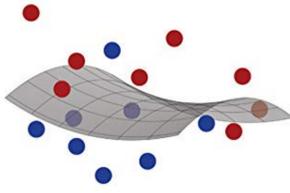


Foundations of
Machine Learning second edition



Mehryar Mohri,
Afshin Rostamizadeh,
and Ameet Talwalkar

[eBooks] Foundations Of Machine Learning, Second Edition (Adaptive Computation And Machine Learning Series)

Recognizing the pretentiousness ways to acquire this ebook **Foundations of Machine Learning, second edition (Adaptive Computation and Machine Learning series)** is additionally useful. You have remained in right site to begin getting this info. get the Foundations of Machine Learning, second edition (Adaptive Computation and Machine Learning series) join that we find the money for here and check out the link.

You could purchase guide Foundations of Machine Learning, second edition (Adaptive Computation and Machine Learning series) or acquire it as soon as feasible. You could speedily download this Foundations of Machine Learning, second edition (Adaptive Computation and Machine Learning series) after getting deal. So, with you require the books swiftly, you can straight acquire it. Its thus unconditionally simple and so fats, isnt it? You have to favor to in this tune

Foundations of Machine Learning-Mehryar Mohri 2012-08-17 This graduate-level textbook introduces fundamental concepts and methods in machine learning. It describes several important modern algorithms, provides the theoretical underpinnings of these algorithms, and illustrates key aspects for their application. The authors aim to present novel theoretical tools and concepts while giving concise proofs even for relatively advanced topics. Foundations of Machine Learning fills the need for a general textbook that also offers theoretical details and an emphasis on proofs. Certain topics that are often treated with insufficient attention are discussed in more detail here; for example, entire chapters are devoted to regression, multi-class classification, and ranking. The first three chapters lay the theoretical foundation for what follows, but each remaining chapter is mostly self-contained. The appendix offers a concise probability review, a short introduction to convex optimization, tools for concentration bounds, and several basic properties of matrices and norms used in the book.The book is intended for graduate students and researchers in machine learning, statistics, and related areas; it can be used either as a textbook or as a reference text for a research seminar.

Foundations of Machine Learning-Mehryar Mohri 2018-12-25 A new edition of a graduate-level machine learning textbook that focuses on the analysis and theory of algorithms. This book is a general introduction to machine learning that can serve as a textbook for graduate students and a reference for researchers. It covers fundamental modern topics in machine learning while providing the theoretical basis and conceptual tools needed for the discussion and justification of algorithms. It also describes several key aspects of the application of these algorithms. The authors aim to present novel theoretical tools and concepts while giving concise proofs even for relatively advanced topics. Foundations of Machine Learning is unique in its focus on the analysis and theory of algorithms. The first four chapters lay the theoretical foundation for what follows; subsequent chapters are mostly self-contained. Topics covered include the Probably Approximately Correct (PAC) learning framework; generalization bounds based on Rademacher complexity and VC-dimension; Support Vector Machines (SVMs); kernel methods; boosting; on-line learning; multi-class classification; ranking; regression; algorithmic stability; dimensionality reduction; learning automata and languages; and reinforcement learning. Each chapter ends with a set of exercises. Appendixes provide additional material including concise probability review. This second edition offers three new chapters, on model selection, maximum entropy models, and conditional entropy models. New material in the appendixes includes a major section on Fenchel duality, expanded coverage of concentration inequalities, and an entirely new entry on information theory. More than half of the exercises are new to this edition.

Fundamentals of Machine Learning for Predictive Data Analytics-John D. Kelleher 2015-07-24 A comprehensive introduction to the most important machine learning approaches used in predictive data analytics, covering both theoretical concepts and practical applications.

Machine Learning Refined-Jeremy Watt 2020-01-29 An intuitive approach to machine learning covering key concepts, real-world applications, and practical Python coding exercises.

Machine Learning-Stephen Marsland 2011-03-23 Traditional books on machine learning can be divided into two groups- those aimed at advanced undergraduates or early postgraduates with reasonable mathematical knowledge and those that are primers on how to code algorithms. The field is ready for a text that not only demonstrates how to use the algorithms that make up machine learning methods, but

Foundations of Deep Reinforcement Learning-Laura Graesser 2019-11-20 The Contemporary Introduction to Deep Reinforcement Learning that Combines Theory and Practice Deep reinforcement learning (deep RL) combines deep learning and reinforcement learning, in which artificial agents learn to solve sequential decision-making problems. In the past decade deep RL has achieved remarkable results on a range of problems, from single and multiplayer games—such as Go, Atari games, and Dota 2—to robotics. Foundations of Deep Reinforcement Learning is an introduction to deep RL that uniquely combines both theory and implementation. It starts with intuition, then carefully explains the theory of deep RL algorithms, discusses implementations in its companion software library SLM Lab, and finishes with the practical details of getting deep RL to work. This guide is ideal for both computer science students and software engineers who are familiar with basic machine learning concepts and have a working understanding of Python. Understand each key aspect of a deep RL problem Explore policy- and value-based algorithms, including REINFORCE, SARSA, DQN, Double DQN, and Prioritized Experience Replay (PER) Delve into combined algorithms, including Actor-Critic and Proximal Policy Optimization (PPO) Understand how algorithms can be parallelized synchronously and asynchronously Run algorithms in SLM Lab and learn the practical implementation details for getting deep RL to work Explore algorithm benchmark results with tuned hyperparameters Understand how deep RL environments are designed Register your book for convenient access to downloads, updates, and/or corrections as they become available. See inside book for details.

A First Course in Machine Learning, Second Edition-Simon Rogers 2016-10-14 "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

Understanding Machine Learning-Shai Shalev-Shwartz 2014-05-19 Introduces machine learning and its algorithmic paradigms, explaining the principles behind automated learning approaches and the considerations underlying their usage.

Mathematics for Machine Learning-Marc Peter Deisenroth 2020-04-23 The fundamental mathematical tools needed to understand machine learning include linear algebra, analytic geometry, matrix decompositions, vector calculus, optimization, probability and statistics. These topics are traditionally taught in disparate courses, making it hard for data science or computer science students, or professionals, to efficiently learn the mathematics. This self-contained textbook bridges the gap between mathematical and machine learning texts, introducing the mathematical concepts with a minimum of prerequisites. It uses these concepts to derive four central machine learning methods: linear regression, principal component analysis, Gaussian mixture models and support vector machines. For students and others with a mathematical background, these derivations provide a starting point to machine learning texts. For those learning the mathematics for the first time, the methods help build intuition and practical experience with applying mathematical concepts. Every chapter includes worked examples and exercises to test understanding. Programming tutorials are offered on the book's web site.

Artificial Intelligence-David L. Poole 2017-09-25 Artificial Intelligence presents a practical guide to AI, including agents, machine learning and problem-solving simple and complex domains.

Introduction to Machine Learning-Ethem Alpaydin 2014-08-29 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.

Reinforcement Learning-Richard S. Sutton 2018-11-13 The significantly expanded and updated new edition of a widely used text on reinforcement learning, one of the most active research areas in artificial intelligence. Reinforcement learning, one of the most active research areas in artificial intelligence, is a computational approach to learning whereby an agent tries to maximize the total amount of reward it receives while interacting with a complex, uncertain environment. In Reinforcement Learning, Richard Sutton and Andrew Barto provide a clear and simple account of the field's key ideas and algorithms. This second edition has been significantly expanded and updated, presenting new topics and updating coverage of other topics. Like the first edition, this second edition focuses on core online learning algorithms, with the more mathematical material set off in shaded boxes. Part I covers as much of reinforcement learning as possible without going beyond the tabular case for which exact solutions can be found. Many algorithms presented in this part are new to the second edition, including UCB, Expected Sarsa, and Double Learning. Part II extends these ideas to function approximation, with new sections on such topics as artificial neural networks and the Fourier basis, and offers expanded treatment of off-policy learning and policy-gradient methods. Part III has new chapters on reinforcement learning's relationships to psychology and neuroscience, as well as an updated case-studies chapter including AlphaGo and AlphaGo Zero, Atari game playing, and IBM Watson's wagering strategy. The final chapter discusses the future societal impacts of reinforcement learning.

Machine Learning-Sergios Theodoridis 2020-02-19 Machine Learning: A Bayesian and Optimization Perspective, 2nd edition, gives a unified perspective on machine learning by covering both pillars of supervised learning, namely regression and classification. The book starts with the basics, including mean square, least squares and maximum likelihood methods, ridge regression, Bayesian decision theory classification, logistic regression, and decision trees. It then progresses to more recent techniques, covering sparse modelling methods, learning in reproducing kernel Hilbert spaces and support vector machines, Bayesian inference with a focus on the EM algorithm and its approximate inference variational versions, Monte Carlo methods, probabilistic graphical models focusing on Bayesian networks, hidden Markov models and particle filtering. Dimensionality reduction and latent variables modelling are also considered in depth. This palette of techniques concludes with an extended chapter on neural networks and deep learning architectures. The book also covers the fundamentals of statistical parameter estimation, Wiener and Kalman filtering, convexity and convex optimization, including a chapter on stochastic approximation and the gradient descent family of algorithms, presenting related online learning techniques as well as concepts and algorithmic versions for distributed optimization. Focusing on the physical reasoning behind the mathematics, without sacrificing rigor, all the various methods and techniques are explained in depth, supported by examples and problems, giving an invaluable resource to the student and researcher for understanding and applying machine learning concepts. Most of the chapters include typical case studies and computer exercises, both in MATLAB and Python. The chapters are written to be as self-contained as possible, making the text suitable for different courses: pattern recognition, statistical/adaptive signal processing, statistical/Bayesian learning, as well as courses on sparse modeling, deep learning, and probabilistic graphical models. New to this edition: Complete rewrite of the chapter on Neural Networks and Deep Learning to reflect the latest advances since the 1st edition. The chapter, starting from the basic perceptron and feed-forward neural networks concepts, now presents an in depth treatment of deep networks, including recent optimization algorithms, batch normalization, regularization techniques such as the dropout method, convolutional neural networks, recurrent neural networks, attention mechanisms, adversarial examples and training, capsule networks and generative architectures, such as restricted Boltzmann machines (RBMs), variational autoencoders and generative adversarial networks (GANs). Expanded treatment of Bayesian learning to include nonparametric Bayesian methods, with a focus on the Chinese restaurant and the Indian buffet processes. Presents the physical reasoning, mathematical modeling and algorithmic implementation of each method Updates on the latest trends, including sparsity, convex analysis and optimization, online distributed algorithms, learning in RKH spaces, Bayesian inference, graphical and hidden Markov models, particle filtering, deep learning, dictionary learning and latent variables modeling Provides case studies on a variety of topics, including protein folding prediction, optical character recognition, text authorship identification, fMRI data analysis, change point detection, hyperspectral image unmixing, target localization, and more

Neural Network Learning-Martin Anthony 2009-08-20 This book describes recent theoretical advances in the study of artificial neural networks. It explores probabilistic models of supervised learning problems, and addresses the key statistical and computational questions. The authors also discuss the computational complexity of neural network learning, describing a variety of hardness results, and outlining two efficient constructive learning algorithms. The book is essentially self-contained, since it introduces the necessary background material on probability, statistics, combinatorics and computational complexity; and it is intended to be accessible to researchers and graduate students in computer science, engineering, and mathematics.

Fundamentals of Deep Learning-Nikhil Buduma 2017-05-25 With the reinvigoration of neural networks in the 2000s, deep learning has become an extremely active area of research, one that's paving the way for modern machine learning. In this practical book, author Nikhil Buduma provides examples and clear explanations to guide you through major concepts of this complicated field. Companies such as Google, Microsoft, and Facebook are actively growing in-house deep-learning teams. For the rest of us, however, deep learning is still a pretty complex and difficult subject to grasp. If you're familiar with Python, and have a background in calculus, along with a basic understanding of machine learning, this book will get you started. Examine the foundations of machine learning and neural networks Learn how to train feed-forward neural networks Use TensorFlow to implement your first neural network Manage problems that arise as you begin to make networks deeper Build neural networks that analyze complex images Perform effective dimensionality reduction using autoencoders Dive deep into sequence analysis to examine language Learn the fundamentals of reinforcement learning

Foundations of Data Science-Avrim Blum 2020-01-31 This book provides an introduction to the mathematical and algorithmic foundations of data science, including machine learning, high-dimensional geometry, and analysis of large networks. Topics include the counterintuitive nature of data in high dimensions, important linear algebraic techniques such as singular value decomposition, the theory of random walks and Markov chains, the fundamentals of and important algorithms for machine learning, algorithms and analysis for clustering, probabilistic models for large networks, representation learning including topic modelling and non-negative matrix

factorization, wavelets and compressed sensing. Important probabilistic techniques are developed including the law of large numbers, tail inequalities, analysis of random projections, generalization guarantees in machine learning, and moment methods for analysis of phase transitions in large random graphs. Additionally, important structural and complexity measures are discussed such as matrix norms and VC-dimension. This book is suitable for both undergraduate and graduate courses in the design and analysis of algorithms for data.

Machine Learning-Kevin P. Murphy 2012-08-24 A comprehensive introduction to machine learning that uses probabilistic models and inference as a unifying approach. Today's Web-enabled deluge of electronic data calls for automated methods of data analysis. Machine learning provides these, developing methods that can automatically detect patterns in data and then use the uncovered patterns to predict future data. This textbook offers a comprehensive and self-contained introduction to the field of machine learning, based on a unified, probabilistic approach. The coverage combines breadth and depth, offering necessary background material on such topics as probability, optimization, and linear algebra as well as discussion of recent developments in the field, including conditional random fields, L1 regularization, and deep learning. The book is written in an informal, accessible style, complete with pseudo-code for the most important algorithms. All topics are copiously illustrated with color images and worked examples drawn from such application domains as biology, text processing, computer vision, and robotics. Rather than providing a cookbook of different heuristic methods, the book stresses a principled model-based approach, often using the language of graphical models to specify models in a concise and intuitive way. Almost all the models described have been implemented in a MATLAB software package—PMTK (probabilistic modeling toolkit)—that is freely available online. The book is suitable for upper-level undergraduates with an introductory-level college math background and beginning graduate students.

Elements of Causal Inference-Jonas Peters 2017-11-22 The mathematization of causality is a relatively recent development, and has become increasingly important in data science and machine learning. This book offers a self-contained and concise introduction to causal models and how to learn them from data. After explaining the need for causal models and discussing some of the principles underlying causal inference, the book teaches readers how to use causal models: how to compute intervention distributions, how to infer causal models from observational and interventional data, and how causal ideas could be exploited for classical machine learning problems. All of these topics are discussed first in terms of two variables and then in the more general multivariate case. The bivariate case turns out to be a particularly hard problem for causal learning because there are no conditional independences as used by classical methods for solving multivariate cases. The authors consider analyzing statistical asymmetries between cause and effect to be highly instructive, and they report on their decade of intensive research into this problem. The book is accessible to readers with a background in machine learning or statistics, and can be used in graduate courses or as a reference for researchers. The text includes code snippets that can be copied and pasted, exercises, and an appendix with a summary of the most important technical concepts.

Machine Learning For Dummies-John Paul Mueller 2021-02-09 One of Mark Cuban's top reads for better understanding A.I. (inc.com, 2021) Your comprehensive entry-level guide to machine learning While machine learning expertise doesn't quite mean you can create your own Turing Test-proof android—as in the movie Ex Machina—it is a form of artificial intelligence and one of the most exciting technological means of identifying opportunities and solving problems fast and on a large scale. Anyone who masters the principles of machine learning is mastering a big part of our tech future and opening up incredible new directions in careers that include fraud detection, optimizing search results, serving real-time ads, credit-scoring, building accurate and sophisticated pricing models—and way, way more. Unlike most machine learning books, the fully updated 2nd Edition of Machine Learning For Dummies doesn't assume you have years of experience using programming languages such as Python (R source is also included in a downloadable form with comments and explanations), but lets you in on the ground floor, covering the entry-level materials that will get you up and running building models you need to perform practical tasks. It takes a look at the underlying—and fascinating—math principles that power machine learning but also shows that you don't need to be a math whiz to build fun new tools and apply them to your work and study. Understand the history of AI and machine learning Work with Python 3.8 and TensorFlow 2.x (and R as a download) Build and test your own models Use the latest datasets, rather than the worn out data found in other books Apply machine learning to real problems Whether you want to learn for college or to enhance your business or career performance, this friendly beginner's guide is your best introduction to machine learning, allowing you to become quickly confident using this amazing and fast-developing technology that's impacting lives for the better all over the world.

Data Mining and Machine Learning-Mohammed J. Zaki 2020-01-31 New to the second edition of this advanced text are several chapters on regression, including neural networks and deep learning.

Introduction to Natural Language Processing-Jacob Eisenstein 2019-10-01 A survey of computational methods for understanding, generating, and manipulating human language, which offers a synthesis of classical representations and algorithms with contemporary machine learning techniques. This textbook provides a technical perspective on natural language processing—methods for building computer software that understands, generates, and manipulates human language. It emphasizes contemporary data-driven approaches, focusing on techniques from supervised and unsupervised machine learning. The first section establishes a foundation in machine learning by building a set of tools that will be used throughout the book and applying them to word-based textual analysis. The second section introduces structured representations of language, including sequences, trees, and graphs. The third section explores different approaches to the representation and analysis of linguistic meaning, ranging from formal logic to neural word embeddings. The final section offers chapter-length treatments of three transformative applications of natural language processing: information extraction, machine translation, and text generation. End-of-chapter exercises include both paper-and-pencil analysis and software implementation. The text synthesizes and distills a broad and diverse research literature, linking contemporary machine learning techniques with the field's linguistic and computational foundations. It is suitable for use in advanced undergraduate and graduate-level courses and as a reference for software engineers and data scientists. Readers should have a background in computer programming and college-level mathematics. After mastering the material presented, students will have the technical skill to build and analyze novel natural language processing systems and to understand the latest research in the field.

Deep Learning-Ian Goodfellow 2016-11-10 An introduction to a broad range of topics in deep learning, covering mathematical and conceptual background, deep learning techniques used in industry, and research perspectives. "Written by three experts in the field, Deep Learning is the only comprehensive book on the subject." —Elon Musk, cochair of OpenAI; cofounder and CEO of Tesla and SpaceX Deep learning is a form of machine learning that enables computers to learn from experience and understand the world in terms of a hierarchy of concepts. Because the computer gathers knowledge from experience, there is no need for a human computer operator to formally specify all the knowledge that the computer needs. The hierarchy of concepts allows the computer to learn complicated concepts by building them out of simpler ones; a graph of these hierarchies would be many layers deep. This book introduces a broad range of topics in deep learning. The text offers mathematical and conceptual background, covering relevant concepts in linear algebra, probability theory and information theory, numerical computation, and machine learning. It describes deep learning techniques used by practitioners in industry, including deep feedforward networks, regularization, optimization algorithms, convolutional networks, sequence modeling, and practical methodology; and it surveys such applications as natural language processing, speech recognition, computer vision, online recommendation systems, bioinformatics, and videogames. Finally, the book offers research perspectives, covering such theoretical topics as linear factor models, autoencoders, representation learning, structured probabilistic models, Monte Carlo methods, the partition function, approximate inference, and deep generative models. Deep Learning can be used by undergraduate or graduate students planning careers in either industry or research, and by software engineers who want to begin using deep learning in their products or platforms. A website offers supplementary material for both readers and instructors.

Artificial Intelligence-Richard E. Neapolitan 2018-03-12 The first edition of this popular textbook, Contemporary Artificial Intelligence, provided an accessible and student friendly introduction to AI. This fully revised and expanded update, Artificial Intelligence: With an Introduction to Machine Learning, Second Edition, retains the same accessibility and problem-solving approach, while providing new material and methods. The book is divided into five sections that focus on the most useful techniques that have emerged from AI. The first section of the book covers logic-based methods, while the second section focuses on probability-based methods. Emergent intelligence is featured in the third section and explores evolutionary computation and methods based on swarm intelligence. The newest section comes next and provides a detailed overview of neural networks and deep learning. The final section of the book focuses on natural language understanding. Suitable for undergraduate and beginning graduate students, this class-tested textbook provides students and other readers with key AI methods and algorithms for solving challenging problems involving systems that behave intelligently in specialized domains such as medical and software diagnostics, financial decision making, speech and text recognition, genetic analysis, and more.

Boosting-Robert E. Schapire 2012-05-18 An accessible introduction and essential reference for an approach to machine learning that creates highly accurate prediction rules by combining many weak and inaccurate ones. Boosting is an approach to machine learning based on the idea of creating a highly accurate predictor by combining many weak and inaccurate "rules of thumb." A remarkably rich theory has evolved around boosting, with connections to a range of topics, including statistics, game theory, convex optimization, and information geometry. Boosting algorithms have also enjoyed practical success in such fields as biology, vision, and speech processing. At various times in its history, boosting has been perceived as mysterious, controversial, even paradoxical. This book, written by the inventors of the method, brings together, organizes, simplifies, and substantially extends two decades of research on boosting, presenting both theory and applications in a way that is accessible to readers from diverse backgrounds while also providing an authoritative reference for advanced researchers. With its introductory treatment of all material and its inclusion of exercises in every chapter, the book is appropriate for course use as well. The book begins with a general introduction to machine learning algorithms and their analysis; then explores the core theory of boosting, especially its ability to generalize; examines some of the myriad other theoretical viewpoints that help to explain and understand boosting; provides practical extensions of boosting for more complex learning problems; and finally presents a number of advanced theoretical topics. Numerous applications and practical illustrations are offered throughout.

Fundamentals of Machine Learning-Thomas Trappenberg 2019-11-28 Interest in machine learning is exploding worldwide, both in research and for industrial applications. Machine learning is fast becoming a fundamental part of everyday life. This book is a brief introduction to this area - exploring its importance in a range of many disciplines, from science to engineering, and even its broader impact on our society. The book is written in a style that strikes a balance between brevity of explanation, rigorous mathematical argument, and outlines principle ideas. At the same time, it provides a comprehensive overview of a variety of methods and their application within this field. This includes an introduction to Bayesian approaches to modeling, as well as deep learning. Writing small programs to apply machine learning techniques is made easy by high level programming systems, and this book shows examples in Python with the machine learning libraries 'sklearn' and 'Keras'. The first four chapters concentrate on the practical side of applying machine learning techniques. The following four chapters discuss more fundamental concepts that includes their formulation in a probabilistic context. This is followed by two more chapters on advanced models, that of recurrent neural networks and that of reinforcement learning. The book closes with a brief discussion on the impact of machine learning and AI on our society. Fundamentals of Machine Learning provides a brief and accessible introduction to this rapidly growing field, one that will appeal to students and researchers across computer science and computational neuroscience, as well as the broader cognitive sciences.

Foundations of Knowledge Acquisition-Susan Chipman 1993 The two volumes of Foundations of Knowledge Acquisition document the recent progress of basic research in knowledge acquisition sponsored by the Office of Naval Research. This volume is subtitled Cognitive Models of Complex Learning, and there is a companion volume, subtitled Machine Learning. Funding was provided by a five-year Accelerated Research Initiative (ARI), and made possible significant advances in the scientific understanding of how machines and humans can acquire new knowledge so as to exhibit improved problem-solving behavior. Knowledge acquisition, as pursued under the ARI, was a coordinated research thrust into both machine learning and the human learning. Chapters in Cognitive Models of Complex Learning thus include summaries of work by cognitive scientists who do computational modeling of human learning. In fact, an accomplishment of research previously sponsored by ONR's Cognitive Science Program gave insight into the knowledge and skills that distinguish human novices from human experts in various domains; the cognitive interest in the ARI was then to characterize how the transition from novice to expert actually takes place. Chapters particularly relevant to that concern are those written by Anderson, Kieras, Marshall, Ohlsson, and VanLehn. Significant progress in machine learning is reported along in a variety of fronts in the companion volume, Machine Learning, also published by Kluwer Academic Publishers. Included is work in analogical reasoning; induction and discovery; explanation-based learning; learning by competition, using genetic algorithms; learning within natural language systems; theoretical limitations, learning in Soar, a proposed general architecture for intelligent systems; and case-based reasoning. These volumes of Foundations of Knowledge Acquisition are excellent reference sources by bringing together descriptions of recent and ongoing research at the forefront of progress in one the most challenging arenas of artificial intelligence and cognitive science. In addition, contributing authors comment on exciting future directions for research.

High-Dimensional Statistics-Martin J. Wainwright 2019-02-21 A coherent introductory text from a groundbreaking researcher, focusing on clarity and motivation to build intuition and understanding.

The Mathematical Foundations of Learning Machines-Nils J. Nilsson 1990 Neural networks research is unified by contributions from computer science, electrical engineering, physics, statistics, cognitive science and neuroscience. Author Nilsson is recognized for his presentation of intuitive geometric and statistical theories. Annotation copyrighted by Book News, Inc., Portland, OR

Introducing Machine Learning-Dino Esposito 2020-02-05 Master machine learning concepts and develop real-world solutions Machine learning offers immense opportunities, and Introducing Machine Learning delivers practical knowledge to make the most of them. Dino and Francesco Esposito start with a quick overview of the foundations of artificial intelligence and the basic steps of any machine learning project. Next, they introduce Microsoft's powerful ML.NET library, including capabilities for data processing, training, and evaluation. They present families of algorithms that can be trained to solve real-life problems, as well as deep learning techniques utilizing neural networks. The authors conclude by introducing valuable runtime services available through the Azure cloud platform and consider the long-term business vision for machine learning. · 14-time Microsoft MVP Dino Esposito and Francesco Esposito help you · Explore what's known about how humans learn and how intelligent software is built · Discover which problems machine learning can address · Understand the machine learning pipeline: the steps leading to a deliverable model · Use AutoML to automatically select the best pipeline for any problem and dataset · Master ML.NET, implement its pipeline, and apply its tasks and algorithms · Explore the mathematical foundations of machine learning · Make predictions, improve decision-making, and apply probabilistic methods · Group data via classification and clustering · Learn the fundamentals of deep learning, including neural network design · Leverage AI cloud services to build better real-world solutions faster About This Book · For professionals who want to build machine learning applications: both developers who need data science skills and data scientists who need relevant programming skills · Includes examples of machine learning coding scenarios built using the ML.NET library

Statistical Machine Learning-Richard Golden 2020-06-24 The recent rapid growth in the variety and complexity of new machine learning architectures requires the development of improved methods for designing, analyzing, evaluating, and communicating machine learning technologies. Statistical Machine Learning: A Unified Framework provides students, engineers, and scientists with tools from mathematical statistics and nonlinear optimization theory to become experts in the field of machine learning. In particular, the material in this text directly supports the mathematical analysis and design of old, new, and not-yet-invented nonlinear high-dimensional machine learning algorithms. Features: Unified empirical risk minimization framework supports rigorous mathematical analyses of widely used supervised, unsupervised, and reinforcement machine learning algorithms Matrix calculus methods for supporting machine learning analysis and design applications Explicit conditions for ensuring convergence of adaptive, batch, minibatch, MCEM, and MCMC learning algorithms that minimize both unimodal and multimodal objective functions Explicit conditions for characterizing asymptotic properties of M-estimators and model selection criteria such as AIC and BIC in the presence of observable

model misspecification This advanced text is suitable for graduate students or highly motivated undergraduate students in statistics, computer science, electrical engineering, and applied mathematics. The text is self-contained and only assumes knowledge of lower-division linear algebra and upper-division probability theory. Students, professional engineers, and multidisciplinary scientists possessing these minimal prerequisites will find this text challenging yet accessible. About the Author: Richard M. Golden (Ph.D., M.S.E.E., B.S.E.E.) is Professor of Cognitive Science and Participating Faculty Member in Electrical Engineering at the University of Texas at Dallas. Dr. Golden has published articles and given talks at scientific conferences on a wide range of topics in the fields of both statistics and machine learning over the past three decades. His long-term research interests include identifying conditions for the convergence of deterministic and stochastic machine learning algorithms and investigating estimation and inference in the presence of possibly misspecified probability models.

Artificial Intelligence and Machine Learning Fundamentals-Zsolt Nagy 2018-12-12 Create AI applications in Python and lay the foundations for your career in data science Key Features Practical examples that explain key machine learning algorithms Explore neural networks in detail with interesting examples Master core AI concepts with engaging activities Book Description Machine learning and neural networks are pillars on which you can build intelligent applications. Artificial Intelligence and Machine Learning Fundamentals begins by introducing you to Python and discussing AI search algorithms. You will cover in-depth mathematical topics, such as regression and classification, illustrated by Python examples. As you make your way through the book, you will progress to advanced AI techniques and concepts, and work on real-life datasets to form decision trees and clusters. You will be introduced to neural networks, a powerful tool based on Moore's law. By the end of this book, you will be confident when it comes to building your own AI applications with your newly acquired skills! What you will learn Understand the importance, principles, and fields of AI Implement basic artificial intelligence concepts with Python Apply regression and classification concepts to real-world problems Perform predictive analysis using decision trees and random forests Carry out clustering using the k-means and mean shift algorithms Understand the fundamentals of deep learning via practical examples Who this book is for Artificial Intelligence and Machine Learning Fundamentals is for software developers and data scientists who want to enrich their projects with machine learning. You do not need any prior experience in AI. However, it's recommended that you have knowledge of high school-level mathematics and at least one programming language (preferably Python).

Introduction to Machine Learning with Python-Andreas C. Müller 2016-09-26 Machine learning has become an integral part of many commercial applications and research projects, but this field is not exclusive to large companies with extensive research teams. If you use Python, even as a beginner, this book will teach you practical ways to build your own machine learning solutions. With all the data available today, machine learning applications are limited only by your imagination. You'll learn the steps necessary to create a successful machine-learning application with Python and the scikit-learn library. Authors Andreas Müller and Sarah Guido focus on the practical aspects of using machine learning algorithms, rather than the math behind them. Familiarity with the NumPy and matplotlib libraries will help you get even more from this book. With this book, you'll learn: Fundamental concepts and applications of machine learning Advantages and shortcomings of widely used machine learning algorithms How to represent data processed by machine learning, including which data aspects to focus on Advanced methods for model evaluation and parameter tuning The concept of pipelines for chaining models and encapsulating your workflow Methods for working with text data, including text-specific processing techniques Suggestions for improving your machine learning and data science skills

The Elements of Statistical Learning-Trevor Hastie 2013-11-11 During the past decade there has been an explosion in computation and information technology. With it have come vast amounts of data in a variety of fields such as medicine, biology, finance, and marketing. The challenge of understanding these data has led to the development of new tools in the field of statistics, and spawned new areas such as data mining, machine learning, and bioinformatics. Many of these tools have common underpinnings but are often expressed with different terminology. This book describes the important ideas in these areas in a common conceptual framework. While the approach is statistical, the emphasis is on concepts rather than mathematics. Many examples are given, with a liberal use of color graphics. It should be a valuable resource for statisticians and anyone interested in data mining in science or industry. The book's coverage is broad, from supervised learning (prediction) to unsupervised learning. The many topics include neural networks, support vector machines, classification trees and boosting—the first comprehensive treatment of this topic in any book. This major new edition features many topics not covered in the original, including graphical models, random forests, ensemble methods, least angle regression & path algorithms for the lasso, non-negative matrix factorization, and spectral clustering. There is also a chapter on methods for “wide” data (p bigger than n), including multiple testing and false discovery rates. Trevor Hastie, Robert Tibshirani, and Jerome Friedman are professors of statistics at Stanford University. They are prominent researchers in this area: Hastie and Tibshirani developed generalized additive models and wrote a popular book of that title. Hastie co-developed much of the statistical modeling software and environment in R/S-PLUS and invented principal curves and surfaces. Tibshirani proposed the lasso and is co-author of the very successful An Introduction to the Bootstrap. Friedman is the co-inventor of many data-mining tools including CART, MARS, projection pursuit and gradient boosting.

Deep Learning for Coders with fastai and PyTorch-Jeremy Howard 2020-06-29 Deep learning is often viewed as the exclusive domain of math PhDs and big tech companies. But as this hands-on guide demonstrates, programmers comfortable with Python can achieve impressive results in deep learning with little math background, small amounts of data, and minimal code. How? With fastai, the first library to provide a consistent interface to the most frequently used deep learning applications. Authors Jeremy Howard and Sylvain Gugger, the creators of fastai, show you how to train a model on a wide range of tasks using fastai and PyTorch. You'll also dive progressively further into deep learning theory to gain a complete understanding of the algorithms behind the scenes. Train models in computer vision, natural language processing, tabular data, and collaborative filtering Learn the latest deep learning techniques that matter most in practice Improve accuracy, speed, and reliability by understanding how deep learning models work Discover how to turn your models into web applications Implement deep learning algorithms from scratch Consider the ethical implications of your work Gain insight from the foreword by PyTorch cofounder, Soumith Chintala

Practical Deep Learning for Cloud, Mobile, and Edge-Anirudh Koul 2019-10-14 Whether you're a software engineer aspiring to enter the world of deep learning, a veteran data scientist, or a hobbyist with a simple dream of making the next viral AI app, you might have wondered where to begin. This step-by-step guide teaches you how to build practical deep learning applications for the cloud, mobile, browsers, and edge devices using a hands-on approach. Relying on years of industry experience transforming deep learning research into award-winning applications, Anirudh Koul, Siddha Ganju, and Meher Kasam guide you through the process of converting an idea into something that people in the real world can use. Train, tune, and deploy computer vision models with Keras, TensorFlow, Core ML, and TensorFlow Lite Develop AI for a range of devices including Raspberry Pi, Jetson Nano, and Google Coral Explore fun projects, from Silicon Valley's Not Hotdog app to 40+ industry case studies Simulate an autonomous car in a video game environment and build a miniature version with reinforcement learning Use transfer learning to train models in minutes Discover 50+ practical tips for maximizing model accuracy and speed, debugging, and scaling to millions of users

Imbalanced Learning-Haibo He 2013-06-07 The first book of its kind to review the current status and future direction of the exciting new branch of machine learning/data mining called imbalanced learning Imbalanced learning focuses on how an intelligent system can learn when it is provided with imbalanced data. Solving imbalanced learning problems is critical in numerous data-intensive networked systems, including surveillance, security, Internet, finance, biomedical, defense, and more. Due to the inherent complex characteristics of imbalanced data sets, learning from such data requires new understandings, principles, algorithms, and tools to transform vast amounts of raw data efficiently into information and knowledge representation. The first comprehensive look at this new branch of machine learning, this book offers a critical review of the problem of imbalanced learning, covering the state of the art in techniques, principles, and real-world applications. Featuring contributions from experts in both academia and industry, Imbalanced Learning: Foundations, Algorithms, and Applications provides chapter coverage on: Foundations of Imbalanced Learning Imbalanced Datasets: From Sampling to Classifiers Ensemble Methods for Class Imbalance Learning Class Imbalance Learning Methods for Support Vector Machines Class Imbalance and Active Learning Nonstationary Stream Data Learning with Imbalanced Class Distribution Assessment Metrics for Imbalanced Learning Imbalanced Learning: Foundations, Algorithms, and Applications will help scientists and engineers learn how to tackle the problem of learning from imbalanced datasets, and gain insight into current developments in the field as well as future research directions.

Algorithmic Aspects of Machine Learning-Ankur Moitra 2018-09-27 Introduces cutting-edge research on machine learning theory and practice, providing an accessible, modern algorithmic toolkit.

Deep Reinforcement Learning Hands-On-Maxim Lapan 2020-01-31 New edition of the bestselling guide to deep reinforcement learning and how it's used to solve complex real-world problems. Revised and expanded to include multi-agent methods, discrete optimization, RL in robotics, advanced exploration techniques, and more Key Features Second edition of the bestselling introduction to deep reinforcement learning, expanded with six new chapters Learn advanced exploration techniques including noisy networks, pseudo-count, and network distillation methods Apply RL methods to cheap hardware robotics platforms Book Description Deep Reinforcement Learning Hands-On, Second Edition is an updated and expanded version of the bestselling guide to the very latest reinforcement learning (RL) tools and techniques. It provides you with an introduction to the fundamentals of RL, along with the hands-on ability to code intelligent learning agents to perform a range of practical tasks. With six new chapters devoted to a variety of up-to-the-minute developments in RL, including discrete optimization (solving the Rubik's Cube), multi-agent methods, Microsoft's TextWorld environment, advanced exploration techniques, and more, you will come away from this book with a deep understanding of the latest innovations in this emerging field. In addition, you will gain actionable insights into such topic areas as deep Q-networks, policy gradient methods, continuous control problems, and highly scalable, non-gradient methods. You will also discover how to build a real hardware robot trained with RL for less than \$100 and solve the Pong environment in just 30 minutes of training using step-by-step code optimization. In short, Deep Reinforcement Learning Hands-On, Second Edition, is your companion to navigating the exciting complexities of RL as it helps you attain experience and knowledge through real-world examples. What you will learn Understand the deep learning context of RL and implement complex deep learning models Evaluate RL methods including cross-entropy, DQN, actor-critic, TRPO, PPO, DDPG, D4PG, and others Build a practical hardware robot trained with RL methods for less than \$100 Discover Microsoft's TextWorld environment, which is an interactive fiction games platform Use discrete optimization in RL to solve a Rubik's Cube Teach your agent to play Connect 4 using AlphaGo Zero Explore the very latest deep RL research on topics including AI chatbots Discover advanced exploration techniques, including noisy networks and network distillation techniques Who this book is for Some fluency in Python is assumed. Sound understanding of the fundamentals of deep learning will be helpful. This book is an introduction to deep RL and requires no background in RL.

Deep Learning with Python-Francois Chollet 2017-10-28 Summary Deep Learning with Python introduces the field of deep learning using the Python language and the powerful Keras library. Written by Keras creator and Google AI researcher François Chollet, this book builds your understanding through intuitive explanations and practical examples. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Machine learning has made remarkable progress in recent years. We went from near-unusable speech and image recognition, to near-human accuracy. We went from machines that couldn't beat a serious Go player, to defeating a world champion. Behind this progress is deep learning—a combination of engineering advances, best practices, and theory that enables a wealth of previously impossible smart applications. About the Book Deep Learning with Python introduces the field of deep learning using the Python language and the powerful Keras library. Written by Keras creator and Google AI researcher François Chollet, this book builds your understanding through intuitive explanations and practical examples. You'll explore challenging concepts and practice with applications in computer vision, natural-language processing, and generative models. By the time you finish, you'll have the knowledge and hands-on skills to apply deep learning in your own projects. What's Inside Deep learning from first principles Setting up your own deep-learning environment Image-classification models Deep learning for text and sequences Neural style transfer, text generation, and image generation About the Reader Readers need intermediate Python skills. No previous experience with Keras, TensorFlow, or machine learning is required. About the Author François Chollet works on deep learning at Google in Mountain View, CA. He is the creator of the Keras deep-learning library, as well as a contributor to the TensorFlow machine-learning framework. He also does deep-learning research, with a focus on computer vision and the application of machine learning to formal reasoning. His papers have been published at major conferences in the field, including the Conference on Computer Vision and Pattern Recognition (CVPR), the Conference and Workshop on Neural Information Processing Systems (NIPS), the International Conference on Learning Representations (ICLR), and others. Table of Contents PART 1 - FUNDAMENTALS OF DEEP LEARNING What is deep learning? Before we begin: the mathematical building blocks of neural networks Getting started with neural networks Fundamentals of machine learning PART 2 - DEEP LEARNING IN PRACTICE Deep learning for computer vision Deep learning for text and sequences Advanced deep-learning best practices Generative deep learning Conclusions appendix A - Installing Keras and its dependencies on Ubuntu appendix B - Running Jupyter notebooks on an EC2 GPU instance

Learning from Data-Vladimir Cherkassky 2007-09-10 An interdisciplinary framework for learning methodologies—covering statistics, neural networks, and fuzzy logic, this book provides a unified treatment of the principles and methods for learning dependencies from data. It establishes a general conceptual framework in which various learning methods from statistics, neural networks, and fuzzy logic can be applied—showing that a few fundamental principles underlie most new methods being proposed today in statistics, engineering, and computer science. Complete with over one hundred illustrations, case studies, and examples making this an invaluable text.