superfluous. The scientific content is just enough to explain the equations, and the mathematical content is not redundant with the main text. In the chapter on bifurcations, for instance, pitchfork and Hopf bifurcations appear only in the main text and in the application, respectively. I particularly enjoyed the applications involving tides, orbital mission design, band gaps, and the KdV equation.

Analytical methods are well explained in most cases, but a few topics may be treated too briefly. One such topic is the connection between linear stability and matrix eigenvalues. Another is the integration of first-order separable equations, which is presented as symbol manipulation of differentials with no mention of the chain rule. An addition I would have found valuable is an explanation of the basic idea behind numerically integrating ODEs, even just Euler's method for initial value problems. This would give students something to picture despite not understanding Chebfun's more sophisticated workings. Some mathematics is referenced that may be unknown to students, such as the equivalence of norms on $\mathbb{R}^n$, or the regarding of functions as elements in a vector space, and a few terms are used without being explained (advection-diffusion, flow field, Jordan block, dissipative). Fortunately most of these gaps would not be hard to fill in during lectures.

The exercises that close each chapter are interesting and original, and they touch on important ideas beyond the scope of the text. Many are multipart questions with an exploratory flavor. Most chapters have enough exercises, although a few have only three or four. There are no exercises for rote practice of analytical methods like separation of variables or integrating factors. A few such exercises would be useful but can be obtained from numerous sources, unlike the questions the authors have crafted.

In short, *Exploring ODEs* is certainly good for independent study, and I expect it will be very good for an introductory ODE course as well. The Chebfun software is a testament to the maturity of methods for computing particular solutions, and its potential for numerically driven pedagogy is not limited to the scope of *Exploring ODEs*. One can imagine similar use of Chebfun in courses on dynamical systems or PDEs. Perhaps books on these topics will follow.

DAVID GOLUSKIN
University of Victoria

First-Order Methods in Optimization. *By Amir Beck.* SIAM, Mathematical Optimization Society, Philadelphia, 2017. $97.00. xii+475 pp., softcover. MOS-SIAM Series on Optimization. Vol. 25. ISBN 978-1-611974-98-0. https://doi.org/10.1137/1.9781611974997.

First-Order Methods in Optimization, as the title suggests, deals with (sub)gradient-based numerical algorithms for continuous optimization problems, where nonsmooth and convex problems are the center of attention. With the rise of *big data*, first-order methods in convex (and most recently nonconvex) optimization have experienced a huge renaissance. Hence a book like this one, in which different theoretical tools and an abundance of numerical methods are presented in a unified and self-contained manner, was sorely needed. Given his contributions and expertise in this research area, Amir Beck is an ideal person to have written it. The book is equally valuable to both researchers and practitioners. It can also be used as a textbook; as a matter of fact I am currently using parts of it to teach a course on convex optimization. The book is subdivided into fifteen chapters: The first seven chapters provide the theoretical tools from convex analysis on which the subsequent study relies. Naturally, a special emphasis is put on subgradients and proximal operators since these are the main building blocks for the optimization methods to be studied later. Moreover, the role of L-smoothness (i.e., L-Lipschitzness of the gradients) and strong convexity are highlighted, as these are fundamental properties that yield desirable convergence rates of numerical optimization methods.

The remaining eight chapters (Chapters 8–15) are devoted to optimization methods for solving (nonsmooth) optimization problems with different underlying structures, like additive composite or (explicitly) constrained problems. Methods such as (projected) subgradient, mirror descent, proximal gradient, conditional gradient, and ADMM are studied. These chapters are largely independent of each other (and hence can be read accordingly) and only build on the tools from the first seven chapters. The convergence results and their technique of proof are state of the art, containing (at the day of release) most current results. The book explicitly does *not* discuss termination criteria.

Many of the algorithms studied are implemented in a free MATLAB toolbox (available from www.siam.org/books/mo25) that nicely complements the book.

The style of the book is very clear and concise and, I should think, also accessible to the uninitiated reader and practitioners with working knowledge in linear algebra, multivariate differential calculus, and some basic topological notions in finite dimension. No space is wasted in the presentation of the theoretical tools, and the weighting of the different topics is clearly motivated by their importance for the subsequent study. The algorithmic part of the book definitely puts an emphasis on a unified, clean, and rigorous convergence analysis rather than on implementation details and numerical issues. Applications are usually presented after the convergence analysis has been carried out and are not used to motivate the algorithmic study but rather to illustrate the versatility of the methods. This part of the book is only slightly selective in terms of the algorithms studied, focusing on the most current ones, and not treating cutting-plane or bundle methods. It is, however, very comprehensive as to how the algorithms are studied, including dual counterparts, acceleration schemes, and specially tailored variants in the strongly convex case and stochastic versions.

I would particularly like to mention the crystal clear presentation of the accelerated proximal gradient method (FISTA), which I think is a particular gem. It is, in my opinion, better and clearer than most attempts made to explain only the special case of accelerated gradient descent. However, a reference to Nesterov's acceleration scheme (on which FISTA relies) could have been placed earlier, in Chapter 10, rather than only in the appendix.

A standing assumption almost throughout the second part of the book is that the smooth part of the objective functions is actually L-smooth. Clearly, this is not always needed to establish mere convergence; however, for the sake of uniformity and with the focus on convergence rates, this is more than justified.

The author's choice not to deal with termination criteria for the methods presented is a deliberate choice, and even though as a reader I would have liked to see some discussion in this regard, it is absolutely acceptable to say that this is beyond the scope and purpose of the book.

All in all, I am a big fan of the book and I strongly recommend it to everybody working in continuous optimization, machine learning, image processing, or related areas.

TIM HOHEISEL
McGill University

Mathematical Modeling. *By Christo Eck, Harald Garcke, and Peter Knabner*. Springer, Cham, 2017. $44.99. xv+509 pp., softcover. ISBN 978-3-319-55160-9.

Mathematical Modeling provides a broad overview of classical mathematical models primarily from physics and engineering targeted to advanced undergraduates and graduate students from mathematics backgrounds. In addition to introducing mathematical models and the relevant scientific background, the book also provides a survey of traditional methods of applied mathematics as needed to solve the model problems. This book is the English version of the previously published book in German.

Mathematical modeling is an extremely broad topic in terms of both mathematical tools and potential application areas. This makes modeling courses challenging to teach, as many students have widely varying backgrounds and expectations in terms of what should be covered in the course. Mathematical tools that may be used in a modeling course include applications of differential equations, networks, graph theory, differential geometry, numerical analysis, discrete dynamic systems, and