Additional Exercises For Convex Optimization Boyd Solutions

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

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

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 book is an abridged version of the two volumes "Convex Analysis and Minimization Algorithms I and II" (Grundlehren der mathematischen Wissenschaften Vol. 305 and 306). It presents an introduction to the basic concepts in convex analysis and a study of convex minimization problems (with an emphasis on numerical algorithms). The "backbone" of bot volumes was extracted, some material deleted which was deemed too advanced for an introduction, or too closely attached to numerical algorithms. Some exercises were included and finally the index has been considerably enriched, making it an excellent choice for the purpose of learning and teaching.

Praise for the Third Edition ". . . guides and leads the reader through the learning path . . . [e]xamples are stated very clearly and the results are presented with attention to detail." —MAA Reviews Fully updated to reflect new developments in the field, the Fourth Edition of Introduction to Optimization fills the need for accessible treatment of optimization theory and methods with an emphasis on engineering design. Basic definitions and notations are provided in addition to the related fundamental background for linear algebra, geometry, and calculus. This new edition explores the essential topics of unconstrained optimization problems, linear programming problems, and nonlinear constrained optimization. The authors also present an optimization perspective on global search methods and include discussions on genetic algorithms, particle swarm optimization, and the simulated annealing algorithm. Featuring an elementary introduction to artificial neural networks, convex optimization, and multi-objective optimization, the Fourth Edition also offers: A new chapter on integer programming Expanded coverage of one-dimensional methods Updated and expanded sections on linear matrix inequalities Numerous new exercises at the end of each chapter MATLAB exercises and drill problems to reinforce the discussed theory and algorithms Numerous diagrams and figures that complement the written presentation of key concepts MATLAB M-files for implementation of the discussed theory and algorithms (available via the book's website) Introduction to Optimization, Fourth Edition is an ideal textbook for courses on optimization theory and methods. In addition, the book is a useful reference for professionals in mathematics, operations research, electrical engineering, economics, statistics, and business.

Foundations of Data Science

First-order and Stochastic Optimization Methods for Machine Learning

Optimization Methods in Finance

Convex Optimization of Power Systems

Interior-point Polynomial Algorithms in Convex Programming

Convex Optimization in Normed Spaces

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

The primary goal of this book is to provide a self-contained, comprehensive study of the main ?rst-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.

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.

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

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.

An Introduction to Optimization

Analysis, Algorithms, and Engineering Applications

Convex Analysis for Optimization

Optimization for Decision Making

First-Order Methods in Optimization

Theory, Algorithms, and Applications with MATLAB

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

Discover New Methods for Dealing with High-Dimensional Data A sparse statistical model has only a small number of nonzero parameters or weights; therefore, it is much easier to estimate and interpret than a dense model. Statistical Learning with Sparsity: The Lasso and Generalizations presents methods that exploit sparsity to help recover the underlying signal in a set of data. Top experts in this rapidly evolving field, the authors describe the lasso for linear regression and a simple coordinate descent algorithm for its computation. They discuss the application of l1 penalties to generalized linear models and support vector machines, cover generalized penalties such as the elastic net and group lasso, and review numerical methods for optimization. They also present statistical inference methods for fitted (lasso) models, including the bootstrap, Bayesian methods, and recently developed approaches. In addition, the book examines matrix decomposition, sparse multivariate analysis, graphical models, and compressed sensing. It concludes with a survey of theoretical results for the lasso. In this age of big data, the number of features measured on a person or object can be large and might be larger than the number of observations. This book shows how the sparsity assumption allows us to tackle these problems and extract useful and reproducible patterns from big datasets. Data analysts, computer scientists, and theorists will appreciate this thorough and up-to-date treatment of sparse statistical modeling.

It was in the middle of the 1980s, when the seminal paper by Kar markar opened a new epoch in nonlinear optimization. The importance of this paper, containing a new polynomial-time algorithm for linear op timization problems, was not only in its complexity bound. At that time, the most surprising feature of this algorithm was that the theoretical pre diction of its high efficiency was supported by excellent computational results. This unusual fact dramatically changed the style and direc tions of the research in nonlinear optimization. Thereafter it became more and more common that the new methods were provided with a complexity analysis, which was considered a better justification of their efficiency than computational experiments. In a new rapidly develop ing field, which got the name "polynomial-time interior-point methods", such a justification was obligatory. Afteralmost fifteen years of intensive research, the main results of this development started to appear in monographs [12, 14, 16, 17, 18, 19]. Approximately at that time the author was asked to prepare a new course on nonlinear optimization for graduate students. The idea was to create a course which would reflect the new developments in the field. Actually, this was a major challenge. At the time only the theory of interior-point methods for linear optimization was polished enough to be explained to students. The general theory of self-concordant functions had appeared in print only once in the form of research monograph [12].

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.

From its origins in the minimization of integral functionals, the notion of variations has evolved greatly in connection with applications in optimization, equilibrium, and control. This book develops a unified framework and provides a detailed exposition of variational geometry and subdifferential calculus in their current forms beyond classical and convex analysis. Also covered are set-convergence, set-valued mappings, epi-convergence, duality, and normal integrands.

Variational Analysis

Numerical Optimization

Convex Optimization Theory

A Gentle Introduction to Optimization

Lectures on Modern Convex Optimization

Introduction to Nonlinear Optimization

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

This book covers not only foundational materials but also the most recent progresses made during the past few years on the area of machine learning algorithms. In spite of the intensive research and development in this area, there does not exist a systematic treatment to introduce the fundamental concepts and recent progresses on machine learning algorithms, especially on those based on stochastic optimization methods, randomized algorithms, nonconvex optimization, distributed and online learning, and projection free methods. This book will benefit the broad audience in the area of machine learning, artificial intelligence and mathematical programming community by presenting these recent developments in a tutorial style, starting from the basic building blocks to the most carefully designed and complicated algorithms for machine learning.

Explore the theoretical foundations and real-world power system applications of convex programming In Mathematical Programming for Power System Operation with Applications in Python, Professor Alejandro Garces delivers a comprehensive overview of power system operations models with a focus on convex optimization models and their implementation in Python. Divided into two parts, the book begins with a theoretical analysis of convex optimization models before moving on to related applications in power systems operations. The author eschews concepts of topology and functional analysis found in more mathematically oriented books in favor of a more natural approach. Using this perspective, he presents recent applications of convex optimization in power system operations problems. Mathematical Programming for Power System Operation with Applications in Python uses Python and CVXPY as tools to solve power system optimization problems and includes models that can be solved with the presented framework. The book also includes: A thorough introduction to power system operation, including economic and environmental dispatch, optimal power flow, and hosting capacity Comprehensive explorations of the mathematical background of power system operation, including quadratic forms and norms and the basic theory of optimization Practical discussions of convex functions and convex sets, including affine and linear spaces, politopes, balls, and ellipsoids In-depth examinations of convex optimization, including global optimums, and first and second order conditions Perfect for undergraduate students with some knowledge in power systems analysis, generation, or distribution, Mathematical Programming for Power System Operation with Applications in Python is also an ideal resource for graduate students and engineers practicing in the area of power system optimization.

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

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 Basic Course

From Theory to Applications in Python
Vectors, Matrices, and Least Squares
Multi-Period Trading Via Convex Optimization
Ant Colony Optimization

Linear programming (LP), modeling, and optimization are very much the fundamentals of OR, and no academic program is complete without them. No matter how highly developed one's LP skills are, however, if a fine appreciation for modeling isn't developed to make the best use of those skills, then the truly 'best solutions' are often not realized, and efforts go wasted. Katta Murty studied LP with George Dantzig, the father of linear programming, and has written the graduate-level solution to that problem. While maintaining the rigorous LP instruction required, Murty's new book is unique in his focus on developing modeling skills to support valid decision making for complex real world problems. He describes the approach as 'intelligent modeling and decision making' to emphasize the importance of employing the best expression of actual problems and then applying the most computationally effective and efficient solution technique for that model.

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

Examines in detail those topics in convex geometry that are concerned with Euclidean space Enriched by numerous examples, illustrations, and exercises, with a good bibliography and index Requires only a basic knowledge of geometry, linear algebra, analysis, topology, and measure theory Can be used for graduates courses or seminars in convex geometry, geometric and convex combinatorics, and convex analysis and optimization

This book has grown out of lectures and courses given at Linköping University, Sweden, over a period of 15 years. It gives an introductory treatment of problems and methods of structural optimization. The three basic classes of geometrical - timization problems of mechanical structures, i. e. , size, shape and topology op- mization, are treated. The focus is on concrete numerical solution methods for d- crete and (?nite element) discretized linear elastic structures. The style is explicit and practical: mathematical proofs are provided when arguments can be kept e- mentary but are otherwise only cited, while implementation details are frequently provided. Moreover, since the text has an emphasis on geometrical design problems, where the design is represented by continuously varying—frequently very many— variables, so-called ?rst order methods are central to the treatment. These methods are based on sensitivity analysis, i. e. , on establishing ?rst order derivatives for - jectives and constraints. The classical ?rst order methods that we emphasize are CONLIN and MMA, which are based on explicit, convex and separable appro- mations. It should be remarked that the classical and frequently used so-called op- mality criteria method is also of this kind. It may also be noted in this context that zero order methods such as response surface methods, surrogate models, neural n- works, genetic algorithms, etc. , essentially apply to different types of problems than the ones treated here and should be presented elsewhere.

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).

Introduction to Applied Linear Algebra
Finite Dimensional Convexity and Optimization
Foundations, Algorithms, and Applications
The Lasso and Generalizations
Machine Learning Refined
Selected Topics in Convex Geometry

This book provides an introduction to the mathematical and algorithmic foundations of data science, including machine learning, high-dimensional geometry, and analysis of large networks. Topics include the counterintuitive nature of data in high dimensions, important linear algebraic techniques such as singular value decomposition, the theory of random walks and Markov chains, the fundamentals of and important algorithms for machine learning, algorithms and analysis for clustering, probabilistic models for large networks, representation learning including topic modelling and non-negative matrix factorization, wavelets and compressed sensing. Important probabilistic techniques are developed including the law of large numbers, tail inequalities, analysis of random projections, generalization guarantees in machine learning, and moment methods for analysis of phase transitions in large random graphs. Additionally, important structural and complexity measures are discussed such as matrix norms and VC-dimension. This book is suitable for both undergraduate and graduate courses in the design and analysis of algorithms for data.

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

This book discusses convex analysis, the basic underlying structure of argumentation in economic theory. Convex analysis is also common to the optimization of problems encountered in many applications. The text is aimed at senior undergraduate students, graduate students, and specialists of mathematical programming who are undertaking research into applied mathematics and economics. The text consists of a systematic development in eight chapters, and contains exercises. The book is appropriate as a class text or for self-study.

This monograph collects in one place the basic definitions, 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.

This work is intended to serve as a guide for graduate students and researchers who wish to get acquainted with the main theoretical and practical tools for the numerical minimization of convex functions on Hilbert spaces. Therefore, it contains the main tools that are necessary to conduct independent research on the topic. It is also a concise, easy-to-follow and self-contained textbook, which may be useful for any researcher working on related fields, as well as teachers giving graduate-level courses on the topic. It will contain a thorough revision of the extant literature including both classical and state-of-the-art references.

An Introduction to Structural Optimization
Introduction to Online Convex Optimization
Convex Optimization
Mathematics for Machine Learning
Optimization Models
Convex Optimization with Computational Errors

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

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

The fundamental mathematical tools needed to understand machine learning include linear algebra, analytic geometry, matrix decompositions, vector calculus, optimization, probability and statistics. These topics are traditionally taught in disparate courses, making it hard for data science or computer science students, or professionals, to efficiently learn the mathematics. This self-contained textbook bridges the gap between mathematical and machine learning texts, introducing the mathematical concepts with a minimum of prerequisites. It uses these concepts to derive four central machine learning methods: linear regression, principal component analysis, Gaussian mixture models and support vector machines. For students and others with a mathematical background, these derivations provide a starting point to machine learning texts. For those learning the mathematics for the first time, the methods help build intuition and practical experience with applying mathematical concepts. Every chapter includes worked examples and exercises to test understanding. Programming tutorials are offered on the book's web site.

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

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

Convex Analysis and Nonlinear Optimization
Proximal Algorithms
Convex Optimization & Euclidean Distance Geometry
Linear Matrix Inequalities in System and Control Theory
Convex Analysis and Optimization
Fundamentals of Convex Analysis

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

Convex OptimizationCambridge University Press

The book is devoted to the study of approximate solutions of optimization problems in the presence of computational errors. It contains a number of results on the convergence behavior of algorithms in a Hilbert space, which are known as important tools for solving optimization problems. The research presented in the book is the continuation and the further development of the author's (c) 2016 book Numerical Optimization with Computational Errors, Springer 2016. Both books study the algorithms taking into account computational errors which are always present in practice. The main goal is, for a known computational error, to find out what an approximate solution can be obtained and how many iterates one needs for this. The main difference between this new book and the 2016 book is that in this present book the discussion takes into consideration the fact that for every algorithm, its iteration consists of several steps and that computational errors for different steps are generally, different. This fact, which was not taken into account in the previous book, is indeed important in practice. For example, the subgradient projection algorithm consists of two steps. The first step is a calculation of a subgradient of the objective function while in the second one we calculate a projection on the feasible set. In each of these two steps there is a computational error and these two computational errors are different in general. It may happen that the feasible set is simple and the objective function is complicated. As a result, the computational error, made when one calculates the projection, is essentially smaller than the computational error of the calculation of the subgradient. Clearly, an opposite case is possible too. Another feature of this book is a study of a number of important algorithms which appeared recently in the literature and which are not discussed in the previous book. This monograph contains 12 chapters. Chapter 1 is an introduction. In Chapter 2 we study the subgradient projection algorithm for minimization of convex and nonsmooth functions. We generalize the results of [NOCE] and establish results which has no prototype in [NOCE]. In Chapter 3 we analyze the mirror descent algorithm for minimization of convex and nonsmooth functions, under the presence of computational errors. For this algorithm each iteration consists of two steps. The first step is a calculation of a subgradient of the objective function while in the second one we solve an auxiliary minimization problem on the set of feasible points. In each of these two steps there is a computational error. We generalize the results of [NOCE] and establish results which has no prototype in [NOCE]. In Chapter 4 we analyze the projected gradient algorithm with a smooth objective function under the presence of computational errors. In Chapter 5 we consider an algorithm, which is an extension of the projection gradient algorithm used for solving linear inverse problems arising in signal/image processing. In Chapter 6 we study continuous subgradient method and continuous subgradient projection algorithm for minimization of convex nonsmooth functions and for computing the saddle points of convex-concave functions, under the presence of computational errors. All the results of this chapter has no prototype in [NOCE]. In Chapters 7-12 we analyze several algorithms under the presence of computational errors which were not considered in [NOCE]. Again, each step of an iteration has a computational errors and we take into account that these errors are, in general, different. An optimization problems with a composite objective function is studied in Chapter 7. A zero-sum game with two-players is considered in Chapter 8. A predicted decrease approximation-based method is used in Chapter 9 for constrained convex optimization. Chapter 10 is devoted to minimization of quasiconvex functions. Minimization of sharp weakly convex functions is discussed in Chapter 11. Chapter 12 is devoted to a generalized projected subgradient method for minimization of a convex function over a set which is not necessarily convex. The book is of interest for researchers and engineers working in optimization. It also can be useful in preparation courses for graduate students. The main feature of the book which appeals specifically to this audience is the study of the influence of computational errors for several important optimization algorithms. The book is of interest for experts in applications of optimization to engineering and economics.

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

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.

Semidefinite Optimization and Convex Algebraic Geometry
Theory and Examples
Linear and Quadratic Models
Mathematical Programming for Power Systems Operation
Introductory Lectures on Convex Optimization
Optimization in Practice with MATLAB

*A modern, up-to-date introduction to optimization theory andmethods This authoritative book serves as an introductory text tooptimization at the senior undergraduate and beginning graduatelevels. With consistently accessible and elementary treatment ofall topics, An Introduction to Optimization, Second Edition helpsstudents build a solid working knowledge of the field, includingunconstrained optimization, linear programming, and constrainedoptimization. Supplemented with more than one hundred tables and illustrations,an extensive bibliography, and numerous worked examples toillustrate both theory and algorithms, this book alsoprovides: * A review of the required mathematical background material * A mathematical discussion at a level accessible to MBA andbusiness students * A treatment of both linear and nonlinear programming * An introduction to recent developments, including neuralnetworks, genetic algorithms, and interior-point methods * A chapter on the use of descent algorithms for the training offeedforward neural networks * Exercise problems after every chapter, many new to thisedition * MATLAB(r) exercises and examples * Accompanying Instructor's Solutions Manual available onrequest An Introduction to Optimization, Second Edition helps studentsprepare for the advanced topics and technological developments thatlie ahead. It is also a useful book for researchers andprofessionals in mathematics, electrical engineering, economics,statistics, and business. An Instructor's Manual presenting detailed solutions to all theproblems in the book is available from the Wiley editorialdepartment.*

Statistical Learning with Sparsity
A Unified Approach
Theory, Methods and Examples