

Foundations Of Algorithms

A concise and self-contained introduction to causal inference, increasingly important in data science and machine learning. 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.

A new and refreshingly different approach to presenting the foundations of statistical algorithms, *Foundations of Statistical Algorithms: With References to R Packages* reviews the historical development of basic algorithms to illuminate the evolution of today's more powerful statistical algorithms. It emphasizes recurring themes in all statistical algorithms, including computation, assessment and verification, iteration, intuition, randomness, repetition and parallelization, and scalability. Unique in scope, the book reviews the upcoming challenge of scaling many of the established techniques to very large data sets and delves into systematic verification by demonstrating how to derive general classes of worst case inputs and emphasizing the importance of testing over a large number of different inputs. Broadly accessible, the book offers examples, exercises, and selected solutions in each chapter as well as access to a supplementary website. After working through the material covered in the book, readers should not only understand current algorithms but also gain a deeper understanding of how algorithms are constructed, how to evaluate new algorithms, which recurring principles are used to tackle some of the tough problems statistical programmers face, and how to take an idea for a new method and turn it into something practically useful.

For a long time, human beings have dreamed of a virtual world where it is possible to interact with synthetic entities as if they were real. It has been shown that the ability to touch virtual objects increases the sense of presence in virtual

environments. This book provides an authoritative overview of state-of-the-art haptic rendering algorithms. Michael Goodrich and Roberto Tamassia, authors of the successful, *Data Structures and Algorithms in Java, 2/e*, have written *Algorithm Engineering*, a text designed to provide a comprehensive introduction to the design, implementation and analysis of computer algorithms and data structures from a modern perspective. This book offers theoretical analysis techniques as well as algorithmic design patterns and experimental methods for the engineering of algorithms. Market: Computer Scientists; Programmers.

Stochastic local search (SLS) algorithms are among the most prominent and successful techniques for solving computationally difficult problems. Offering a systematic treatment of SLS algorithms, this book examines the general concepts and specific instances of SLS algorithms and considers their development, analysis and application. *Foundations of Algorithms, Fifth Edition* offers a well-balanced presentation of algorithm design, complexity analysis of algorithms, and computational complexity. Ideal for any computer science students with a background in college algebra and discrete structures, the text presents mathematical concepts using standard English and simple notation to maximize accessibility and user-friendliness. Concrete examples, appendices reviewing essential mathematical concepts, and a student-focused approach reinforce theoretical explanations and promote learning and retention. C++ and Java pseudocode help students better understand complex algorithms. A chapter on numerical algorithms includes a review of basic number theory, Euclid's Algorithm for finding the greatest common divisor, a review of modular arithmetic, an algorithm for solving modular linear equations, an algorithm for computing modular powers, and the new polynomial-time algorithm for determining whether a number is prime. The revised and updated Fifth Edition features an all-new chapter on genetic algorithms and genetic programming, including approximate solutions to the traveling salesperson problem, an algorithm for an artificial ant that navigates along a trail of food, and an application to financial trading. With fully updated exercises and examples throughout and improved instructor resources including complete solutions, an Instructor's Manual and PowerPoint lecture outlines, *Foundations of Algorithms* is an essential text for undergraduate and graduate courses in the design and analysis of algorithms. Key features include:

- The only text of its kind with a chapter on genetic algorithms
- Use of C++ and Java pseudocode to help students better understand complex algorithms
- No calculus background required
- Numerous clear and student-friendly examples throughout the text
- Fully updated exercises and examples throughout
- Improved instructor resources, including complete solutions, an Instructor's Manual, and PowerPoint lecture outlines

Discrete Mathematics has permeated the whole of mathematics so much so it has now come to be taught even at the high school level. This book presents the basics of Discrete Mathematics and its applications to day-to-day problems in

several areas. This book is intended for undergraduate students of Computer Science, Mathematics and Engineering. A number of examples have been given to enhance the understanding of concepts. The programming languages used are Pascal and C.

Formal Design Theory (PDT) is a mathematical theory of design. The main goal of PDT is to develop a domain independent core model of the design process. The book focuses the reader's attention on the process by which ideas originate and are developed into workable products. In developing PDT, we have been striving toward what has been expressed by the distinguished scholar Simon (1969): that "the science of design is possible and some day we will be able to talk in terms of well-established theories and practices. " The book is divided into five interrelated parts. The conceptual approach is presented first (Part I); followed by the theoretical foundations of PDT (Part II), and from which the algorithmic and pragmatic implications are deduced (Part III). Finally, detailed case-studies illustrate the theory and the methods of the design process (Part IV), and additional practical considerations are evaluated (Part V). The generic nature of the concepts, theory and methods are validated by examples from a variety of disciplines. FDT explores issues such as: algebraic representation of design artifacts, idealized design process cycle, and computational analysis and measurement of design process complexity and quality. FDT's axioms convey the assumptions of the theory about the nature of artifacts, and potential modifications of the artifacts in achieving desired goals or functionality. By being able to state these axioms explicitly, it is possible to derive theorems and corollaries, as well as to develop specific analytical and constructive methodologies.

To truly understand how the Internet and Web are organized and function requires knowledge of mathematics and computation theory. Mathematical and Algorithmic Foundations of the Internet introduces the concepts and methods upon which computer networks rely and explores their applications to the Internet and Web. The book offers a unique approach to mathematical and algorithmic concepts, demonstrating their universality by presenting ideas and examples from various fields, including literature, history, and art. Progressing from fundamental concepts to more specific topics and applications, the text covers computational complexity and randomness, networks and graphs, parallel and distributed computing, and search engines. While the mathematical treatment is rigorous, it is presented at a level that can be grasped by readers with an elementary mathematical background. The authors also present a lighter side to this complex subject by illustrating how many of the mathematical concepts have counterparts in everyday life. The book provides in-depth coverage of the mathematical prerequisites and assembles a complete presentation of how computer networks function. It is a useful resource for anyone interested in the inner functioning, design, and organization of the Internet.

Intro Computer Science (CS0)

The objectives of this book are to provide a solid foundation for the theory of algorithms. The book is intended to serve as a text for

a core upper division undergraduate course in the design and analysis of algorithms.

Data Structures & Theory of Computation

This book probes the stable marriage problem and its variants as a rich source of problems and ideas that illustrate both the design and analysis of efficient algorithms. It covers the most recent structural and algorithmic work on stable matching problems, simplifies and unifies many earlier proofs, strengthens several earlier results, and presents new results and more efficient algorithms. The authors develop the structure of the set of stable matchings in the stable marriage problem in a more general and algebraic context than has been done previously; they discuss the problem's structure in terms of rings of sets, which allows many of the most useful features to be seen as features of a more general set of problems. The relationship between the structure of the stable marriage problem and the more general stable roommates problem is demonstrated, revealing many commonalities. The results the authors obtain provide an algorithmic response to the practical, and political, problems created by the asymmetry inherent in the Gale Shapley solutions, leading to alternative methods and better compromises than are provided by the Gale Shapley method. And, in contrast to Donald Knuth's earlier work which primarily focused on the application of mathematics to the analysis of algorithms, this book illustrates the productive and almost inseparable relationship between mathematical insight and the design of efficient algorithms. Dan Gusfield is Associate Professor of Computer Science at the University of California, Davis. Robert W. Irving is Senior Lecturer in Computing Science at the University of Glasgow. The Stable Marriage Problem is included in the Foundations of Computing Series, edited by Michael Garey and Albert Meyer.

Algorithms and Data Structures for External Memory describes several useful paradigms for the design and implementation of efficient external memory (EM) algorithms and data structures. The problem domains considered include sorting, permuting, FFT, scientific computing, computational geometry, graphs, databases, geographic information systems, and text and string processing. This book is devoted to the most difficult part of concurrent programming, namely synchronization concepts, techniques and principles when the cooperating entities are asynchronous, communicate through a shared memory, and may experience failures. Synchronization is no longer a set of tricks but, due to research results in recent decades, it relies today on sane scientific foundations as explained in this book. In this book the author explains synchronization and the implementation of concurrent objects, presenting in a uniform and comprehensive way the major theoretical and practical results of the past 30 years. Among the key features of the book are a new look at lock-based synchronization (mutual exclusion, semaphores, monitors, path expressions); an introduction to the atomicity consistency criterion and its properties and a specific chapter on transactional memory; an introduction to mutex-freedom and associated progress conditions such as obstruction-freedom and wait-freedom; a presentation of Lamport's hierarchy of safe, regular and atomic registers and associated wait-free constructions; a description of numerous wait-free constructions of concurrent objects (queues, stacks, weak counters, snapshot objects, renaming objects, etc.); a presentation of the computability power of concurrent objects including the notions of universal construction, consensus number and the associated Herlihy's hierarchy; and a survey of failure detector-based constructions of consensus objects. The book is

suitable for advanced undergraduate students and graduate students in computer science or computer engineering, graduate students in mathematics interested in the foundations of process synchronization, and practitioners and engineers who need to produce correct concurrent software. The reader should have a basic knowledge of algorithms and operating systems. 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.

The problem of privacy-preserving data analysis has a long history spanning multiple disciplines. As electronic data about individuals becomes increasingly detailed, and as technology enables ever more powerful collection and curation of these data, the need increases for a robust, meaningful, and mathematically rigorous definition of privacy, together with a computationally rich class of algorithms that satisfy this definition. Differential Privacy is such a definition. *The Algorithmic Foundations of Differential Privacy* starts out by motivating and discussing the meaning of differential privacy, and proceeds to explore the fundamental techniques for achieving differential privacy, and the application of these techniques in creative combinations, using the query-release problem as an ongoing example. A key point is that, by rethinking the computational goal, one can often obtain far better results than would be achieved by methodically replacing each step of a non-private computation with a differentially private implementation. Despite some powerful computational results, there are still fundamental limitations. Virtually all the algorithms discussed herein maintain differential privacy against adversaries of arbitrary computational power -- certain algorithms are computationally intensive, others are efficient. Computational complexity for the adversary and the algorithm are both discussed. The monograph then turns from fundamentals to applications other than query-release, discussing differentially private methods for mechanism design and machine learning. The vast majority of the literature on differentially private algorithms considers a single, static, database that is subject to many analyses. Differential privacy in other models, including distributed databases and

computations on data streams, is discussed. The Algorithmic Foundations of Differential Privacy is meant as a thorough introduction to the problems and techniques of differential privacy, and is an invaluable reference for anyone with an interest in the topic.

Computational finance is an interdisciplinary field which joins financial mathematics, stochastics, numerics and scientific computing. Its task is to estimate as accurately and efficiently as possible the risks that financial instruments generate. This volume consists of a series of cutting-edge surveys of recent developments in the field written by leading international experts. These make the subject accessible to a wide readership in academia and financial businesses. The book consists of 13 chapters divided into 3 parts: foundations, algorithms and applications. Besides surveys of existing results, the book contains many new previously unpublished results.

Foundations of Algorithms Using C++ Pseudocode, Third Edition offers a well-balanced presentation on designing algorithms, complexity analysis of algorithms, and computational complexity. The volume is accessible to mainstream computer science students who have a background in college algebra and discrete structures. To support their approach, the authors present mathematical concepts using standard English and a simpler notation than is found in most texts. A review of essential mathematical concepts is presented in three appendices. The authors also reinforce the explanations with numerous concrete examples to help students grasp theoretical concepts.

This book is Part I of the fourth edition of Robert Sedgwick and Kevin Wayne's Algorithms , the leading textbook on algorithms today, widely used in colleges and universities worldwide. Part I contains Chapters 1 through 3 of the book. The fourth edition of Algorithms surveys the most important computer algorithms currently in use and provides a full treatment of data structures and algorithms for sorting, searching, graph processing, and string processing -- including fifty algorithms every programmer should know. In this edition, new Java implementations are written in an accessible modular programming style, where all of the code is exposed to the reader and ready to use. The algorithms in this book represent a body of knowledge developed over the last 50 years that has become indispensable, not just for professional programmers and computer science students but for any student with interests in science, mathematics, and engineering, not to mention students who use computation in the liberal arts. The companion web site, algs4.cs.princeton.edu contains An online synopsis Full Java implementations Test data Exercises and answers Dynamic visualizations Lecture slides Programming assignments with checklists Links to related material The MOOC related to this book is accessible via the "Online Course" link at algs4.cs.princeton.edu. The course offers more than 100 video lecture segments that are integrated with the text, extensive online assessments, and the large-scale discussion forums that have proven so valuable. Offered each fall and spring, this course regularly attracts tens of thousands of registrants. Robert Sedgwick and Kevin Wayne are developing a modern approach to disseminating knowledge that fully embraces technology, enabling people all around the world to discover new ways of learning and teaching. By integrating their textbook, online content, and MOOC, all at the state of the art, they have built a unique resource that greatly expands the breadth and depth of the educational experience.

This book offers a well-balanced presentation on designing algorithms, complexity analysis of algorithms, and computational complexity that is accessible to mainstream computer science students who have a background in college algebra and discrete structures.

In this second book of what will be a four-volume series, the authors present, in a mathematically rigorous way, the essential foundations of both the theory and practice of algorithms, approximation, and optimization—essential topics in modern applied and computational mathematics. This material is the introductory framework upon which algorithm analysis, optimization, probability, statistics, machine learning, and control theory are built. This text gives a unified treatment of several topics that do not usually appear together: the theory and analysis of algorithms for mathematicians and data science students; probability and its applications; the theory and applications of approximation, including Fourier series, wavelets, and polynomial approximation; and the theory and practice of optimization, including dynamic optimization. When used in concert with the free supplemental lab materials, Foundations of Applied Mathematics, Volume 2: Algorithms, Approximation, Optimization teaches not only the theory but also the computational practice of modern mathematical methods. Exercises and examples build upon each other in a way that continually reinforces previous ideas, allowing students to retain learned concepts while achieving a greater depth. The mathematically rigorous lab content guides students to technical proficiency and answers the age-old question “When am I going to use this?” This textbook is geared toward advanced undergraduate and beginning graduate students in mathematics, data science, and machine learning.

Foundations of Algorithms, Fifth Edition offers a well-balanced presentation of algorithm design, complexity analysis of algorithms, and computational complexity. Ideal for any computer science students with a background in college algebra and discrete structures, the text presents mathematical concepts using standard English and simple notation to maximize accessibility and user-friendliness. Concrete examples, appendices reviewing essential mathematical concepts, and a student-focused approach reinforce theoretical explanations and promote learning and retention. C++ and Java pseudocode help students better understand complex algorithms. A chapter on numerical algorithms includes a review of basic number theory, Euclid's Algorithm for finding the greatest common divisor, a review of modular arithmetic, an algorithm for solving modular linear equations, an algorithm for computing modular powers, and the new polynomial-time algorithm for determining whether a number is prime. The revised and updated Fifth Edition features an all-new chapter on genetic algorithms and genetic programming, including approximate solutions to the traveling salesperson problem, an algorithm for an artificial ant that navigates along a trail of food, and an application to financial trading. With fully updated exercises and examples throughout and improved instructor resources including complete solutions, an Instructor's Manual and PowerPoint lecture outlines, Foundations of Algorithms is an essential text for undergraduate and graduate courses in the design and analysis of algorithms. Key features include: The only text of its kind with a chapter on genetic algorithms Use of C++ and Java pseudocode to help students better understand complex algorithms No calculus background required Numerous clear and student-friendly examples throughout the text Fully updated exercises and examples

throughout Improved instructor resources, including complete solutions, an Instructor s Manual, and PowerPoint lecture outlines" This treatment focuses on the analysis and algebra underlying the workings of convexity and duality and necessary/sufficient local/global optimality conditions for unconstrained and constrained optimization problems. 2015 edition.

Volume 1.

This book reviews the algorithms for processing geometric data, with a practical focus on important techniques not covered by traditional courses on computer vision and computer graphics. Features: presents an overview of the underlying mathematical theory, covering vector spaces, metric space, affine spaces, differential geometry, and finite difference methods for derivatives and differential equations; reviews geometry representations, including polygonal meshes, splines, and subdivision surfaces; examines techniques for computing curvature from polygonal meshes; describes algorithms for mesh smoothing, mesh parametrization, and mesh optimization and simplification; discusses point location databases and convex hulls of point sets; investigates the reconstruction of triangle meshes from point clouds, including methods for registration of point clouds and surface reconstruction; provides additional material at a supplementary website; includes self-study exercises throughout the text.

Introduces machine learning and its algorithmic paradigms, explaining the principles behind automated learning approaches and the considerations underlying their usage.

This book presents extensive research on two main problems in robotics: the path planning problem and the multi-robot task allocation problem. It is the first book to provide a comprehensive solution for using these techniques in large-scale environments containing randomly scattered obstacles. The research conducted resulted in tangible results both in theory and in practice. For path planning, new algorithms for large-scale problems are devised and implemented and integrated into the Robot Operating System (ROS). The book also discusses the parallelism advantage of cloud computing techniques to solve the path planning problem, and, for multi-robot task allocation, it addresses the task assignment problem and the multiple traveling salesman problem for mobile robots applications. In addition, four new algorithms have been devised to investigate the cooperation issues with extensive simulations and comparative performance evaluation. The algorithms are implemented and simulated in MATLAB and Webots.

This is a comprehensive overview of the basics of fuzzy control, which also brings together some recent research results in soft computing, in particular fuzzy logic using genetic algorithms and neural networks. This book offers researchers not only a solid background but also a snapshot of the current state of the art in this field.

Foundations of Algorithms Jones & Bartlett Learning

Goals of the Book Over the last thirty years there has been a revolution in diagnostic radiology as a result

of the emergence of computerized tomography (CT), which is the process of obtaining the density distribution within the human body from multiple x-ray projections. Since an enormous variety of possible density values may occur in the body, a large number of projections are necessary to ensure the accurate reconstruction of their distribution. There are other situations in which we desire to reconstruct an object from its projections, but in which we know that the object to be reconstructed has only a small number of possible values. For example, a large fraction of objects scanned in industrial CT (for the purpose of nondestructive testing or reverse engineering) are made of a single material and so the ideal reconstruction should contain only two values: zero for air and the value associated with the material composing the object. Similar assumptions may even be made for some specific medical applications; for example, in angiography of the heart chambers the value is either zero (indicating the absence of dye) or the value associated with the dye in the chamber. Another example arises in the electron microscopy of biological macromolecules, where we may assume that the object to be reconstructed is composed of ice, protein, and RNA. One can also apply electron microscopy to determine the presence or absence of atoms in crystalline structures, which is again a two-valued situation.

The advent of multi-core architectures and cloud-computing has brought parallel programming into the mainstream of software development. Unfortunately, writing scalable parallel programs using traditional lock-based synchronization primitives is well known to be a hard, time-consuming and error-prone task, mastered by only a minority of specialized programmers. Building on the familiar abstraction of atomic transactions, Transactional Memory (TM) promises to free programmers from the complexity of conventional synchronization schemes, simplifying the development and verification of concurrent programs, enhancing code reliability, and boosting productivity. Over the last decade TM has been subject to intense research on a broad range of aspects including hardware and operating systems support, language integration, as well as algorithms and theoretical foundations. On the industrial side, the major players of the software and hardware markets have been up-front in the research and development of prototypical products providing support for TM systems. This has recently led to the introduction of hardware TM implementations on mainstream commercial microprocessors and to the integration of TM support for the world's leading open source compiler. In such a vast interdisciplinary domain, the Euro-TM COST Action (IC1001) has served as a catalyzer and a bridge for the various research communities looking at disparate, yet subtly interconnected, aspects of TM. This book emerged from the idea having Euro-TM experts compile recent results in the TM area in a single and consistent volume. Contributions have been carefully selected and revised to provide a broad coverage of several fundamental issues associated with the design and implementation of TM systems, including their theoretical underpinnings and algorithmic foundations, programming language integration and verification tools, hardware supports, distributed TM systems, self-tuning mechanisms, as well

as lessons learnt from building complex TM-based applications.

This book presents a systematic approach to analyze nature-inspired algorithms. Beginning with an introduction to optimization methods and algorithms, this book moves on to provide a unified framework of mathematical analysis for convergence and stability. Specific nature-inspired algorithms include: swarm intelligence, ant colony optimization, particle swarm optimization, bee-inspired algorithms, bat algorithm, firefly algorithm, and cuckoo search. Algorithms are analyzed from a wide spectrum of theories and frameworks to offer insight to the main characteristics of algorithms and understand how and why they work for solving optimization problems. In-depth mathematical analyses are carried out for different perspectives, including complexity theory, fixed point theory, dynamical systems, self-organization, Bayesian framework, Markov chain framework, filter theory, statistical learning, and statistical measures. Students and researchers in optimization, operations research, artificial intelligence, data mining, machine learning, computer science, and management sciences will see the pros and cons of a variety of algorithms through detailed examples and a comparison of algorithms.

An up-to-date, self-contained introduction to a state-of-the-art machine learning approach, Ensemble Methods: Foundations and Algorithms shows how these accurate methods are used in real-world tasks. It gives you the necessary groundwork to carry out further research in this evolving field. After presenting background and terminology, the book covers the main algorithms and theories, including Boosting, Bagging, Random Forest, averaging and voting schemes, the Stacking method, mixture of experts, and diversity measures. It also discusses multiclass extension, noise tolerance, error-ambiguity and bias-variance decompositions, and recent progress in information theoretic diversity. Moving on to more advanced topics, the author explains how to achieve better performance through ensemble pruning and how to generate better clustering results by combining multiple clusterings. In addition, he describes developments of ensemble methods in semi-supervised learning, active learning, cost-sensitive learning, class-imbalance learning, and comprehensibility enhancement.

This is a central topic in any computer science curriculum. To distinguish this textbook from others, the author considers probabilistic methods as being fundamental for the construction of simple and efficient algorithms, and in each chapter at least one problem is solved using a randomized algorithm. Data structures are discussed to the extent needed for the implementation of the algorithms. The specific algorithms examined were chosen because of their wide field of application. This book originates from lectures for undergraduate and graduate students. The text assumes experience in programming algorithms, especially with elementary data structures such as chained lists, queues, and stacks. It also assumes familiarity with mathematical methods, although the author summarizes some basic notations and results from probability theory and related mathematical terminology in the appendices. He includes many examples to explain the individual steps of the algorithms, and he concludes each chapter with numerous exercises.

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.

A systematic survey of many of these recent results on Gossip network algorithms.

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.

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.

Providing a unique approach to machine learning, this text contains fresh and intuitive, yet rigorous, descriptions of all fundamental concepts necessary to conduct research, build products, tinker, and play. By prioritizing geometric intuition, algorithmic thinking, and practical real world applications in disciplines including computer vision, natural language processing, economics, neuroscience, recommender systems, physics, and biology, this text provides readers with both a lucid understanding of foundational material as well as the practical tools needed to solve real-world problems. With in-depth Python and MATLAB/OCTAVE-based computational exercises and a complete treatment of cutting edge numerical optimization techniques, this is an essential resource for students and an ideal reference for researchers and practitioners working in machine learning, computer science, electrical engineering, signal processing, and numerical optimization.

[Copyright: e7839b73033c90deb1d949082be85af0](#)