

Technion – Israel Institute of Technology
Faculty of Industrial Engineering and Management
Minerva Optimization Center
Technion City, Haifa 32000, Israel
Fax 972-4-8235194

LECTURES ON MODERN CONVEX OPTIMIZATION
ANALYSIS, ALGORITHMS, ENGINEERING APPLICATIONS

Aharon Ben-Tal and Arkadi Nemirovski

Ben-Tal, A., and Nemirovski, A. *Lectures on Modern Convex Optimization: Analysis, Algorithms and Engineering Applications*, MPS-SIAM Series on Optimization, SIAM, Philadelphia, 2001

Copyright 2000, Aharon Ben-Tal and Arkadi Nemirovski

Preface

The goals. To make decisions optimally is a basic desire of a human being. Whenever the situation and the objectives can be described quantitatively, this desire can be satisfied, to some extent, by using mathematical tools, specifically those provided by Optimization theory and algorithms. For our purposes, a sufficiently general mathematical setting of an optimization problem is offered by *Mathematical Programming*:

$$\begin{aligned} & \text{minimize} && f_0(x) \\ & \text{subject to} && \\ & && f_i(x) \leq b_i, \quad i = 1, \dots, m, \\ & && x \in X \subset \mathbf{R}^n. \end{aligned} \tag{P}$$

In this problem, we are given an *objective* $f_0(x)$ and finitely many *functional constraints* $f_i(x)$, $i = 1, \dots, m$, which are real-valued functions of n -dimensional *design vector* x varying in a given domain X . Our goal is to minimize the objective over the *feasible set* of the problem – the set which is cut off the domain X by the system of inequalities $f_i(x) \leq b_i$, $i = 1, \dots, m$.

In typical engineering applications, the design vector specifies a decision to be made (e.g., the physical sizes of the bars in a truss-like structure), the domain X is the set of “meaningful” values of the design vector, and the functional constraints represent design specifications – restrictions (physical, technical, financial) on certain characteristics of the decision.

The last decade has witnessed a major progress in Optimization, especially in the area of convex programming. Powerful modeling languages and database technology extended our abilities to model large scale complex real-life problems; progress in complexity theory improved our understanding of the advantages of certain algorithms, and the limitations of others, and resulted in developing efficient interior point algorithms for a wide family of convex programs. Combined with the dramatic improvement of computers, the above progress enables us today to solve problems which were considered out of reach for Optimization just ten years ago.

Regrettably, this promising state of affairs is yet unrealized, and consequently not utilized, by potential end-users (engineers, managers, etc.). This book aims at presenting modern optimization, combining it with applications (mainly from engineering) and thus helping to bridge the gap between researchers and practitioners. This ultimate goal determines the approach we are undertaking, and dictates several specific targets which we are trying to achieve.

Theoretically speaking, what modern Optimization can solve “well” are *convex optimization problems*. In essence, the two-decade-long investigation of complexity issues in Optimization can be summarized in the following conclusion:

(!) *From the viewpoint of the numerical processing of problems (P), there exists a “solvable case” – the one of convex optimization problems, those where the domain X is a closed convex subset of \mathbf{R}^n , and the objective $f_0(x)$ and the functional constraints $f_i(x)$, $i = 1, \dots, m$, are convex functions on X .*

Under minimal additional “computability assumptions” (which are satisfied in basically all applications), a convex optimization problem is “computationally tractable” – the computational effort required to solve the problem to a given accuracy “grows moderately” with the dimensions of the problem and the required number of accuracy digits.

In contrast to this, general-type non-convex problems are too difficult for numerical solution; the computational effort required to solve such a problem, by the best numerical methods known so far, grows prohibitively fast with the dimensions of the problem and the number of accuracy digits. Moreover, there are serious theoretical reasons to conjecture that this is an intrinsic feature of non-convex problems rather than a drawback of the existing optimization techniques.

As an example, consider the pair of optimization problems (A) and (B) below. The first is the nonconvex problem

$$\begin{aligned} & \text{maximize} && \sum_{i=1}^n x_i \\ & \text{subject to} && \\ & && x_i^2 - x_i = 0, \quad i = 1, \dots, n; \\ & && x_i x_j = 0 \quad \forall (i, j) \in \Gamma, \end{aligned} \tag{A}$$

Γ being a given set of pairs (i, j) of indices i, j . This is a fundamental combinatorial problem of computing the *stability number* of a graph. It arises, e.g., in the following *channel communication problem*:

There is an alphabet of n letters a_i , $i = 1, 2, \dots, n$, say, the 256 usual bytes. A letter a_i can be sent through a communication channel; when passing through it, it either remains unchanged, or can be converted to another letter a_j due to transmission errors. The errors are assumed to be “symmetric” (if a_i can be converted into a_j , then a_j can be converted into a_i as well), and Γ is the set of (indices of) those pairs of letters which can be converted from one into another. Assume that we are interested in a “non-confusing” communication, where the addressee either gets a correct letter, or is able to conclude that a transmission error has occurred, but never misreads the input letter. In this case we should restrict the sub-alphabet S we use to be *independent*, meaning that no two distinct letters from S can be converted from one to another by the channel. In order to get from the channel “its most”, we would like to use a independent sub-alphabet of maximal cardinality. It turns out that the optimal value in (A) is exactly the cardinality of such a maximal independent sub-alphabet.

The second problem is the convex program

$$\begin{aligned} & \text{minimize} && x_0 \\ & \text{subject to} && \\ & && \lambda_{\min} \begin{pmatrix} x_1 & & & x_1^\ell \\ & \ddots & & \vdots \\ & & x_m & x_m^\ell \\ x_1^\ell & \cdots & x_m^\ell & x_0 \end{pmatrix} \geq 0, \quad \ell = 1, \dots, k, \\ & && \sum_{j=1}^m a_j x_j^\ell = b^\ell, \quad \ell = 1, \dots, k, \\ & && \sum_{j=1}^m x_j = 1, \end{aligned} \tag{B}$$

where $\lambda_{\min}(A)$ denotes the minimum eigenvalue of a symmetric matrix A . This problem originates from design of a *truss* (a mechanical construction built from thin elastic bars, like an electric mast, a bridge or the Eifel Tower) capable to withstand k nonsimultaneous loads.

Looking at the analytical forms of (A) and (B), it seems that the first problem is easier than the second: the constraints in (A) are simple explicit quadratic equations, while the constraints in (B) involve much more complicated functions

of the design variables – the eigenvalues of certain matrices depending on the design vector. The truth, however, is that the first problem is in a sense “as difficult as an optimization problem can be”, and the worst-case computational effort to solve it within absolute inaccuracy 0.5 is about 2^n operations for all known optimization methods. For $n = 256$ (“alphabet of bytes”), the complexity $2^n \approx 10^{77}$ is, for all practical purposes, the same as $+\infty$. In contrast to this, the second problem is quite “computationally tractable”. E.g., for $k = 6$ (6 loading scenarios) and $m = 100$ (a 100-bar construction) the problem has 701 variables (2.7 times the number of variables in the “byte” version of (A)); however, it can be reliably solved within 6 accuracy digits in a couple of minutes. The dramatic difference in the computational effort required to solve (A) and (B) is due to the fact that (A) is a non-convex optimization problem, while (B) is convex.

The above discussion explains the words “Convex Programming” in the title of our book. We now explain the word “modern”. The emphasis in the book is on *well-structured convex problems* such as Linear, Conic Quadratic and Semidefinite Programming. These areas are the ones most affected by the recent progress in Optimization, areas where we possess well-developed techniques for building large-scale models, analyzing them theoretically (“on paper”) and processing them numerically. Except for Linear Programming, these areas basically did not exist 10 years ago, so that most of the users who could potentially benefit from recent developments are not aware of their existence, let alone their usefulness. Enlarging the scope of classical optimization (LP, general Nonlinear Programming) by introduction of CQP and SDP, new application areas were created. Examples include: semidefinite relaxations of difficult combinatorial problems, Linear Matrix Inequality-based techniques in Control, Mathematical Programming with uncertain data, etc. These new applications create synergies between Optimization, Computer Science and Engineering, with potentially far-reaching consequences on our ability to solve complex real-life decision problems.

At this point, we want to make a remark addressing experts in optimization rather than “general” readers. The history of Convex Programming, as far as applied aspects are concerned, started with the invention of Linear Programming (G. Dantzig, circa 1948). LPs possess the simplest possible and transparent analytical structure, which from the very beginning was heavily exploited both theoretically (the completely algorithmic LP Duality Theory) and computationally (the Simplex method). With subsequent nonlinear extensions, the focus was shifted, basically “in one giant step”, from the simplest possible linear structure to the most general one, where all we know about the entities occurring in (P) is that they are called f_i , $i = 0, 1, \dots, m$, and X , that they are convex, and that f_i are 0/1/2/... times continuously differentiable. At the theoretical level, the focus on this general setting yielded very deep results in Convex Analysis (Lagrange duality, Karush-Kuhn-Tucker optimality conditions, etc.); however, “the price” of these really fundamental achievements was a lack of an *algorithmic* content of the subject. For example, the Lagrange dual of a general-type convex program (P) is something which exists and possesses very nice properties; this “something”, however, normally cannot be written down explicitly (in a sharp contrast with the “fully algorithmic” LP duality). At the computational level, the “emphasis on generality” resulted in general-purpose “near-sighted” (and thus slow, as compared to LP algorithms) optimization methods, those utilizing purely local information on the problem.

To some extent, recent trends (the last decade or so) in Convex Optimization stem from realizing that there is something “in between” the relatively narrow Linear Programming and the completely unstructured universe of Convex Programming; this “something” are “well-structured” generic convex optimization problems like Conic Quadratic and Semidef-

inite Programming. Needless to say, interest in “special cases” is not a complete novelty for our area (recall Linearly Constrained Convex Quadratic and Geometric Programming); what is indeed a novelty, is the recent emphasis on well-structured generic problems, along with outstanding achievements resulting from this emphasis. The most remarkable of these achievements is, of course, the “interior point revolution”, which has extended dramatically our abilities to process convex programs numerically, at the same time creating a completely new, challenging and attractive area of theoretical research. The emphasis on well-structured “special cases” has, however, another, a less evident (perhaps not less important) consequence which can be summarized as follows:

- *When restricted to “nice” generic convex programs, like LP, CQP and SDP, Convex Analysis becomes “an algorithmic calculus” – basically as algorithmic as in the LP case. For example, the SDP Duality is as “explicit and symmetric” as the LP one. In fact, the same can be said about all other basic operations of Convex Analysis, from the simplest (like taking intersections and affine images of convex sets, or sums of convex functions) to the more sophisticated ones (like passing from a convex set to its polar or from a convex function to its Legendre transformation). Whenever the “operands” of such a construction can be represented, in a properly defined precise sense, via, say, SDP, the same is true for the resulting entity, and the “SDP representation” of the result is readily given by those of the operands.*

As a result,

- *An instance of a “nice” generic convex problem can be processed, up to a certain point, “on paper” (and in a quite routine fashion). In many cases this allows to obtain important qualitative information on the problem and/or to convert it into an equivalent one, better suited for subsequent numerical processing. For example, applying SDP Duality, one can reduce dramatically the design dimension of the Truss Topology Design problem, and as a result – increase by orders of magnitude the sizes of the TTD problems which can be processed numerically in practice.*

Moreover,

- *“Nice” generic convex problems, like CQP and especially SDP, possess vast “expressive abilities”, which allows to utilize the above advantages in a very wide spectrum of applications, much wider than the one covered by Linear Programming.*

When writing this book, our major concern was to emphasize the issues just raised, and this emphasis is perhaps the most characteristic (and, hopefully, to some extent a novel) feature of the book.

Restricting ourselves to “well-structured” convex optimization problems, it becomes logical to skip a significant part of what is traditionally thought of as “the theory and algorithms of Mathematical Programming”. Readers interested in the Gradient Descent, Quasi-Newton methods, and even Sequential Quadratic Programming, are kindly advised to use the excellent existing textbooks on these important subjects; our book should be thought of as a self-contained complement to, and not as an extended version of, these textbooks. We even have “dared” to omit the Karush-Kuhn-Tucker optimality conditions in their standard form, since they are too general to be “algorithmic”; the role of the KKT conditions in our exposition is played by their particular (and truly “algorithmic”) case, expressed by the so called Conic Duality Theorem.

The book is primarily addressed to potential users (mainly, engineers). Consequently, *our emphasis is on building and processing instructive engineering models rather than on describing the details of optimization algorithms.* The underlying motivation is twofold. First, we are trying to convince an engineer that Optimization has indeed something to do with his/her

professional interests. Second, we believe that a crucial necessary condition for successful practical applications of optimization is the understanding of what is desirable and what should be avoided at the *modeling* stage. Thus, important questions to be addressed are:

(a) *What are optimization models which can be efficiently processed* (to certain point – “on paper” and then - on a computer),

(b) *How one can convert* (provided that it is possible) a “*seemingly bad*” *initial description of a problem into a “tractable and well-structured” optimization model.*

We believe that the best way to provide a reader with a relevant insight is to present, along with general concepts and techniques, many applications of these concepts and techniques. As about optimization algorithms, we believe that their presentation in a user-oriented book should be as non-technical as possible (to drive a car, no expertise in engines is necessary, otherwise there is something bad with cars...) The part of the book devoted to algorithms presents the Ellipsoid method (due to its simplicity, combined with its capability to answer affirmatively the fundamental question of whether Convex Programming is “computationally tractable”) and an overview of polynomial time interior point methods for Linear, Conic Quadratic and Semidefinite Programming.

In spite of the fact that the book is user-oriented, it is a mathematical book. The goal we try to achieve is to demonstrate that when processing “meaningful” *mathematical* models by *rigorous* mathematical methods (and not by their “engineering surrogates”), one can obtain results which have meaningful and instructive engineering interpretation. Whether we have reached this goal – this is another story, and here the judgment is upon the reader.

Last, but not least, a reader should keep in mind that what follows are *Lecture Notes*; *our intention was to highlight those issues which we find most interesting and instructive*, rather than to present a complete overview of Convex Optimization. Consequently, we are ready to take blame for being boring, or unclear, or focusing on issues of a minor importance, in any material which we *include* in the book, however, we do not accept “money back” requests based on claims that something (however important) is *not* included. Along with “immunity w.r.t. what is absent”, a “Lecture Notes type” book provides us with some other privileges, like a style which is a bit more vivid compared to the academic standards, and a rather short list of bibliography references (embedded as footnotes in the body of the text). In this latter respect, a reader should be aware that the fact that if a statement (even a “non-common-knowledge” one) appears in the text without a reference, this does *not* mean that we are claiming for the authorship; it merely reflects the fact that our focus is on the state of Convex Optimization rather than on its history.

Audience and prerequisites. Formally, readers are supposed to know basic facts from Linear Algebra and Analysis – those presented in standard undergraduate mathematical courses for engineers. As far as optimization-related areas are concerned, readers are assumed to know the definitions of a convex set and a convex function and to have heard (no more than that!) what is a Mathematical Programming problem. What is highly desirable, is a possession of the basic elements of mathematical culture.

The exercises accompanying each lecture form a significant part of the book. They are organized in small groups, each devoted to a specific topic somehow related to the corresponding lecture. Typically, the task in an exercise is to prove something. Most of the exercises are not too easy... The results stated by the exercises are used in the subsequent parts of the book in the same way as the statements presented and proved in the main body of the book; consequently,

a reader is kindly asked if not to do, then at least to read carefully all exercises. The order of exercises is of primary importance: in many cases preceding exercises contain information/hints necessary to succeed in the subsequent ones.

Acknowledgments. A quite significant part of applications discussed in our book is taken from the papers of Prof. Stephen Boyd from Stanford University and his colleagues. We are greatly indebted to Prof. Boyd for providing us with access to this material and stimulating discussions. Applications related to Structural Design were developed in tight collaboration with Prof. Jochem Zowe and Dr. Michael Kočvara from Erlangen University. Parts of the book were written when the authors were visiting the Statistics and Operations Research Department of the Technical University of Delft, and we are thankful to our hosts Prof. Kees Roos and Prof. Tamas Terlaky.

Aharon Ben-Tal and Arkadi Nemirovski

August 2000, Haifa, Israel

Contents

1	Linear Programming	13
1.1	Linear programming: basic notions	13
1.2	Example: Tschebyshev approximation and its applications	14
1.2.1	Best uniform approximation	14
1.2.2	Application: synthesis of filters	15
1.2.3	Filter synthesis revisited	16
1.2.4	Synthesis of arrays of antennae	19
1.3	Duality in Linear Programming	22
1.3.1	Certificates for solvability and insolvability	23
1.3.2	Dual to an LP program: the origin	26
1.3.3	The LP Duality Theorem	28
1.3.4	Illustration: the problem dual to the Tschebyshev approximation problem	29
1.3.5	Application: Truss Topology Design	31
1.4	Exercises to Lecture 1	38
1.4.1	Around uniform approximation	38
1.4.2	Around the Theorem on Alternative	41
1.4.3	Proof of the Homogeneous Farkas Lemma	43
1.4.4	The Helley Theorem	46
1.4.5	How many bars are needed in an optimal truss?	49
2	From Linear Programming to Conic Programming	51
2.1	Orderings of \mathbf{R}^m and convex cones	51
2.2	“Conic programming” – what is it?	54
2.3	Conic Duality	57
2.3.1	Geometry of the primal and the dual problems	60
2.4	The Conic Duality Theorem	64
2.4.1	Is something wrong with conic duality?	67
2.4.2	Consequences of the Conic Duality Theorem	68
2.4.3	Robust solvability status	70
2.5	Conic Duality revisited	73
2.6	Exercises to Lecture 2	78
2.6.1	Around cones	78
2.6.2	Around conic problems	81
2.6.3	Feasible and level sets of conic problems	82

3	Conic Quadratic Programming	83
3.1	Conic Quadratic problems: preliminaries	83
3.2	Examples of conic quadratic problems	85
3.2.1	Best linear approximation of complex-valued functions	85
3.2.2	Contact problems with static friction	86
3.3	What can be expressed via conic quadratic constraints?	89
3.3.1	More examples of CQ-representable functions/sets	105
3.4	More applications	110
3.4.1	Tschebyshev approximation in relative scale	110
3.4.2	Robust Linear Programming	111
3.4.3	Truss Topology Design	119
3.5	Exercises to Lecture 3	128
3.5.1	Optimal control in discrete time linear dynamic system	128
3.5.2	Around CQR's	129
3.5.3	Whether Conic Quadratic Programming does exist?	134
4	Semidefinite Programming	137
4.1	Semidefinite cone and Semidefinite programs	137
4.1.1	Preliminaries	137
4.2	What can be expressed via LMIs?	141
4.3	Applications, I: Combinatorics	156
4.3.1	Shor's Semidefinite Relaxation scheme	157
4.3.2	Stability number, Shannon capacity and Lovasz capacity of a graph	160
4.3.3	The MAXCUT problem	165
4.3.4	Extensions	168
4.3.5	\mathcal{S} -Lemma	170
4.4	Applications, II: Stability Analysis	173
4.4.1	Dynamic stability in Mechanics	173
4.4.2	Lyapunov stability analysis/synthesis	175
4.4.3	Interval Stability Analysis/Synthesis	183
4.5	Applications, III: Robust Quadratic Programming	193
4.6	Applications, IV: Synthesis of filters and antennae arrays	200
4.7	Applications, V: Design of chips	208
4.7.1	Building the model	209
4.7.2	Wire sizing	215
4.8	Applications, VI: Structural design	216
4.8.1	Building a model	216
4.8.2	The standard case	221
4.8.3	Semidefinite reformulation of the Standard SSD problem	224
4.8.4	From primal to dual	230
4.8.5	Back to primal	234
4.8.6	Explicit forms of the Standard Truss and Shape problems	237
4.9	Applications, VII: Extremal ellipsoids	240
4.9.1	Ellipsoidal approximations of unions/intersections of ellipsoids	245
4.9.2	Approximating sums of ellipsoids	247
4.10	Exercises to Lecture 4	258
4.10.1	Around positive semidefiniteness, eigenvalues and \succeq -ordering	258

4.10.2	SD representations of epigraphs of convex polynomials	272
4.10.3	Around the Lovasz capacity number and semidefinite relaxations of combinatorial problems	274
4.10.4	Around Lyapunov Stability Analysis	279
4.10.5	Around the \mathcal{S} -Lemma	281
4.10.6	Around Antenna Synthesis	300
4.10.7	Around ellipsoidal approximations	303
5	Computational tractability of convex programs	313
5.1	Numerical solution of optimization programs – preliminaries	313
5.1.1	Mathematical Programming programs	313
5.1.2	Solution methods and efficiency	314
5.2	“Black box”-represented convex programs	319
5.3	Polynomial solvability of Convex Programming	329
5.3.1	Polynomial solvability of Convex Programming	335
5.4	Difficult problems and NP-completeness	338
5.4.1	CCT – a quick introduction	339
5.4.2	From the Real Arithmetic Complexity theory to the CCT and back . . .	343
6	Interior point polynomial time methods for LP, CQP and SDP	353
6.1	Motivation	353
6.1.1	Interior point methods	354
6.2	The Newton method and the Interior penalty scheme	355
6.2.1	Unconstrained minimization and the Newton method	355
6.2.2	Classical interior penalty scheme: the construction	356
6.2.3	Classical interior penalty scheme: the drawbacks	357
6.2.4	But...	358
6.3	Interior point methods for LP, CQP, and SDP: building blocks	359
6.3.1	Canonical cones and canonical barriers	360
6.3.2	Elementary properties of canonical barriers	362
6.4	Primal-dual pair of problems and the primal-dual central path	364
6.4.1	The problem(s)	364
6.4.2	The central path(s)	365
6.5	Tracing the central path	371
6.5.1	The path-following scheme	371
6.5.2	Speed of path-tracing	373
6.5.3	The primal and the dual path-following methods	377
6.5.4	The SDP case	379
6.6	Complexity bounds for LP, CQP, SDP	394
6.6.1	Complexity of LP	394
6.6.2	Complexity of CQP	395
6.6.3	Semidefinite Programming	395
6.7	Concluding remarks	396
6.8	Exercises to Lecture 6	398
6.8.1	Around canonical barriers	398
6.8.2	Scalings of canonical cones	398
6.8.3	The Dikin ellipsoid	401

6.8.4	More on canonical barriers	403
6.8.5	Around the primal path-following method	404
6.8.6	Infeasible start path-following method	406
7	Solutions to selected exercises	415
7.1	Exercises to Lecture 1	415
7.2	Exercises to Lecture 2	417
7.3	Exercises to Lecture 3	421
7.4	Exercises to Lecture 4	428
7.5	Exercises to Lecture 6	442