


The book was found

Machine Learning: The Art And Science Of Algorithms That Make Sense Of Data



Synopsis

As one of the most comprehensive machine learning texts around, this book does justice to the field's incredible richness, but without losing sight of the unifying principles. Peter Flach's clear, example-based approach begins by discussing how a spam filter works, which gives an immediate introduction to machine learning in action, with a minimum of technical fuss. Flach provides case studies of increasing complexity and variety with well-chosen examples and illustrations throughout. He covers a wide range of logical, geometric and statistical models and state-of-the-art topics such as matrix factorisation and ROC analysis. Particular attention is paid to the central role played by features. The use of established terminology is balanced with the introduction of new and useful concepts, and summaries of relevant background material are provided with pointers for revision if necessary. These features ensure Machine Learning will set a new standard as an introductory textbook.

Book Information

File Size: 15319 KB

Print Length: 409 pages

Simultaneous Device Usage: Up to 4 simultaneous devices, per publisher limits

Publisher: Cambridge University Press; 1 edition (September 20, 2012)

Publication Date: September 20, 2012

Sold by: Digital Services LLC

Language: English

ASIN: B00P82T41Q

Text-to-Speech: Not enabled

X-Ray for Textbooks: Enabled

Word Wise: Not Enabled

Lending: Not Enabled

Enhanced Typesetting: Not Enabled

Best Sellers Rank: #109,341 Paid in Kindle Store (See Top 100 Paid in Kindle Store) #19

in Books > Computers & Technology > Computer Science > AI & Machine Learning > Computer Vision & Pattern Recognition #657 in Books > Computers & Technology > Networking & Cloud Computing #1095 in Kindle Store > Kindle eBooks > Computers & Technology

Customer Reviews

I'm probably to be considered a "very advanced amateur" or "informal professional" in machine

learning techniques. I never studied them in school, but I presently make a living in part by coding up new ones and coming up with ML solutions for commercial problems. Confronting the field for the first time, I wondered how people learned the stuff. Most of the introductory texts just covered neural nets ... and very very badly. Neural nets are still useful, and probably the most mature of machine learning techniques, but throwing them at a beginner without context is a recipe for confusion and dismay. This text, by contrast, barely mentions them, and puts them in their proper context for the beginner. The right way to think about machine learning is starting with *very* basic statistical techniques and probability theory, and building up from there into simple classification and scoring systems, and then on to the rest of the field. The author of this text does it the right way. One of the difficulties of didactic texts in the subject is ... machine learning is a very diverse field. All kinds of gizmos are helpful, and there isn't an obvious taxonomy, as there is in, say, linear time series models. The author takes a very high level view; breaking the field down into geometric, probabilistic and "logical" models. I believe this to be original, and a very powerful way of looking at things for the beginner. The progression is well thought out, and each chapter comes with a useful summary and references (one of which has already proved helpful to me) for further reading.

In my Advanced Statistics class one of the text books was *The Elements of Statistical Learning: Data Mining, Inference, and Prediction, Second Edition*. I found some of the chapters in this book heavy going and I had to read them several times and ask the professor lots of questions before I understood the material. In contrast, *Machine Learning* by Peter Flach is a very well written, very gentle introduction to machine learning algorithms. Prof. Flach writes that he spent four years writing this book and it shows in the care with which the material is presented. The mathematics used is algebra, exponents, summations, products and a bit of linear algebra. There are only a few places where derivatives are used (as it turns out, basic linear algebra can be used to describe many machine learning algorithms). The level of the *Machine Learning* makes it appropriate for an undergraduate Machine Learning course. *Machine Learning* covers most of the core algorithms in machine learning. Of necessity what is provided is an overview of topics like linear regression and linear classifiers like Support Vector Machines. These are topics that are covered in depth in books like *Applied Regression Analysis* and *An Introduction to Support Vector Machines and Other Kernel-based Learning Methods*.

Over a decade ago, Peter Flach of Bristol University wrote a paper on the topic of "On the state of the art in machine learning: A personal review" in which he reviewed several, then recent books,

related to developments in machine learning. This included Pat Langley's Elements of Machine Learning (Morgan Kaufmann), Tom Mitchell's Machine Learning (McGraw-Hill), and Data Mining: Practical Machine Learning Tools and Techniques with Java Implementations by Ian Witten and Eibe Frank (Morgan Kaufman) among many others. Dr. Flach mentioned Michael Berry and Gordon Linoff's Data Mining Techniques for Marketing, Sales, and Customer Support (John Wiley) for its excellent writing style citing the paragraph below and commending "I wish that all computer science textbooks were written like this." "People often find it hard to understand why the training set and test set are tainted once they have been used to build a model. An analogy may help: Imagine yourself back in the 5th grade. The class is taking a spelling test. Suppose that, at the end of the test period, the teacher asks you to estimate your own grade on the quiz by marking the words you got wrong. You will give yourself a very good grade, but your spelling will not improve. If, at the beginning of the period, you thought there should be an \hat{e} at the end of tomato , nothing will have happened to change your mind when you grade your paper. No new data has entered the system. You need a test set! Now, imagine that at the end of the test the teacher allows you to look at the papers of several neighbors before grading your own. If they all agree that tomato has no final \hat{e} , you may decide to mark your own answer wrong. If the teacher gives the same quiz tomorrow, you will do better. But how much better?

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

Machine Learning: The Art and Science of Algorithms that Make Sense of Data
Machine Learning with Spark - Tackle Big Data with Powerful Spark Machine Learning Algorithms
Analytics: Data Science, Data Analysis and Predictive Analytics for Business (Algorithms, Business Intelligence, Statistical Analysis, Decision Analysis, Business Analytics, Data Mining, Big Data)
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)
A collection of Advanced Data Science and Machine Learning Interview Questions Solved in Python and Spark (II): Hands-on Big Data and Machine ... Programming Interview Questions) (Volume 7)
Data Analytics: What Every Business Must Know About Big Data And Data Science (Data Analytics for Business, Predictive Analysis, Big Data)

Unsupervised Machine Learning in Python: Master Data Science and Machine Learning with Cluster Analysis, Gaussian Mixture Models, and Principal Components Analysis 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) Deep Learning: Recurrent Neural Networks in Python: LSTM, GRU, and more RNN machine learning architectures in Python and Theano (Machine Learning in Python) 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 ... 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) Java: Artificial Intelligence; Made Easy, w/ Java Programming; Learn to Create your * Problem Solving * Algorithms! TODAY! w/ Machine Learning & Data Structures (Artificial Intelligence Series) Javascript Artificial Intelligence: Made Easy, w/ Essential Programming; Create your * Problem Solving * Algorithms! TODAY! w/ Machine Learning & Data Structures (Artificial Intelligence Series) Swift: Programming, Master's Handbook; A TRUE Beginner's Guide! Problem Solving, Code, Data Science, Data Structures & Algorithms (Code like a PRO in ... engineering, r programming, iOS development) Ruby: Programming, Master's Handbook: A TRUE Beginner's Guide! Problem Solving, Code, Data Science, Data Structures & Algorithms (Code like a PRO in ... web design, tech, perl, ajax, swift, python,)

[Dmca](#)