

Preface

This book provides an introduction to the applications, theory, and algorithms of linear and nonlinear optimization. The emphasis is on practical aspects—modern algorithms, as well as the influence of theory on the interpretation of solutions or on the design of software. Two important goals of this book are to present linear and nonlinear optimization in an integrated setting, and to incorporate up-to-date interior-point methods in linear and nonlinear optimization.

As an illustration of this unified approach, almost every algorithm in this book is presented in the form of a General Optimization Algorithm. This algorithm has two major steps: an optimality test, and a step that improves the estimate of the solution. This framework is general enough to encompass the simplex method and various interior-point methods for linear programming, as well as Newton's method and active-set methods for nonlinear optimization. The optimality test in this algorithm motivates the discussion of optimality conditions for a variety of problems. The step procedure motivates the discussion of feasible directions (for constrained problems) and Newton's method and its variants (for nonlinear problems).

In general, there is an attempt to develop the material from a small number of basic concepts, emphasizing the interrelationships among the many topics. Our hope is that, by emphasizing a few fundamental principles, it will be easier to understand and assimilate the vast panorama of linear and nonlinear optimization.

We have attempted to make accessible a number of topics that are not often found in textbooks. Within linear programming, we have emphasized the importance of sparse matrices on the design of algorithms, described computational techniques used in sophisticated software packages, and derived the primal-dual interior-point method together with the predictor-corrector technique. Within nonlinear optimization, we have included discussions of truncated-Newton methods for large problems, convergence theory for trust-region methods, filter methods, and techniques for alleviating the ill-conditioning in barrier methods. We hope that the book serves as a useful introduction to research papers in these areas.

The book was designed for use in courses and course sequences that discuss both linear and nonlinear optimization. We have used consistent approaches when discussing the two topics, often using the same terminology and notation in order to emphasize the similarities between the two topics. However, it can also be used in traditional (and separate) courses in Linear Programming and Nonlinear Optimization—in fact, that is the way we use it in the courses that we teach. At the end of this preface are chapter descriptions and course outlines indicating these possibilities.

We have also used the book for more advanced courses. The later chapters (and the later sections within chapters) contain a great deal of material that would be difficult to cover in an introductory course. The Notes at the ends of many sections contain pointers to research papers and other references, and it would be straightforward to use such materials to supplement the book.

The book is divided into four parts plus appendices. Part I (Basics) contains material that might be used in a number of different topics. It is not intended that all of this material be presented in the classroom. Some of it might be irrelevant (as the sample course outlines illustrate). In other cases, material might be familiar to the students from other courses, or simple enough to be assigned as a reading exercise. The material in Part I could also be taught in stages, as it is needed. In a course on Nonlinear Optimization, for example, Chapter 4 (Representation of Linear Constraints) could be delayed until after Part III (Unconstrained Optimization). Our intention in designing Part I was to make the book as flexible as possible, and instructors should feel free to exploit this flexibility.

Part II (Linear Programming) and Part III (Unconstrained Optimization) are independent of each other. Either one could be taught or read before the other. In addition, it is not necessary to cover Part II before going on to Part IV (Nonlinear Optimization), although the material in Part IV will benefit from an understanding of Linear Programming. The material in the appendices may already be familiar. If not, it could either be presented in class or left for students to read independently.

Many sections in the book can be omitted without interrupting the flow of the discussions (detailed information on this is given below). Proofs of theorems and lemmas can similarly be omitted. Roughly speaking, it is possible to skip later sections within a chapter and later chapters within a part and move on to later chapters in the book. The book was organized in this way so that it would be accessible to a wider audience, as well as to increase its flexibility.

Many of the exercises are computational. In some cases, pencil-and-paper techniques would suffice, but the use of a computer is recommended. We have not specified how the computer might be used, and we leave this up to the instructor. In courses with an emphasis on modeling, a specialized linear or nonlinear optimization package might be appropriate. In other courses, the students might be asked to program algorithms themselves. We leave these decisions up to the instructor. Some information about software packages can be found in Appendix C. In addition, some exercises depend on auxiliary data sets that can be found on the web site for the book:

<http://www.siam.org/books/ot108>

In our own classes, we use the MATLAB[®] software package for class demonstrations and homework assignments. It allows us to demonstrate a great many techniques easily, and it allows students to program individual algorithms without much difficulty. It also includes (in its toolboxes) prepared algorithms for many of the optimization problems that we discuss.

We have gone to considerable effort to ensure the accuracy of the material in this book. Even so, we expect that some errors remain. For this reason, we have set up an online page for errata. It can be obtained at the book Web site.

Using This Book

This book is designed to be flexible. It can be read and taught in many different ways.

The material in the appendices can be taught as needed, or left to the students to read independently. Also, all formally identified proofs can be omitted.

Part II (Linear Programming) and Part III (Unconstrained Optimization) are independent of each other. Part II does not assume any knowledge of Calculus. Part IV (Nonlinear Optimization) does not assume that Part II has been read (with the exception of Section 14.4.1).

The only “essential” chapters in Part II are Chapters 4 (Geometry of Linear Programming), 5 (The Simplex Method), and 6 (Duality). The only “essential” chapter in Part III is Chapter 11 (Basics of Unconstrained Optimization). The other chapters can be skipped.

We now describe the chapters individually, pointing out various ways they can be used. The sample course outlines that follow indicate how chapters might be selected to construct individual courses (based on a 15-week semester).¹

Part I: Basics

- *Chapter 1: Optimization Models.* This chapter is self-contained and describes a variety of optimization models. Sections 1.3–1.5 are independent of one another. Section 1.6 includes more realistic models and assumes that the reader is familiar with the basic models described in the earlier sections. The subsections of Section 1.6 are independent of one another.
- *Chapter 2: Fundamentals of Optimization.* For Part II, only Sections 2.1–2.4 are needed (and Section 2.3.1 can be omitted). For Parts III and IV the whole chapter is relevant.
- *Chapter 3: Representation of Linear Constraints.* Sections 3.3.2–3.3.4 can be omitted (although Section 3.3.2 is needed for Part IV). This chapter is only relevant to Parts II and IV; it is not needed for Part III.

Part II: Linear Programming

- *Chapter 4: Geometry of Linear Programming.* All sections of this chapter are needed in Part II.
- *Chapter 5: The Simplex Method.* Sections 5.1 and 5.2 are the most important. How the rest of the chapter is used depends on the goals of the instructor, in particular with regard to tableaus. In a number of examples, we use the full simplex tableau to display data for linear programs. Thus, it is necessary to be able to read these tableaus to extract information. This is the only use we make of the tableaus elsewhere in the book. It is not necessary to be able to manipulate these tableaus.

¹Throughout the book, the number of a section or subsection begins with the chapter number. That is, Section 10.3 refers to the third section in Chapter 10, and Section 16.7.2 refers to the second subsection in the seventh section of Chapter 16. Also, a reference to Appendix A.9 refers to the ninth section of Appendix A. A similar system is used for tables, examples, theorems, etc.; Figure 8.10 refers to the tenth figure in Chapter 8, for example. For exercises, however, the chapter number is omitted, e.g., Exercise 4.7 is the seventh exercise in Section 4 of the current chapter (unless another chapter is specified).

- *Chapter 6: Duality and Sensitivity.* Sections 6.1 and 6.2 are the most important. The remaining sections can be skipped, if desired. If taught, we recommend that Sections 6.3–6.5 be taught in order, although Section 6.3 is only used in a minor way in the remaining two sections. It would be possible to stop after any section. Note: The remaining chapters in Part II are independent of each other.
- *Chapter 7: Enhancements of the Simplex Method.* The sections in this chapter are independent of each other. The instructor is free to pick and choose material, with one partial exception: the discussion of the decomposition principle is easier to understand if column generation has already been read.
- *Chapter 8: Network Problems.* In this chapter, the sections must be taught in order. It would be possible to stop after any section.
- *Chapter 9: Computational Complexity of Linear Programming.* The first two sections contain basic material used in Sections 9.3–9.5. Ideally, the remaining sections should be taught in order, although Sections 9.4 and 9.5 are independent of each other. Even if some topics are not of interest, at least the introductory paragraphs of each section should be read. (Section 9.5 requires some knowledge of statistics.)
- *Chapter 10: Interior-Point Methods for Linear Programming.* Sections 10.1 and 10.2 are the most important. The later sections could be skipped but, if taught, Sections 10.4–10.6 should be taught in order. Section 10.4 reviews some fundamental concepts from nonlinear optimization needed in Sections 10.5–10.6.

Part III: Unconstrained Optimization

- *Chapter 11: Basics of Unconstrained Optimization.* We recommend reading all of this chapter (with the exception of the proofs). If desired, either Section 11.5 or Section 11.6 could be omitted, but not both. Chapters 12 and 13 could be omitted. Chapter 13 makes more sense if taught after Chapter 12, but in fact, only Section 13.5 makes explicit use of the material in Chapter 12.
- *Chapter 12: Methods for Unconstrained Optimization.* Sections 12.1–12.3 are the most important. All the remaining sections and subsections can be taught independently of each other.
- *Chapter 13: Low-Storage Methods for Unconstrained Problems.* Once Sections 13.1 and 13.2 have been taught, the remaining sections are independent of each other.

Part IV: Nonlinear Optimization

- *Chapter 14: Optimality Conditions for Constrained Problems.* We recommend reading Sections 14.1–14.6. The rest of the chapter may be omitted. Within Section 14.8, Sections 14.8.3 and 14.8.5 can be taught without teaching the remaining subsections, although Section 14.8.5 depends on Section 14.8.3. (The discussion of nonlinear duality in Section 14.8 is only needed in Sections 16.6–16.8 of Chapter 16.)
- *Chapter 15: Feasible-Point Methods.* We recommend reading Sections 15.1–15.4 (although Section 15.4.1 could be omitted). These sections explain how to solve problems with linear constraints. Sections 15.5–15.7 discuss methods for problems

with nonlinear constraints. Sections 15.5 and 15.6 are independent of each other, but Section 15.7 depends on Section 15.5.

- *Chapter 16: Penalty and Barrier Methods.* We recommend reading Sections 16.1 and 16.2 (although Section 16.2.3 could be omitted). If more of the chapter is covered, then Section 16.3 should be read. Sections 16.4–16.8 are independent of each other. Sections 16.6–16.8 use Section 14.8.3 of Chapter 14.

Changes in the Second Edition

The overall structure of the book has not changed in the new addition, and the major topic areas are the same. However, we have updated certain topics to reflect developments since the first edition appeared. We list the major changes here.

Chapter 1 has been expanded to include examples of more realistic optimization models (Section 1.6). The description of interior-point methods for linear programming has been thoroughly revised and restructured (Chapter 10). The discussion of derivative-free methods has been extensively revised to reflect advances in theory and algorithms (Section 12.5). In Part IV we have added material on filter methods (Section 15.7), nonlinear primal-dual methods (Section 16.7), and semidefinite programming (Section 16.8). In addition, numerous smaller changes have been made throughout the book.

Some material from the first edition has been omitted here. The most notable examples are the chapter on nonlinear least-squares data fitting, and the sections on interior-point methods for convex programming. These topics from the first edition are available at the book Web site (see above for the URL).

Sample Course Outlines

We provide below some sample outlines for courses that might use this book. If a section is listed without mention of subsections, then it is assumed that all the subsections will be taught. If a subsection is specified, then the unmentioned subsections may be omitted.

Proposed Course Outline: Linear Programming

I: Foundations

Chapter 1. Optimization Models

1. Introduction
3. Linear Equations
4. Linear Optimization
7. Optimization Applications
 1. Crew Scheduling and Fleet Scheduling

Chapter 2. Fundamentals of Optimization

1. Introduction
2. Feasibility and Optimality
3. Convexity
4. The General Optimization Algorithm

Chapter 3. Representation of Linear Constraints

1. Basic Concepts
2. Null and Range Spaces
3. Generating Null-Space Matrices
 1. Variable Reduction Method

II: Linear Programming**Chapter 4. Geometry of Linear Programming**

1. Introduction
2. Standard Form
3. Basic Solutions and Extreme Points
4. Representation of Solutions; Optimality

Chapter 5. The Simplex Method

1. Introduction
2. The Simplex Method
3. The Simplex Method (Details)
4. Getting Started—Artificial Variables
 1. The Two-Phase Method
5. Degeneracy and Termination

Chapter 6. Duality and Sensitivity

1. The Dual Problem
2. Duality Theory
3. The Dual Simplex Method
4. Sensitivity

Chapter 7. Enhancements of the Simplex Method

1. Introduction
2. Problems with Upper Bounds
3. Column Generation
5. Representation of the Basis

Chapter 9. Computational Complexity of Linear Programming

1. Introduction
2. Computational Complexity
3. Worst-Case Behavior of the Simplex Method
4. The Ellipsoid Method
5. The Average-Case Behavior of the Simplex Method

Chapter 10. Interior-Point Methods for Linear Programming

1. Introduction
2. The Primal-Dual Interior-Point Method

Proposed Course Outline: Nonlinear Optimization**I: Foundations****Chapter 1. Optimization Models**

1. Introduction

3. Linear Equations
5. Least-Squares Data Fitting
6. Nonlinear Optimization
7. Optimization Applications²
 2. Support Vector Machines
 3. Portfolio Optimization
 4. Intensity Modulated Radiation Treatment Planning
 5. Positron Emission Tomography Image Reconstruction
 6. Shape Optimization

Chapter 2. Fundamentals of Optimization

1. Introduction
2. Feasibility and Optimality
3. Convexity
4. The General Optimization Algorithm
5. Rates of Convergence
6. Taylor Series
7. Newton's Method for Nonlinear Equations

Chapter 3. Representation of Linear Constraints³

1. Basic Concepts
2. Null and Range Spaces
3. Generating Null-Space Matrices
 1. Variable Reduction Method

III: Unconstrained Optimization

Chapter 11. Basics of Unconstrained Optimization

1. Introduction
2. Optimality Conditions
3. Newton's Method for Minimization
4. Guaranteeing Descent
5. Guaranteeing Convergence: Line Search Methods
6. Guaranteeing Convergence: Trust-Region Methods

Chapter 12. Methods for Unconstrained Optimization

1. Introduction
2. Steepest-Descent Method
3. Quasi-Newton Methods

Chapter 13. Low-Storage Methods for Unconstrained Problems

1. Introduction
2. The Conjugate-Gradient Method for Solving Linear Equations
3. Truncated-Newton Methods
4. Nonlinear Conjugate-Gradient Methods
5. Limited-Memory Quasi-Newton Methods

²Not all the applications need be taught.

³The material in Chapter 3 is not needed until Part IV.

IV: Nonlinear Optimization

Chapter 14. Optimality Conditions for Constrained Problems

1. Introduction
2. Optimality Conditions for Linear Equality Constraints
3. The Lagrange Multipliers and the Lagrangian Function
4. Optimality Conditions for Linear Inequality Constraints
5. Optimality Conditions for Nonlinear Constraints
6. Preview of Methods
8. Duality
 3. Wolfe Duality
 5. Duality in Support Vector Machines

Chapter 15. Feasible-Point Methods

1. Introduction
2. Linear Equality Constraints
3. Computing the Lagrange Multipliers
4. Linear Inequality Constraints
5. Sequential Quadratic Programming

Chapter 16. Penalty and Barrier Methods

1. Introduction
2. Classical Penalty and Barrier Methods

Proposed Course Outline: Introduction to Optimization**I: Foundations**

Chapter 1. Optimization Models

1. Introduction
3. Linear Equations
4. Linear Optimization
5. Least-Squares Data Fitting
6. Nonlinear Optimization
7. Optimization Applications⁴

Chapter 2. Fundamentals of Optimization

1. Introduction
2. Feasibility and Optimality
3. Convexity
4. The General Optimization Algorithm
5. Rates of Convergence
6. Taylor Series
7. Newton's Method for Nonlinear Equations

Chapter 3. Representation of Linear Constraints

1. Basic Concepts
2. Null and Range Spaces

⁴Not all the applications need be taught.

3. Generating Null-Space Matrices
 1. Variable Reduction Method

II: Linear Programming

Chapter 4. Geometry of Linear Programming

1. Introduction
2. Standard Form
3. Basic Solutions and Extreme Points
4. Representation of Solutions; Optimality

Chapter 5. The Simplex Method

1. Introduction
2. The Simplex Method
3. The Simplex Method (Details)
4. Getting Started—Artificial Variables
 1. The Two-Phase Method
5. Degeneracy and Termination

Chapter 6. Duality and Sensitivity

1. The Dual Problem
2. Duality Theory
4. Sensitivity

Chapter 8. Network Problems

1. Introduction
2. Basic Concepts and Examples

III: Unconstrained Optimization

Chapter 11. Basics of Unconstrained Optimization

1. Introduction
2. Optimality Conditions
3. Newton's Method for Minimization
4. Guaranteeing Descent
5. Guaranteeing Convergence: Line Search Methods

IV: Nonlinear Optimization

Chapter 14. Optimality Conditions for Constrained Problems

1. Introduction
2. Optimality Conditions for Linear Equality Constraints
3. The Lagrange Multipliers and the Lagrangian Function
4. Optimality Conditions for Linear Inequality Constraints
5. Optimality Conditions for Nonlinear Constraints
6. Preview of Methods

Acknowledgments

We owe a great deal of thanks to the people who have assisted us in preparing this second edition of this book. In particular, we would like to thank the following individuals for reviewing various portions of the manuscript and providing helpful advice and guidance: Erling Andersen, Bob Bixby, Sanjay Mehrotra, Hans Mittelmann, Michael Overton, Virginia Torczon, and Bob Vanderbei. We are especially grateful to Sara Murphy at SIAM for guiding us through the preparation of the manuscript.

Special thanks also to Galina Spivak, whose design for the front cover skillfully conveys, in our minds, the spirit of the book.

We continue to be grateful to those individuals who contributed to the preparation of the first edition. These include: Kurt Anstreicher, John Anzalone, Todd Beltracchi, Dimitri Bertsekas, Bob Bixby, Paul Boggs, Dennis Bricker, Tony Chan, Jessie Cohen, Andrew Conn, Blaine Crowthers, John Dennis, Peter Foellbach, John Forrest, Bob Fourer, Christoph Luitpold Frommel, Saul Gass, David Gay, James Ho, Sharon Holland, Jeffrey Horn, Soonam Kahng, Przemyslaw Kowalik, Michael Lewis, Lorin Lund, Irvin Lustig, Maureen Mackin, Eric Munson, Arkadii Nemirovsky, Florian Potra, Michael Rothkopf, Michael Saunders, David Shanno, Eric Smith, Martin Smith, Pete Stewart, André Tits, Michael Todd, Virginia Torczon, Luis Vicente, Don Wagner, Bing Wang, and Tjalling Ypma.

While preparing the first edition, we received valuable support from the National Science Foundation. We also benefited from the facilities of the National Institute of Standards and Technology and Rice University.

Igor Griva
Stephen G. Nash
Ariela Sofer