


The book was found

Machine Learning For Hackers: Case Studies And Algorithms To Get You Started



Synopsis

If you're an experienced programmer interested in crunching data, this book will get you started with machine learning—a toolkit of algorithms that enables computers to train themselves to automate useful tasks. Authors Drew Conway and John Myles White help you understand machine learning and statistics tools through a series of hands-on case studies, instead of a traditional math-heavy presentation. Each chapter focuses on a specific problem in machine learning, such as classification, prediction, optimization, and recommendation. Using the R programming language, you'll learn how to analyze sample datasets and write simple machine learning algorithms. Machine Learning for Hackers is ideal for programmers from any background, including business, government, and academic research. Develop a naive Bayesian classifier to determine if an email is spam, based only on its text. Use linear regression to predict the number of page views for the top 1,000 websites. Learn optimization techniques by attempting to break a simple letter cipher. Compare and contrast U.S. Senators statistically, based on their voting records. Build a "whom to follow" recommendation system from Twitter data.

Book Information

Paperback: 324 pages

Publisher: O'Reilly; 1st edition (February 25, 2012)

Language: English

ISBN-10: 1449303714

ISBN-13: 978-1449303716

Product Dimensions: 7 x 0.8 x 9.2 inches

Shipping Weight: 1.4 pounds (View shipping rates and policies)

Average Customer Review: 2.9 out of 5 stars 30 customer reviews

Best Sellers Rank: #373,500 in Books (See Top 100 in Books) #44 in Books > Computers & Technology > Computer Science > AI & Machine Learning > Natural Language Processing #74 in Books > Computers & Technology > Computer Science > AI & Machine Learning > Machine Theory #101 in Books > Textbooks > Computer Science > Algorithms

Customer Reviews

Drew Conway is a PhD candidate in Politics at NYU. He studies international relations, conflict, and terrorism using the tools of mathematics, statistics, and computer science in an attempt to gain a deeper understanding of these phenomena. His academic curiosity is informed by his years as an

analyst in the U.S. intelligence and defense communities. John Myles White is a PhD candidate in Psychology at Princeton. He studies pattern recognition, decision-making, and economic behavior using behavioral methods and fMRI. He is particularly interested in anomalies of value assessment.

This book is more of an introduction to R than anything to do with Machine Learning.. as a R introduction it's not bad minus the horibad sample code ... which won't even work if you copy and paste it, my advice find a different book there any many, many more accurate and detailed books on R and machine learning

I'm very disappointed that O'Reilly put there name on this book, it's boring doesn't really explain anything except how to use R to query data. But we all know machine learning is the process used to build that data. So the choice of ar arbitrary data sets to to explain what data is, its simply philosophical bull crap. It feels like this book is writen from the point of view of an idiot a proper title would be R query hack for databas tables. I really want a refund because I feel like I was duped ... and then Once you read the authors intro and that the book was written by someone who obviously has nothing to do with hacking or programming. Then it make sense, I don't see any actual maching learning period or detailed and correct explanation machine learning algorithms. O'R

This book may be good for those with little mathematical or statistical background, but its background information sections are too long and its treatment of ML topics too superficial for the book to be very useful for someone with the requisite background to actually implement the methods described in the book.

I learned a little about the R language and how a lot of scientist of physicist use it the write functions to dynamically create graphs to analyze the data they compiled. It also encompasses a lot of the same math you learn in statistics classes which makes sense. I even ended up downloading an R language reference app from the Apple app store to have an object reference of the language.

As other reviews have noted, this book is R-heavy. And R turns out to be a poor choice for a lot of this data manipulation. As a "hacker" I'd prefer to use python or something similar for a lot of the code and then use R for the stuff that it's good at. Also the kindle formatting of this book is terrible. The code is formatted very poorly (I hope this is a kindle translation problem and not the way that it looks in the book). That poor formatting makes it difficult to read through the code. I'd recommend

anyone who knows a decent programming language should try a different book.

The way this is formatted for Kindle makes it basically unusable. Page numbers in the middle of the text, mangled figures and tables, no chapter breaks and more. I had to buy the book again in paper.

I used this book to teach students about data mining and machine learning with a hands-on approach. I intended it to be used as a book for the students to rely and fall back on. It is not suited well for that purpose. Pros: The book is affordable and nicely written. The authors take great care in making the book useful and entertaining and one can immediately start putting things into practise. Also, the R examples are interesting and by itself motivating. Cons: The book has a couple of very grievous errors, that make me wonder the authors understand the subject matter. This is especially striking in the chapters on PCA and Multidimensional Scaling (which I covered in some depth in the class), but also to a lesser degree in other parts of the book that I have read more thoroughly (like optimization and linear and nonlinear regression). Many errors are not typos or simple mistakes but seem to be proof of a profound misunderstanding of concepts by the authors. I am sorry to be so blunt, but one should not write a book about topics that one is not intimate with. Given that the book is probably quite successful, it propagates error into a community whose members may not have the statistical background to spot the errors immediately. Some methods used in the book are quite hard to understand even for graduate students and to be so nonchalant about the underlying theory can be dangerous. I realize that the book is intended to be superficial with regards to mathematical or conceptual depth, but this combined with some of the presented high-level techniques can easily backfire when people are given the tools, but not the understanding. Especially when the explanations on interpretation are plainly wrong (I am talking about using standard deviations instead of variances, substantive interpretation of methodological artifacts, wrong explanation of R output, etc.). Additionally, certain parts of the book became outdated as soon as the book came out, such as the Google example. Overall, I do not recommend the book. I now only use it as a collection of nice examples and sometimes borrow bits of their R code.

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

Machine Learning for Hackers: Case Studies and Algorithms to Get You Started Machine Learning: For Beginners: Definitive Guide for Neural Networks, Algorithms, Random Forests and Decision Trees Made Simple (Machine Learning, Book 1) Fundamentals of Machine Learning for Predictive Data Analytics: Algorithms, Worked Examples, and Case Studies (MIT Press) Machine Learning: Fundamental Algorithms for Supervised and Unsupervised Learning With Real-World Applications

Algorithms for Reinforcement Learning (Synthesis Lectures on Artificial Intelligence and Machine Learning) The Don't Get Me Started! Toolkit - Workbook and Teacher Answer Key: Strategies for a Culturally-Challenged World (The Don't Get Me Started! Toolkit - Workbook and Teacher Key) (Volume 1) Machine Learning: A Probabilistic Perspective (Adaptive Computation and Machine Learning series) Introduction to Machine Learning (Adaptive Computation and Machine Learning series) Machine Learning: An Algorithmic Perspective, Second Edition (Chapman & Hall/Crc Machine Learning & Pattern Recognition) Beginner's Guide to Free-Motion Quilting: 50+ Visual Tutorials to Get You Started Professional-Quality Results on Your Home Machine To Be a Machine: Adventures Among Cyborgs, Utopians, Hackers, and the Futurists Solving the Modest Problem of Death Numerical Algorithms: Methods for Computer Vision, Machine Learning, and Graphics Genetic Algorithms in Search, Optimization, and Machine Learning Fundamentals of Deep Learning: Designing Next-Generation Machine Intelligence Algorithms BREAD MACHINE COOKBOOK: 120 Most Delicious Bread Machine Recipes (bread, bread bible, bread makers, breakfast, bread machine cookbook, bread baking, bread making, healthy, healthy recipes) Evolutionary Algorithms in Theory and Practice: Evolution Strategies, Evolutionary Programming, Genetic Algorithms Bundle of Algorithms in C++, Parts 1-5: Fundamentals, Data Structures, Sorting, Searching, and Graph Algorithms (3rd Edition) (Pts. 1-5) Practical Algorithms in Pediatric Hematology and Oncology: (Practical Algorithms in Pediatrics. Series Editor: Z. Hochberg) Practical Algorithms in Pediatric Nephrology: (Practical Algorithms in Pediatrics. Series Editor: Z. Hochberg) Practical Algorithms in Pediatric Gastroenterology: (Practical Algorithms in Pediatrics. Series Editor: Z. Hochberg)

[Contact Us](#)

[DMCA](#)

[Privacy](#)

[FAQ & Help](#)