

Preface

For many years all three of us have been interested in, and have tried to make contributions to, derivative-free optimization. Our motivation for writing this book resulted from various circumstances. We had the opportunity to work closely with the leading contributors to the field, including especially John Dennis, Michael Powell, Philippe Toint, and Virginia Torczon. We had some knowledge of various facets of the recent developments in the area, and yet felt there was a need for a unified view, and we hoped thereby to gain a better understanding of the field. In addition we were enticed by the growing number of applications. We also felt very strongly that there was a considerable need for a textbook on derivative-free optimization, especially since the foundations, algorithms, and applications have become significantly enriched in the past decade. Finally, although the subject depends upon much that is true for, and was developed for, optimization with derivatives, the issues that arise are new. The absence of computable derivatives naturally prohibits the use of Taylor models—so common in derivative-based optimization. The fact that typical applications involve expensive function evaluations shifts the emphasis from the cost of the linear algebra, or other contributors to the iteration complexity, to simply the number of function evaluations. Also, the noise in the function values affects the local convergence expectations. Thus, the area is both simpler, in the sense of diminished expectations, and harder, in the sense that one is trying to achieve something with considerably less information. It is definitely fun and challenging and, not incidentally, very useful.

Although we do make considerable effort to give a sense of the current state of the art, we do not attempt to present a comprehensive treatise on all the work in the area. This is in part because we think that the subject is not yet mature enough for such a treatise. For similar reasons our emphasis is on the unconstrained problem, although we include a review on the work done so far in constrained derivative-free optimization. The constrained problem is in fact very important for applications, but theoretical treatment of derivative-free methods for constrained problems is very limited in the literature published to date, and thus, for the present volume at least, we are content to concentrate on the unconstrained case.

The book is meant to be reasonably self-contained and is addressed to readers at the level of a graduate student or a senior undergraduate with a background in calculus, linear algebra, and numerical analysis. Some elementary notions of optimization would be helpful but are not necessary. It is certainly our intent that practitioners would find the material covered to be both accessible and reasonably complete for their needs, whether their emphasis is on the algorithms or the applications. We have also made an effort to include figures and exercises when appropriate. The major aims include giving any interested reader a good idea of the state of the art of derivative-free optimization, with a detailed de-

scription of the basic theory to the extent that the reader can well understand what is needed to ensure convergence, how this affects algorithm design, and what kind of success one can expect and where. Thus, it is certainly our goal that the material be of interest to those who want to do research in the area.

As we state in our introduction, due to the growing sophistication and efficiency of computer simulations as well as of other applications, there is an increasing number of instances where one wishes to perform optimization of a complex system and the derivative information of the resulting objective functions is not available. This book is intended to help the reader to study and select, if necessary, suitable approaches to do exactly that. It is also intended to extract and emphasize the common theoretical features used by modern derivative-free algorithms, as well as highlight the differences.

We would be remiss if we ended our preface without some indicator of what the future holds in this area. There is still much waiting to be discovered. Undoubtedly, researchers and practitioners, perhaps soon, will discover ways to tackle much larger problems, whether through the use of massively parallel architectures or through advances in hardware yet to be realized, or through breakthroughs in the theory, or likely all three. The theory of constrained derivative-free optimization is also likely to advance significantly in the near future. Certainly, and especially because of the broad availability of difficult and important applications, this promises to be an exciting, interesting, and challenging area for many years to come.

Aside from the colleagues we mentioned in the first paragraph of this preface, there are many others we would like to thank who have provided us their valuable feedback in specific parts of the book, including, in particular, Natalia Alexandrov, Charles Audet, and Paul Tseng. We are very grateful to Ana Luísa Custódio for her tireless proofreading of the manuscript. We would also like to thank the reviewers who have refereed the version originally submitted to SIAM, in particular Tim Kelley and Jorge Nocedal, for their many interesting comments and suggestions.