Preface

Optimization is presented in most multivariable calculus courses as an application of the gradient, and while this treatment makes sense for a calculus course, there is much more to the theory of optimization. Moreover, optimization is actually used every day in a way that is much different from what one is led to believe in the typical calculus course. Our world and its societies have for many centuries generated interesting and important optimization problems, and the theory of optimization has grown and developed in response to the challenges presented by these problems. In fact, optimization theory continues to be developed today in response to practical concerns encountered in applications, which makes optimization an ideal topic of study in modern applied mathematics. Through the study of optimization theory, the power and beauty of mathematics can be observed in close connection to interesting and relevant problems of our world.

The goals of this text are to help readers understand how optimization is actually done in practice and, at the same time, to expand their appreciation of the richness of the theory behind the practice. The only prerequisite knowledge for getting the most out of this text is a basic understanding of multivariable calculus. Matrices are one important additional tool for some of this material, but whatever is needed is developed here in a series of "Asides" scattered throughout the text in places where they are relevant. Readers who wish to learn all of these tools at once (or who just need to easily find one or more Aside) can refer to the Appendix, where all the Asides are collected in one place.

One important aspect of practical optimization is the mathematical modeling of real situations in the general format used for optimization. This is an art which is best learned through as much practice as possible, particularly since different real situations often present unique modeling issues. Thus, this book begins in Chapter 1 with the presentation of step-by-step solutions for five prototypical examples which fit the general optimization model. After one example is presented in the first section, there is a section giving four general rules of thumb for modeling. The remaining sections of Chapter 1 present four additional modeling examples with detailed solutions. Readers who wish to focus exclusively on modeling at the start can work carefully through all of Chapter 1 and solve the Modeling Problems presented at its end. When I teach from this text, I move quickly to Chapter 2 and assign several of the Modeling Problems in Chapter 1 to be completed during each week of the course.

Solving real-world problems via optimization models generally involves two steps: constructing an appropriate optimization model of the real-world problem, and then solving the optimization model in order to solve the real-world problem. Model construction can generally be done by hand, but, as will be shown in Chapter 2, model-solving typically requires a computer. The current wide availability of generally useful optimization software means that one can produce candidate solutions to optimization problems after a little training in how to feed the models into a computer. We take advantage of this here by encouraging readers to use software to "solve" the models they construct throughout the text. However, it can be difficult to interpret the computer's results without an understanding of how these results are produced. In some cases, the computer can even produce false "solutions" which an uninformed practitioner would have no reason to doubt. By understanding the material presented here on optimization methods, the reader should at least be prepared to make better informed—and therefore more intelligent—use of results produced by commercial optimization software. Moreover, since most of the optimization methods presented in this book were developed relatively recently, there is room not only to improve them but also to develop new optimization methods. Readers who are intrigued enough by this important and developing area should consider further study in it so that they might actually contribute to its development.

As a general principle here, practical issues are dealt with first as a means of motivating the theory. In particular, this means that optimization methods take a front seat and that various general rules from the theory are presented as they are actually used. There are many different optimization methods tailored to different circumstances; however, the focus in this text is on a few core procedures which are essentially prototypical of derivative-based optimization. Chapter 3 begins with the most fundamental of these methods for optimizing functions, and Chapters 4 and 5 cover some of the adjustments that are typically made to the fundamental method in order to accommodate various practical concerns. Chapter 6 presents two fundamental approaches to optimizing functions under constraints, with practical adjustments covered in Chapter 7. An important special case when the variables are integers is covered in Chapter 8.

Lurking in Chapter 7 is an introduction to linear programming, which traditionally has been treated in a much more central way. The special historical status of linear programming is reflected in the label of "nonlinear programming" given to the rest of optimization. The integrated treatment of linear programming in this book reflects the current view of the subject, as is clear from the following quote from Margaret H. Wright [9]:

> Today, in complete contrast to the era before 1984, researchers view linear and nonlinear programming from a unified perspective; the magnitude of this change can be seen simply by noting that no one would seriously argue today that linear programming is independent of nonlinear programming.

The recent merging of linear and nonlinear programming has occurred because the methods developed in each domain have been applied with great success across the traditional boundary. For instance, the *simplex method* for solving linear programs extends to nonlinear optimization problems as the *working-set method*, and the *logarithmic barrier method* for solving nonlinear optimization problems was specialized by Narendra Karmarkar as an *interior point method* for solving linear programs. In this book, the methods will be developed for general optimization problems (both linear and nonlinear) and applied to linear programs as a special case. This broader development of optimization not only highlights the current state of the subject but also connects much more naturally to calculus.

In addition to the Modeling Problems at the end of Chapter 1, there are four other basic forms of learning-by-doing activities incorporated throughout this text:

- Exercises: These are situated throughout the text precisely where they are relevant. They are meant to be attempted immediately when encountered and are short enough to be used as class questions during a lecture on the material.
- Problems: These are more in-depth activities than the Exercises and typically appear toward the end of the relevant sections. Problems too can be used as in-class student activities, but students will need more time to complete them than they will for Exercises.
- Computational Problems: As their title suggests, these activities depend in some way on the use of a computer. These are typically even longer-term activities than the Problems and are best completed out of class or as part of a devoted computer laboratory.
- Implementations: These involve the implementation of the optimization methods presented in this course and therefore depend heavily on the use of a computer. These are also the longest-term activities in the book and so are best completed out of class, possibly as a collaborative activity. Readers with no programming experience will need some support to complete the Implementations, and *Mathematica* notebooks which greatly simplify the programming can be downloaded from

  http://www.siam.org/books/ot114

This text certainly does not present the complete picture of numerical optimization, nor does it give the historical background of the development of this field. Indeed, one goal here is to stimulate the reader's interest in learning more about numerical optimization, and the bibliography contains several very good resources where the reader can take a next step. In light of this goal, Chapter 9 is devoted to a Research Project through which readers can explore something else in this broad and vibrant field.

Acknowledgments

The content and form of this book have benefited greatly from the careful reading and thoughtful comments of my colleague Steve Wright, as well as the anonymous reviewers coordinated by Sara Murphy at SIAM. My students have also made many important contributions while this text was under development, which I appreciate very much.