

Convex Analysis And Optimization Bertsekas

This book offers a concise and in-depth exposition of specific algorithmic solutions for distributed optimization based control of multi-agent networks and their performance analysis. It synthesizes and analyzes distributed strategies for three collaborative tasks: distributed cooperative optimization, mobile sensor deployment and multi-vehicle formation control. The book integrates miscellaneous ideas and tools from dynamic systems, control theory, graph theory, optimization, game theory and Markov chains to address the particular challenges introduced by such complexities in the environment as topological dynamics, environmental uncertainties, and potential cyber-attack by human adversaries. The book is written for first- or second-year graduate students in a variety of engineering disciplines, including control, robotics, decision-making, optimization and algorithms and with backgrounds in aerospace engineering, computer science, electrical engineering, mechanical engineering and operations research. Researchers in these areas may also find the book useful as a reference.

From the reviews: "The account is quite detailed and is written in a manner that will appeal to analysts and numerical practitioners alike...they contain everything from rigorous proofs to tables of numerical calculations.... one of the strong features of these books...that they are designed not for the expert, but for those who wish to learn the subject matter starting from little or no background...there are numerous examples, and counter-examples, to back up the theory...To my knowledge, no other authors have given such a clear geometric account of convex analysis." "This innovative text is well written, copiously illustrated, and accessible to a wide audience"

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).

???????????????????????????????????? 30

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.

Computational contact mechanics is a broad topic which bringstogether algorithmic, geometrical, optimization and numericalaspects for a robust, fast and accurate treatment of contactproblems. This book covers all the basic ingredients of contact andcomputational contact mechanics: from efficient contact detectionalgorithms and classical optimization methods to new developmentsin contact kinematics and resolution schemes for both sequentialand parallel computer architectures. The book is self-contained andintended for people working on the implementation and improvementof contact algorithms in a finite element software. Using a new tensor algebra, the authors introduce some originalnotions in contact kinematics and extend the classical formulationof contact elements. Some classical and new resolution methods forcontact problems and associated ready-to-implement expressions areprovided. Contents: 1. Introduction to Computational Contact. 2. Geometry in Contact Mechanics. 3. Contact Detection. 4. Formulation of Contact Problems. 5. Numerical Procedures. 6. Numerical Examples. About the Authors Vladislav A. Yastrebov is a postdoctoral-fellow in ComputationalSolid Mechanics at MINES ParisTech in France. His work incomputational contact mechanics was recognized by the CSMA awardand by the Prix Paul Caseau of the French Academy of Technology andElectricité de France.

Convex Analysis may be considered as a refinement of standard calculus, with equalities and approximations replaced by inequalities. As such, it can easily be integrated into a graduate study curriculum. Minimization algorithms, more specifically those adapted to non-differentiable functions, provide an immediate application of convex analysis to various fields related to optimization and operations research. These two topics making up the title of the book, reflect the two origins of the authors, who belong respectively to the academic world and to that of applications. Part I can be used as an introductory textbook (as a basis for courses, or for self-study); Part II continues this at a higher technical level and is addressed more to specialists, collecting results that so far have not appeared in books.

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 self-contained textbook is an informal introduction to optimization through the use of numerous illustrations and applications. The focus is on analytically solving optimization problems with a finite number of continuous variables. In addition, the authors provide introductions to classical and modern numerical methods of optimization and to dynamic optimization. The book's overarching point is that most problems may be solved by the direct application of the theorems of Fermat, Lagrange, and Weierstrass. The authors show how the intuition for each of the theoretical results can be supported by simple geometric figures. They include numerous applications through the use of varied classical and practical problems. Even experts may find some of these applications truly surprising. A basic mathematical knowledge is sufficient to understand the topics covered in this book. More advanced readers, even experts, will be surprised to see how all main results can be grounded on the Fermat-Lagrange theorem. The book can be used for courses on continuous optimization, from introductory to advanced, for any field for which optimization is relevant.

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

This book provides a comprehensive and accessible presentation of algorithms for solving convex optimization problems. It relies on rigorous mathematical analysis, but also aims at an intuitive exposition that makes use of visualization where possible. This is facilitated by the extensive use of analytical and algorithmic concepts of duality, which by nature lend themselves to geometrical interpretation. The book places particular emphasis on modern developments, and their widespread applications in fields such as large-scale resource allocation problems, signal processing, and machine learning. The book is aimed at students, researchers, and practitioners, roughly at the first year graduate level. It is similar in style to the author's 2009 "Convex Optimization Theory" book, but can be read independently. The latter book focuses on convexity theory and optimization duality, while the present book focuses on algorithmic issues. The two books share notation, and together cover the entire finite-dimensional convex optimization methodology. To facilitate readability, the statements of definitions and results of the "theory book" are reproduced without proofs in Appendix B.

Optimality Conditions in Convex Optimization explores an important and central issue in the field of convex optimization: optimality conditions. It brings together the most important and recent results in this area that have been scattered in the literature—notably in the area of convex analysis—essential in developing many of the important results in this book, and not usually found in conventional texts. Unlike other books on convex optimization, which usually discuss algorithms along with some basic theory, the sole focus of this book is on fundamental and advanced convex optimization theory. Although many results presented in the book can also be proved in infinite dimensions, the authors focus on finite dimensions to allow for much deeper results and a better understanding of the structures involved in a convex optimization problem. They address semi-infinite optimization problems; approximate solution concepts of convex optimization problems; and some classes of non-convex problems which can be studied using the tools of convex analysis. They include examples wherever needed, provide details of major results, and discuss proofs of the main results.

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 presents tools and methods for large-scale and distributed optimization. Since many methods in "Big Data" fields rely on solving large-scale optimization problems, often in distributed fashion, this topic has over the last decade emerged to become very important. As well as specific coverage of this active research field, the book serves as a powerful source of information for practitioners as well as theoreticians. Large-Scale and Distributed Optimization is a unique combination of contributions from leading experts in the field, who were speakers at the LCCC Focus Period on Large-Scale and Distributed Optimization, held in Lund, 14th–16th June 2017. A source of information and innovative ideas for current and future research, this book will appeal to researchers, academics, and students who are interested in large-scale optimization.

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

?A research monograph providing a synthesis of old research on the foundations of dynamic programming, with the modern theory of approximate dynamic programming and new research on semicontractive models. It aims at a unified and economical development of the core theory and algorithms of total cost sequential decision problems, based on the strong connections of the subject with fixed point theory. The analysis focuses on the abstract mapping that underlies dynamic programming and defines the mathematical character of the associated problem. The discussion centers on two fundamental properties that this mapping may have: monotonicity and (weighted sup-norm) contraction. It turns out that the nature of the analytical and algorithmic DP theory is determined primarily by the presence or absence of these two properties, and the rest of the problem's structure is largely inconsequential. New research is focused on two areas: 1) The ramifications of these properties in the context of algorithms for approximate dynamic programming, and 2) The new class of semicontractive models, exemplified by stochastic shortest path problems, where some but not all policies are contractive. The 2nd edition aims primarily to amplify the presentation of the semicontractive models of Chapter 3 and Chapter 4 of the first (2013) edition, and to supplement it with a broad spectrum of research results that I obtained and published in journals and reports since the first edition was written (see below). As a result, the size of this material more than doubled, and the size of the book increased by nearly 40%. The book is an excellent supplement to several of our books: Dynamic Programming and Optimal Control (Athena Scientific, 2017), and Neuro-Dynamic Programming (Athena Scientific, 1996).

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.

MIMO Transceiver Design via Majorization Theory presents an up-to-date unified mathematical framework for the design of point-to-point MIMO transceivers with channel state information (CSI) at both sides of the link according to an arbitrary cost function as a measure of the system performance.

????????????????,????????????????,????????????????;????????????????

Recently Geometric Programming has been applied to study a variety of problems in the analysis and design of communication systems from information theory and queuing theory to signal processing and network protocols. Geometric Programming for Communication Systems begins its comprehensive treatment of the subject by providing an in-depth tutorial on the theory, algorithms, and modeling methods of Geometric Programming. It then gives a systematic survey of the applications of Geometric Programming to the study of communication systems. It collects in one place various published results in this area, which are currently scattered in several books and many research papers, as well as to date unpublished results. Geometric Programming for Communication Systems is intended for researchers and students who wish to have a comprehensive starting point for understanding the theory and applications of geometric programming in communication systems.

This is the leading and most up-to-date textbook on the far-ranging algorithmic methodology of Dynamic Programming, which can be used for optimal control, Markovian decision problems, planning and sequential decision making under uncertainty, and discrete/combinatorial optimization. The treatment focuses on basic unifying themes, and conceptual foundations. It illustrates the versatility, power, and generality of the method with many examples and applications from engineering, operations research, and other fields. It also addresses extensively the practical application of the methodology, possibly through the use of approximations, and provides an extensive treatment of the far-reaching methodology of Neuro-Dynamic Programming/Reinforcement Learning. Among its special features, the book 1) provides a unifying framework for sequential decision making, 2) treats simultaneously deterministic and stochastic control problems popular in modern control theory and Markovian decision popular in operations research, 3) develops the theory of deterministic optimal control problems including the Pontryagin Minimum Principle, 4) introduces recent suboptimal control and simulation-based approximation techniques (neuro-dynamic programming), which allow the practical application of dynamic programming to complex problems that involve the dual curse of large dimension and lack of an accurate mathematical model, 5) provides a comprehensive treatment of infinite horizon problems in the second volume, and an introductory treatment in the first volume.

The theory, methods and applications of matrix analysis are presented here in a novel theoretical framework.

This book reviews and discusses recent advances in the development of methods and algorithms for nonlinear optimization and its applications, focusing on the large-dimensional case, the current forefront of much research. Individual chapters, contributed by eminent authorities, provide an up-to-date overview of the field from different and complementary standpoints, including theoretical analysis, algorithmic development, implementation issues and applications.

This textbook provides students, researchers, and engineers in the area of electrical engineering with advanced mathematical optimization methods. Presented in a readable format, this book highlights fundamental concepts of advanced optimization used in electrical engineering. Chapters provide a collection that ranges from simple yet important concepts such as unconstrained optimization to highly advanced topics such as linear matrix inequalities and artificial intelligence-based optimization methodologies. The reader is motivated to engage with the content via numerous application examples of optimization in the area of electrical engineering. The book begins with an extended review of linear algebra that is a prerequisite to mathematical optimization. It then precedes with unconstrained optimization, convex programming, duality, linear matrix inequality, and intelligent optimization methods. This book can be used as the main text in courses such as Engineering Optimization, Convex Engineering Optimization, Advanced Engineering Mathematics and Robust Optimization and will be useful for practicing design engineers in electrical engineering fields. Author provided cases studies and worked examples are included for student and instructor use.

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

The core of this paper is a general set of variational principles for the problems of computing marginal probabilities and modes, applicable to multivariate statistical models in the exponential family.

A unified framework for developing planning and control algorithms for active sensing, with examples of applications for specific sensor technologies. Active sensor systems, increasingly deployed in such applications as unmanned vehicles, mobile robots, and environmental monitoring, are characterized by a high degree of autonomy, reconfigurability, and redundancy. This book is the first to offer a unified framework for the development of planning and control algorithms for active sensing, with examples of applications for a range of specific sensor technologies. The methods presented can be characterized as information-driven because their goal is to optimize the value of information, rather than to optimize traditional guidance and navigation objectives.

This compendium provides a self-contained introduction to mathematical analysis in the field of machine learning and data mining. The mathematical analysis component of the typical mathematical curriculum for computer science students omits these very important ideas and techniques which are indispensable for approaching specialized area of machine learning centered around optimization such as support vector machines, neural networks, various types of regression, feature selection, and clustering. The book is of special interest to researchers and graduate students who will benefit from these application areas discussed in the book.

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>

This treatment focuses on the analysis and algebra underlying the workings of convexity and duality and necessary/sufficient local/global optimality conditions for unconstrained and constrained optimization problems. 2015 edition.

Convex Analysis and Optimization Athena Scientific

Nature-inspired computation and swarm intelligence have become popular and effective tools for solving problems in optimization, computational intelligence, soft computing and data science. Recently, the literature in the field has expanded rapidly, with new algorithms and applications emerging. Nature-Inspired Computation and Swarm Intelligence: Algorithms, Theory and Applications is a timely reference giving a comprehensive review of relevant state-of-the-art developments in algorithms, theory and applications of nature-inspired algorithms and swarm intelligence. It reviews and documents the new developments, focusing on nature-inspired algorithms and their theoretical analysis, as well as providing a guide to their implementation. The book includes case studies of diverse real-world applications, balancing explanation of the theory with practical implementation. Nature-Inspired Computation and Swarm Intelligence: Algorithms, Theory and Applications is suitable for researchers and graduate students in computer science, engineering, data science, and management science, who want a comprehensive review of algorithms, theory and implementation within the fields of nature inspired computation and swarm intelligence. Introduces nature-inspired algorithms and their fundamentals, including: particle swarm optimization, bat algorithm, cuckoo search, firefly algorithm, flower pollination algorithm, differential evolution and genetic algorithms as well as multi-objective optimization algorithms and others Provides a theoretical foundation and analyses of algorithms, including: statistical theory and Markov chain theory on the convergence and stability of algorithms, dynamical system theory, benchmarking of optimization, no-free-lunch theorems, and a generalized mathematical framework Includes a diversity of case studies of real-world applications: feature selection, clustering and classification, tuning of restricted Boltzmann machines, travelling salesman problem, classification of white blood cells, music generation by artificial intelligence, swarm robots, neural networks, engineering designs and others

Traditional network optimization focuses on a single control objective in a network populated by obedient users and limited dispersion of information. However, most of today's networks are large-scale with lack of access to centralized information, consist of users with diverse requirements, and are subject to dynamic changes. These factors naturally motivate a new distributed control paradigm, where the network infrastructure is kept simple and the network control functions are delegated to individual agents which make their decisions independently ("selfishly"). The interaction of multiple independent decision-makers necessitates the use of game theory, including economic notions related to markets and incentives. This monograph studies game theoretic models of resource allocation among selfish agents in networks. The first part of the monograph introduces fundamental game theoretic topics. Emphasis is given to the analysis of dynamics in game theoretic situations, which is crucial for design and control of networked systems. The second part of the monograph applies the game theoretic tools for the analysis of resource allocation in communication networks. We set up a general model of routing in wireline networks, emphasizing the congestion problems caused by delay and packet loss. In particular, we develop a systematic approach to characterizing the inefficiencies of network equilibria, and highlight the effect of autonomous service providers on network performance. We then turn to examining distributed power control in wireless networks. We show that the resulting Nash equilibria can be efficient if the degree of freedom given to end-users is properly designed. Table of Contents: Static Games and Solution Concepts / Game Theory Dynamics / Wireline Network Games / Wireless Network Games / Future Perspectives

This book considers large and challenging multistage decision problems, which can be solved in principle by dynamic programming (DP), but their exact solution is computationally intractable. We discuss solution methods that rely on approximations to produce suboptimal policies with adequate performance. These methods are collectively known by several essentially equivalent names: reinforcement learning, approximate dynamic programming, neuro-dynamic programming. They have been at the forefront of research for the last 25 years, and they underlie, among others, the recent impressive successes of self-learning in the context of games such as chess and Go. Our subject has benefited greatly from the interplay of ideas from optimal control and from artificial intelligence, as it relates to reinforcement learning and simulation-based neural network methods. One of the aims of the book is to explore the common boundary between these two fields and to form a bridge that is accessible by workers with background in either field. Another aim is to organize coherently the broad mosaic of methods that have proved successful in practice while having a solid theoretical and/or logical foundation. This may help researchers and practitioners to find their way through the maze of competing ideas that constitute the current state of the art. This book relates to several of our other books: Neuro-Dynamic Programming (Athena Scientific, 1996), Dynamic Programming and Optimal Control (4th edition, Athena Scientific, 2017), Abstract Dynamic Programming (2nd edition, Athena Scientific, 2018), and Nonlinear Programming (Athena Scientific, 2016). However, the mathematical style of this book is somewhat different. While we provide a rigorous, albeit short, mathematical account of the theory of finite and infinite horizon dynamic programming, and some fundamental approximation methods, we rely more on intuitive explanations and less on proof-based insights. Moreover, our mathematical requirements are quite modest: calculus, a minimal use of matrix-vector algebra, and elementary probability (mathematically complicated arguments involving laws of large numbers and stochastic convergence are bypassed in favor of intuitive explanations).

The book illustrates the methodology with many examples and illustrations, and uses a gradual expository approach, which proceeds along four directions: (a) From exact DP to approximate DP: We first discuss exact DP algorithms, explain why they may be difficult to implement, and then use them as the basis for approximations. (b) From finite horizon to infinite horizon problems: We first discuss finite horizon exact and approximate DP methodologies, which are intuitive and mathematically simple, and then progress to infinite horizon problems. (c) From deterministic to stochastic models: We often discuss separately deterministic and stochastic problems, since deterministic problems are simpler and offer special advantages for some of our methods. (d) From model-based to model-free implementations: We first discuss model-based implementations, and then we identify schemes that can be appropriately modified to work with a simulator. The book is related and supplemented by the companion research monograph Rollout, Policy Iteration, and Distributed Reinforcement Learning (Athena Scientific, 2020), which focuses more closely on several topics related to rollout, approximate policy iteration, multiagent problems, discrete and Bayesian optimization, and distributed computation, which are either discussed in less detail or not covered at all in the present book. The author's website contains class notes, and a series of videolectures and slides from a 2021 course at ASU, which address a selection of topics from both books.

This book is the first easy-to-read text on nonsmooth optimization (NSO, not necessarily differentiable optimization). Solving these kinds of problems plays a critical role in many industrial applications and real-world modeling systems, for example in the context of image denoising, optimal control, neural network training, data mining, economics and computational chemistry and physics. The book covers both the theory and the numerical methods used in NSO and provide an overview of different problems arising in the field. It is organized into three parts: 1. convex and nonconvex analysis and the theory of NSO; 2. test problems and practical applications; 3. a guide to NSO software. The book is ideal for anyone teaching or attending NSO courses. As an accessible introduction to the field, it is also well suited as an independent learning guide for practitioners already familiar with the basics of optimization.

The study of Euclidean distance matrices (EDMs) fundamentally asks what can be known geometrically given only distance information between points in Euclidean space. Each point may represent simply location or, 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 recognition will 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, and we 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 introduced as 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 bodies is 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 explain how 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 affine function relates to convex functions as the hyperplane relates to convex sets. Here, also, pertinent results for multidimensional convex functions are presented that are largely ignored in the literature; tricks and tips for determining their convexity and discerning their geometry, particularly with regard to matrix calculus which remains largely unsystematized when 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 beyond the triangle inequality. We proceed by deriving the fifth Euclidean axiom and then explain why furthering this endeavor is 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 are presented; e.g., we generate a recognizable isotonic map of the United States using only 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 proof relies 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 devise a polyhedral cone required for determining membership of a candidate matrix (in Cayley-Menger form) to the convex cone of Euclidean distance matrices (EDM cone); i.e., 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 between the EDM cone, two positive semidefinite cones, and the ellipsope. We illustrate geometric requirements, in particular, for projection of a candidate matrix on 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, is revealed 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 that is simpler than the Schoenberg criterion. We derive a new concise expression for the EDM cone and its dual involving two subspaces and a positive semidefinite cone. "Semidefinite programming" is reviewed with particular attention to optimality conditions of 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 3×3 symmetric matrices is introduced; a tool for visualizing in 6 dimensions. In "EDM proximity" we explore methods of solution to a few fundamental and prevalent Euclidean distance matrix proximity problems; the problem of finding that Euclidean distance matrix closest to 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 Euclidean projection of a point on a subset of the positive semidefinite cone comprising all positive semidefinite matrices having rank not exceeding a prescribed limit ρ . We explain how this problem is transformed to a convex optimization for any rank ρ .

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

This book provides a comprehensive and accessible presentation of algorithms for solving continuous optimization problems. It relies on rigorous mathematical analysis, but also aims at an intuitive exposition that makes use of visualization where possible. It places particular emphasis on modern developments, and their widespread applications in fields such as large-scale resource allocation problems, signal processing, and machine learning. The 3rd edition brings the book in closer harmony with the companion works Convex Optimization Theory (Athena Scientific, 2009), Convex Optimization Algorithms (Athena Scientific, 2015), Convex Analysis and Optimization (Athena Scientific, 2003), and Network Optimization (Athena Scientific, 1998). These works are complementary in that they deal primarily with convex, possibly nondifferentiable, optimization problems and rely on convex analysis. By contrast the nonlinear programming book focuses primarily on analytical and computational methods for possibly nonconvex differentiable problems. It relies primarily on calculus and variational analysis, yet it still contains a detailed presentation of duality theory and its uses for both convex and nonconvex problems. This online edition contains detailed solutions to all the theoretical book exercises. Among its special features, the book: Provides extensive coverage of iterative optimization methods within a unifying framework Covers in depth duality theory from both a variational and a geometric point of view Provides a detailed treatment of interior point methods for linear programming Includes much new material on a number of topics, such as proximal algorithms, alternating direction methods of multipliers, and conic programming Focuses on large-scale optimization topics of much current interest, such as first order methods, incremental methods, and distributed asynchronous computation, and their applications in machine learning, signal processing, neural network training, and big data applications Includes a large number of examples and exercises Was developed through extensive classroom use in first-year graduate courses

The purpose of this book is to develop in greater depth some of the methods from the author's Reinforcement Learning and Optimal Control recently published textbook (Athena Scientific, 2019). In particular, we present new research, relating to systems involving multiple agents, partitioned architectures, and distributed asynchronous computation. We pay special attention to the contexts of dynamic programming/policy iteration and control theory/model predictive control. We also discuss in some detail the application of the methodology to challenging discrete/combinatorial optimization problems, such as routing, scheduling, assignment, and mixed integer programming, including the use of neural network approximations within these contexts. The book focuses on the fundamental idea of policy iteration, i.e., start from some policy, and successively generate one or more improved policies. If just one improved policy is generated, this is called rollout, which, based on broad and consistent computational experience, appears to be one of the most versatile and reliable of all reinforcement learning methods. In this book, rollout algorithms are developed for both discrete deterministic and stochastic DP problems, and the development of distributed implementations in both multiagent and multiprocessor settings, aiming to take advantage of parallelism. Approximate policy iteration is more ambitious than rollout, but it is a strictly off-line method, and it is generally far more computationally intensive. This motivates the use of parallel and distributed computation. One of the purposes of the monograph is to discuss distributed (possibly asynchronous) methods that relate to rollout and policy iteration, both in the context of an exact and an approximate implementation involving neural networks or other approximation architectures. Much of the new research is inspired by the remarkable AlphaZero chess program, where policy iteration, value and policy networks, approximate lookahead minimization, and parallel computation all play an important role.

Performance optimization is vital in the design and operation of modern engineering systems, including communications, manufacturing, robotics, and logistics. Most engineering systems are too complicated to model, or the system parameters cannot be easily identified, so learning techniques have to be applied. This book provides a unified framework based on a sensitivity point of view. It also introduces new approaches and proposes new research topics within this sensitivity-based framework. This new perspective on a popular topic is presented by a well respected expert in the field.

[Copyright: b35dfe310262793518670df870506556](https://www.amazon.com/dp/B035DFE310)