


The book was found

Machine Learning: An Algorithmic Perspective, Second Edition (Chapman & Hall/Crc Machine Learning & Pattern Recognition)



Synopsis

A Proven, Hands-On Approach for Students without a Strong Statistical Foundation Since the best-selling first edition was published, there have been several prominent developments in the field of machine learning, including the increasing work on the statistical interpretations of machine learning algorithms. Unfortunately, computer science students without a strong statistical background often find it hard to get started in this area. Remedy this deficiency, Machine Learning: An Algorithmic Perspective, Second Edition helps students understand the algorithms of machine learning. It puts them on a path toward mastering the relevant mathematics and statistics as well as the necessary programming and experimentation. New to the Second Edition Two new chapters on deep belief networks and Gaussian processes Reorganization of the chapters to make a more natural flow of content Revision of the support vector machine material, including a simple implementation for experiments New material on random forests, the perceptron convergence theorem, accuracy methods, and conjugate gradient optimization for the multi-layer perceptron Additional discussions of the Kalman and particle filters Improved code, including better use of naming conventions in Python Suitable for both an introductory one-semester course and more advanced courses, the text strongly encourages students to practice with the code. Each chapter includes detailed examples along with further reading and problems. All of the code used to create the examples is available on the author's website.

Book Information

Series: Chapman & Hall/Crc Machine Learning & Pattern Recognition

Hardcover: 457 pages

Publisher: Chapman and Hall/CRC; 2 edition (October 8, 2014)

Language: English

ISBN-10: 1466583282

ISBN-13: 978-1466583283

Product Dimensions: 1.2 x 7 x 10 inches

Shipping Weight: 2.2 pounds (View shipping rates and policies)

Average Customer Review: 4.0 out of 5 stars 42 customer reviews

Best Sellers Rank: #43,839 in Books (See Top 100 in Books) #13 in Books > Textbooks > Computer Science > Algorithms #31 in Books > Computers & Technology > Programming > Algorithms #35 in Books > Computers & Technology > Databases & Big Data > Data Mining

Customer Reviews

"I thought the first edition was hands down, one of the best texts covering applied machine learning from a Python perspective. I still consider this to be the case. The text, already extremely broad in scope, has been expanded to cover some very relevant modern topics | I highly recommend this text to anyone who wants to learn machine learning | I particularly recommend it to those students who have followed along from more of a statistical learning perspective (Ng, Hastie, Tibshirani) and are looking to broaden their knowledge of applications. The updated text is very timely, covering topics that are very popular right now and have little coverage in existing texts in this area." •Intelligent Trading Tech blog, April 2015 "The book's emphasis on algorithms distinguishes it from other books on machine learning (ML). This is further highlighted by the extensive use of Python code to implement the algorithms. ... The topics chosen do reflect the current research areas in ML, and the book can be recommended to those wishing to gain an understanding of the current state of the field." •J. P. E. Hodgson, Computing Reviews, March 27, 2015 "I have been using this textbook for an undergraduate machine learning class for several years. Some of the best features of this book are the inclusion of Python code in the text (not just on a website), explanation of what the code does, and, in some cases, partial numerical run-throughs of the code. This helps students understand the algorithms better than high-level descriptions and equations alone and eliminates many sources of ambiguity and misunderstanding." •Daniel Kifer "This book will equip and engage students with its well-organised and -presented material. In each chapter, they will find thorough explanations, figures illustrating the discussed concepts and techniques, lots of programming (Python) and worked examples, practice questions, further readings, and a support website. The book will also be useful to professionals who can quickly inform and refresh their memory and knowledge of how machine learning works and what are the fundamental approaches and methods used in this area. As a whole, it provides an essential source for machine learning methodologies and techniques, how they work, and what are their application areas." •Ivan Jordanov, University of Portsmouth, UK Praise for the First Edition:" | liberally illustrated with many programming examples, using Python. It includes a basic primer on Python and has an accompanying website. It has excellent breadth and is comprehensive in terms of the topics it covers, both in terms of methods and in terms of concepts and theory. | I think the author has succeeded in his aim: the book provides an accessible introduction to machine learning. It would be excellent as a first exposure to the subject, and would put the various ideas in context |" •David J. Hand, International Statistical Review (2010), 78 "If you are interested in learning enough AI to understand the sort of new techniques being introduced into Web 2 applications, then this is a good place to start. | it covers the subject matter of many an introductory course on AI and it has references to

the source material and further reading but it is written in a fairly casual style. Overall it works and much of the mathematics is explained in ways that make it fairly clear what is going on [1]. This is a suitable introduction to AI if you are studying the subject on your own and it would make a good course text for an introduction and overview of AI." -I-Programmer, November 2009

Stephen Marsland is a professor of scientific computing and the postgraduate director of the School of Engineering and Advanced Technology (SEAT) at Massey University. His research interests in mathematical computing include shape spaces, Euler equations, machine learning, and algorithms. He received a PhD from Manchester University

For those diving into machine learning who are rusty at math or not a math expert this is a solid, understandable book on the topic. It covers a wide variety of machine learning algorithms, and while it does include some math, the math isn't the primary and only focus like other books on the topic. The math sections are less involved, giving the formulas and basic information the various algorithms are based on, but most importantly it's accompanied by easier to understand explanations and pseudo code/actual code implementations. I certainly hope the author continues to update and maintain this book over time - I've shown this book to coworkers who are also less math inclined and they liked the way the book was written and were interested in picking up their own copies.

This book has sat in my Wish List for several years. I always put off buying it - until I actually needed a machine learning book. Pros: Unlike too many technical books, the author of this book is not trying to display his brilliance. This is a nuts and bolts book - not a cookbook, but not a Springer-style book of axioms, either. Reading this book reminded me of school. There seem to be two kinds of professors - one kind tends to prefer examples and problems that display some sort of first principle or other basic and fundamental problem. These professors are good at teaching new concepts, but sometimes fail at teaching the practicalities. The other kind of professor favors practical problems and how-do-I-do-this-for-real issues. This book sits firmly in the latter camp. Cons: As mentioned above, this book is not a cookbook, and yet is also not rigorous. For such a practical book, I would have wanted more pseudocode and algorithms. For example, near the end of the book, the author goes over Kalman filters and particle filters. He gives one algorithm for a Kalman filter (which no treatment to the different kinds and uses of Kalman filters), and only a slight description of a particle filter (and no pseudocode). To be fair, books and volumes are to be

had on the subject. But, I was left wanting more. This book is a good introduction to how to use various aspects and techniques of machine learning. If you are looking for mathematical rigour, look elsewhere. If you are looking for cookbook-style algorithms, use this book as a supplement. If you want a practical overview of machine learning methods before setting on your course, buy this book. Recommended, but with qualification.

I am updating my review of this book because apparently in my first review I didn't do a very good job. This made the review less than useful. I will try to do a better job this time. If it still isn't helpful let me know and I will try again. Like the title says, this book takes an algorithmic approach to teaching machine learning - as opposed to an applied or example based approach. The expectation is that you would get a tutorial on all the main algorithms rather than how to put various algorithms together to solve a particular problem in, say, fraud detection. The Contents reveal the algorithm basis: 1. Introduction (types of machine learning, why you would want to do it in the first place and a quick introduction to supervised learning) 2. Preliminaries (Key ideas about the problem of overfitting and the what I consider the most important topic: how to test and know when you have a program that has learned something other than the noise). Here the author also covers some ideas about the role of probability. Calling it "turning data into probabilities" is a bit odd, but that's really what we do. Early on he gets the key ideas of the ROC curve out of the way - something many texts just gloss over. I think the secret to understanding machine learning is understanding the idea behind the bias-variance trade-off (it is also handled very well in *The Elements of Statistical Learning: Data Mining, Inference, and Prediction, Second Edition* (Springer Series in Statistics) which I used to teach a class and read before I read this book. 3. Coverage of Artificial Neural Networks starting with the perceptron and why you would want to go beyond linear discriminators 4. The multilayer ANN 5. Radial Basis Functions and Splines - this is interesting because, Andrew Ng presents a linear regression as the most basic learning algorithm in his Coursera course which means all of the fitting methods, even when not used for classification are relevant. 6. Is a section on dimensional reduction - feature selection and other methods like PCA and even factor analysis (most people stop with PCA which I personally think is a mistake, because you can accidentally end up keeping the features with all the noise and throwing out the meaningful linear combinations). 7. This is a cool section not seen in basic books on probabilistic methods - sure everyone teaches k-NN, but this one has a nice discussion of Gaussian mixture models 8. Talks about the support vector machine. Most people don't get introduced to the idea that ANN and SVMs are actually very similar - they are both large margin classifiers and so knowing something about

SVMs will help, even if you end up with some other large margin classifier (with or without kernels)9. This section talks about search and optimization. The way Ng teaches machine learning, you always begin with the error surface, take the derivative and then search for a minimum in the learning function. You quit teaching when you have minimized the error on the training set without out driving the error too high on the validation set - so in a way all these approaches are optimization methods.10. A whole section of genetic algorithms (which I jumped to first) a very clear explanation and a good example that really ran so I could see what was going on.11. Reinforcement learning12. Learning with trees - CART trees end the chapter something everyone working in this area should know something about. He saves random forests for the next section (where I suppose it really belongs)13. This section is on bagging an boosting and then compares the idea of a collection of weak learners (like stubby random trees) as a powerful tool - the idea behind random forests.14. Unsupervised learning. People tend to focus on supervised learning for a very good reason, but there are lots of examples where the cost of putting a label on a data example is too high, so an unsupervised method is a good call.15. Coverage of Markov Chain methods (MCMC) - again this does not get covered in every applied book.16. Graphical models - the Bayesian Network and Probabilistic Network Models along with the hidden Markov models17. Deep belief networks18. Gaussian process regression and classificationThe book concludes with an appendix on Python - getting started etc. I don't think this is quite enough Python unless you are already pretty familiar with the language. A critic of my first review suggested that I just bashed R and didn't talk about the book - not a completely unfair statement. R keeps data in data frames and Python is much more list and directory based. Data frames and collections are related and there are ways to do list comprehension in both languages. and Python has a data frame package to make using R-like constructs easier if you happen to be coming from R and like them (the package is called Pandas) Both are good languages, but I will stand by my original statement that R is a statistical language at its core. Many of the packages are written in C so they are fast (like the ones written for Python). It has been my experience that the open source tools for R development are just what the commentator said: they are adequate. In my humble opinion, R-Studio has a lot of catching up to do to be as good as professional tools like the JetBrains stuff (PyCharm). Look at MatLab compared to Octave. At least the community version of PyCharm is free. R-Studio is not fast, and the dirty secret of R is that everything you do has to fit in memory at once, so you have to be very careful with memory management or you will run out and R will crash - it happens to me everyday. Almost all of the ML methods are statistically based so R and all the books (like An Introduction to Statistical Learning: with Applications in R (Springer Texts in Statistics)) are totally brilliant. But if you want to see what it

is under the hood, I suggest you look at Advanced R (Chapman & Hall/CRC The R Series). This will give you a deep dive on the internals. Compare it to Python Scripting for Computational Science (Texts in Computational Science and Engineering) and make the call for yourself. I have used both R and Python for both prototyping advanced algorithms and putting code live in production. What tipped the scale for me was the productivity. Now that the data science community has started building state-of-art tools for Python, (not to say anything negative about the statistics community who put all of machine learning on a solid footing), I prefer a rapid development language with good tools, test-first frameworks and solid software engineering practices as part of the culture. The book reviewed here allows you to learn almost all of the algorithms used for machine learning and in the end you will be able to produce fast, readable, testable code.

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

Machine Learning: An Algorithmic Perspective, Second Edition (Chapman & Hall/Crc Machine Learning & Pattern Recognition) Statistics and Data Analysis for Microarrays Using R and Bioconductor, Second Edition (Chapman & Hall/CRC Mathematical and Computational Biology) Linear Models with R, Second Edition (Chapman & Hall/CRC Texts in Statistical Science) Generalized Linear Models, Second Edition (Chapman & Hall/CRC Monographs on Statistics & Applied Probability) Modeling and Analysis of Stochastic Systems, Second Edition (Chapman & Hall/CRC Texts in Statistical Science) Introduction to Scientific Programming and Simulation Using R, Second Edition (Chapman & Hall/CRC The R Series) Introduction to Modern Cryptography, Second Edition (Chapman & Hall/CRC Cryptography and Network Security Series) Measure and Integral: An Introduction to Real Analysis, Second Edition (Chapman & Hall/CRC Pure and Applied Mathematics) Topological Vector Spaces, Second Edition (Chapman & Hall/CRC Pure and Applied Mathematics) Environmental and Ecological Statistics with R, Second Edition (Chapman & Hall/CRC Applied Environmental Statistics) Draw in Perspective: Step by Step, Learn Easily How to Draw in Perspective (Drawing in Perspective, Perspective Drawing, How to Draw 3D, Drawing 3D, Learn to Draw 3D, Learn to Draw in Perspective) Modeling and Analysis of Stochastic Systems, Third Edition (Chapman & Hall/CRC Texts in Statistical Science) Introduction to Set Theory, Third Edition, Revised and Expanded (Chapman & Hall/CRC Pure and Applied Mathematics) A Concise Introduction to Pure Mathematics, Fourth Edition (Chapman Hall/CRC Mathematics Series) Computational Statistics Handbook with MATLAB, Third Edition (Chapman & Hall/CRC Computer Science & Data Analysis) Modelling Survival Data in Medical Research, Third Edition (Chapman & Hall/CRC Texts in Statistical Science) An Introduction to Systems Biology: Design Principles of Biological Circuits (Chapman & Hall/CRC Mathematical and Computational Biology) Introduction to

Scientific Programming and Simulation Using R (Chapman & Hall/CRC The R Series) Design of Experiments: An Introduction Based on Linear Models (Chapman & Hall/CRC Texts in Statistical Science) Introduction to Computational Biology: Maps, Sequences and Genomes (Chapman & Hall/CRC Interdisciplinary Statistics)

[Contact Us](#)

[DMCA](#)

[Privacy](#)

[FAQ & Help](#)