, machine learning, and bioinformatics. Many of these tools have common underpinnings but are often expressed with different terminology. This book describes the important ideas in these areas in a common conceptual framework. While the approach is statistical, the emphasis is on concepts rather than mathematics. Many examples are given, with a liberal use of color graphics. It is a valuable resource for statisticians and anyone interested in data mining in science or industry. The book's coverage is broad, from supervised learning (prediction) to unsupervised learning. The many topics include neural networks, support vector machines, classification trees and boosting---the first comprehensive treatment of this topic in any book. This major new edition features many topics not covered in the original, including graphical models, random forests, ensemble methods, least angle regression & path algorithms for the lasso, non-negative matrix factorization, and spectral clustering. There is also a chapter on methods for "wide" data (p bigger than n), including multiple testing and false discovery rates.

Book Information

Series: Springer Series in Statistics

Hardcover: 745 pages

Publisher: Springer; 2nd ed. 2009. Corr. 7th printing 2013 edition (April 12, 2011)

Language: English

ISBN-10: 0387848576

ISBN-13: 978-0387848570

Product Dimensions: 1.2 x 6.2 x 9.2 inches

Shipping Weight: 3 pounds (View shipping rates and policies)

Average Customer Review: 4.1 out of 5 stars See all reviews (75 customer reviews)

Best Sellers Rank: #3,481 in Books (See Top 100 in Books) #1 in Books > Computers & Technology > Computer Science > Bioinformatics #1 in Books > Textbooks > Computer Science > Artificial Intelligence #1 in Books > Computers & Technology > Computer Science > AI & Machine Learning > Intelligence & Semantics

Customer Reviews

This review is written from the perspective of a programmer who has sometimes had the chance to choose, hire, and work with algorithms and the mathematician/statisticians that love them in order to get things done for startup companies. I don't know if this review will be as helpful to professional mathematicians, statisticians, or computer scientists. The good news is, this is pretty much the most important book you are going to read in the space. It will tie everything together for you in a way that I haven't seen any other book attempt. The bad news is you're going to have to work for it. If you just need to use a tool for a single task this book won't be worth it; think of it as a way to train yourself in the fundamentals of the space, but don't expect a recipe book. Get something in the "using R" series for that. When it came out in 2001 my sense of machine learning was of a jumbled set of recipes that tended to work in some cases. This book showed me how the statistical concepts of bias, variance, smoothing and complexity cut across both fields of traditional statistics and inference and the machine learning algorithms made possible by cheaper cpus. Chapters 2-5 are worth the price of the book by themselves for their overview of learning, linear methods, and how those methods can be adopted for non-linear basis functions. The hard parts: First, don't bother reading this book if you aren't willing to learn at least the basics of linear algebra first. Skim the second and third chapters to get a sense for how rusty your linear algebra is and then come back when you're ready. Second, you really really want to use the SQRRR technique with this book. Having that glimpse of where you are going really helps guide your understanding when you dig in for real. Third, I wish I had known of R when I first read this; I recommend using it along with some sample data sets to follow along with the text so the concepts become skills not just abstract relationships to forget. It would probably be worth the extra time, and I wish I had known to do that then. Fourth, if you are reading this on your own time while making a living, don't expect to finish the book in a month or two.

I have been using The Elements of Statistical Learning for years, so it is finally time to try and review it. The Elements of Statistical Learning is a comprehensive mathematical treatment of machine learning from a statistical perspective. This means you get good derivations of popular methods such as support vector machines, random forests, and graphical models; but each is developed only after the appropriate (and wrongly considered less sexy) statistical framework has already been derived (linear models, kernel smoothing, ensembles, and so on). In addition to having excellent and correct mathematical derivations of important algorithms The Elements of Statistical Learning is fairly unique in that it actually uses the math to accomplish big things. My favorite examples come from Chapter 3 "Linear Methods for Regression." The standard treatments of these

methods depend heavily on respectful memorization of regurgitation of original iterative procedure definitions of the various regression methods. In such a standard formulation two regression methods are different if they have superficially different steps or if different citation/priority histories. The Elements of Statistical Learning instead derives the stopping conditions of each method and considers methods the same if they generate the same solution (regardless of how they claim they do it) and compares consequences and results of different methods. This hard use of isomorphism allows amazing results such as Figure 3.15 (which shows how Least Angle Regression differs from Lasso regression, not just in algorithm description or history: but by picking different models from the same data) and section 3.5.2 (which can separate Partial Least Squares' design CLAIM of fixing the x-dominance found in principle components analysis from how effective it actually is as fixing such problems). The biggest issue is who is the book for? This is a mathy book emphasizing deep understanding over mere implementation. Unlike some lesser machine learning books the math is not there for appearances or mere intimidating typesetting: it is there to allow the authors to organize many methods into a smaller number of consistent themes. So I would say the book is for researchers and machine algorithm developers. If you have a specific issue that is making inference difficult you may find the solution in this book. This is good for researchers but probably off-putting for tinkers (as this book likely has methods superior to their current favorite new idea). The interested student will also benefit from this book, the derivations are done well so you learn a lot by working through them. Finally- don't buy the kindle version, but the print book. This book is satisfying deep reading and you will want the advantages of the printed page (and 's issues in conversion are certainly not the authors' fault).

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

The Elements of Statistical Learning: Data Mining, Inference, and Prediction, Second Edition (Springer Series in Statistics) Analytics: Data Science, Data Analysis and Predictive Analytics for Business (Algorithms, Business Intelligence, Statistical Analysis, Decision Analysis, Business Analytics, Data Mining, Big Data) Statistical Methods for Dynamic Treatment Regimes: Reinforcement Learning, Causal Inference, and Personalized Medicine (Statistics for Biology and Health) Data Analytics: Practical Data Analysis and Statistical Guide to Transform and Evolve Any Business. Leveraging the Power of Data Analytics, Data ... (Hacking Freedom and Data Driven) (Volume 2) RapidMiner: Data Mining Use Cases and Business Analytics Applications (Chapman & Hall/CRC Data Mining and Knowledge Discovery Series) An Introduction to Statistical Learning: with Applications in R (Springer Texts in Statistics) Statistics and Data Analysis for Financial Engineering: with R examples (Springer Texts in Statistics) A First Course in Bayesian Statistical

Methods (Springer Texts in Statistics) Data Analytics: What Every Business Must Know About Big Data And Data Science (Data Analytics for Business, Predictive Analysis, Big Data) Web Data Mining: Exploring Hyperlinks, Contents, and Usage Data (Data-Centric Systems and Applications) Data Science for Business: What You Need to Know about Data Mining and Data-Analytic Thinking Mining the Social Web: Data Mining Facebook, Twitter, LinkedIn, Google+, GitHub, and More Healthcare Data Analytics (Chapman & Hall/CRC Data Mining and Knowledge Discovery Series) Statistical Learning with Sparsity: The Lasso and Generalizations (Chapman & Hall/CRC Monographs on Statistics & Applied Probability) Data Analysis and Data Mining using Microsoft Business Intelligence Tools: Excel 2010, Access 2010, and Report Builder 3.0 with SQL Server Exploratory Data Mining and Data Cleaning Time Series: Theory and Methods (Springer Series in Statistics) Yellowcake Towns - Uranium Mining Communities in the American West (Mining the American West) Bitcoin Mining: The Bitcoin Beginner's Guide (Proven, Step-By-Step Guide To Making Money With Bitcoins) (Bitcoin Mining, Online Business, Investing for ... Beginner, Bitcoin Guide, Bitcoin Trading) Information Theory, Inference and Learning Algorithms

[Dmca](#)