Rogers / Girolami

A First Course in Machine Learning

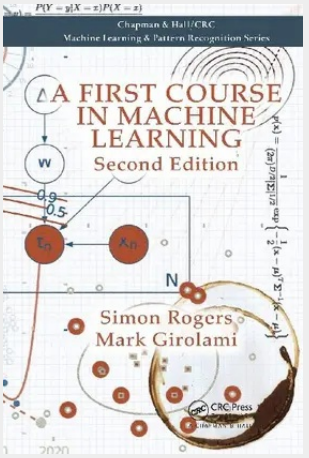

"A First Course in Machine Learning by Simon Rogers and Mark Girolami is the best introductory book for ML currently available. It combines rigor and precision with accessibility, starts from a detailed explanation of the basic foundations of Bayesian analysis in the simplest of settings, and goes all 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

The new edition of this popular, undergraduate textbook has been revised and updated to reflect current growth areas in Machine Learning. The new edition includes three new chapters with more detailed discussion of Markov Chain Monte Carlo techniques, Classification and Regression with Gaussian Processes, and Dirichlet Process models.

68,60 €
68,60 € (zzgl. MwSt.)

Lieferfrist: bis zu 10 Tage

Artikelnummer: 9780367574642
Medium: Buch
ISBN: 978-0-367-57464-2
Verlag: Chapman and Hall/CRC
Erscheinungstermin: 30.06.2020
Sprache(n): Englisch
Auflage: 2. Auflage 2020
Serie: Chapman & Hall/CRC Machine Learning & Pattern Recognition
Produktform: Kartoniert
Gewicht: 646 g
Seiten: 428
Format (B x H): 156 x 234 mm