

Convex Optimization Stephen Boyd

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>

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.

Important text examines most significant algorithms for optimizing large systems and clarifying relations between optimization procedures. Initial chapter on linear and nonlinear programming provide the foundation for the rest of the book. Appendixes.

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.

Optimization Models

From Theory to Implementation

Algorithms for Optimization

Primal-dual Interior-Point Methods

Convex Analysis and Nonlinear Optimization

Convex optimization problems arise frequently in many different fields. This book provides a comprehensive introduction to the subject, and shows in detail how such problems can be solved numerically with great efficiency. The book begins with the basic elements of convex sets and functions, and then describes various classes of convex optimization problems. Duality and approximation techniques are then covered, as are statistical estimation techniques. Various geometrical problems are then presented, and there is detailed discussion of unconstrained and constrained minimization problems, and interior-point methods. The focus of the book is on recognizing convex optimization problems and then finding the most appropriate technique for solving them. It contains many worked examples and homework exercises and will appeal to students, researchers and practitioners in fields such as engineering, computer science,

mathematics, statistics, finance and economics. Most global optimization literature focuses on theory. This book, however, contains descriptions of new implementations of general-purpose or problem-specific global optimization algorithms. It discusses existing software packages from which the entire community can learn. The contributors are experts in the discipline of actually getting global optimization to work, and the book provides a source of ideas for people needing to implement global optimization software.

A PRACTICAL GUIDE TO OPTIMIZATION PROBLEMS WITH DISCRETE OR INTEGER VARIABLES, REVISED AND UPDATED The revised second edition of *Integer Programming* explains in clear and simple terms how to construct custom-made algorithms or use existing commercial software to obtain optimal or near-optimal solutions for a variety of real-world problems. The second edition also includes information on the remarkable progress in the development of mixed integer programming solvers in the 22 years since the first edition of the book appeared. The updated text includes information on the most recent developments in the field such as the much improved preprocessing/presolving and the many new ideas for primal heuristics included in

the solvers. The result has been a speed-up of several orders of magnitude. The other major change reflected in the text is the widespread use of decomposition algorithms, in particular column generation (branch-(cut)-and-price) and Benders' decomposition. The revised second edition: Contains new developments on column generation Offers a new chapter on Benders' algorithm Includes expanded information on preprocessing, heuristics, and branch-and-cut Presents several basic and extended formulations, for example for fixed cost network flows Also touches on and briefly introduces topics such as non-bipartite matching, the complexity of extended formulations or a good linear program for the implementation of lift-and-project Written for students of integer/mathematical programming in operations research, mathematics, engineering, or computer science, Integer Programming offers an updated edition of the basic text that reflects the most recent developments in the field. 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.

Limits of Performance

Lectures on Modern Convex Optimization

Convex Optimization in Signal Processing and Communications

Selected Applications of Convex Optimization

Recent Advances in Learning and Control

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

Convexity is a simple idea that manifests itself in a surprising variety of places. This fertile field has an immensely rich structure and numerous applications. Barvinok demonstrates that simplicity, intuitive appeal, and the universality of applications make teaching (and learning) convexity a gratifying experience. The book will benefit both teacher and student: It is easy to understand, entertaining to the reader, and includes many exercises that vary in degree of difficulty. Overall, the author demonstrates the power of a few simple unifying principles in a variety of pure and applied problems. The

prerequisites are minimal amounts of linear algebra, analysis, and elementary topology, plus basic computational skills. Portions of the book could be used by advanced undergraduates. As a whole, it is designed for graduate students interested in mathematical methods, computer science, electrical engineering, and operations research. The book will also be of interest to research mathematicians, who will find some results that are recent, some that are new, and many known results that are discussed from a new perspective. This volume is composed of invited papers on learning and control. The contents form the proceedings of a workshop held in January 2008, in Hyderabad that honored the 60th birthday of Doctor Mathukumalli Vidyasagar. The 14 papers, written by international specialists in the field, cover a variety of interests within the broader field of learning and control. The diversity of the research provides a comprehensive overview of a field of great interest to control and system theorists.

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 appendixes, making it ideal for use as a graduate text.

*Introduction to Nonlinear Optimization
Optimization Theory for Large Systems
Problem Complexity and Method
Efficiency in Optimization
Convex Analysis*

Vectors, Matrices, and Least Squares

The primary goal of this book is to provide a self-contained, comprehensive study of the main first-order methods that are frequently used in solving large-scale problems. First-order methods exploit information on values and

gradients/subgradients (but not Hessians) of the functions composing the model under consideration. With the increase in the number of applications that can be modeled as large or even huge-scale optimization problems, there has been a revived interest in using simple methods that require low iteration cost as well as low memory storage. The author has gathered, reorganized, and synthesized (in a unified manner) many results that are currently scattered throughout the literature, many of which cannot be typically found in optimization books. *First-Order Methods in Optimization* offers comprehensive study of first-order methods with the theoretical foundations; provides plentiful examples and illustrations; emphasizes rates of convergence and complexity analysis of the main first-order methods used to solve large-scale problems; and covers both variables and functional decomposition methods.

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

Leading experts provide the theoretical underpinnings of the subject plus tutorials on a wide range of applications, from automatic code generation to robust broadband beamforming. Emphasis on cutting-edge research and formulating problems in convex form make this an ideal textbook for advanced graduate courses and a useful self-study guide.

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.

Linear Controller Design

Convex Optimization of Power Systems

Global Optimization

Analysis, Algorithms, and Engineering Applications

The Pia Contingency

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

This Fourth Edition introduces the latest theory and applications in optimization. It emphasizes constrained optimization, beginning with a substantial treatment of linear programming and then proceeding to convex analysis, network flows, integer programming, quadratic programming, and convex optimization. Readers will discover a host of practical business applications as well as non-business applications. Topics are clearly developed with many

numerical examples worked out in detail. Specific examples and concrete algorithms precede more abstract topics. With its focus on solving practical problems, the book features free C programs to implement the major algorithms covered, including the two-phase simplex method, primal-dual simplex method, path-following interior-point method, and homogeneous self-dual methods. In addition, the author provides online JAVA applets that illustrate various pivot rules and variants of the simplex method, both for linear programming and for network flows. These C programs and JAVA tools can be found on the book's website. The website also includes new online instructional tools and exercises.

A mathematically rigorous guide to convex optimization for power systems engineering.

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

Numerical Optimization

Integer Programming

Convex Optimization South Asia Edition

Semidefinite Optimization and Convex Algebraic Geometry

This book focuses on the applications of convex optimization and highlights several topics, including support vector machines, parameter estimation, norm approximation and regularization, semi-definite programming problems, convex relaxation, and geometric problems. All derivation processes are presented in detail to aid in comprehension. The book offers concrete

guidance, helping readers recognize and formulate convex optimization problems they might encounter in practice.

In this book the authors reduce a wide variety of problems arising in system and control theory to a handful of convex and quasiconvex optimization problems that involve linear matrix inequalities. These optimization problems can be solved using recently developed numerical algorithms that not only are polynomial-time but also work very well in practice; the reduction therefore can be considered a solution to the original problems. This book opens up an important new research area in which convex optimization is combined with system and control theory, resulting in the solution of a large number of previously unsolved problems.

A comprehensive introduction to optimization with a focus on practical algorithms for the design of engineering systems. This book offers a comprehensive introduction to optimization with a focus on practical algorithms. The book approaches optimization from an engineering perspective, where the objective is to design a system that optimizes a set of metrics subject to constraints. Readers will learn about computational approaches for a range of

challenges, including searching high-dimensional spaces, handling problems where there are multiple competing objectives, and accommodating uncertainty in the metrics. Figures, examples, and exercises convey the intuition behind the mathematical approaches. The text provides concrete implementations in the Julia programming language. Topics covered include derivatives and their generalization to multiple dimensions; local descent and first- and second-order methods that inform local descent; stochastic methods, which introduce randomness into the optimization process; linear constrained optimization, when both the objective function and the constraints are linear; surrogate models, probabilistic surrogate models, and using probabilistic surrogate models to guide optimization; optimization under uncertainty; uncertainty propagation; expression optimization; and multidisciplinary design optimization. Appendixes offer an introduction to the Julia language, test functions for evaluating algorithm performance, and mathematical concepts used in the derivation and analysis of the optimization methods discussed in the text. The book can be used by advanced

undergraduates and graduate students in mathematics, statistics, computer science, any engineering field, (including electrical engineering and aerospace engineering), and operations research, and as a reference for professionals.

The new edition of this book presents a comprehensive and up-to-date description of the most effective methods in continuous optimization. It responds to the growing interest in optimization in engineering, science, and business by focusing on methods best suited to practical problems. This edition has been thoroughly updated throughout. There are new chapters on nonlinear interior methods and derivative-free methods for optimization, both of which are widely used in practice and are the focus of much current research. Because of the emphasis on practical methods, as well as the extensive illustrations and exercises, the book is accessible to a wide audience.

Convex Optimization Theory

Statistical Inference Via Convex

Optimization

Introduction to Online Convex Optimization

Lectures on Convex Optimization

(PMS-28)

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.

In the past decade, primal-dual algorithms have emerged as the most important and useful algorithms from the interior-point class. This book presents the major primal-dual algorithms for linear programming in straightforward terms. A thorough description of the theoretical properties of these methods is given, as are a discussion of practical and computational aspects and a summary of current software. This is an excellent, timely, and well-written work. The major primal-dual algorithms covered in this book are path-following algorithms (short- and long-step, predictor-corrector), potential-reduction algorithms, and

infeasible-interior-point algorithms. A unified treatment of superlinear convergence, finite termination, and detection of infeasible problems is presented. Issues relevant to practical implementation are also discussed, including sparse linear algebra and a complete specification of Mehrotra's predictor-corrector algorithm. Also treated are extensions of primal-dual algorithms to more general problems such as monotone complementarity, semidefinite programming, and general convex programming problems. Surveys the theory and history of the alternating direction method of multipliers, and discusses its applications to a wide variety of statistical and machine learning problems of recent interest, including the lasso, sparse logistic regression, basis pursuit, covariance selection, support vector machines, and many others.

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

Foundations and Extensions

Optimization in Practice with MATLAB

Introduction to Applied Linear Algebra

Convex Optimization China Edition

Linear Programming

A new habitable world was discovered in the universe. The people of Earth could get there but only as a space colony. The real question was who would go? The atheists of Earth agreed to go to explore the new planet but only if they could form a new society, a society free from all religious beliefs. The untold secret goal of the group was even darker than they told their everyone, darker than anyone could have imagined. As soon as the space colony ship went past the point of no

return, the leaders announced a harsher system that was originally proposed. Anyone caught worshipping anything was immediately put to death. Unbeknown to the leaders of the original colony, a small group of Believers imbedded themselves into the colony. They were visited by an angel of God who told them to go. They were told that they didn't have to go but if they didn't billions of souls would be lost for all of eternity. Some of the group that were visited by the angel went on the mission, some did not. The new planet was named PIA, an abbreviation for Planet of Intergalactic Atheists. The story picks up several generations after the original colonists landed. The grandson of the founder of the atheist group was in charge of the planet just as his father had been and his father's father before him. Stephen Steele was the grandson of one the original Believers that had imbedded themselves onto the space colony. The Believers spread out across the planet when they first arrived so as to keep safe. They formed small cell groups as well. That way if one group got caught the others would stay safe. His Grandparents migrated to the mountains. They lived in a small isolated village they named named Ash, Ash was very far away from the Darwin, the capital of the planet. Kristin Knope was born to be a prostitute just as her mother and her mother's mother was. Kristin never met her father as even her mother didn't know who he was. Even if she did know it

would not have mattered much. Most children on Pia were raised in state run orphanages. If either of their parents stayed around it was usually the mother. Child were considered throw aways, especially children born into the lower class and Kristin was the lowest of the lowest class. She was born into the Hopeless class. On Pia the manual laborers were necessary at first, not so much later on. The laborers soon became useless. As more and more machinery arrived from Earth they were tossed aside. Soon an entire class of people moved into the gutters of Pia. They became the Hopeless class. Stephen Steele's parents and grandparents were part of the original team of Believers that developed their computer cyber systems. They were necessary to run the planet. They were considered valuable and given much freedom in the early days of Pia. They also helped create the Pianet, their world wide web. Fresh out of law school, Stephen Steele came across Kristin Knope on the streets of downtown Darwin. He had barely heard of a prostitution when he met Kristin due to his isolated upbringing. When he recognized what she was he determine to help her. He sent her to a charity that helped prostitutes reeducate themselves. It was years before they met again. By that time Kristin had a degree as a teacher and was teaching Kindergartners at Simpka elementary. Due to unfortunate circumstances Kristin lost her job and had to get work where

ever she could. She also had a degree in paralegal studies and got a job in Stephen Steele's law office. They didn't recognise each other at first from their long ago chance encounter but they soon remembered each other. Stephen's real mission in life was the same as his parents and grandparents, to tell people about God and let them decide for themselves whether or not to follow him. The Supreme Chancellor of Pia says that he is against all religions but in truth he is not. In fact he is being controlled by one. A mysterious being that suddenly appears.

This monograph presents the main complexity theorems in convex optimization and their corresponding algorithms. It begins with the fundamental theory of black-box optimization and proceeds to guide the reader through recent advances in structural optimization and stochastic optimization. The presentation of black-box optimization, strongly influenced by the seminal book by Nesterov, includes the analysis of cutting plane methods, as well as (accelerated) gradient descent schemes. Special attention is also given to non-Euclidean settings (relevant algorithms include Frank-Wolfe, mirror descent, and dual averaging), and discussing their relevance in machine learning. The text provides a gentle introduction to structural optimization with FISTA (to optimize a sum of a smooth and a simple non-smooth term), saddle-

point mirror prox (Nemirovski's alternative to Nesterov's smoothing), and a concise description of interior point methods. In stochastic optimization it discusses stochastic gradient descent, mini-batches, random coordinate descent, and sublinear algorithms. It also briefly touches upon convex relaxation of combinatorial problems and the use of randomness to round solutions, as well as random walks based methods.

Available for the first time in paperback, R. Tyrrell Rockafellar's classic study presents readers with a coherent branch of nonlinear mathematical analysis that is especially suited to the study of optimization problems. Rockafellar's theory differs from classical analysis in that differentiability assumptions are replaced by convexity assumptions. The topics treated in this volume include: systems of inequalities, the minimum or maximum of a convex function over a convex set, Lagrange multipliers, minimax theorems and duality, as well as basic results about the structure of convex sets and the continuity and differentiability of convex functions and saddle- functions. This book has firmly established a new and vital area not only for pure mathematics but also for applications to economics and engineering. A sound knowledge of linear algebra and introductory real analysis should provide readers with sufficient background for this book. There is also a guide

for the reader who may be using the book as an introduction, indicating which parts are essential and which may be skipped on a first reading. 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).

Distributed Optimization and Statistical Learning Via the Alternating Direction Method of Multipliers

Convex Optimization

First-Order Methods in Optimization

Convex Analysis and Optimization

Linear Matrix Inequalities in System and Control Theory

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.

This book provides a comprehensive, modern introduction to convex optimization, a field that is becoming increasingly important in applied mathematics, economics and finance, engineering, and computer science, notably in data science and machine learning. Written by a leading expert in the field, this book includes recent advances in the

algorithmic theory of convex optimization, naturally complementing the existing literature. It contains a unified and rigorous presentation of the acceleration techniques for minimization schemes of first- and second-order. It provides readers with a full treatment of the smoothing technique, which has tremendously extended the abilities of gradient-type methods. Several powerful approaches in structural optimization, including optimization in relative scale and polynomial-time interior-point methods, are also discussed in detail. Researchers in theoretical optimization as well as professionals working on optimization problems will find this book very useful. It presents many successful examples of how to develop very fast specialized minimization algorithms. Based on the author's lectures, it can naturally serve as the basis for introductory and advanced courses in convex optimization for students in engineering, economics, computer science and mathematics.

Proximal Algorithms

Algorithms and Complexity

Theory, Algorithms, and Applications with MATLAB

A Basic Course

Optimization for Machine Learning