


The book was found

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



Synopsis

This book describes the important ideas in a variety of fields such as medicine, biology, finance, and marketing 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 colour 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 factorisation, 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 edition (2016)

Language: English

ISBN-10: 0387848576

ISBN-13: 978-0387848570

Product Dimensions: 1.2 x 6.2 x 9.2 inches

Shipping Weight: 3.3 pounds (View shipping rates and policies)

Average Customer Review: 4.1 out of 5 stars 132 customer reviews

Best Sellers Rank: #3,544 in Books (See Top 100 in Books) #1 in Books > Computers & Technology > Computer Science > Bioinformatics #3 in Books > Textbooks > Computer Science > Artificial Intelligence #3 in Books > Textbooks > Computer Science > Database Storage & Design

Customer Reviews

From the reviews: "Like the first edition, the current one is a welcome edition to researchers and academicians equally. Almost all of the chapters are revised. The Material is nicely reorganized and repackaged, with the general layout being the same as that of the first edition. If you bought the first edition, I suggest that you buy the second editon for maximum effect, and if you haven't, then I still strongly recommend you have this book at

your desk. Is it a good investment, statistically speaking!" (Book Review Editor, *Technometrics*, August 2009, VOL. 51, NO. 3) From the reviews of the second edition: "This second edition pays tribute to the many developments in recent years in this field, and new material was added to several existing chapters as well as four new chapters. These additions make this book worthwhile to obtain. In general this is a well written book which gives a good overview on statistical learning and can be recommended to everyone interested in this field. The book is so comprehensive that it offers material for several courses." (Klaus Nordhausen, *International Statistical Review*, Vol. 77 (3), 2009) "The second edition features about 200 pages of substantial new additions in the form of four new chapters, as well as various complements to existing chapters. The book may also be of interest to a theoretically inclined reader looking for an entry point to the area and wanting to get an initial understanding of which mathematical issues are relevant in relation to practice. This is a welcome update to an already fine book, which will surely reinforce its status as a reference." (Gilles Blanchard, *Mathematical Reviews*, Issue 2012 d) "The book would be ideal for statistics graduate students. This book really is the standard in the field, referenced in most papers and books on the subject, and it is easy to see why. The book is very well written, with informative graphics on almost every other page. It looks great and inviting. You can flip the book open to any page, read a sentence or two and be hooked for the next hour or so." (Peter Rabinovitch, *The Mathematical Association of America*, May, 2012)

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 and 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. Trevor Hastie, Robert Tibshirani, and Jerome Friedman are professors of statistics at Stanford University. They are prominent researchers in this area: Hastie and Tibshirani developed generalized additive models and wrote a popular book of that title. Hastie co-developed much of the statistical modeling software and environment in R/S-PLUS and invented principal curves and surfaces. Tibshirani proposed the lasso and is co-author of the very successful *An Introduction to the Bootstrap*. Friedman is the co-inventor of many data-mining tools including CART, MARS, projection pursuit and gradient boosting.

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 or 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).

A bunch of pages were missing and replaced by a bunch of repeated pages. The problem is the editor (Springer). The contents are superb, and it is an excellent reference book (unless what you are looking for ends up in some of the missing pages, as it happened to me last week).

Anyone working in data mining and predictive analytics or in applied statistics in other areas should have this on their shelf. It is well written and very comprehensive. Not light reading but not a pure math stat book either. It helped bring me up to date on many of the recent developments in machine learning and applied statistics. An excellent reference.

I would recommend this book to those who need to use machine learning. It is great as a reference book. It's compact style means that it will most benefit those with a background of linear algebra (matrices) and some calculus. It may not be the best book to start learning these techniques from scratch.

Here are my impressions about the book.* The organization of book is clear and concise. No too lengthy and fuzzy discussions. Method and technique explanations are clear though sometimes bit hard to follow.* Liberal usage of colors. It helps to make drawings more attractive and to show more info. May be too bright for some.* My parts like SVM and LDA, PCA.... they seem to be explained in quite difficult way. Though it is clear enough and self-sufficient* There are many references and pointers to original works and related fields. Big plus here.* Work and results on real sample data. I particularly enjoy it! It helps to better understand each method and its +/-This is a good book for those who use statistics for data mining. It seems that there is more accent on probabilistic (frequentist) field. Though a reader is expected to have good enough level in maths to grasp the concepts quickly. May not be appropriate for the beginners in the field.

great product thanks

This is THE book to read/study-from/have-on-your-shelf-as-a-reference if you are serious about data science

This book is beautiful. It is definitely not an introductory book, but the effort is worth it. It should probably be read along with an open Wikipedia tab, for all the non-statistics wizards out there. In summary, this book takes effort to read, but it is a priceless jewel of knowledge in its field.

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

The Elements of Statistical Learning: Data Mining, Inference, and Prediction, Second Edition (Springer Series in Statistics) All of Statistics: A Concise Course in Statistical Inference (Springer Texts 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) Big Data For Business: Your Comprehensive Guide to Understand Data Science, Data Analytics and Data Mining to Boost More Growth and Improve Business - Data Analytics Book, Series 2 Discovering Knowledge in Data: An Introduction to Data Mining (Wiley Series on Methods and Applications in Data Mining) Analytics: Business Intelligence, Algorithms and Statistical Analysis (Predictive Analytics, Data Visualization, Data Analytics, Business Analytics, Decision Analysis, Big Data, Statistical Analysis) Statistics, Data Mining, and Machine Learning in Astronomy: A Practical Python Guide for the Analysis of Survey Data (Princeton Series in Modern Observational Astronomy) Data Mining: Practical Machine Learning Tools and Techniques, Second Edition (Morgan Kaufmann Series in Data Management Systems) Statistical Inference Based on Divergence Measures (Statistics: A Series of Textbooks and Monographs) Data Analytics: What Every Business Must Know About Big Data And Data Science (Data Analytics for Business, Predictive Analysis, Big Data Book 1) Data Analytics: Applicable Data Analysis to Advance Any Business Using the Power of Data Driven Analytics (Big Data Analytics, Data Science, Business Intelligence Book 6) Data Mining, Fourth Edition: Practical Machine Learning Tools and Techniques (Morgan Kaufmann Series in Data Management Systems) Data Mining: Practical Machine Learning Tools and Techniques, Third Edition (Morgan Kaufmann Series in Data Management Systems) Statistics and Data Analysis for Financial Engineering: with R examples (Springer Texts in Statistics) Data Science for Business: What You Need to Know about Data Mining and Data-Analytic Thinking An Introduction to Statistical Learning: with Applications in R (Springer Texts in Statistics) Data Mining: Practical Machine Learning Tools and Techniques (Morgan Kaufmann

Series in Data Management Systems) Statistics for People Who (Think They) Hate Statistics
(Salkind, Statistics for People Who(Think They Hate Statistics(Without CD)) An Introduction to
Probability and Statistical Inference, Second Edition Data Mining: Concepts and Techniques, Third
Edition (The Morgan Kaufmann Series in Data Management Systems)

[Contact Us](#)

[DMCA](#)

[Privacy](#)

[FAQ & Help](#)