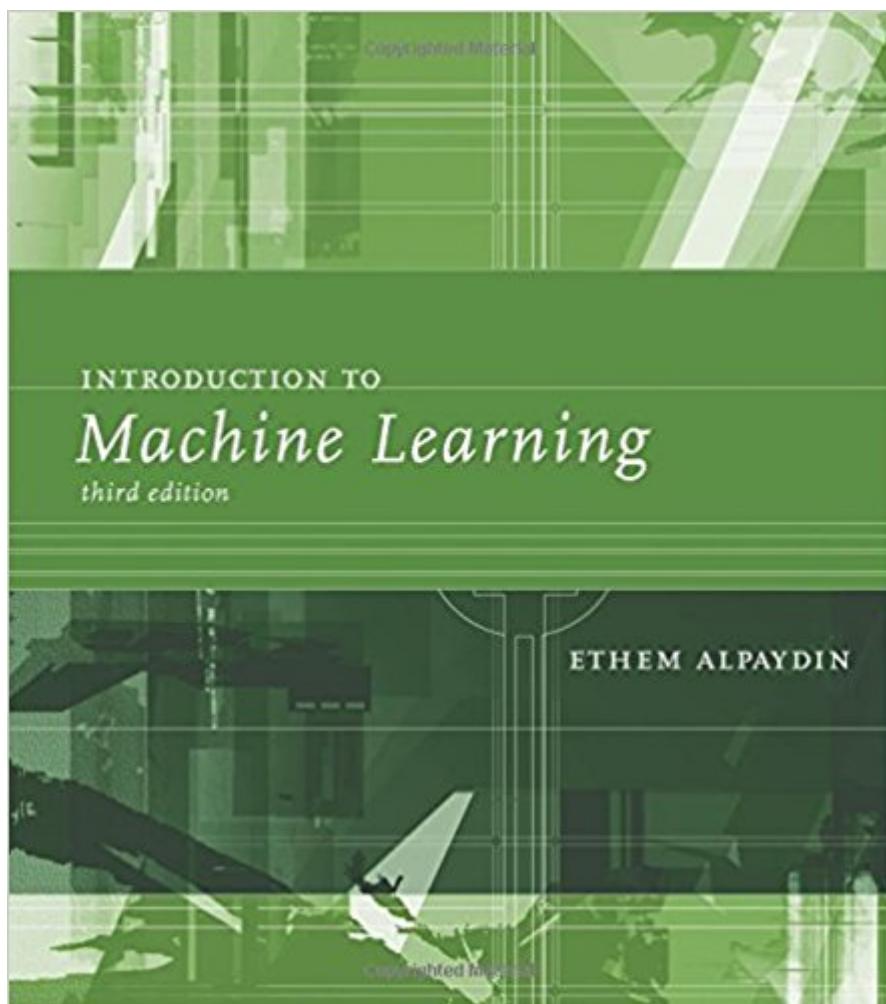


The book was found

Introduction To Machine Learning (Adaptive Computation And Machine Learning Series)



Synopsis

The goal of machine learning is to program computers to use example data or past experience to solve a given problem. Many successful applications of machine learning exist already, including systems that analyze past sales data to predict customer behavior, optimize robot behavior so that a task can be completed using minimum resources, and extract knowledge from bioinformatics data.

Introduction to Machine Learning is a comprehensive textbook on the subject, covering a broad array of topics not usually included in introductory machine learning texts. Subjects include supervised learning; Bayesian decision theory; parametric, semi-parametric, and nonparametric methods; multivariate analysis; hidden Markov models; reinforcement learning; kernel machines; graphical models; Bayesian estimation; and statistical testing. Machine learning is rapidly becoming a skill that computer science students must master before graduation. The third edition of Introduction to Machine Learning reflects this shift, with added support for beginners, including selected solutions for exercises and additional example data sets (with code available online). Other substantial changes include discussions of outlier detection; ranking algorithms for perceptrons and support vector machines; matrix decomposition and spectral methods; distance estimation; new kernel algorithms; deep learning in multilayered perceptrons; and the nonparametric approach to Bayesian methods. All learning algorithms are explained so that students can easily move from the equations in the book to a computer program. The book can be used by both advanced undergraduates and graduate students. It will also be of interest to professionals who are concerned with the application of machine learning methods.

Book Information

Series: Adaptive Computation and Machine Learning series

Hardcover: 640 pages

Publisher: The MIT Press; third edition edition (August 22, 2014)

Language: English

ISBN-10: 0262028182

ISBN-13: 978-0262028189

Product Dimensions: 8 x 0.9 x 9 inches

Shipping Weight: 2.8 pounds

Average Customer Review: 3.4 out of 5 stars 8 customer reviews

Best Sellers Rank: #356,781 in Books (See Top 100 in Books) #71 in Books > Computers & Technology > Computer Science > AI & Machine Learning > Machine Theory #111 in Books >

Customer Reviews

Ethem Alpaydin's Introduction to Machine Learning provides a nice blending of the topical coverage of machine learning (À la Tom Mitchell) with formal probabilistic foundations (À la Christopher Bishop). This newly updated version now introduces some of the most recent and important topics in machine learning (e.g., spectral methods, deep learning, and learning to rank) to students and researchers of this critically important and expanding field. (John W. Sheppard, Professor of Computer Science, Montana State University) I have used Introduction to Machine Learning for several years in my graduate Machine Learning course. The book provides an ideal balance of theory and practice, and with this third edition, extends coverage to many new state-of-the-art algorithms. I look forward to using this edition in my next Machine Learning course. (Larry Holder, Professor of Electrical Engineering and Computer Science, Washington State University) This volume is both a complete and accessible introduction to the machine learning world. This is a 'Swiss Army knife' book for this rapidly evolving subject. Although intended as an introduction, it will be useful not only for students but for any professional looking for a comprehensive book in this field. Newcomers will find clearly explained concepts and experts will find a source for new references and ideas. (Hilario Gómez-Moreno, IEEE Senior Member, University of Alcalá, Spain)

Ethem Alpaydin is a Professor in the Department of Computer Engineering at Bogaziçi University, Istanbul.

I hope one day to meet Alpaydin in person to thank him profusely for this book (and earlier versions). He unpacks the major concepts of machine learning in a manner that makes it very easy to follow. I probably have 3 copies of the earlier edition. I bought this one, and am very pleased with the updates - specifically related to neural networks and deep learning. If you're running around in this domain - this book is crucial.

I understand ML very well, and I find this text nearly impossible to penetrate. Formulas are reduced to their most rudimentary forms. Sure it is impressive that the author obviously has a good grasp on the topic, but there are virtually no explanations behind the math. This book was written just to show

off, not to teach. Definitely the most pompous book on ML I've ever seen.

This is a very challenging textbook that is best used as a part of a graduate machine learning class, as a reference or for a very dedicated self-learner. I recommend the edition published by MIT or the ebook from the MIT site. The focus of this book is application and the essentials of theory are covered very briefly in the beginning. Sometimes more exposition is needed to understand topics, however, this book has been the single most important to my work as a data scientist and it is always on my desk. It took me about one year to get from cover to cover. Alpaydin provides comprehensive coverage on the most common machine learning techniques, starting from a probabilistic perspective and continuing to discriminant models. Bayesian analysis, dimensionality reduction, support vector machines (kernel machines), and unsupervised learning are covered in detail, along with techniques applicable to image recognition, NLP, and AI. The end of each chapter includes a bibliography which helps for deeper dives into specific topics. The notation is simple and fairly consistent throughout the book. Since the field has become so large no textbook on machine learning can stand alone. Classes in calculus, linear algebra, probability and statistics are recommended first, but you can pick this up on the fly when going through the book. Before reading this book, it is helpful to go through an applied approach such as *Hands-On Machine Learning with Scikit-Learn and TensorFlow* (Géron, 2017). I recommend *Deep Learning* (Goodfellow et al, 2015) as a continuation to the chapters on multilayer perceptrons. A deeper exploration of theory is provided in texts such as *Learning from Data* (Abu Mostafa, 2012), *Foundations of Machine Learning* (Mohri et al, 2012), and *Foundations of Data Science* (Blum et al, 2016). The PyMC3 documentation is a good companion for the Bayesian sections and the Scikit-learn documentation helps with the content as well. I hope that Alpaydin releases a fourth edition, however, as it stands I highly recommend this text.

I took this book out from a university library because I needed to learn some ML and they were out of any book I'd heard of. It is utterly horrible. Very simple concepts explained in such a way as to make them confusing, with strange notation, terrible figures, and very little English analysis or context. I only read the first thirty pages but oh man were they bad. I stopped shortly after the completely garbled and internally inconsistent definition of Vapnik-Chervonenkis dimension (page 27). I have some background in ML, a pretty strong math background, and a PhD in biophysics, and there is nothing I hate more than simple things I already understand or could easily understand explained so as to be incomprehensible.

Agreed 'written to show off, not to teach', inconsistent notation, got very very confused...

good for beginners

Very nice and clear content. Thanks....

Item as described. Fast shipment.

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

Introduction to Machine Learning (Adaptive Computation and Machine Learning series) Machine Learning: A Probabilistic Perspective (Adaptive Computation and Machine Learning series) Reinforcement Learning with Python: An Introduction (Adaptive Computation and Machine Learning series) Reinforcement Learning: An Introduction (Adaptive Computation and Machine Learning) Deep Learning (Adaptive Computation and Machine Learning series) Semi-Supervised Learning (Adaptive Computation and Machine Learning series) Probabilistic Graphical Models: Principles and Techniques (Adaptive Computation and Machine Learning series) IntAR, Interventions Adaptive Reuse, Volume 03; Adaptive Reuse in Emerging Economies Human Computation (Synthesis Lectures on Artificial Intelligence and Machine Learning) Machine Learning: For Beginners: Definitive Guide for Neural Networks, Algorithms, Random Forests and Decision Trees Made Simple (Machine Learning, Book 1) Machine Learning: An Algorithmic Perspective, Second Edition (Chapman & Hall/Crc Machine Learning & Pattern Recognition) 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) Introduction to Linear Optimization (Athena Scientific Series in Optimization and Neural Computation, 6) Introduction to Practical Peridynamics: Computational Solid Mechanics Without Stress and Strain (Frontier Research in Computation and Mechanics of Materials) Number Talks: Whole Number Computation, Grades K-5: A Multimedia Professional Learning Resource A Course in Derivative Securities: Introduction to Theory and Computation (Springer Finance) Introduction to Automata Theory, Languages, and Computation (3rd Edition) Introduction to Automata Theory, Languages, and Computation (2nd Edition) Introduction to Computation and Programming Using Python: With Application to Understanding Data (MIT Press) Medical Terminology & Anatomy for ICD-10 Coding - Text and Elsevier Adaptive Learning Package, 3e

[Contact Us](#)

[DMCA](#)

[Privacy](#)

[FAQ & Help](#)