Additional Exercises For Convex Optimization Solutions

A textbook for an undergraduate course in mathematical programming for students with a knowledge of elementary real analysis, linear algebra, and classical linear programming (simple techniques). Focuses on the computation and characterization of global optima of nonlinear functions, rather than the locally optimal solutions addressed by most books on optimization. Incorporates the theoretical, algorithmic, and computational advances of the past three decades that help solve globally multi-extreme problems in the mathematical modeling of real world systems. Annotation copyright by Book News, Inc., Portland, OR

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.

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.
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.
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 on-line 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
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.
The new edition of this book presents a comprehensive and up-to-date description of the most effective methods in continuous optimization. It responds to the growing interest in optimization in engineering, science, and business by focusing on methods best suited to practical problems. This edition has been thoroughly updated throughout. There are new chapters on nonlinear interior methods and derivative-free methods for optimization, both of which are widely used in practice and are the focus of much current research. Because of the emphasis on practical methods, as well as the extensive illustrations and exercises, the book is accessible to a wide audience.
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.
A comprehensive introduction to the tools, techniques and applications of convex optimization.
An insightful, concise, and rigorous treatment of the basic theory of convex sets and functions in finite dimensions, and the analytical/geometrical foundations of convex optimization and duality theory. Convexity theory is first developed in a simple accessible manner, using easily visualized proofs. Then the focus shifts to a transparent geometrical line of analysis to develop the fundamental duality between descriptions of convex functions in terms of points, and in terms of hyperplanes. Finally, convexity theory and abstract duality are applied to problems of constrained optimization, Fenchel and conic duality, and game theory to develop the sharpest possible duality results within a highly visual geometric framework. This on-line version of the book, includes an extensive set of theoretical problems with detailed high-quality solutions, which significantly extend the range and value of the book. The book may be used as a text for a theoretical convex optimization course; the author has taught several variants of such a course at MIT and elsewhere over the last ten years. It may also be used as a supplementary source for nonlinear programming classes, and as a theoretical foundation for classes focused on convex optimization models (rather than theory). It is an excellent supplement to several of our books: Convex Optimization Algorithms (Athena Scientific, 2015), Nonlinear Programming (Athena Scientific, 2017), Network Optimization(Athena Scientific, 1998), Introduction to Linear Optimization (Athena Scientific, 1997), and Network Flows and Monotropic Optimization (Athena Scientific, 1998).

This 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 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.

This textbook is a completely revised, updated, and expanded English edition of the important Analyse fonctionnelle (1983). In addition, it contains a wealth of problems and exercises (with solutions) to guide the reader. Uniquely, this book presents in a coherent, concise and unified way the main results from functional analysis together with the main results from the theory of partial differential equations (PDEs). Although there are many books on functional analysis and many on PDEs, this is the first to cover both of these closely connected topics. Since the French book was first published, it has been translated into Spanish, Italian, Japanese, Korean, Romanian, Greek and Chinese. The English edition makes a welcome addition to this list.

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.

Distills key concepts from linear algebra, geometry, matrices, calculus, optimization, probability and statistics that are used in machine learning.

Fully describes optimization methods that are currently most valuable in solving real-life problems. Since optimization has applications in almost every branch of science and technology, the text emphasizes their practical aspects in conjunction with the heuristics useful in making them perform more reliably and efficiently. To this end, it presents comparative numerical studies to give readers a feel for possibile applications and to illustrate the problems in assessing evidence. Also provides theoretical background which provides insights into how methods are derived. This edition offers revised coverage of basic theory and standard techniques, with updated discussions of line search methods, Newton and quasi-Newton methods, and conjugate direction methods, as well as a comprehensive treatment of restricted step or trust region methods not commonly found in the literature. Also includes recent developments in hybrid methods for nonlinear least squares; an extended discussion of linear programming, with new methods for stable updating of LU factors; and a completely new section on network programming. Chapters include computer subroutines, worked examples, and study questions.

The power grid can be considered one of twentieth-century engineering?s greatest achievements, and as grids and populations grow, robustness is a factor that planners must take into account. Power grid robustness is a complex problem for two reasons:

the underlying physics is mathematically complex, and modeling is complicated by lack of accurate data. This book sheds light on this complex problem by introducing the engineering details of power grid operations from the basic to the detailed; describing how to use optimization and stochastic modeling, with special focus on the modeling of cascading failures and robustness; providing numerical examples that show ?how things work?; and detailing the application of a number of optimization theories to power grids.÷

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.

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.

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

The goal of this book is to present the main ideas and techniques in the field of continuous smooth and nonsmooth optimization. Starting with the case of differentiable data and the classical results on constrained optimization problems, and continuing with the topic of nonsmooth objects involved in optimization theory, the book concentrates on both theoretical and practical aspects of this field. This book prepares those who are engaged in research by giving repeated insights into ideas that are subsequently dealt with and illustrated in detail.

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.

Non-convex Optimization for Machine Learning takes an in-depth look at the basics of non-convex optimization with applications to machine learning. It introduces the rich literature in this area, as well as equips the reader with the tools and techniques needed to apply and analyze simple but powerful procedures for non-convex problems. Non-convex Optimization for Machine Learning is as self-contained as possible while not losing focus of the main topic of non-convex optimization techniques. The monograph initiates the discussion with entire chapters devoted to presenting a tutorial-like treatment of basic concepts in convex analysis and optimization, as well as their non-convex counterparts. The monograph concludes with a look at four interesting applications in the areas of machine learning and signal processing, and exploring how the non-convex optimization techniques introduced earlier can be used to solve these problems. The monograph also contains, for each of the topics discussed, exercises and figures designed to engage the reader, as well as extensive bibliographic notes pointing towards classical works and recent advances. Non-convex Optimization for Machine Learning can be used for a semester-length course on the basics of non-convex optimization with applications to machine learning. On the other hand, it is also possible to cherry pick individual portions, such the chapter on sparse recovery, or the EM algorithm, for inclusion in a broader course. Several courses such as those in machine learning, optimization, and signal processing may benefit from the inclusion of such topics.

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

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

A practical guide to problem solving using MATLAB. Designed to complement a taught course introducing MATLAB but ideally suited for any beginner. This book provides a brief tour of some of the tasks that MATLAB is perfectly suited to instead of focusing on any particular topic. Providing instruction, guidance and a large supply of exercises, this book is meant to stimulate problem-solving skills rather than provide an in-depth knowledge of the MATLAB language.

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

This Fourth Edition introduces the latest theory and applications in optimization. It emphasizes constrained optimization, beginning with a substantial treatment of linear programming and then proceeding to convex analysis, network flows, integer programming, quadratic programming, and convex optimization. Readers will discover a host of practical business applications as well as non-business applications. Topics are clearly developed with many numerical examples worked out in detail. Specific examples and concrete algorithms precede more abstract topics. With its focus on solving practical problems, the book features free C programs to implement the major algorithms covered, including the two-phase simplex method, primal-dual simplex method, path-following interior-point method, and homogeneous self-dual methods. In addition, the author provides online JAVA applets that illustrate various pivot rules and variants of the simplex method, both for linear programming and for network flows. These C programs and JAVA tools can be found on the book's website. The website also includes new online instructional tools and exercises.

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.

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.

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

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

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.

Functional analysis owes much of its early impetus to problems that arise in the calculus of variations. In turn, the methods developed there have been applied to optimal control, an area that also requires new tools, such as nonsmooth analysis. This self-contained textbook gives a complete course on all these topics. It is written by a leading specialist who is also a noted expositor. This book provides a thorough introduction to functional analysis and includes many novel elements as well as the standard topics. A short course on nonsmooth analysis and geometry completes the first half of the book whilst the second half concerns the calculus of variations and optimal control. The author provides a comprehensive course on these subjects, from their inception through to the present. A notable feature is the inclusion of recent, unifying developments on regularity, multiplier rules, and the Pontryagin maximum principle, which appear here for the first time in a textbook. Other major themes include existence and Hamilton-Jacobi methods. The many substantial examples, and the more than three hundred exercises, treat such topics as viscosity solutions, nonsmooth Lagrangians, the logarithmic Sobolev inequality, periodic trajectories, and systems theory. They also touch lightly upon several fields of application: mechanics, economics, resources, finance, control engineering. Functional Analysis, Calculus of Variations and Optimal Control is intended to support several different courses at the first-year or second-year graduate level, on functional analysis, on the calculus of variations and optimal control, or on some combination. For this reason, it has been organized with customization in mind. The text also has considerable value as a reference. Besides its advanced results in the calculus of variations and optimal control, its polished presentation of certain other topics (for example convex analysis, measurable selections, metric regularity, and nonsmooth analysis) will be appreciated by researchers in these and related fields.

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.

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.

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

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.