

Preface

This book began as a set of notes that I developed while teaching a graduate course on computational methods and uncertainty quantification (UQ) for inverse problems at the University of Montana (UM) in 2012, 2014, 2016, and 2017. I had three goals in mind as I wrote the manuscript. First, I wanted the book to start with a self-contained introduction to computational inverse problems that would be accessible to the wide range of students that I see in my graduate courses at UM: from beginning graduate students to more advanced Ph.D. students in pure and applied math, computer science, and statistics. Second, I wanted the book to end with Markov chain Monte Carlo (MCMC) sampling methods for UQ in inverse problems, my primary research focus in recent years. And third, I wanted the manuscript to be a cohesive whole, proceeding from the introductory computational inverse problems material in the first half of the book to the UQ for inverse problems material in the second half of the book, without large jumps in the progression of ideas. In other words, I wrote the book to be learned from.

Chapters 1 to 3 contain the introduction to computational inverse problems, with a focus on linear problems. In Chapter 1, the singular value decomposition (SVD) is used to illustrate the fundamental properties of inverse problems and to explore the statistical properties of the least squares solution. In Chapter 2, motivated by observations made in Chapter 1, regularization is implemented via the “filtered SVD,” yielding regularized solutions of inverse problems with better statistical properties than the least squares solution. Also in Chapter 2, several standard techniques for choosing the regularization parameter are presented. In Chapter 3, the methods developed in Chapters 1 and 2, which were tested on one-dimensional examples, are extended to two-dimensional test cases. Moreover, we introduce the conjugate gradient (CG) iteration for use on the two-dimensional test cases in which one cannot compute an SVD, as is the case for computed tomography (CT).

In Chapters 4 to 6, the focus shifts to Bayesian methods and UQ for inverse problems. The Bayesian point of view interfaces naturally with the introductory chapters on computational inverse problems, making for a smooth transition between the two parts of the book. In Chapter 4, we discuss prior modeling using Markov random fields, focusing mainly on Gaussian Markov random fields (GMRFs) and their relationship to differential operator–based regularization, but we also present Laplace Markov random fields and their relationship to total variation regularization. We then discuss methods for computing the maximum a posteriori (MAP) estimator, which is synonymous with the regularized solution in computational inverse problems.

The focus in Chapters 5 and 6 switches to sampling from the probability distributions that arise in Bayesian inverse problems. In Chapter 5, we begin with sampling from large-scale Gaussian probability densities and then turn our attention to various MCMC methods for sampling from the non-Gaussian posterior density function that arises in a hierarchical Bayesian approach to inverse problems. Finally, in Chapter 6, we

extend the methods presented in Chapter 5—which were developed for linear inverse problems with Gaussian priors—to nonlinear inverse problems and/or non-Gaussian priors. The MCMC methods presented in Chapter 6 mirror those presented in Chapter 5, although different algorithms, and innovations, are needed to account for the nonlinearity in the inverse problem, or prior, considered.

Readers not well versed in computational inverse problems should carefully read the material in Chapters 1 to 3, while those with more experience, and whose primary goal is to learn UQ, can move quickly through these chapters and focus more attention on Chapters 4 to 6. It is possible to cover all six chapters in a one semester course provided that the instructor moves quickly and skips some material. The book could also be used as a blueprint for a two-semester course, although additional material would need to be brought in. Finally, each chapter contains a large number of exercises meant to be accessible to graduate students.

Accompanying each chapter is a collection of MATLAB codes that implement the algorithms and generate the figures in the book. The codes are referenced within the manuscript, and many of the exercises require their use or modification. They can be downloaded at

www.siam.org/books/cs19

For instructions on how to run the codes, see `README.txt` in the base directory; my intention is that the codes will run on standard MATLAB, without any toolboxes.

I would like to thank my colleagues at UM for the opportunity to teach four courses using the various iterations of the book over the past seven years. I would also like to thank Per Christian Hansen, who brought me to the Danish Technical University in December 2017 to give a one-week short course from the book, as this forced me to work on and improve many of the book's exercises. A special thanks goes out to the students in those courses for their invaluable feedback and for providing inspiration. And finally, I would like to thank Marko Laine of the Finnish Meteorological Institute for allowing me to use and include his codes for performing MCMC chain convergence diagnostics.

Johnathan M. Bardsley
Missoula, Montana