development techniques that are its subject. Particular attention is given to the development of large-scale software. The third part of the book is devoted to efficiency with respect to both time and memory. The BLAS (basic linear algebra subprograms) provide one illustration of the interaction of numerical algorithms and computer architecture. This is developed further by considering how level 3 BLAS are used in LAPACK to solve linear systems efficiently. Fortran programmers accustomed to static memory will especially benefit from the treatment of dynamic memory management in this part as they move to Fortran 90. A fourth part describes sources of scientific software and tools for managing revisions of large software projects and the like. To make points about how issues may be handled differently in the various programming languages, there are examples throughout in C, C++, Fortran 90, Java, MATLAB, etc. A reader who is not familiar with either C++ or Fortran 90 will not be able to appreciate fully some sections of the book. Indeed, the last part presents two medium-sized examples of numerical software development. One is the development in Fortran 90 of software for cubic splines, and the other is the development in C++ of software for solving linear systems by a multigrid method.

I found this book to be quite readable and, with its good advice on a wide range of topics, I learned a lot, so I recommend it warmly to readers with an interest in scientific software. In reading the book I found myself wondering whether I could use it in the classroom. It does not correspond to any of the courses being taught at present in the mathematics or computer science departments here. However, the authors have used portions of the book in several courses and I look forward to doing the same in a numerical analysis course.

L. F. SHAMPINE
Southern Methodist University

Deblurring Images: Matrices, Spectra, and Filtering. *By Per Christian Hansen, James G. Nagy, and Dianne P. O'Leary.* SIAM, Philadelphia, PA, 2006. $63.00. xiv+130 pp., softcover. ISBN 978-0-898716-18-4.

When I was asked to do this review or suggest another reviewer, my first inclination was to pass it on to other members of our team who are still working on very similar problems. But then I reconsidered. I thought that perhaps it would be fun and, more to the point, useful to the readers to do this review as a panel discussion between several of our team members. Very briefly, our team, the Data Driven Modeling and Analysis (DDMA) team, is spread over several groups at Los Alamos National Laboratory and several universities such as Michigan, Duke, and Rice. The team members are mathematical and computational scientists with specialties ranging from wavelets and signal processing to geometric measure theory to inverse problems and imaging and physics. (See http://ddma.lanl.gov for more information.)

Four DDMA team members agreed to do this with me. Rick Chartrand (RC) is an analyst who began life working in spaces of holomorphic functions on the unit disk but now does image analysis and compressed sensing research. Brendt E. Wohlberg (BEW) is an electrical engineer with expertise in wavelets and signal processing who works on algorithms for data modeling and imaging. William K. Allard (WKA) is very well known for his work in geometric measure theory, has worked on difficult numerical problems, and is now working in image data analysis from a geometric analysis viewpoint. Matthew J. Sottile (MJS) is a computer scientist known for his work in parallel architectures and parallel computing performance assessment, now spending significant effort on algorithms for image and data analysis. I (KRV) have worked on inverse problems, dynamical systems, and data modeling and analysis and am now focused on geometric measure theory and geometric analysis with a view to data challenges. In the review that follows, I will use initials to indicate the speaker.

Panel Review.

KRV: My overall impression of the book is very good. In particular, I think this book is quite ideal for a problem-focused path into numerical linear algebra. I think that an interesting course could be designed using this text for the first semester, with side trips to various related topics, with a second semester devoted to a text like Golub and Van Loan [2] or Demmel [1], filling in various topics needed to round out the course. The resulting course would give a nicely motivated and focused introduction to numerical linear algebra for those who are interested in using that subject in a wide range of computational contexts.

BEW: I agree. The book is clearly written and coherently structured. The restriction to the linear algebraic mathematical formulation, as explained by the authors in the preface, seems to have been very successful in making the book more accessible. Similarly, the careful attention to computational issues is likely to be very useful for the target audience. Despite the introductory level of the book, the authors have, commendably, included coverage of a number of more advanced topics, such as regularization parameter choice methods.

KRV: I thought that the "Very Important Points" (VIPs), "Pointers," and "Challenges" are very helpful.

MJS: Yes, I also liked the style of bringing key points out on their own as VIPs, and providing "Pointers" at the appropriate places to indicate important information relevant to implementation and using MATLAB. For a relative newcomer to this field of work, I found the book to be highly accessible to people who do not work with the theoretical foundations of the work on a frequent basis. The smooth presentation of the theory on which the deblurring techniques are built was quite nice—very little was assumed about the knowledge of the reader, yet the book avoided dwelling on excessive background material and proceeded at a reasonable pace.

The book does a good job of laying the groundwork for the methods and showing how the theory connects to practical implementation in MATLAB. I was also impressed by the care taken to inform the reader of implementation considerations that are often glossed over or ignored in other books. For example, the discussion of fft2 and ifft2 implementations on page 41 does a good job of explaining the wrinkles that should be understood when translating this particular method to concrete code.

WKA: As one whose knowledge of image analysis is deep in a few areas but rather spotty in most, I found the book to be an extremely readable presentation of the ideas behind and the practice of deblurring. Particularly useful and instructive are the many examples which one can study very easily using the MATLAB routines which are supplied. In my opinion, the availability of these examples and the skillful way they are employed in the text is the main strength of the book. Indeed, I intend to go through the book again and run most, if not all, of the MATLAB examples.

KRV: What were the quibbles you all had with the text?

BEW: I do have some minor quibbles:

- Most chapters follow a natural progression from problem statement, to mathematical formulation, to computational details of the solution. Chapter 2 is somewhat out of place as a tutorial on MATLAB functions, and would have been more appropriate as an appendix.
- Section 7.2 recommends solving Tikhonov problems on the least squares formulation, using a QR or SVD decomposition. The advice against solution using the normal equations overlooks the possibility of operators being available as matrix-vector product functions rather than explicit matrices. In this case QR or SVD decompositions are not possible, and the normal equations are indeed the correct approach.

RC: My observations, indexed by position in the text, are:

p. 2 Dark purple text is barely distinguishable from black, and a poor choice of a color for emphasizing important terms.

p. 3 Clearly the authors and the printer aren't on the same page as far as color goes: Figure 1.1 has purple where gray should be.

p. 8 Restriction to linear blur is understandable, but limits the applicability of the book, especially at LANL.

p. 24 Restriction to spatial invariant blur is less understandable, and should be made more clear. Most books aren't read front to back.

Chap. 4 The authors were following a nice expository path up to this point; one would have expected the discussion of the elements of the blurring model to lead into a description of the deblurring process. Instead, they suddenly bog down into details of special computations with particular kinds of matrices, which would have been better placed after one could see how it would be useful.

Chap. 5 This chapter opens, "The previous chapter shows that it is easy to solve noise-free problems that have structure." Not having paid attention to every detail of the matrix calculations, I missed any such mention entirely. It only came up on p. 50 under "Kronecker Product Matrices." Anything that assumes the reader is reading from cover to cover makes a book much less useful.

p. 56 About TSVD, they say "in spite of its simplicity, this method can work quite well," but the one example given seems pretty poor.

p. 76 Then again, looking at Figure 6.2, maybe that's the best that can be done, which seems disappointing.

p. 80 In the discrepancy principle, why not just use $\tau = 1$?

p. 88 Now it's blue that comes out purple, in a discussion of RBG. I wonder whether this is the fault of the printer or the publisher, the latter being responsible either way.

p. 96 The authors don't see fit to explain why p should be at least one, even though the narrative speaks in favor of smaller p.

p. 99 In my copy, this page could benefit from the techniques in the book, although the blur seems to be worse down on the page, meaning the spatial invariance assumption is not met. There are other pages like this.

Finally, the last 3 sections are too brief to be useful.

MJS: Like Rick, I found the linear blur formulation to be a bit restrictive in general practical applications. I agree with the choice the authors made, though, as this clear presentation of the restricted case will provide a firm foundation on which a reader can then build an understanding of how more complex blur models can be dealt with. Often books can become inaccessible to people new to an area by becoming complex in the interest of generality, which can lead to frustration.

On the other hand, a discussion (possibly appropriate in the final chapter) on how to step beyond the linear blurring model would have been appreciated. It would have been nice if topics relevant for addressing, say, camera jitter in photographs were at least introduced with pointers to sources of information on more complex sources of blur.

I would also agree with Brendt that Chapter 2 would be more appropriate as an appendix for those not familiar with the basics of MATLAB and the Image Processing Toolbox. Chapters 1 and 3 are natural neighbors and would flow well into each other. The presence of Chapter 2 disturbs this flow if one attempts to read the book in page order without skipping any material.

I found Chapter 4 to be a bit awkward to read, as I frequently found myself paging back and forth to make sure I was keeping track of the BCCB/BTTB/BTHB/BHHB terminology correctly. I don't think this is easily avoided by the authors, but readers should take time early in the chapter to ensure that they fully understand what each of these means so they can follow the remainder of the chapter without too much reference to the definitions on pages 37 and 38.

Finally, Chapter 7 is also different from the other chapters, but I do not find it to be useless. The chapter really should simply be titled "Advanced Topics," with no mention of the subtopics as none are covered in any great detail. I was a bit disappointed that each topic was simply presented very rapidly before a new topic was immediately picked up and discussed in equally brief terms. Fortunately, the chapter does come with a rich set of pointers to literature that the reader can refer to in order to explore these areas.

WKA: I also have a few complaints. In places the matrix algebra gets a little too pedantic. In other places, the impression is given that the reader should understand the result of a computation when in fact no justification whatsoever is given for the way things turn out. An example of this is the picture in the middle of page 5. While it is obvious to anyone that this picture is useless except as an example of how things may go wrong, no indication whatsoever is given as to why this is so. Like Matt, I was unable to remember what BCCB/BTTB/BTHB/BHHB meant. Perhaps a short glossary of acronyms and other symbols used in the text would be useful, if only so that one could read the book without starting at the beginning and without tedious backtracking fishing for the meaning of some symbol.

KRV: I think that the comment in Figure 7.5 can't be right: "The top image, which was computed with $p = 2$, has sharp edges . . . " since the Dirichlet regularization leads to a heat kernel in the gradient (Euler–Lagrange equation). This in turns always blurs edges. The only explanation can be that the α was too small or the method was far from convergence.

Like Rick, I was disappointed with results demonstrated. But this is a minor point since I believe the main point of the text is to be a very clear, useful path into numerical linear algebra using a particular imaging inverse problem as motivation.

KRV: Any final comments?

MJS: Like Rick, I also had to laugh that due to some printing error, the later pages in my text were perfect examples of spatially-dependent blur—although quite likely not due to the authors.

BEW: I think this book would be an excellent choice as a textbook for an undergraduate electrical engineering course.

RC: The book best meets the expectations as defined by the series title, "Fundamentals of Algorithms." The book discusses fairly thoroughly the fundamentals on one type of approach for one type of deblurring problem.

WKA: The book is accessible to anyone with a modest knowledge of matrix algebra and maybe a little Fourier analysis. In particular, it is accessible to good undergraduates with a technical background. I believe I could hand the book to undergraduates and be confident they would get from the book the basics of the subject without much help from me.

KRV: The large degree of agreement on the high merits of the text under review is in itself quite noteworthy: this group of five people are each quite opinionated and, even though we collaborate well, we are never given to agreement for the sake of agreement! Closing anecdote: When my colleague Tom Asaki and I first started working on inverse problems at Los Alamos a few years ago, we discovered smart scientists who had not yet computed the SVD of their linear operators and in some cases hadn't even heard of the SVD. I don't know how many times we explained the simple (yet very important) contents of pages 9–11 to people who needed to understand linear inverse problems. In fact, for some of these expensive inverse problems, we found that a combination of various methods (SVD, TV, etc.), which are explained in this book, worked very well. All this has led me to the impression that scientists, and not just those who work in the computational sciences, should be taught the SVD at about the same time they learn to walk. The text under review makes this task (and others like it) easier.

End of Panel Review.

To a very large degree, we all found the book clearly written and in line with the obvious objectives of the authors. In particular, we believe that this book will be great to learn and teach from; thin books are good, and this thin book will be sure to be appreciated by students and professors alike.

REFERENCES

[1] J. W. Demmel, *Applied Numerical Linear Algebra*, SIAM, Philadelphia, 1997.

[2] G. H. Golub and C. F. Van Loan, *Matrix Computations*, Johns Hopkins University Press, Baltimore, MD, 1996.

Kevin R. Vixie
Los Alamos National Laboratory