

Convex Optimization

Stephen Boyd

*Department of Electrical Engineering
Stanford University*

Lieven Vandenberghe

*Electrical Engineering Department
University of California, Los Angeles*



CAMBRIDGE
UNIVERSITY PRESS

PUBLISHED BY THE PRESS SYNDICATE OF THE UNIVERSITY OF CAMBRIDGE
The Pitt Building, Trumpington Street, Cambridge, United Kingdom

CAMBRIDGE UNIVERSITY PRESS

The Edinburgh Building, Cambridge, CB2 2RU, UK

40 West 20th Street, New York, NY 10011-4211, USA

477 Williamstown Road, Port Melbourne, VIC 3207, Australia

Ruiz de Alarcón 13, 28014 Madrid, Spain

Dock House, The Waterfront, Cape Town 8001, South Africa

<http://www.cambridge.org>

© Cambridge University Press 2004

This book is in copyright. Subject to statutory exception
and to the provisions of relevant collective licensing agreements,
no reproduction of any part may take place without
the written permission of Cambridge University Press.

First published 2004

Reprinted with corrections 2004

Printed in the United Kingdom at the University Press, Cambridge

Typeset in Computer Modern Roman using L^AT_EX

A catalogue record for this book is available from the British Library

Library of Congress Cataloguing-in-Publication data

Boyd, Stephen P.

Convex Optimization / Stephen Boyd & Lieven Vandenberghe

p. cm.

Includes bibliographical references and index.

ISBN 0 521 83378 7

1. Mathematical optimization. 2. Convex functions. I. Vandenberghe, Lieven. II.
Title.

QA402.5.B69 2004

519.6-dc22 2003063284

For

Anna, Nicholas, and Nora

Daniël and Margriet

Contents

Preface	xi
1 Introduction	1
1.1 Mathematical optimization	1
1.2 Least-squares and linear programming	4
1.3 Convex optimization	7
1.4 Nonlinear optimization	9
1.5 Outline	11
1.6 Notation	14
Bibliography	16
I Theory	19
2 Convex sets	21
2.1 Affine and convex sets	21
2.2 Some important examples	27
2.3 Operations that preserve convexity	35
2.4 Generalized inequalities	43
2.5 Separating and supporting hyperplanes	46
2.6 Dual cones and generalized inequalities	51
Bibliography	59
Exercises	60
3 Convex functions	67
3.1 Basic properties and examples	67
3.2 Operations that preserve convexity	79
3.3 The conjugate function	90
3.4 Quasiconvex functions	95
3.5 Log-concave and log-convex functions	104
3.6 Convexity with respect to generalized inequalities	108
Bibliography	112
Exercises	113

4	Convex optimization problems	127
4.1	Optimization problems	127
4.2	Convex optimization	136
4.3	Linear optimization problems	146
4.4	Quadratic optimization problems	152
4.5	Geometric programming	160
4.6	Generalized inequality constraints	167
4.7	Vector optimization	174
	Bibliography	188
	Exercises	189
5	Duality	215
5.1	The Lagrange dual function	215
5.2	The Lagrange dual problem	223
5.3	Geometric interpretation	232
5.4	Saddle-point interpretation	237
5.5	Optimality conditions	241
5.6	Perturbation and sensitivity analysis	249
5.7	Examples	253
5.8	Theorems of alternatives	258
5.9	Generalized inequalities	264
	Bibliography	272
	Exercises	273
II	Applications	289
6	Approximation and fitting	291
6.1	Norm approximation	291
6.2	Least-norm problems	302
6.3	Regularized approximation	305
6.4	Robust approximation	318
6.5	Function fitting and interpolation	324
	Bibliography	343
	Exercises	344
7	Statistical estimation	351
7.1	Parametric distribution estimation	351
7.2	Nonparametric distribution estimation	359
7.3	Optimal detector design and hypothesis testing	364
7.4	Chebyshev and Chernoff bounds	374
7.5	Experiment design	384
	Bibliography	392
	Exercises	393

8	Geometric problems	397
8.1	Projection on a set	397
8.2	Distance between sets	402
8.3	Euclidean distance and angle problems	405
8.4	Extremal volume ellipsoids	410
8.5	Centering	416
8.6	Classification	422
8.7	Placement and location	432
8.8	Floor planning	438
	Bibliography	446
	Exercises	447
III	Algorithms	455
9	Unconstrained minimization	457
9.1	Unconstrained minimization problems	457
9.2	Descent methods	463
9.3	Gradient descent method	466
9.4	Steepest descent method	475
9.5	Newton's method	484
9.6	Self-concordance	496
9.7	Implementation	508
	Bibliography	513
	Exercises	514
10	Equality constrained minimization	521
10.1	Equality constrained minimization problems	521
10.2	Newton's method with equality constraints	525
10.3	Infeasible start Newton method	531
10.4	Implementation	542
	Bibliography	556
	Exercises	557
11	Interior-point methods	561
11.1	Inequality constrained minimization problems	561
11.2	Logarithmic barrier function and central path	562
11.3	The barrier method	568
11.4	Feasibility and phase I methods	579
11.5	Complexity analysis via self-concordance	585
11.6	Problems with generalized inequalities	596
11.7	Primal-dual interior-point methods	609
11.8	Implementation	615
	Bibliography	621
	Exercises	623

Appendices	631
A Mathematical background	633
A.1 Norms	633
A.2 Analysis	637
A.3 Functions	639
A.4 Derivatives	640
A.5 Linear algebra	645
Bibliography	652
B Problems involving two quadratic functions	653
B.1 Single constraint quadratic optimization	653
B.2 The S-procedure	655
B.3 The field of values of two symmetric matrices	656
B.4 Proofs of the strong duality results	657
Bibliography	659
C Numerical linear algebra background	661
C.1 Matrix structure and algorithm complexity	661
C.2 Solving linear equations with factored matrices	664
C.3 LU, Cholesky, and LDL ^T factorization	668
C.4 Block elimination and Schur complements	672
C.5 Solving underdetermined linear equations	681
Bibliography	684
References	685
Notation	697
Index	701

Preface

This book is about *convex optimization*, a special class of mathematical optimization problems, which includes least-squares and linear programming problems. It is well known that least-squares and linear programming problems have a fairly complete theory, arise in a variety of applications, and can be solved numerically very efficiently. The basic point of this book is that the same can be said for the larger class of convex optimization problems.

While the mathematics of convex optimization has been studied for about a century, several related recent developments have stimulated new interest in the topic. The first is the recognition that interior-point methods, developed in the 1980s to solve linear programming problems, can be used to solve convex optimization problems as well. These new methods allow us to solve certain new classes of convex optimization problems, such as semidefinite programs and second-order cone programs, almost as easily as linear programs.

The second development is the discovery that convex optimization problems (beyond least-squares and linear programs) are more prevalent in practice than was previously thought. Since 1990 many applications have been discovered in areas such as automatic control systems, estimation and signal processing, communications and networks, electronic circuit design, data analysis and modeling, statistics, and finance. Convex optimization has also found wide application in combinatorial optimization and global optimization, where it is used to find bounds on the optimal value, as well as approximate solutions. We believe that many other applications of convex optimization are still waiting to be discovered.

There are great advantages to recognizing or formulating a problem as a convex optimization problem. The most basic advantage is that the problem can then be solved, very reliably and efficiently, using interior-point methods or other special methods for convex optimization. These solution methods are reliable enough to be embedded in a computer-aided design or analysis tool, or even a real-time reactive or automatic control system. There are also theoretical or conceptual advantages of formulating a problem as a convex optimization problem. The associated dual problem, for example, often has an interesting interpretation in terms of the original problem, and sometimes leads to an efficient or distributed method for solving it.

We think that convex optimization is an important enough topic that everyone who uses computational mathematics should know at least a little bit about it. In our opinion, convex optimization is a natural next topic after advanced linear algebra (topics like least-squares, singular values), and linear programming.

Goal of this book

For many general purpose optimization methods, the typical approach is to just try out the method on the problem to be solved. The full benefits of convex optimization, in contrast, only come when the problem is known ahead of time to be convex. Of course, many optimization problems are not convex, and it can be difficult to recognize the ones that are, or to reformulate a problem so that it is convex.

Our main goal is to help the reader develop a working knowledge of convex optimization, i.e., to develop the skills and background needed to recognize, formulate, and solve convex optimization problems.

Developing a working knowledge of convex optimization can be mathematically demanding, especially for the reader interested primarily in applications. In our experience (mostly with graduate students in electrical engineering and computer science), the investment often pays off well, and sometimes very well.

There are several books on linear programming, and general nonlinear programming, that focus on problem formulation, modeling, and applications. Several other books cover the theory of convex optimization, or interior-point methods and their complexity analysis. This book is meant to be something in between, a book on general convex optimization that focuses on problem formulation and modeling.

We should also mention what this book is *not*. It is not a text primarily about convex analysis, or the mathematics of convex optimization; several existing texts cover these topics well. Nor is the book a survey of algorithms for convex optimization. Instead we have chosen just a few good algorithms, and describe only simple, stylized versions of them (which, however, do work well in practice). We make no attempt to cover the most recent state of the art in interior-point (or other) methods for solving convex problems. Our coverage of numerical implementation issues is also highly simplified, but we feel that it is adequate for the potential user to develop working implementations, and we do cover, in some detail, techniques for exploiting structure to improve the efficiency of the methods. We also do not cover, in more than a simplified way, the complexity theory of the algorithms we describe. We do, however, give an introduction to the important ideas of self-concordance and complexity analysis for interior-point methods.

Audience

This book is meant for the researcher, scientist, or engineer who uses mathematical optimization, or more generally, computational mathematics. This includes, naturally, those working directly in optimization and operations research, and also many others who use optimization, in fields like computer science, economics, finance, statistics, data mining, and many fields of science and engineering. Our primary focus is on the latter group, the potential *users* of convex optimization, and not the (less numerous) experts in the field of convex optimization.

The only background required of the reader is a good knowledge of advanced calculus and linear algebra. If the reader has seen basic mathematical analysis (*e.g.*, norms, convergence, elementary topology), and basic probability theory, he or she should be able to follow every argument and discussion in the book. We hope that

readers who have not seen analysis and probability, however, can still get all of the essential ideas and important points. Prior exposure to numerical computing or optimization is not needed, since we develop all of the needed material from these areas in the text or appendices.

Using this book in courses

We hope that this book will be useful as the primary or alternate textbook for several types of courses. Since 1995 we have been using drafts of this book for graduate courses on linear, nonlinear, and convex optimization (with engineering applications) at Stanford and UCLA. We are able to cover most of the material, though not in detail, in a one quarter graduate course. A one semester course allows for a more leisurely pace, more applications, more detailed treatment of theory, and perhaps a short student project. A two quarter sequence allows an expanded treatment of the more basic topics such as linear and quadratic programming (which are very useful for the applications oriented student), or a more substantial student project.

This book can also be used as a reference or alternate text for a more traditional course on linear and nonlinear optimization, or a course on control systems (or other applications area), that includes some coverage of convex optimization. As the secondary text in a more theoretically oriented course on convex optimization, it can be used as a source of simple practical examples.

Acknowledgments

We have been developing the material for this book for almost a decade. Over the years we have benefited from feedback and suggestions from many people, including our own graduate students, students in our courses, and our colleagues at Stanford, UCLA, and elsewhere. Unfortunately, space limitations and shoddy record keeping do not allow us to name everyone who has contributed. However, we wish to particularly thank A. Aggarwal, V. Balakrishnan, A. Bernard, B. Bray, R. Cottle, A. d'Aspremont, J. Dahl, J. Dattorro, D. Donoho, J. Doyle, L. El Ghaoui, P. Glynn, M. Grant, A. Hansson, T. Hastie, A. Lewis, M. Lobo, Z.-Q. Luo, M. Mesbahi, W. Naylor, P. Parrilo, I. Pressman, R. Tibshirani, B. Van Roy, L. Xiao, and Y. Ye. We are grateful to J. Jalden and A. d'Aspremont, who contributed to the time-frequency analysis example in §6.5.4, and the consumer preference bounding example in §6.5.5, respectively. P. Parrilo suggested exercises 4.4 and 4.56.

We want to single out two others for special acknowledgment. Arkadi Nemirovski incited our original interest in convex optimization, and encouraged us to write this book. We also want to thank Kishan Baheti for playing a critical role in the development of this book. In 1994 he encouraged us to apply for a National Science Foundation combined research and curriculum development grant, on convex optimization with engineering applications, and this book is a direct (if delayed) consequence.

Stephen Boyd
Lieven Vandenberghe

Stanford, California
Los Angeles, California

July 2003