

Additional Exercises Convex Optimization Solution Boyd

An insightful, concise, and rigorous treatment of the basic theory of convex sets and functions in finite dimensions, and the analytical/geometrical foundations of convex optimization and duality theory. Convexity theory is first developed in a simple accessible manner, using easily visualized proofs. Then the focus shifts to a transparent geometrical line of analysis to develop the fundamental duality between descriptions of convex functions in terms of points, and in terms of hyperplanes. Finally, convexity theory and abstract duality are applied to problems of constrained optimization, Fenchel and conic duality, and game theory to develop the sharpest possible duality results within a highly visual geometric framework. This on-line version of the book, includes an extensive set of theoretical problems with detailed high-quality solutions, which significantly extend the range and value of the book. The book may be used as a text for a theoretical convex optimization course; the author has taught several variants of such a course at MIT and elsewhere over the last ten years. It may also be used as a supplementary source for nonlinear programming classes, and as a theoretical foundation for classes focused on convex optimization models (rather than theory). It is an excellent supplement to several of our books: Convex Optimization Algorithms (Athena Scientific, 2015), Nonlinear Programming (Athena Scientific, 2017), Network Optimization (Athena Scientific, 1998), Introduction to Linear Optimization (Athena Scientific, 1997), and Network Flows and Monotropic Optimization (Athena Scientific, 1998).

This accessible textbook demonstrates how to recognize, simplify, model and solve optimization problems - and apply these principles to new projects.

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.

Proximal Algorithms discusses proximal operators and proximal algorithms, and illustrates their applicability to standard and distributed convex optimization in general and many applications of recent interest in particular. Much like Newton's method is a standard tool for solving unconstrained smooth optimization problems of modest size, proximal algorithms can be viewed as an analogous tool for nonsmooth, constrained, large-scale, or distributed versions of these problems. They are very generally applicable, but are especially well-suited to problems of substantial recent interest involving large or high-dimensional datasets. Proximal methods sit at a higher level of abstraction than classical algorithms like Newton's method: the base operation is evaluating the proximal operator of a function, which itself involves solving a small convex optimization problem. These subproblems, which generalize the problem of projecting a point onto a convex set, often admit closed-form solutions or can be solved very quickly with standard or simple specialized methods. Proximal Algorithms discusses different interpretations of proximal operators and algorithms, looks at their connections to many other topics in optimization and applied mathematics, surveys some popular algorithms, and provides a large number of examples of proximal operators that commonly arise in practice.

Convex Optimization in Normed Spaces

Convexity and Well-Posed Problems

Statistical Inference via Convex Optimization

Optimization for Machine Learning

From Theory to Algorithms

Understanding Machine Learning

Convex optimization has an increasing impact on many areas of mathematics, applied sciences, and practical applications. It is now being taught at many universities and being used by researchers of different fields. As convex analysis is the mathematical f

Optimization is an important tool used in decision science and for the analysis of physical systems used in engineering. One can trace its roots to the Calculus of Variations and the work of Euler and Lagrange. This natural and reasonable approach to mathematical programming covers numerical methods for finite-dimensional optimization problems. It begins with very simple ideas progressing through more complicated concepts, concentrating on methods for both unconstrained and constrained optimization.

A uniquely pedagogical, insightful, and rigorous treatment of the analytical/geometrical foundations of optimization. The book provides a comprehensive development of convexity theory, and its rich applications in optimization, including duality, minimax/saddle point theory, Lagrange multipliers, and Lagrangian relaxation/nondifferentiable optimization. It is an excellent supplement to several of our books: Convex Optimization Theory (Athena Scientific, 2009), Convex Optimization Algorithms (Athena Scientific, 2015), Nonlinear Programming (Athena Scientific, 2016), Network Optimization (Athena Scientific, 1998), and Introduction to Linear Optimization (Athena Scientific, 1997). Aside from a thorough account of convex analysis and optimization, the book aims to restructure the theory of the subject, by introducing several novel unifying lines of analysis, including: 1) A unified development of minimax theory and constrained optimization duality as special cases of duality between two simple geometrical problems. 2) A unified development of conditions for existence of solutions of convex optimization problems, conditions for the minimax equality to hold, and conditions for the absence of a duality gap in constrained optimization. 3) A unification of the major constraint qualifications allowing the use of Lagrange multipliers for nonconvex constrained optimization, using the notion of constraint pseudonormality and an enhanced form of the Fritz John necessary optimality conditions. Among its features the book: a) Develops rigorously and comprehensively the theory of convex sets and functions, in the classical tradition of Fenchel and Rockafellar b) Provides a geometric, highly visual treatment of convex and nonconvex optimization problems, including existence of solutions, optimality conditions, Lagrange multipliers, and duality c) Includes an insightful and comprehensive presentation of minimax theory and zero sum games, and its connection with duality d) Describes dual optimization, the associated computational methods, including the novel incremental subgradient methods, and applications in linear, quadratic, and integer programming e) Contains many examples, illustrations, and exercises with complete solutions (about 200 pages) posted at the publisher's web site <http://www.athenasc.com/convexity.html>

In the last few years, Algorithms for Convex Optimization have revolutionized algorithm design, both for discrete and continuous optimization problems. For problems like maximum flow, maximum matching, and submodular function minimization, the fastest algorithms involve essential methods such as gradient descent, mirror descent, interior point methods, and ellipsoid methods. The goal of this self-contained book is to enable researchers and professionals in computer science, data science, and machine learning to gain an in-depth understanding of these algorithms. The text emphasizes how to derive key algorithms for convex optimization from first principles and how to establish precise running time bounds. This modern text explains the success of these algorithms in problems of discrete optimization, as well as how these methods have significantly pushed the state of the art of convex optimization itself.

A Textbook

Analysis, Algorithms, and Engineering Applications

Introduction to Applied Linear Algebra

Semidefinite Optimization and Convex Algebraic Geometry

Convex Optimization

Mathematics for Machine Learning

This book provides the foundations of the theory of nonlinear optimization as well as some related algorithms and presents a variety of applications from diverse areas of applied sciences. The author combines three pillars of optimization—theoretical and algorithmic foundation, familiarity with various applications, and the ability to apply the theory and algorithms on actual problems—and rigorously and gradually builds the connection between theory, algorithms, applications, and implementation. Readers will find more than 170 theoretical, algorithmic, and numerical exercises that deepen and enhance the reader's understanding of the topics. The author includes offers several subjects not typically found in optimization books—for example, optimality conditions in sparsity-constrained optimization, hidden convexity, and total least squares. The book also offers a large number of applications discussed theoretically and algorithmically, such as circle fitting, Chebyshev center, the Fermat-Weber problem, denoising, clustering, total least squares, and orthogonal regression and theoretical and algorithmic topics demonstrated by the MATLAB toolbox CVX and a package of m-files that is posted on the book's web site.

This textbook offers graduate students a concise introduction to the classic notions of convex optimization. Written in a highly accessible style and including numerous examples and illustrations, it presents everything readers need to know about convexity and convex optimization. The book introduces a systematic three-step method for doing everything, which can be summarized as "conify, work, deconify". It starts with the concept of convex sets, their primal description, constructions, topological properties and dual description, and then moves on to convex functions and the fundamental principles of convex optimization and their use in the complete analysis of convex optimization problems by means of a systematic four-step method. Lastly, it includes chapters on alternative formulations of optimality conditions and on illustrations of their use. "The author deals with the delicate subjects in a precise yet light-minded spirit... For experts in the field, this book not only offers a unifying view, but also opens a door to new discoveries in convexity and optimization...perfectly suited for classroom teaching." Shuzhong Zhang, Professor of Industrial and Systems Engineering, University of Minnesota

Optimization models play an increasingly important role in financial decisions. This is the first textbook devoted to explaining how recent advances in optimization models, methods and software can be applied to solve problems in computational finance more efficiently and accurately. Chapters discussing the theory and efficient solution methods for all major classes of optimization problems alternate with chapters illustrating their use in modeling problems of mathematical finance. The reader is guided through topics such as volatility estimation, portfolio optimization problems and constructing an index fund, using techniques such as nonlinear optimization models, quadratic programming formulations and integer programming models respectively. The book is based on Master's courses in financial engineering and comes with worked examples, exercises and case studies. It will be welcomed by applied mathematicians, operational researchers and others who work in mathematical and computational finance and who are seeking a text for self-learning or for use with courses.

Optimization is a rich and thriving mathematical discipline, and the underlying theory of current computational optimization techniques grows ever more sophisticated. This book aims to provide a concise, accessible account of convex analysis and its applications and extensions, for a broad audience. Each section concludes with an often extensive set of optional exercises. This new edition adds material on semismooth optimization, as well as several new proofs.

Data Mining: Concepts and Techniques

An Introduction to Optimization

Convex Analysis and Optimization

Linear Algebra and Optimization for Machine Learning

An Easy Path to Convex Analysis and Applications

A comprehensive introduction to the tools, techniques and applications of convex optimization.

Flexible graduate textbook that introduces the applications, theory, and algorithms of linear and nonlinear optimization in a clear succinct style, supported by numerous examples and exercises. It introduces important realistic applications and explains how optimization can address them.

"This textbook is designed to accompany a one- or two-semester course for advanced undergraduates or beginning graduate students in computer science and applied mathematics. - It gives an excellent introduction to the probabilistic techniques and paradigms used in the development of probabilistic algorithms and analyses. - It assumes only an elementary background in discrete mathematics and gives a rigorous yet accessible treatment of the material, with numerous examples and applications."--Jacket.

Noncooperative Game Theory is aimed at students interested in using game theory as a design methodology for solving problems in engineering and computer science. João Hespanha shows that such design challenges can be analyzed through game theoretical perspectives that help to pinpoint each problem's essence: Who are the players? What are their goals? Will the solution to "the game" solve the original design problem? Using the fundamentals of game theory, Hespanha explores these issues and more. The use of game theory in technology design is a recent development arising from the intrinsic limitations of classical optimization-based designs. In optimization, one attempts to find values for parameters that minimize suitably defined criteria—such as monetary cost, energy consumption, or heat generated. However, in most engineering applications, there is always some uncertainty as to how the selected parameters will affect the final objective. Through a sequential and easy-to-understand discussion, Hespanha examines how to make sure that the selection leads to acceptable performance, even in the presence of uncertainty—the unforgiving variable that can wreck engineering designs. Hespanha looks at such standard topics as zero-sum, non-zero-sum, and dynamics games and includes a MATLAB guide to coding. Noncooperative Game Theory offers students a fresh way of approaching engineering and computer science applications. An introduction to game theory applications for students of engineering and computer science Materials presented sequentially and in an easy-to-understand fashion Topics explore zero-sum, non-zero-sum, and dynamics games MATLAB commands are included

Optimization in Practice with MATLAB

Statistical Learning with Sparsity

A Unified Approach

Second Edition

Vectors, Matrices, and Least Squares

Optimization Models

This textbook is designed for students and industry practitioners for a first course in optimization integrating MATLAB® software.

Optimization is an essential technique for solving problems in areas as diverse as accounting, computer science and engineering. Assuming only basic linear algebra and with a clear focus on the fundamental concepts, this textbook is the perfect starting point for first- and second-year undergraduate students from a wide range of backgrounds and with varying levels of ability. Modern, real-world examples motivate the theory throughout. The authors keep the text as concise and focused as possible, with more advanced material treated separately or in starred exercises. Chapters are self-contained so that instructors and students can adapt the material to suit their own needs and a wide selection of over 140 exercises gives readers the opportunity to try out the skills they gain in each section. Solutions are available for instructors. The book also provides suggestions for further reading to help students take the next step to more advanced material.

This authoritative book draws on the latest research to explore the interplay of high-dimensional statistics with optimization. Through an accessible analysis of fundamental problems of hypothesis testing and signal recovery, Anatoli Juditsky and Arkadi Nemirovski show how convex optimization theory can be used to devise and analyze near-optimal statistical inferences. Statistical Inference via Convex Optimization is an essential resource for optimization specialists who are new to statistics and its applications, and for data scientists who want to improve their optimization methods. Juditsky and Nemirovski provide the first systematic treatment of the statistical techniques that have arisen from advances in the theory of optimization. They focus on four well-known statistical problems—sparse recovery, hypothesis testing, and recovery from indirect observations of both signals and functions of signals—demonstrating how they can be solved more efficiently as convex optimization problems. The emphasis throughout is on achieving the best possible statistical performance. The construction of inference routines and the quantification of their statistical performance are given by efficient computation rather than by analytical derivation typical of more conventional statistical approaches. In addition to being computation-friendly, the methods described in this book enable practitioners to handle numerous situations too difficult for closed analytical form analysis, such as composite hypothesis testing and signal recovery in inverse problems. Statistical Inference via Convex Optimization features exercises with solutions along with extensive appendices, making it ideal for use as a graduate text.

This book serves as a reference for a self-contained course on online convex optimization and the convex optimization approach to machine learning for the educated graduate student in computer science/electrical engineering/ operations research/statistics and related fields. An ideal reference.

Linear and Nonlinear Optimization

Numerical Optimization

Introduction to Nonlinear Optimization

How to Think About Algorithms

Theory, Algorithms, and Applications with MATLAB

Introductory Course

Targeting readers with backgrounds in economics, Intermediate Financial Theory, Third Edition includes new material on the asset pricing implications of behavioral finance perspectives, recent developments in portfolio choice, derivatives-risk neutral pricing research, and implications of the 2008 financial crisis. Each chapter concludes with questions, and for the first time a freely accessible website presents complementary and supplementary material for every chapter. Known for its rigor and intuition, Intermediate Financial Theory is perfect for those who need basic training in financial theory and those looking for a user-friendly introduction to advanced theory. Completely updated edition of classic textbook that fills a gap between MBA- and PhD-level texts Focuses on clear explanations of key concepts and requires limited mathematical prerequisites Online solutions manual available Updates include new structure emphasizing the distinction between the equilibrium and the arbitrage perspectives on valuation and pricing, and a new chapter on asset management for the long-term investor

This book provides a complete and comprehensive reference/guide to Pyomo (Python Optimization Modeling Objects) for both beginning and advanced modelers, including students at the undergraduate and graduate levels, academic researchers, and practitioners. The text illustrates the breadth of the modeling and analysis capabilities that are supported by the software and support of complex real-world

applications. Pyomo is an open source software package for formulating and solving large-scale optimization and operations research problems. The text begins with a tutorial on simple linear and integer programming models. A detailed reference of Pyomo's modeling components is illustrated with extensive examples, including a discussion of how to load data from data sources like spreadsheets and databases. Chapters describing advanced modeling capabilities for nonlinear and stochastic optimization are also included. The Pyomo software provides familiar modeling features within Python, a powerful dynamic programming language that has a very clear, readable syntax and intuitive object orientation. Pyomo includes Python classes for defining sparse sets, parameters, and variables, which can be used to formulate algebraic expressions that define objectives and constraints. Moreover, Pyomo can be used from a command-line interface and within Python's interactive command environment, which makes it easy to create Pyomo models, apply a variety of optimizers, and examine solutions. The software supports a different modeling approach than commercial AML (Algebraic Modeling Languages) tools, and is designed for flexibility, extensibility, portability, and maintainability but also maintains the central ideas in modern AMLs.

An up-to-date account of the interplay between optimization and machine learning, accessible to students and researchers in both communities. The interplay between optimization and machine learning is one of the most important developments in modern computational science. Optimization formulations and methods are proving to be vital in designing algorithms to extract essential knowledge from huge volumes of data. Machine learning, however, is not simply a consumer of optimization technology but a rapidly evolving field that is itself generating new optimization ideas. This book captures the state of the art of the interaction between optimization and machine learning in a way that is accessible to researchers in both fields. Optimization approaches have enjoyed prominence in machine learning because of their wide applicability and attractive theoretical properties. The increasing complexity, size, and variety of today's machine learning models call for the reassessment of existing assumptions. This book starts the process of reassessment. It describes the resurgence in novel contexts of established frameworks such as first-order methods, stochastic approximations, convex relaxations, interior-point methods, and proximal methods. It also devotes attention to newer themes such as regularized optimization, robust optimization, gradient and subgradient methods, splitting techniques, and second-order methods. Many of these techniques draw inspiration from other fields, including operations research, theoretical computer science, and subfields of optimization. The book will enrich the ongoing cross-fertilization between the machine learning community and these other fields, and within the broader optimization community.

This reference text, now in its second edition, offers a modern unifying presentation of three basic areas of nonlinear analysis: convex analysis, monotone operator theory, and the fixed point theory of nonexpansive operators. Taking a unique comprehensive approach, the theory is developed from the ground up, with the rich connections and interactions between the areas as the central focus, and it is illustrated by a large number of examples. The Hilbert space setting of the material offers a wide range of applications while avoiding the technical difficulties of general Banach spaces. The authors have also drawn upon recent advances and modern tools to simplify the proofs of key results making the book more accessible to a broader range of scholars and users. Combining a strong emphasis on applications with exceptionally lucid writing and an abundance of exercises, this text is of great value to a large audience including pure and applied mathematicians as well as researchers in engineering, data science, machine learning, physics, decision sciences, economics, and inverse problems. The second edition of Convex Analysis and Monotone Operator Theory in Hilbert Spaces greatly expands on the first edition, containing over 140 pages of new material, over 270 new results, and more than 100 new exercises. It features a new chapter on proximity operators including two sections on proximity operators of matrix functions, in addition to several new sections distributed throughout the original chapters. Many existing results have been improved, and the list of references has been updated. Heinz H. Bauschke is a Full Professor of Mathematics at the Kelowna campus of the University of British Columbia, Canada. Patrick L. Combettes, IEEE Fellow, was on the faculty of the City University of New York and of Université Pierre et Marie Curie – Paris 6 before joining North Carolina State University as a Distinguished Professor of Mathematics in 2016.

Convex Analysis for Optimization

Optimization Methods in Finance

A Gentle Introduction to Optimization

Introduction to OnLine Convex Optimization

Convex Optimization Theory

Interior-point Polynomial Algorithms in Convex Programming

This book provides easy access to the basic principles and methods for solving constrained and unconstrained convex optimization problems. Included are sections that cover: basic methods for solving constrained and unconstrained optimization problems with differentiable objective functions; convex sets and their properties; convex functions and their properties and generalizations; and basic principles of sub-differential calculus and convex programming problems. Convex Optimization provides detailed proofs for most of the results presented in the book and also includes many figures and exercises for a better understanding of the material. Exercises are given at the end of each chapter, with solutions and hints to selected exercises given at the end of the book. Undergraduate and graduate students, researchers in different disciplines, as well as practitioners will all benefit from this accessible approach to convex optimization methods.

An overview of the rapidly growing field of ant colony optimization that describes theoretical findings, the major algorithms, and current applications. The complex social behaviors of ants have been much studied by science, and computer scientists are now finding that these behavior patterns can provide models for solving difficult combinatorial optimization problems. The attempt to develop algorithms inspired by one aspect of ant behavior, the ability to find what computer scientists would call shortest paths, has become the field of ant colony optimization (ACO), the most successful and widely recognized algorithmic technique based on ant behavior. This book presents an overview of this rapidly growing field, from its theoretical inception to practical applications, including descriptions of many available ACO algorithms and their uses. The book first describes the translation of observed ant behavior into working optimization algorithms. The ant colony metaheuristic is then introduced and viewed in the general context of combinatorial optimization. This is followed by a detailed description and guide to all major ACO algorithms and a report on current theoretical findings. The book surveys ACO applications now in use, including routing, assignment, scheduling, subset, machine learning, and bioinformatics problems. AntNet, an ACO algorithm designed for the network routing problem, is described in detail. The authors conclude by summarizing the progress in the field and outlining future research directions. Each chapter ends with bibliographic material, bullet points setting out important ideas covered in the chapter, and exercises. Ant Colony Optimization will be of interest to academic and industry researchers, graduate students, and practitioners who wish to learn how to implement ACO algorithms.

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

Here is a book devoted to well-structured and thus efficiently solvable convex optimization problems, with emphasis on conic quadratic and semidefinite programming. The authors present the basic theory underlying these problems as well as their numerous applications in engineering, including synthesis of filters, Lyapunov stability analysis, and structural design. The authors also discuss the complexity issues and provide an overview of the basic theory of state-of-the-art polynomial time interior point methods for linear, conic quadratic, and semidefinite programming. The book's focus on well-structured convex problems in conic form allows for unified theoretical and algorithmical treatment of a wide spectrum of important optimization problems arising in applications.

Convex Analysis and Nonlinear Optimization

Theory and Examples

Lectures on Modern Convex Optimization

Probability and Computing

Statistical Inference Via Convex Optimization

Convex Analysis and Monotone Operator Theory in Hilbert Spaces

A groundbreaking introduction to vectors, matrices, and least squares for engineering applications, offering a wealth of practical examples.

This book deals mainly with the study of convex functions and their behavior from the point of view of stability with respect to perturbations. We shall consider convex functions from the most modern point of view: a function is defined to be convex whenever its epigraph, the set of the points lying above the graph, is a convex set. Thus many of its properties can be seen also as properties of a certain convex set related to it. Moreover, we shall consider extended real valued functions, i. e. , functions taking possibly the values $+\infty$ and $-\infty$. The reason for considering the value $+\infty$ is the powerful device of including the constraint set of a constrained minimum problem into the objective function itself (by redefining it as $+\infty$ outside the constraint set). Except for trivial cases, the minimum value must be taken at a point where the function is not $+\infty$, hence at a point in the constraint set. And the value $-\infty$ is allowed because useful operations, such as the inf-convolution, can give rise to functions valued $-\infty$ even when the primitive objects are real valued. Observe that defining the objective function to be $+\infty$ outside the closed constraint set preserves lower semicontinuity, which is the pivotal and minimal continuity assumption one needs when dealing with minimum problems. Variational calculus is usually based on derivatives.

Discover New Methods for Dealing with High-Dimensional Data A sparse statistical model has only a small number of nonzero parameters or weights; therefore, it is much easier to estimate and interpret than a dense model. Statistical Learning with Sparsity: The Lasso and Generalizations presents methods that exploit sparsity to help recover the underlying signal in a set of data. Top experts in this rapidly evolving field, the authors describe the lasso for linear regression and a simple coordinate descent algorithm for its computation. They discuss the application of l_1 penalties to generalized linear models and support vector machines, cover generalized penalties such as the elastic net and group lasso, and review numerical methods for optimization. They also present statistical inference methods for fitted (lasso) models, including the bootstrap, Bayesian methods, and recently developed approaches. In addition, the book examines matrix decomposition, sparse multivariate analysis, graphical models, and compressed sensing. It concludes with a survey of theoretical results for the lasso. In this age of big data, the number of features measured on a person or object can be large and might be larger than the number of observations. This book shows how the sparsity assumption allows us to tackle these problems and extract useful and reproducible patterns from big datasets. Data analysts, computer scientists, and theorists will appreciate this thorough and up-to-date treatment of sparse statistical modeling.

An accessible introduction to convex algebraic geometry and semidefinite optimization. For graduate students and researchers in mathematics and computer science.

Theory, Methods and Examples

Noncooperative Game Theory

Pyomo – Optimization Modeling in Python

The Lasso and Generalizations

Proximal Algorithms

An Introduction for Engineers and Computer Scientists

Data Mining: Concepts and Techniques provides the concepts and techniques in processing gathered data or information, which will be used in various applications. Specifically, it explains data mining and the tools used in discovering knowledge from the collected data. This book is referred as the knowledge discovery from data (KDD). It focuses on the feasibility, usefulness, effectiveness, and scalability of techniques of large data sets. After describing data mining, this edition explains the methods of knowing, preprocessing, processing, and warehousing data. It then presents information about data warehouses, online analytical processing (OLAP), and data cube technology. Then, the methods involved in mining frequent patterns, associations, and correlations for large data sets are described. The book details the methods for data classification and introduces the concepts and methods for data clustering. The remaining chapters discuss the outlier detection and the trends, applications, and research frontiers in data mining. This book is intended for Computer Science students, application developers, business professionals, and researchers who seek information on data mining. Presents dozens of algorithms and implementation examples, all in pseudo-code and suitable for use in real-world, large-scale data mining projects Addresses advanced topics such as mining object-relational databases, spatial databases, multimedia databases, time-series databases, text databases, the World Wide Web, and applications in several fields Provides a comprehensive, practical look at the concepts and techniques you need to get the most out of your data

This textbook introduces linear algebra and optimization in the context of machine learning. Examples and exercises are provided throughout this text book together with access to a solution's manual. This textbook targets graduate level students and professors in computer science, mathematics and data science. Advanced undergraduate students can also use this textbook. The chapters for this textbook are organized as follows: 1. Linear algebra and its applications: The chapters focus on the basics of linear algebra together with their common applications to singular value decomposition, matrix factorization, similarity matrices (kernel methods), and graph analysis. Numerous machine learning applications have been used as examples, such as spectral clustering, kernel-based classification, and outlier detection. The tight integration of linear algebra methods with examples from machine learning differentiates this book from generic volumes on linear algebra. The focus is clearly on the most relevant aspects of linear algebra for machine learning and to teach readers how to apply these concepts. 2. Optimization and its applications: Much of machine learning is posed as an optimization problem in which we try to maximize the accuracy of regression and classification models. The “parent problem” of optimization-centric machine learning is least-squares regression. Interestingly, this problem arises in both linear algebra and optimization, and is one of the key connecting problems of the two fields. Least-squares regression is also the starting point for support vector machines, logistic regression, and recommender systems. Furthermore, the methods for dimensionality reduction and matrix factorization also require the development of optimization methods. A general view of optimization in computational graphs is discussed together with its applications to back propagation in neural networks. A frequent challenge faced by beginners in machine learning is the extensive background required in linear algebra and optimization. One problem is that the existing linear algebra and optimization courses are not specific to machine learning; therefore, one would typically have to complete more course material than is necessary to pick up machine learning. Furthermore, certain types of ideas and tricks from optimization and linear algebra recur more frequently in machine learning than other application-centric settings. Therefore, there is significant value in developing a view of linear algebra and optimization that is better suited to the specific perspective of machine learning.

This textbook, for second- or third-year students of computer science, presents insights, notations, and analogies to help them describe and think about algorithms like an expert, without grinding through lots of formal proof. Solutions to many problems are provided to let students check their progress, while class-tested PowerPoint slides are on the web for anyone running the course. By looking at both the big picture and easy step-by-step methods for developing algorithms, the author guides students around the common pitfalls. He stresses paradigms such as loop invariants and recursion to unify a huge range of algorithms into a few meta-algorithms. The book fosters a deeper understanding of how and why each algorithm works. These insights are presented in a careful and clear way, helping students to think abstractly and preparing them for creating their own innovative ways to solve problems.

Specialists working in the areas of optimization, mathematical programming, or control theory will find this book invaluable for studying interior-point methods for linear and quadratic programming, polynomial-time methods for nonlinear convex programming, and efficient computational methods for control problems and variational inequalities. A background in linear algebra and mathematical programming is necessary to understand the book. The detailed proofs and lack of “numerical examples” might suggest that the book is of limited value to the reader interested in the practical aspects of convex optimization, but nothing could be further from the truth. An entire chapter is devoted to potential reduction methods precisely because of their great efficiency in practice.

Ant Colony Optimization

Randomized Algorithms and Probabilistic Analysis

Algorithms for Convex Optimization

Nonlinear Optimization

Intermediate Financial Theory

This textbook on nonlinear optimization focuses on model building, real world problems, and applications of optimization models to natural and social sciences. Organized into two parts, this book may be used as a primary text for courses on convex optimization and non-convex methods are well illustrated and all chapters contain compelling exercises. The exercises emphasize fundamental theoretical results on optimality and duality theorems, numerical methods with or without constraints, and derivative-free optimization. Selected solutions are given. methods are highlighted to help students comprehend methods and techniques.

Convex OptimizationCambridge University Press

This work is intended to serve as a guide for graduate students and researchers who wish to get acquainted with the main theoretical and practical tools for the numerical minimization of convex functions on Hilbert spaces. Therefore, it contains the main tools that are necessary also a concise, easy-to-follow and self-contained textbook, which may be useful for any researcher working on related fields, as well as teachers giving graduate-level courses on the topic. It will contain a thorough revision of the extant literature including both classical and state-of-the-art. Praise for the Third Edition “. . . guides and leads the reader through the learning path. . . [e]xamples are stated very clearly and the results are presented with attention to detail.” —MAA Reviews Fully updated to reflect new developments in the field, the Fourth Edition of Introduction to Optimization Theory and Methods with an emphasis on engineering design. Basic definitions and notations are provided in addition to the related fundamental background for linear algebra, geometry, and calculus. This new edition explores the essential topics of unconstrained optimization, programming problems, and nonlinear constrained optimization. The authors also present an optimization perspective on global search methods and include discussions on genetic algorithms, particle swarm optimization, and the simulated annealing algorithm. Featuring an element of convex optimization, and multi-objective optimization, the Fourth Edition also offers: A new chapter on integer programming Expanded coverage of one-dimensional methods Updated and expanded sections on linear matrix inequalities Numerous new exercises at the end of each chapter to reinforce the discussed theory and algorithms Numerous diagrams and figures that complement the written presentation of key concepts MATLAB M-files for implementation of the discussed theory and algorithms (available via the book's website) Introduction to Optimization, Theory and Methods. In addition, the book is a useful reference for professionals in mathematics, operations research, electrical engineering, economics, statistics, and business.