

els (e.g., Ising model and the hard core gas models), the antiferromagnetic model and infinite graph models, among others. Although the list of applications of perfect sampling is growing, few fall within Bayesian computation, especially when compared with the huge impact MCMC has had on the analysis of complex Bayesian models. A possible companion to these chapters is the review article of Craiu and Meng (2005) that would complement the book with more emphasis on the statistical applications of CFTP, read-once CFTP and Fill's algorithm along with a particular focus on coupling techniques and variance reduction tricks that can be easily implemented for these procedures.

Chapter 7 is devoted to perfect simulation for Spatial Point Processes and the subsequent chapters move beyond the use of Markov chains for perfect simulation. Chapters 8 and 9 review the randomness recycler and advanced Accept/Reject, respectively. I found intriguing the concept of a Bernoulli factory and its use for generating solutions for general stochastic

Statistical Learning with Sparsity: The Lasso and Generalizations. Trevor Hastie, Robert Tibshirani, and Martin Wainwright. Boca Raton, FL: Chapman & Hall/CRC Press, 2015, xv + 351 pp., \$98.95 (H), ISBN: 978-1-49-871216-3.

The authors of some of the seminal articles on exploiting sparsity have written a book that describes the many applications of sparse methods and recent research results in the area. Exploiting sparsity in real-world problems has proven to be very effective in a wide range of application areas such as recommendation systems and signal reconstruction (compressed sensing), to name just two. This book introduces concepts with minimal jargon and is written assuming that the reader has only a basic background in optimization. However, understanding the advanced material toward the end of each chapter requires

more knowledge of optimization techniques. The book includes all the major branches of statistical learning. For each topic, the authors first give a concise introduction of the basic problem, evaluate conventional methods, pointing out their deficiencies, and then introduce a method based on sparsity. Thus, the book has the potential to be the standard textbook on the topic.

Each chapter includes comprehensive references and notes describing the evolution of its methods. For many the authors have developed toy problems to illustrate how the lasso-based solution depends on tuning parameters. The corresponding plots and other figures are helpful in developing the intuition behind why the lasso methods work. Additional detail is provided as exercises. Thus, the exercises complete the description of the methods in each chapter and are not just instruments for assessing how well the reader has understood the material. The complexity of the exercises makes the book suitable to be used as a textbook for an advanced graduate course. However, the book does not include detailed case studies of applications or code that can be directly applied to datasets. In the context of a formal course, the instructor could provide additional material to cover applications and thus provide a mix of both theoretical foundations and practice.

Chapter 2 describes the core principles behind the lasso and its properties. The authors present the advantages of exploiting sparsity while ensuring that the solutions can be efficiently computed, a recurring theme in the book. The authors draw parallels with other statistical approaches and in particular show how sparse constraints can be interpreted from a Bayesian perspective. This chapter includes an overview of the theory and its historical development. In Chapter 3, the authors give the formulation for different generalized linear models, including for special cases that arise in practice such as when observations are sequences of different lengths. A real-world but easy-to-follow example is provided for each generalized linear model. In

Chapter 4, the authors generalize the lasso principle to different penalty models. Again, the authors present real-world problems where such generalization is useful. Chapter 5 is an overview of optimization methods that are applicable to the lasso and its variants and includes a runtime comparison of two methods. This overview will be useful to the practitioner who needs guidance in choosing an appropriate optimization algorithm. Chapter 6 presents techniques to compute the statistical strength of the variables from the lasso and includes both traditional methods such as the bootstrap and newer research directions.

The remaining chapters present how the sparsity property can be applied to several standard problems with matrices, multivariate analysis, and graphical models. Matrix completion methods based on exploiting the sparsity principle are described in Chapter 7. The chapter includes a description of the popular Netflix challenge and the winning solution. Thus, this chapter can also be used independently to augment reading on recommender systems. Chapter 8 covers sparse variants of principal components analysis, correlation analysis, linear discriminant analysis, and clustering. Multiple ways of using sparsity for graphical model structure learning are presented in Chapter 9, with a particular emphasis on Gaussian graphical models.

The parallel development of compressed sensing techniques is covered in Chapter 10 along with sparse signal approximation methods. The book concludes with a set of theoretical results on bounds on the magnitude of error in the lasso solution and proofs for these. Figures are provided to illustrate these results in small problems and this helps in developing an appreciation of the theory even without an in-depth understanding of the theorems.

Anand Panangadan
California State University, Fullerton