Convex Optimization Stephen Boyd

The purpose of this annual series, Applied and Computational Control, Signals, and Circuits, is to keep abreast of the fast-paced developments in computational mathematics and scientific computing and their increasing use by researchers and engineers in control, signals, and circuits. The series is dedicated to fostering effective communication between mathematicians, computer scientists, computational scientists, software engineers, theorists, and practicing engineers. This interdisciplinary scope is meant to blend areas of mathematics (such as linear algebra, operator theory, and certain branches of analysis) and computational mathematics (numerical linear algebra, numerical differential equations, large scale and parallel matrix computations, numerical optimization) with control and systems theory, signal and image processing, and circuit analysis and design. The disciplines mentioned above have long enjoyed a natural synergy. There are distinguished journals in the fields of control and systems the ory, as well as signal processing and circuit theory, which publish high quality papers on mathematical and engineering aspects of these areas; however, articles on their computational and applications aspects appear only sporadically. At the same time, there has been tremendous recent growth and development of computational mathematics, scientific comput ing, and mathematical software, and the resulting sophisticated techniques are being gradually adapted by engineers, software designers, and other scientists to the needs of those applied disciplines.

Convex optimization is widely used, in many fields, but is nearly always constrained to problems solved in a few minutes or seconds, and even then, nearly always with a human in the loop. The advent of parser-solvers has made convex optimization simpler and more accessible, and greatly increased the number of people using convex optimization. Most current applications, however, are for the design of systems or analysis of data. It is possible to use convex optimization for real-time or embedded applications, where the optimization solver is a part of a larger system. Here, the optimization algorithm must find solutions much faster than a generic solver, and often has a hard, real-time deadline. Use in embedded applications additionally means that the solver cannot fail, and must be robust even in the presence of relatively poor quality data. For ease of embedding, the solver should be simple, and have minimal dependencies on external libraries. Convex optimization has been successfully applied in such settings in the past. However, they have usually necessitated a custom, hand-written solver. This requires signficant time and expertise, and has been a major factor preventing the adoption of convex optimization in embedded applications. This work describes the implementation and use of a prototype code generator for convex optimization, CVXGEN, that creates high-speed solvers automatically. Using the principles of disciplined convex programming, CVXGEN allows the user to describe an optimization problem in a convenient, high-level language, then receive code for compilation into an extremely fast, robust, embeddable solver.

??????????????????????????????,??,???????????????

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

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

Convex OptimizationCambridge University Press

The core research conducted over the period of this grant focused on distributed and networked optimization in a variety of practical settings, including design (such as, for circuits), parameter estimation and fault identification (for integrated vehicle health management), network flow control, control of wireless and sensor networks, methods for general purpose control, and power optimization in processors. The research explored a variety of applications of distributed optimization in which local decisions are coordinated to produce a good overall solution. These problems arise in many areas, such as monitoring the health of a complex system, coordinating the design of a complex system, or fusing together information from disparate sources. The researchers exploited mathematical properties of the underlying problems to be solved (specifically, convexity) to develop effective methods. Following a brief description of the core research conducted during this grant, the report presents a bibliography of 56 archival publications in which Stephen Boyd is a co-author and other achievements realized during the course of the grant: 8 student doctoral dissertations, 4 software programs, honors and awards, and 7 plenary/keynote talks.

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.

Artificial Intelligence continues to be one of the most exciting and fast-developing fields of computer science. This book presents the 177 long papers and 123 short papers accepted for ECAI 2016, the latest edition of the biennial European Conference on Artificial Intelligence, Europe’s premier venue for presenting scientific results in AI. The conference was held in The Hague, the Netherlands, from August 29 to September 2, 2016. ECAI 2016 also incorporated the conference on Prestigious Applications of Intelligent Systems (PAIS) 2016, and the Starting AI Researcher Symposium (STAIRS). The papers from PAIS are included in this volume; the papers from STAIRS are published in a separate volume in the Frontiers in Artificial Intelligence and Applications (FAIA) series. Organized by the European Association for Artificial Intelligence (EurAI) and the Benelux Association for Artificial Intelligence (BNVKI), the

ECAI conference provides an opportunity for researchers to present and hear about the very best research in contemporary AI. This proceedings will be of interest to all those seeking an overview of the very latest innovations and developments in this field.

This book focuses on recent research in modern optimization and its implications in control and data analysis. This book is a collection of papers from the conference “Optimization and Its Applications in Control and Data Science” dedicated to Professor Boris T. Polyak, which was held in Moscow, Russia on May 13-15, 2015. This book reflects developments in theory and applications rooted by Professor Polyak’s fundamental contributions to constrained and unconstrained optimization, differentiable and nonsmooth functions, control theory and approximation. Each paper focuses on techniques for solving complex optimization problems in different application areas and recent developments in optimization theory and methods. Open problems in optimization, game theory and control theory are included in this collection which will interest engineers and researchers working with efficient algorithms and software for solving optimization problems in market and data analysis. Theoreticians in operations research, applied mathematics, algorithm design, artificial intelligence, machine learning, and software engineering will find this book useful and graduate students will find the state-of-the-art research valuable.

This groundbreaking textbook combines straightforward explanations with a wealth of practical examples to offer an innovative approach to teaching linear algebra. Requiring no prior knowledge of the subject, it covers the aspects of linear algebra – vectors, matrices, and least squares – that are needed for engineering applications, discussing examples across data science, machine learning and artificial intelligence, signal and image processing, tomography, navigation, control, and finance. The numerous practical exercises throughout allow students to test their understanding and translate their knowledge into solving real-world problems, with lecture slides, additional computational exercises in Julia and MATLAB, and data sets accompanying the book online at https://web.stanford.edu/~boyd/vmls/. Suitable for both one-semester and one-quarter courses, as well as self-study, this self-contained text provides beginning students with the foundation they need to progress to more advanced study.

This is a comprehensive book on robust portfolio optimization, which includes up-to-date developments and will interest readers looking for advanced material on portfolio optimization. The book will also attract introductory-level readers because it begins by reviewing the foundations of portfolio optimization. The material in this book emphasizes applications in equity portfolio management and includes MATLAB codes that can assist readers of all levels in implementing robust models. The book aims to help the reader fully understand formulations, performances, and properties of robust portfolios. Application in the equity market is described throughout the book and the implementation of robust models is explained in detail with example code.

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.

??,????????????????????

??????????????????,???????????????????????,????????????????????????????

?????????;???????????????????

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 Analysis is an emerging calculus of inequalities while Convex Optimization is its application. Analysis is the domain of the mathematician while Optimization belongs to the engineer. In layman's terms, the mathematical science of Optimization is a study of how to make good choices when confronted with conflicting requirements and demands. The qualifier Convex means: when an optimal solution is found, then it is guaranteed to be a best solution; there is no better choice. As any convex optimization problem has geometric interpretation, this book is about convex geometry (with particular attention to distance geometry) and nonconvex, combinatorial, and geometrical problems that can be relaxed or transformed into convexity. A virtual flood of new applications follows by epiphany that many problems, presumed nonconvex, can be so transformed. This is a BLACK & WHITE paperback. A hardcover with full color interior, as originally conceived, is available at lulu.com/spotlight/dattorro

The study of Euclidean distance matrices (EDMs) fundamentally asks what can be known geometrically given onlydistance information between points in Euclidean space. Each point may represent simply locationor, abstractly, any entity expressible as a vector in finite-dimensional Euclidean space.The answer to the question posed is that very much can be known about the points;the mathematics of this combined study of geometry and optimization is rich and deep.Throughout we cite beacons of historical accomplishment.The application of EDMs has already proven invaluable in discerning biological molecular conformation.The emerging practice of localization in wireless sensor networks, the global positioning system (GPS), and distance-based pattern recognitionwill certainly simplify and benefit from this theory.We study the pervasive convex Euclidean bodies and their various representations.In particular, we make convex polyhedra, cones, and dual cones more visceral through illustration, andwe study the geometric relation of polyhedral cones to nonorthogonal bases biorthogonal expansion.We explain conversion between halfspace- and vertex-descriptions of convex cones,we provide formulae for determining dual cones,and we show how classic alternative systems of linear inequalities or linear matrix inequalities and optimality conditions can be explained by generalized inequalities in terms of convex cones and their duals.The conic analogue to linear independence, called conic independence, is introducedas a new tool in the study of classical cone theory; the logical next step in the progression:linear, affine, conic.Any convex optimization problem has geometric interpretation.This is a powerful attraction: the ability to visualize geometry of an optimization problem.We provide tools to make visualization easier.The concept of faces, extreme points, and extreme directions of convex Euclidean bodiesis explained here, crucial to understanding convex optimization.The convex cone of positive semidefinite matrices, in particular, is studied in depth.We mathematically interpret, for example,its inverse image under affine transformation, and we explainhow higher-rank subsets of its boundary united with its interior are convex.The Chapter on "Geometry of convex functions",observes analogies between convex sets and functions:The set of all vector-valued convex functions is a closed convex cone.Included among the examples in this chapter, we show how the real affinefunction relates to convex functions as the hyperplane relates to convex sets.Here, also, pertinent results formultidimensional convex functions are presented that are largely ignored in the literature;tricks and tips for determining their convexityand discerning their geometry, particularly with regard to matrix calculus which remains largely unsystematizedwhen compared with the traditional practice of ordinary calculus.Consequently, we collect some results of matrix differentiation in the appendices.The Euclidean distance matrix (EDM) is studied,its properties and relationship to both positive semidefinite and Gram matrices.We relate the EDM to the four classical axioms of the Euclidean metric;thereby, observing the existence of an infinity of axioms of the Euclidean metric beyondthe triangle inequality. We proceed byderiving the fifth Euclidean axiom and then explain why furthering this endeavoris inefficient because the ensuing criteria (while describing polyhedra)grow linearly in complexity and number.Some geometrical problems solvable via EDMs,EDM problems posed as convex optimization, and methods of solution arepresented;\eg, we generate a recognizable isotonic map of the United States usingonly comparative distance information (no distance information, only distance inequalities).We offer a new proof of the classic Schoenberg criterion, that determines whether a candidate matrix is an EDM. Our proofrelies on fundamental geometry; assuming, any EDM must correspond to a list of points contained in some polyhedron(possibly at its vertices) and vice versa.It is not widely known that the Schoenberg criterion implies nonnegativity of the EDM entries; proved here.We characterize the eigenvalues of an EDM matrix and then devisea polyhedral cone required for determining membership of a candidate matrix(in Cayley-Menger form) to the convex cone of Euclidean distance matrices (EDM cone); \ie,a candidate is an EDM if and only if its eigenspectrum belongs to a spectral cone for EDM^N.We will see spectral cones are not unique.In the chapter "EDM cone", we explain the geometric relationship betweenthe EDM cone, two positive semidefinite cones, and the elliptope.We illustrate geometric requirements, in particular, for projection of a candidate matrixon a positive semidefinite cone that establish its membership to the EDM cone. The faces of the EDM cone are described,but still open is the question whether all its faces are exposed as they are for the positive semidefinite cone.The classic Schoenberg criterion, relating EDM and positive semidefinite cones, isrevealed to be a discretized membership relation (a generalized inequality, a new Farkas''''''-like lemma)between the EDM cone and its ordinary dual. A matrix criterion for membership to the dual EDM cone is derived thatis simpler than the Schoenberg criterion.We derive a new concise expression for the EDM cone and its dual involvingtwo subspaces and a positive semidefinite cone."Semidefinite programming" is reviewedwith particular attention to optimality conditionsof prototypical primal and dual conic programs,their interplay, and the perturbation method of rank reduction of optimal solutions(extant but not well-known).We show how to solve a ubiquitous platonic combinatorial optimization problem from linear algebra(the optimal Boolean solution x to Ax=b)via semidefinite program relaxation.A three-dimensional polyhedral analogue for the positive semidefinite cone of 3X3 symmetricmatrices is introduced; a tool for visualizing in 6 dimensions.In "EDM proximity"we explore methods of solution to a few fundamental and prevalentEuclidean distance matrix proximity problems; the problem of finding that Euclidean distance matrix closestto a given matrix in the Euclidean sense.We pay particular attention to the problem when compounded with rank minimization.We offer a new geometrical proof of a famous result discovered by Eckart \& Young in 1936 regarding Euclideanprojection of a point on a subset of the positive semidefinite cone comprising all positive semidefinite matriceshaving rank not exceeding a prescribed limit rho.We explain how this problem is transformed to a convex optimization for any rank rho.

This book is derived from lecture notes for a course on Fourier analysis for engineering and science students at the advanced undergraduate or beginning graduate level. Beyond teaching specific topics and techniques—all of which are important in many areas of engineering and science—the author's goal is to help engineering and science students cultivate more advanced mathematical know-how and increase confidence in learning and using mathematics, as well as appreciate the coherence of the subject. He promises the readers a little magic on every page. The section headings are all recognizable to mathematicians, but the arrangement and emphasis are directed toward students from other disciplines. The material also serves as a foundation for advanced courses in signal processing and imaging. There are over 200 problems, many of which are oriented to applications, and a number use standard software. An unusual feature for courses meant for engineers is a more detailed and accessible treatment of distributions and the generalized Fourier transform. There is also more coverage of higher-dimensional phenomena than is found in most books at this level.

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.

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.

This monograph collects in one place the basic de?nitions, a careful description of the model, and discussion of how convex optimization can be used in multi-period trading, all in a common notation and framework.

??? ??? C#????????????????????????????????????? ???Fortran?MATLAB?C++?Java ????????????? ??? ???

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.

Although LMI has emerged as a powerful tool with applications across the major domains of systems and control, there has been a need for a textbook that provides an accessible introduction to LMIs in control systems analysis and design. Filling this need, LMIs in Control Systems: Analysis, Design and Applications focuses on the basic analysis and d

We consider a basic model of multi-period trading, which can be used to evaluate the performance of a trading strategy. We describe a framework for single-period optimization, where the trades in each period are found by solving a convex optimization problem that trades off expected return, risk, transaction cost and holding cost such as the borrowing cost for shorting assets. We then describe a multi-period version of the trading method, where optimization is used to plan a sequence of trades, with only the first one executed, using estimates of future quantities that are unknown when the trades are chosen. The single period method traces back to Markowitz; the multi-period methods trace back to model predictive control. Our contribution is to describe the single-period and multi-period methods in one simple framework, giving a clear description of the development and the approximations made. In this paper, we do not address a critical component in a trading algorithm, the predictions or forecasts of future quantities. The methods we describe in this paper can be thought of as good ways to exploit predictions, no matter how they are made. We have also developed a

companion open-source software library that implements many of the ideas and methods described in the paper.
??????????????????????????????????????K-T????????????????????????????????. ????:??
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.
This book constitutes the refereed proceedings of the 7th International Symposium on Self-Stabilizing Systems, SSS 2005, held in Barcelona, Spain, in October 2005. The 15 revised full papers presented were carefully reviewed and selected from 33 submissions. The papers address classical topics of self-stabilization, prevailing extensions to the field, such as snap-stabilization, code stabilization, self-stabilization with either dynamic, faulty or Byzantine components, or deal with applications of self-stabilization, either related to operating systems, security, or mobile and ad hoc networks.
An accessible introduction to convex algebraic geometry and semidefinite optimization. For graduate students and researchers in mathematics and computer science.
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.
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.
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.