


The book was found

A First Course In Machine Learning, Second Edition



Synopsis

"A First Course in Machine Learning by Simon Rogers and Mark Girolami is the best introductory book for ML currently available. It combines rigor and precision with accessibility, starts from a detailed explanation of the basic foundations of Bayesian analysis in the simplest of settings, and goes all the way to the frontiers of the subject such as infinite mixture models, GPs, and MCMC." •Devdatt Dubhashi, Professor, Department of Computer Science and Engineering, Chalmers University, Sweden "This textbook manages to be easier to read than other comparable books in the subject while retaining all the rigorous treatment needed. The new chapters put it at the forefront of the field by covering topics that have become mainstream in machine learning over the last decade." •Daniel Barbara, George Mason University, Fairfax, Virginia, USA "The new edition of A First Course in Machine Learning by Rogers and Girolami is an excellent introduction to the use of statistical methods in machine learning. The book introduces concepts such as mathematical modeling, inference, and prediction, providing 'just in time' the essential background on linear algebra, calculus, and probability theory that the reader needs to understand these concepts." •Daniel Ortiz-Arroyo, Associate Professor, Aalborg University Esbjerg, Denmark "I was impressed by how closely the material aligns with the needs of an introductory course on machine learning, which is its greatest strength. Overall, this is a pragmatic and helpful book, which is well-aligned to the needs of an introductory course and one that I will be looking at for my own students in coming months." •David Clifton, University of Oxford, UK "The first edition of this book was already an excellent introductory text on machine learning for an advanced undergraduate or taught masters level course, or indeed for anybody who wants to learn about an interesting and important field of computer science. The additional chapters of advanced material on Gaussian process, MCMC and mixture modeling provide an ideal basis for practical projects, without disturbing the very clear and readable exposition of the basics contained in the first part of the book." •Gavin Cawley, Senior Lecturer, School of Computing Sciences, University of East Anglia, UK "This book could be used for junior/senior undergraduate students or first-year graduate students, as well as individuals who want to explore the field of machine learning. The book introduces not only the concepts but the underlying ideas on algorithm implementation from a critical thinking perspective." •Guangzhi Qu, Oakland University, Rochester, Michigan, USA

Book Information

Paperback: 427 pages

Publisher: Chapman and Hall/CRC; 2 edition (August 15, 2016)

Language: English

ISBN-10: 1498738486

ISBN-13: 978-1498738484

Product Dimensions: 9.6 x 1.1 x 6.5 inches

Shipping Weight: 1.7 pounds (View shipping rates and policies)

Average Customer Review: Be the first to review this item

Best Sellers Rank: #739,333 in Books (See Top 100 in Books) #109 in Books > Computers & Technology > Computer Science > AI & Machine Learning > Machine Theory #422 in Books > Computers & Technology > Databases & Big Data > Data Mining #668 in Books > Business & Money > Education & Reference > Statistics

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

Bioinformatics: The Machine Learning Approach, Second Edition (Adaptive Computation and

Machine Learning) Deep Learning: Recurrent Neural Networks in Python: LSTM, GRU, and more
RNN machine learning architectures in Python and Theano (Machine Learning in Python)

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 First Course in Machine Learning, Second Edition

Python: Python Programming Course: Learn the Crash Course to Learning the Basics of Python
(Python Programming, Python Programming Course, Python Beginners Course) Machine Learning:

A Probabilistic Perspective (Adaptive Computation and Machine Learning series) Unsupervised

Machine Learning in Python: Master Data Science and Machine Learning with Cluster Analysis,
Gaussian Mixture Models, and Principal Components Analysis Machine Learning with Spark -

Tackle Big Data with Powerful Spark Machine Learning Algorithms Foundations of Machine

Learning (Adaptive Computation and Machine Learning series) Introduction to Machine Learning

(Adaptive Computation and Machine Learning series) Gaussian Processes for Machine Learning

(Adaptive Computation and Machine Learning series) First-Time Machine Applique: Learning to

Machine Applique in Nine Easy Lessons 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) Machine Learning with R - Second Edition Machine Learning with R - Second Edition - Deliver Data Insights with R and Predictive Analytics Introduction to Statistical Relational Learning (Adaptive Computation and Machine Learning series) Reinforcement Learning: An Introduction (Adaptive Computation and Machine Learning series)

[Dmca](#)