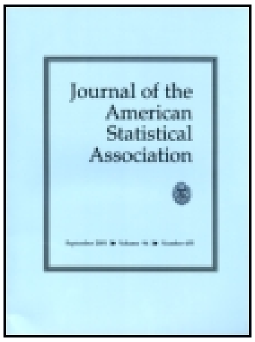

Journal of the American Statistical Association

ISSN: 0162-1459 (Print) 1537-274X (Online) Journal homepage: http://www.tandfonline.com/loi/uasa20

Book Reviews

To cite this article: (2016) Book Reviews, Journal of the American Statistical Association, 111:514, 912-919, DOI: 10.1080/01621459.2016.1200851

To link to this article: http://dx.doi.org/10.1080/01621459.2016.1200851

Published online: 18 Aug 2016.

JOURNAL OF THE AMERICAN STATISTICAL ASSOCIATION
2016, VOL. 111, NO. 514, 912–919, Book Reviews
http://dx.doi.org/10.1080/01621459.2016.1200851

Design and Analysis of Cross-Over Trials (3rd ed.).

Byron Jones and Michael G. Kenward. Boca Raton, FL: Chapman & Hall/CRC Press, 2014, xxvi + 412 pp., $99.95 (H), ISBN: 978-1-4398-6142-4.

As in the previous two editions, this edition offers a comprehensive coverage on the design and analysis of cross-over trials. With several major noteworthy updates, it will assist statisticians to conveniently tackle practical issues that arise in a cross-over trial. First, this new edition incorporates the newly available SAS procedures `proc glimmix` for categorical data analysis and `proc mcmc` for Bayesian analysis. Second, it includes updates on the use of period-dependent baseline covariates and the analysis of very small trials. Third, Chapter 4 employs the R package `Crossover`, which was specifically developed to accompany this book, to facilitate the design of cross-over trials. The package `Crossover` provides a graphical user interface (GUI) and can calculate the efficiencies and variances of a design that is either user-defined or selected from a built-in (and expanding) catalog. It also allows the users to search for an optimal design under three different choices for correlation structure of the within-subject errors and eight different models for carry-over effect. The most substantial update is the addition of seven new chapters (Chapters 8–14) in the form of short case studies. These real-world examples cover a wide range of issues and solutions above and beyond what is commonly encountered in a cross-over trial and significantly broaden the book. Specifically, Chapter 8 presents an example for fitting the Emax dose–response model with `proc nlmixed`, testing for noninferiority, and choosing the design for the last period based on interim results in a phase I trial. Chapter 9 deals with the comparison of several candidate dose–response models using the R package `DoseFinding`. Chapter 10 illustrates the computation of conditional power and sample size modification at the interim analysis from a phase III trial. Chapter 11 focuses on sample size calculation and reestimation based on a blinded estimate of the within-subject variance at the interim analysis from a phase II proof of concept trial. Together Chapters 12–14 give detailed discussions on sample size reestimation in bioequivalence studies. In summary, the third edition of Design and Analysis of Cross-Over Trials remains an outstanding reference for statisticians who work on cross-over trials, whether occasionally or frequently.

Haiying Chen
Wake Forest School of Medicine

Graphical Data Analysis with R.

Antony Unwin. Boca Raton, FL: Chapman & Hall/CRC Press, 2015, xiii + 296 pp., $69.95 (H), ISBN: 978-1-49-871523-2.

This is a slender volume that covers a remarkable amount of ground. The main emphasis is on how to make use of graphics to understand a dataset. Instructions on how to use R to achieve these graphics are inserted without weighing down the text, by the device of supplying the few lines of code, above each graph, which are needed to create them. The code is available on the book's web page. There is an extensive selection of exercises at the end of each chapter. The book would be accessible to a reader about to take their first introductory statistics class or embark on their first data analysis, but it is also important reading for instructors of data analysis courses and people involved in consulting and statistical practice. It could be used as the main textbook for an independent class on graphical data analysis.

What is in the Book?

Material is organized into an introduction about the purpose and principles of graphical data analysis, a summary of the literature, with the main content then bringing the reader through the use of graphics for simple to increasingly complex problems, and finishes with a capstone of case studies, notes on technology and philosophy. The primary subject matter chapters include

- Single variable plots, continuous and categorical types (chapters 3 and 4).
- Multiple variable plots, looking for relationships (chapters 5–7).
- Approaching and getting started with a dataset (chapters 8 and 9).
- Strategies for making useful plots (chapter 10)—may be the most important chapter in the book.
- Plotting temporal data (chapter 11), which harks back to Unwin's early research in statistical graphics.

The introductory chapter certainly sets the scene as its title promises: convincing the reader through a few examples that there are many seldom-discussed aspects of graphics that we really do need to think about. It starts with an example on skiing speeds, where a limited understanding of the dataset will lead to an inaccurate interpretation of a histogram. It also draws our attention to the fact that the importance and implications of aspects of a graphic, such as outliers, can strongly depend on the subject of the measurements. Thus, an identical graphic could attract a very different commentary, depending on what data are plotted. The examples in this chapter introduce many of the issues covered in more depth in later chapters. Perhaps more interesting, because it is less discussed in the general statistics literature, is the use of graphics to examine the quality of the data. For example, by employing very small bin widths in a histogram, a version of the famous Galton dataset of heights, available in R, is disclosed to have been reverse engineered from a tabular summary. There is further discussion of data quality in Chapter 9. Although "it is not always possible to find a test which can help you assess a feature you have discovered in a graphic" (p. 5) it is made clear that graphics are not a substitute for other analysis, and virtually every chapter has a section on modeling and testing relevant to the data whose graphics are under discussion in that chapter.

Chapter 2 provides the historical and literature context for graphical data analysis, including current web resources, an understanding of the availability of data, and some explanation of the choice of R. One advantage is "the integrated access to R's extensive range of statistical models and tools." (p. 22) Many of the well-known datasets available in R packages are used in the text explanations and exercises, and a further package of datasets has been created for the book, for particular topics when the existing data-sets are inadequate, for example, Hertzsprung–Russell diagram of stars magnitude and temperature (p. 97).

Chapters 3–7 work us through the toolbox of single to multivariate plots. Basic plots such as bar charts, pie charts, histograms, and scatterplots are discussed along with less commonly used displays such as parallel coordinate plots, mosaic plots, doubledecker plots, and fluctuation diagrams. Very useful tidbits in these chapters are the sections on choice of display, summary of a variety of functions available to make the plots in R, and nuances of construction. These five chapters fill more than half the pages of the book.

Chapters 8 and 9 lead the reader through obtaining an overview of a dataset and dealing with practical hurdles such as handling missing values.

Chapter 10 explains how to make comparisons with graphics. "At the heart of quantitative reasoning is a single question: compared to what?" (p. 199, attributed to Tufte). We make comparisons between subsets, old data, different populations, sources, against a standard—it is the basis of statistical thinking. Making effective comparisons is difficult, and this material elucidates the process.

Chapter 11 contains a brief treatment of time series data. It would have been nice to see a chapter on spatial data following this, because it is also an area where Antony Unwin has made a lot of research contributions.

Putting it all together is the focus of the case study in Chapter 12. This chapter is too brief! There is only one case study and then many more potential case studies are left to the reader to work through themselves as exercises. This is the chance for the reader to discover what they have learned from the earlier chapters in terms of solving problems using graphical data analysis.

There is one chapter focusing on graphics with R just before the final summary chapter. Chapter 13 has a discussion of what to aim for in a graphic, as well as supplying some understanding of how R and its packages work, and the advantages and disadvantages of various different graphics systems that are available in R. Many aspects of appearance that require control in R are also discussed, for example, the effect on the appearance of a graph of both size and aspect ratio.

From a Teacher's Perspective

This is potentially useful as a textbook for a course primarily focused on exploring data visually. There are several of these courses offered at different institutions around the world, for example, Monash University's ETX2250 Data Visualization and Analytics, or Carnegie Mellon University's 36-721, Statistical Graphics & Visualization. There are likely to be more courses along these lines as more data science programs emerge. To be successful using the book for teaching the instructor would need to supplement the material with additional R instruction, from books such as Winston Chang's (2013) *R Graphics Cookbook* and Garret Grolemund's (2014) *Hands-On Programming with R*. It may also be possible to use Naomi Robbins' (2013) book *Creating More Effective Graphs* with its sage advice and Excel examples in an introductory few weeks of such a class. For working on current data challenges, it would also be ideal to incorporate new data for projects or longer assignments.

From a Consulting Statistician's Perspective

W. Edwards Deming said that without data you are just another person with an opinion. Today, when we are flooded with complicated and large datasets, I would dare to say that without a plot you are just a person missing a convincing argument (potentially for a future confirmatory analysis). This book is a great reference book for a researcher or a consultant to get inspiration about different ways of exploring the features in the analyzed data. In addition, the book increases the awareness of the observers' perception of the data displayed in graphs with different graphical choices.

Concluding Remarks

The book emphasizes exploratory data analysis, as opposed to presentation, and recommends that we do not need to be bound to use only one plot, but to make and show many different displays of the same dataset. This is an important point. So often, we try to draw one optimal, overloaded, plot that actually hampers learning what the data have to say. It is relatively easy to generate many displays, and even include many additional displays in electronic supplementary material in journal articles, and this is emphasized in the ensemble graphics material in chapter 13.

The approach to presenting R code is just one example of very careful organization of the content of the book, allowing it to supply a broad range of ideas without rendering the book heavy. Other proof of clever organization includes well-targeted use of example datasets and the occasional use of succinctly written and well-presented lists containing useful commentary. Examples of these are what features might be visible in a scatterplot (p. 77) and a possible strategy for outliers (p. 192). In a similar way, chapters end with a summary of the main points, making it easier to refind examples and ideas.

This book comprises the career work of Antony Unwin, on developing and using data plots to analyze and solve problems. The author is well known for his work in interactive graphics. The software that has emerged from his group during and after their studies with Unwin include DiamondFast (for time series), MANET (Missings are now equally treated), Klimt (decision trees), Cassatt (parallel coordinates), Mondrian (multivariate data and maps), iplots (interactive graphics in R), and RCloud (running large data analyses). Most of these softwares provide interactive graphics tools for exploring data, so it would have been nice to see something about interactive graphics in this book, but the omission is understandable. The technology for seamlessly integrating interactive graphics and modeling is not available yet, so this is something for the future. All of the above-mentioned packages are stand-alone and do not integrate easily with R. In this book, Antony Unwin adapts to the current technology taking currently available graphics software (primarily `ggplot2` in R) and shows how to make good static plots of data to solve problems. The book fills a much needed gap in the current literature and will be a valuable resource for some time.

One of the reviewers, Dianne Cook, also reviewed a prepublication version of this book for the publisher.

References

Chang, W. (2013), *R Graphics Cookbook*, Sebastopol, CA: O'Reilly Media. [914]
Grolemund, G. (2014), *Hands-On Programming With R*, Sebastopol, CA: O'Reilly Media. [914]
Robbins, N. (2013), *Creating More Effective Graphs*, Chart House Press. [914]

Dianne Cook, Jill Wright, and Julia Polak
Monash University

Mean Field Simulation for Monte Carlo Integration.

Pierre Del Moral. Boca Raton, FL: Chapman & Hall/CRC Press, 2013, xlvii + 578 pp., $115.95 (H), ISBN: 978-1-4665-0405-9.

In the last 50 years we have witnessed a rapid growth in the development of probabilistic methods for approximating the evolution of dynamical systems. A class of these probabilistic methods uses sets of particles with an evolution that is (possibly) random and which interact with each other. Such methods have been the focus of attention of mathematicians, statisticians, engineers, computer scientists, physicists, computational biologists, etc. Depending on the area of application, the component particles can be interpreted as individuals, random variables, samples, realizations, candidate solutions, etc., while the corresponding methods are known under various names, including genetic algorithms, sequential Monte Carlo methods, particle filters, interacting particle models, etc. Most often than not, the interaction is of a *mean field* type. This means that each particle interacts with the others through the occupation measure (empirical distribution) of the entire set of particles, typically through the value of a particular statistics averaged over particles' spatial positions.

The literature covering this class of methods is highly fragmented. Many of these methods are analyzed and developed virtually independently in each area of application or community of interest. The monograph *Mean Field Simulation for Monte Carlo Integration* aims to address all the areas of application of mean field simulation in an attempt to provide an integrated approach of both the theoretical and the methodological developments governing these classes of probabilistic methods. The book is divided into four parts and, in the following, I will describe briefly their contents.

The first part contains an extended justification of the importance (if one was needed!) and the applicability of mean field particle approximations and also a description of the corresponding theory of mean field simulation in an abbreviated form. As the author strongly recommends, this part should be compulsory reading. The salient parts of the work are discussed and a collection of selected theoretical results are included that will guide both practitioners and theoreticians through the material.

The second part of the book introduces a general class of evolution systems that are amenable for approximation via mean field interaction particle systems. Such systems are termed Feynman–Kac models. The author presents in the second part a large sample of examples both in continuous time and in discrete time, including nonlinear filtering problems, particle absorption models, spatial branching processes, etc. The evolution of these systems in discrete time can be modeled via a sequence $(\eta_n)_{n\geq 0}$ of probability distributions satisfying a nonlinear equation of the form

$$\eta_{n+1} = \Phi_n \eta_n. \tag{1}$$

In (1), Φ_n is the composition of two operators

$$\Phi_n = \Psi_n \circ M_n. \tag{2}$$

The operator Ψ_n is termed a Boltzmann–Gibbs transformation. For any probability measure μ, this transformation returns a probability measure $\Psi_n(\mu)$ absolutely continuous with respect to μ. The density of $\Psi_n(\mu)$ with respect to μ is a given function g_n. In other words, for any set A belonging to the σ-algebra on which μ is defined, we have

$$\Psi_n(\mu)\,(A) := \frac{\int_A g_n(x)\,\mu(dx)}{\int g_n(x)\,\mu(dx)} = \frac{1}{c_\mu}\int_A g_n(x)\,\mu(dx). \tag{3}$$

In (3), c_μ is the normalization constant $c_\mu = \int g_n(x)\,\mu(dx)$. The second transformation in (2) is a standard Markov operator. The nonlinearity of the operator Φ_n occurs as a result of the normalization procedure incorporated in the Boltzmann–Gibbs transformation Ψ_n. Both the analysis of Feynman–Kac models and the corresponding methodology to approximate them, that is, the mean field Monte Carlo method have to cope with this difficulty.

In the second part of the book no less than four mathematically equivalent particle interpretations are described: spatial branching models, sequential Monte Carlo methods, interacting Markov chain Monte Carlo algorithms, and mean field interacting particle models. Each of these particle interpretations involves a sequence $(\eta_n^N)_{n\geq 0}$ of probability distributions, where, depending on the field of application or on the type of scientist that studies them, N is the number of particles, samples, realizations, etc., whose empirical distribution is η_n^N. The sequence $(\eta_n^N)_{n\geq 0}$ satisfies a nonlinear equation of the form

$$\eta_{n+1}^N = \Phi_n^N \eta_n^N. \tag{4}$$

In (4), Φ_n^N is a random (Monte Carlo) approximation of Φ_n, which improves as N increases. The approximations are asymptotically consistent in the sense that

$$\lim_{N\to\infty} \eta_n^N = \eta_n$$

for all $n \geq 0$.

The second part concludes with a chapter on nonlinear evolution of intensity measure models highly relevant to applications involving multiple-object filtering (multi-target tracking).

The third part of the book includes two substantial application domains for mean field particle approximations. The first entails particle absorption models including Markov chains with restricted movements, polymer models, self-avoiding walks. It also contains applications to approximations of the largest eigenvalue of Schrödinger operators and quasi-invariant measures. The second domain represents applications to signal processing and control systems: linear and nonlinear filtering, parameter estimation in hidden Markov chains, or optimal stopping problems.

The fourth part of the book (Chapters 9 to 17) is the largest in the book. It is here that the author proceeds to develop a comprehensive theory for mean field simulation. Chapter 9 revisits the concept of a Feynman–Kac model. This is an elegant self-contained chapter that could be the focus of a short introduction to both the topic and some basic convergence results. More details about the modularity of the book can be found below. Chapter 10 covers a general class of mean field models, some illustrative example, and a short fluctuation analysis section. Chapter 11 covers nonasymptotic and concentration results, maximal, and Cramer–Chernov inequalities for empirical processes.

The author discusses Feynman–Kac and intensity measure semigroups and particle density profiles in Chapters 12–14. The last three chapters of the book cover genealogical tree models, approximation methods for the normalizing constants, and backward particle Markov models.

The book serves both research and pedagogic needs. On the one hand, it serves as a monograph covering an extensive arrays of theoretical and methodological results. On the other hand, it is cleverly built in a modular manner (which unavoidably requires repetition of some of the material) so that it can serve as a textbook to teach the subject at different levels and for different audiences. The author gives detailed instructions on the pedagogical usage of the book in the preface.

The monograph is a welcome continuation of the author's successful monograph (Del Moral 2004). It is certain to initiate and facilitate the interaction among the communities whose research revolves around interaction, more precisely around interacting particle systems. All in all it is a must read to both theoreticians and practitioners interested in probabilistic methods for approximating the evolution of dynamical systems.

Reference

Del Moral, P. (2004), *Feynman-Kac Formulae. Genealogical and Interacting Particle Systems With Applications. Probability and its Applications*, New York: Springer-Verlag. [916]

Dan Crisan
Department of Mathematics, Imperial College London

Measuring Statistical Evidence Using Relative Belief.

Michael Evans. Boca Raton, FL: Chapman & Hall/CRC Press, 2015, xvii + 232 pp., $89.95 (H), ISBN: 978-1-48-224279-9.

This book presents a comprehensive survey of the literature on measuring statistical evidence based on relative beliefs. General conceptual discussions on statistical evidence have a long and rich history including, for instance, the works by Birnbaum (1962), Good (1985), Royall (2000), and Thompson (2007). This book integrates recent developments into a coursebook format with a pedagogical introduction to the area. The author is a main contributor in the field and the book systematically assembles material conveying his findings and his views on statistics as a whole. Emphasizing the inherent subjectivity of statistical analyses, he strives to establish a "gold standard" for a statistical analysis. The author's strategy is to perform statistical inference based on updating relative beliefs, is in a certain sense embedded into the Bayesian paradigm, and exploits the hope that the data at least are objective.

From my perspective, there are two main parts of the book. The first three chapters provide the readers with a basic introduction to the foundations of probability and statistics viewed through the lens of the author's school of thought. The author points out shortcomings of popular methodologies and argues that finite parameter spaces are sufficient for nearly any real-world application of statistics. This part summarizes the major positions taken on probability and statistics and prepares the reader for the second part. I found this to be an enjoyable read. Many illustrative examples reveal intriguing paradoxes in statistical theories, some of them are well-known and complemented with a broad informative discussion and others are less obvious.

At the same time, I felt slightly indoctrinated with the authors' perspectives: the main messages are frequently repeated and the entire text seems to be designed to fortify his point of view. Basically probabilities are interpreted to express degrees of beliefs and the ingredients of a statistical analysis, namely, a model and a prior, should be investigated according to the principal of empirical criticism. Most importantly, this means that subjectively chosen ingredients should not be contradicted by the objective data.

The second part of the book is devoted to the measurement of statistical evidence based on relative beliefs. Chapter 4 focuses on measuring evidence and the assessment of the strength of the evidence. Given that the model and prior are properly chosen, the evidence is deduced by a change in beliefs induced by the observation of the data. At the heart of the analysis is the relative belief ratio that is, for events with positive mass, defined as a quotient of conditional and unconditional probabilities. For some considered event, a relative belief ratio larger than one reflects an increase of belief whereas a ratio smaller than one gives evidence against the event. First, this principle of evidence does not apply in case of continuous densities. Also, ratios for different sets are not directly comparable. Suitable regularity conditions allow, however, a generalization of the concept by considering a ratio of conditional and unconditional densities. It is convincingly demonstrated that a few desirable properties for a measure of evidence naturally lead to the relative belief ratio. As a device to quantify the strength of evidence certain p-values are suggested. In these central passages, the exposition becomes more detailed. An increased number of theorems and proofs bear witness to the more condense mathematical content and rigor.

The developed statistical theory includes the least relative surprise estimator that is the maximizer of relative belief and in its nature close to the classical concept of M-estimation. Hypotheses assessment is established and carefully separated from standard statistical testing approaches. The parameter space is finite. Possible extensions are addressed and require certain regularity conditions; the author takes a firm stance in professing that this does not cause serious restrictions. The theorems that I found most appealing include those regarding asymptotic large sample theory, consistency, and certain optimality results.

The recommended recipe to conduct a statistical analysis can be summarized as follows:

- Presuppose properly collected data (objective element);
- Choose a (subjective) model and judge it against the data;
- If the model passes the check, choose a prior and judge it against the data;
- If both are adequate, acquire statistical evidence based on changes of relative beliefs.

Advice on the first steps, choosing and checking the model and the prior, is provided in the fifth chapter. It is admitted that to some extent models, typically determined by a compromise to fit the data reasonably well while preserving enough simplicity to be useful for statistical inference, can be selected according to the preferences of the statistician in the context of a specific problem. However, gross violations of the stylized facts should be avoided. The theory to evaluate the model by Evans and Jang (2010) is discussed. For choosing proper priors, an elicitation process is prescribed. Only in case that the described tests for prior-data conflicts react is the prior rejected.

Promising a new "gold standard" to measure statistical evidence the author raises high expectations. Having read his book and pondered it, the main findings and the stated messages are more modest. My impression is that the discussed assessment of evidence still leaves some freedom to the statistician. Even more, accepting a degree of subjectivity these beliefs take the role of an important instrument of the process. The method is not completely opposed to other traditional approaches—but rather attempts to combine the best features of several. With some distance from this research area, I can say that reading the book has influenced my general view on statistics without, however, changing my everyday work with statistics. Probably, most high-level statistical analyses, including those that do not formally follow this mode of inference, are aware of and proceed in accordance with the so-called principle of empirical criticism.

Though the first (overview) part of the book does not presuppose much background knowledge, my feeling is that a sound grasp of the subject will help the reader come away with a solid understanding and interpretation of this tour. On the other hand, the main theoretical part does not require many technical prerequisites and is accessible to students having a basic background. I am not sure of the intended audience, but I guess the author's goal is to disseminate the theory to a broad readership; in this he is successful. Personally, I have not yet decided if I will use the book for student seminars. It certainly does a good job of gathering works on measuring statistical evidence, especially contributions from the authors' own work including Evans and Jang (2010) and Evans and Jang (2011) among many others, hitherto unavailable.

References

Birnbaum, A. (1962), "On the Foundations of Statistical Inference" (with discussion), *Journal of the American Statistical Association*, 57, 269–306, 307–326. [916]

Evans, M., and Jang, G. H. (2010), "Invariant P-values for Model Checking," *The Annals of Statistics*, 38, 512–525. [917]

——— (2011), "Weak Informativity and the Information in One Prior Relative to Another," *Statistical Science*, 26, 423–439. [917]

Good, I. J. (1985), "Weight of Evidence: A Brief Survey" (with discussion), in *Bayesian Statistics, Proceedings of the Second Valencia International Meeting*, eds. J. M. Bernardo, M. H. DeGroot, D. V. Lindley and A. F. M. Smith, Amsterdam: North-Holland Publishing Co., pp. 249–269. [916]

Royall, R. (2000), "On the Probability of Observing Misleading Statistical Evidence" (with discussion), *Journal of the American Statistical Association*, 95, 760–780. [916]

Thompson, B. (2007), *The Nature of Statistical Evidence. Lecture Notes in Statistics* (Vol. 189), New York: Springer. [916]

Markus Bibinger
University of Marburg

Robust Methods for Data Reduction. Alessio Farcomeni and Luca Greco. Boca Raton, FL: Chapman & Hall/CRC Press, 2015, xxvii + 269 pp., $99.95 (H), ISBN: 978-1-46-659062-5.

Technological advances in data collection and assimilation techniques over the last decade have resulted in the modern data deluge, in which large-scale datasets are increasingly being collected in business, media, medicine, and science. These massive datasets typically contain either a large number of observations, a large number of variables (dimensions), or both. Data-reduction techniques are used to extract relevant information from these massive pools of data in tractable forms that can be subsequently visualized and subjected to further statistical analysis. This well-established sub-topic of multivariate data analysis has become a vibrant area of current research as traditional data-reduction methods need to be enhanced to deal with big datasets. Also, these massive volumes of data are usually unstructured and contain varying levels of noise. Robust statistical methods should be used to analyze them, because, when applied to noisy datasets, methods that are heavily influenced by outliers may lead to misleading results. Robust methods have a long history in statistics, but are seldom used in practice because they can be complicated and complex.

Robust Methods for Data Reduction makes it easy for practitioners of big-data analytics to conduct robust and efficient data reduction. It is a timely topic in which recently prescribed algorithms and methodological research findings are properly assimilated and presented in a lucid fashion. The book serves as a good introductory book that motivates and teaches the art of developing robust frameworks for synthesis and reduction of large, complex datasets.

The first two chapters of the book provide a terse, gradual introduction to basic robustness concepts and classical multivariate data-reduction methods. These two chapters do not require any prior subject knowledge and provide an ideal starting point for beginners. The authors have taken care to clearly explain the statistical notions and mathematical definitions. The most appealing aspect of this book is that all of the concepts and algorithms described are inspired by real-data examples. All of the methods presented in this book are accompanied by extensive codes and exhaustive documentation on how to implement them in the R computing environment. Readers can download the data and the computer code used in the book from the publisher's webpage. The codes, along with demonstrations provided in the chapters, can be used as a step-by-step, hands-on, self-learning experience. Besides the implementation of methods in the R computing environment, the book also offers informative references to associated toolboxes in other mathematical and statistical software.

The remainder of the book is divided into two sections. The first, presented in Chapters 3–6, describes dimension-reduction techniques that involve mapping multivariate datasets into lower-dimensional manifolds ideal for inference and visualization. The second section describes methods that cluster observations into groups and quantify the information spread across the observations to form fewer groups with distinct heterogenous characteristics. These methods are presented in Chapters 7–11 and are called sample-reduction algorithms because they split the sample into a few representative groups.

A fundamental dimensional-reduction technique, principal component analysis (PCA), is discussed in detail in Chapter 3. The classical nonrobust approach is illustrated first. Thereafter, the authors present several robust adaptations of the traditional technique and demonstrate outlier-detection strategies based on PCA analysis. Chapter 4 is entirely devoted to robust sparse PCA and its associated algorithms. Unlike PCA, in which the major directions of variation are explained by linear combinations of all variables, the assumption in sparse PCA is that only a smaller subset of the variables contributes to significant variability directions. Thus, sparse PCA transforms the data into a lower dimensional manifold in such a way that the reduced space can be explained by the effects of a smaller tractable subset of variables. In this context, sparse PCA conducts variable selection while reducing data dimension. Chapter 4 explains robust sparse PCA algorithms and their implementational details. Data-adaptive strategies to detect sparsity levels are discussed, and the applicability of the robust sparse PCA method is demonstrated in real data problems. Chapters 5 and 6 describe the usage of canonical correlation analysis and factor analysis to reducing data dimension robustly.

The sample-reduction section begins in Chapter 7 with a primer on traditional clustering criteria and algorithms. This chapter describes two major classes of clustering algorithms, (a) distance-based and (b) model-based. Robust adaptations of the methods in the two classes appear in Chapters 8 and 9, respectively. Many real-world data-reduction problems (Tanay, Sharan, and Shamir 2005) need the observations (in rows) as well as the variables (in columns) to be simultaneously clustered. Robust bi-clustering algorithms for solving these kinds of problems are presented in Chapter 10. In Chapter 11, the problem of supervised sample reduction is analyzed via discriminant analysis.

The authors point out from the beginning that their objective is not to provide an all-inclusive exposition of the large literature on data reduction, but to emphasize the applicability and usage of robust data-reduction methods by discussing few popular methods in details. They have done an impressive job in this respect. The book is concise and has a very readable style. Through figures, computer code, and data examples, it successfully communicates the applicability of some technically complicated methods. The book follows a logical structure; the discussion in the chapters is well supported by relevant illustrations. However, low-rank models (Markovsky 2011), which are very popular nowadays for robust data reduction, do not receive considerable coverage. The authors, however, have provided fairly comprehensive bibliographical notes that point to the important results of the field, and interested readers can follow these resources to learn more about methods not described in this book.

In summary, *Robust Methods for Data Reduction* is a useful book that will greatly help practitioners and analysts to understand, appreciate, and implement correct data-reduction techniques. The collection of data examples and the pedagogical writing style make it an ideal text for instructors aiming to quickly train students on proper data-reduction techniques. The book does not contain end-of-chapter exercises and is mainly suitable for graduate level data analytics courses, though parts of it could be included in an undergraduate applied statistics course. This book will be particularly useful for courses with R labs. It is bound to find a wide and enduring readership and will be a valuable addition to the library of any data scientist.

References

Markovsky, I. (2011), *Low Rank Approximation: Algorithms, Implementation, Applications*, London: Springer Science & Business Media. [918]

Tanay, A., Sharan, R., and Shamir, R. (2005), "Biclustering Algorithms: A Survey," *Handbook of Computational Molecular Biology*, 9.1-20, 122–124. [918]

Gourab Mukherjee
University of Southern California

Theory of Factorial Design: Single- and Multi-Stratum Experiments. Ching-Shui Cheng. Boca Raton, FL: Chapman & Hall/CRC Press, 2013, xvi + 393 pp., $109.95 (H), ISBN: 978-1-46-650557-5.

The field of experimental design aims to help practitioners collect their data in a more efficient manner, or more specifically, run their experiments more effectively. There are many good textbooks in this area: the classical ones of the early 50's (e.g., Cochran and Cox 1957) focused more on agricultural experimentation; the later ones of the late 70's (e.g., Box, Hunter, and Hunter 1978) focused more on industrial experimentation, and the recent ones (e.g., Santner, Williams, and Notz 2003; Fang, Li, and Sudjianto 2006) focused more on computer experiments. There are also some theoretical approaches, notably on optimal design (e.g., Pukelsheim 1993) and combinatorics (e.g., Street and Street 1987). This (Cheng's) book is clearly one of the very first about design of experiment from a multi-stratum approach. I expect that it will prove to be an influential design book for years to come. The closest book is probably Bailey (2008), but Cheng's book is much more completed and updated.

In the preface, the author states that "*The objective of this book is to provide a rigorous, systematic, and up-to-date treatment of the theoretical aspects of this subject.*" and goes on to say "*A theory of orthogonal block structure that goes back to John Nelder provides a unifying framework for the design and analysis of multi-stratum experiments. One feature of the present book is to present this elegant and general theory which, once understood, is simple to use, and can be applied to various structures of experimental units in a unified and systematic way.*" This more or less indicates the uniqueness of this beautifully presented book. Some topics have never appeared in any other book and the author has produced elegant mathematics accompanied with lucid explanations.

The book is organized into 15 chapters. Chapter 1 provides an overview, which is very helpful for those who decide to carefully go through the entire book, or to use this book as a textbook. Chapters 2–5 are devoted to background material. It took me a while to get used to the notation used in the very brief Chapter 2 that contains a brief review of linear models that are used for the analysis and design criteria in later chapters. Chapter 3 (again very brief) introduces basic design concepts on randomization and blocking. Chapter 4 is rather unique in the way it makes use of proportional frequencies, first introduced in Chapter 2, to build a foundation for future chapters. The Hasse diagram is used here to illustrate block structure. I personally think Chapter 4's title "Factors" is somehow noninformative, however. Chapter 5 is the key component of the book. According to the author (p. 12), "*it [mainly Theorem 5.1] is to present a unified treatment of the analyses of three classes of orthogonal designs (completely randomized designs, complete block designs, and Latin square design) ... It is also a key result for developing a general theory of orthogonal designs for experiments with more complicated block structure.*" It is probably fair to say that Theorem 5.1 (p. 52) is the key theorem for the entire book. (Another key theorem is Theorem 13.2.) Theorem 5.1 is a powerful result and covers design construction for random effects, including fixed effect designs as special cases.

Chapters 6 and 7 are devoted to (full) factorial designs. Chapter 6 introduces treatment factorial designs through orthogonal polynomials, finite Euclidean geometry, and Abelian groups. Chapter 7 presents complete factorial designs in incomplete block, row-column layout, including split-plot and strip-plot designs. The concept of "design keys" is cleverly applied here.

Chapters 8–11 are devoted to fractional factorial designs. The combinatorial structure of orthogonal arrays is introduced in Chapter 8 and a brief connection to (recently popularized) computer experiments is discussed at the end of this chapter. Chapter 9 presents the regular fractional factorial designs, but uses a different set of notation than those used in previous chapters. Chapter 10 gives some recent advances in fractional factorial designs (mainly based on minimum aberration criterion). Chapter 11 focuses on Two-Level Resolution IV designs. While the content is good, many recent results have been missed, especially those pertaining to designs not of the foldover type; it seems to me that this topic is too narrow for a full chapter. The claim (p. 12) of "*a unifying framework for the design*" does not seem to fit these chapters particularly well.

Chapters 12–14 are devoted to factorial designs with more complicated block structures, namely, the multi-stratum designs. They connect nicely with Chapters 5–7 in terms of their content, approach, and notation. They are unique in that they are not included in other design books. Perhaps this is why they are mentioned in the title "Theory of Factorial Design: Single- and Multi-Stratum Experiments." Chapter 12 begins with a simple block structure. Chapter 13 then provides a general theory for multi-stratum complete factorial designs and details "design keys" in Section 13.6. Chapter 14 provides a rather comprehensive study on the construction of multi-stratum fractional factorial designs.

Chapter 15 is a relatively self-contained presentation of nonregular designs. Some recent interesting and important advances in this area are surveyed. As noted by the author that (p. xvi) "*The research on nonregular designs is still very active and expands rapidly. It deserves another volume.*"

I believe that this excellent book will soon become a must read for researchers and educators in experimental design. It could serve as a great reference or textbook for a high-level design course. Such a high-level design class, however, may be too specialized for most statistics education programs. For practitioners this book may be too theoretical: it would take some time to become familiar with its notation and concepts. I hope that this is a groundless concern from a person who has long known and respected the author. As claimed by the author, "*...once understood, (the notation) is simple to use, and can be applied to various structures of experimental units in a unified and systematic way.*" It would be ideal to implement the important concepts and theorems in this book into software so that they can be readily used for practitioners who do not typically have much theoretical background in design.

References

Bailey, R. A. (2008), *Design of Comparative Experiments*, Cambridge, UK: Cambridge University Press. [918]

Box, G. E. P., Hunter, W. G., and Hunter, J. S. (1978), *Statistics for Experimenters*, New York: Wiley. [918]

Cochran, W. G., and Cox, G. M. (1957), *Experimental Designs*, New York: Wiley. [918]

Fang, K. T., Li, Runze, and Sudjianto, A. (2006), *Design and Modeling for Computer Experiments*, Boca Raton, FL: Chapman and Hall/CRC. [918]

Pukelsheim, F. (1993), *Optimal Design of Experiments*, Philadelphia, PA: SIAM. [918]

Santner, T. J., Williams, B. J., and Notz, W. I. (2003), *The Design and Analysis of Computer Experiments*, New York: Springer. [918]

Street, A. P., and Street, D. (1987), *Combinatorics of Experimental Design*, Oxford, UK: Oxford Science Publications. [918]

Dennis Lin
Penn State University