

**Excerpts from Preface to
*Introduction to Stochastic Search and Optimization:
Estimation, Simulation, and Control*
by J. C. Spall**

© John Wiley and Sons, Inc., 2003

THE SUBJECT

Individuals and organizations are often faced with making tradeoffs in order to achieve desirable outcomes. Choosing these tradeoffs in the “best” way is the essence of the search and optimization problem. Mathematical algorithms for search and optimization play a large role in finding the best options in many problems in engineering, business, medicine, and the physical and social sciences. *Stochastic* search and optimization pertains to problems where there is random noise in the measurements provided to the algorithm and/or there is injected (Monte Carlo) randomness in the algorithm itself. A countless number of stochastic search and optimization problems arise in industry, academia, and government.

SUMMARY OF THE BOOK

Introduction to Stochastic Search and Optimization provides a broad survey of many of the most important methods in stochastic search and optimization. These include random search, recursive least squares, stochastic approximation, simulated annealing, evolutionary computation (including genetic algorithms), and reinforcement learning. Also included is a discussion of closely related subjects such as multiple statistical comparisons, model selection, simulation-based optimization, Markov chain Monte Carlo, and experimental design. These subjects are covered in 17 chapters. Each chapter ends with some concluding remarks, including comments on the historical perspective and on the nexus with other topics in the book. Five appendices review essential background information in multivariate analysis, matrix theory, statistical testing, probability theory, pseudorandom number generation, and Markov chains. All chapters and appendices include exercises. The book concludes with an extensive list of references.

Broadly speaking, the first part of the book (Chapters 1 – 12) is devoted to a description of core algorithms for stochastic search and optimization. The second part (Chapters 13–17) discusses some closely related subjects in modeling, simulation, and estimation, including some important applications of the algorithms in Chapters 1–12.

PREREQUISITES AND INTENDED AUDIENCE

Readers should have a working knowledge of basic probability and statistics, together with knowledge of multivariate calculus and basic matrix algebra. Previous optimization experience is helpful, but not required. With the exception of some of the material on probabilistic convergence (Section C.2 in Appendix C), random number generation (Appendix D), and possibly Markov processes (Appendix E), the appendices are largely a review of the prerequisite subject matter.

The book may serve as either a reference book for researchers and practitioners or as a textbook, the latter use being supported by exercises at the end of every chapter and appendix.

The relatively modest prerequisites are intended to make the book broadly accessible, including to those who will actually *use* the methods in practical settings.

PHILOSOPHY OF THE BOOK

As any author must do, I had to make many choices about overall philosophy and topics to include or exclude. Aside from the inevitable personal biases, my choices were guided by the need to keep the book from reaching biblical proportions and the desire to make the book accessible to industrial practitioners and master's level students in engineering and other areas.

This book focuses on methods that have a solid theoretical foundation and that have a track record of effectiveness in a *broad range* of practical applications. Without such restrictions, there are an almost limitless number of methods that could be discussed (in particular, there are a huge number of highly specialized and/or ad hoc methods discussed in the literature). Even with these restrictions, there are a very large number of possible algorithms. Although the methods here are among the most popular in the field, the coverage is not encyclopedic.

Rather than simply present various stochastic search and optimization algorithms as a collection of distinct techniques, the book compares and contrasts the algorithms within a broader context of stochastic methods. From this perspective, the book demonstrates that relatively simple algorithms may perform as well as—or better than—other “exotic” algorithms that may have received a disproportionate amount of attention. A correct choice of algorithm may lead to significant cost savings through the production of a better solution to the problem and through savings in the effort required for implementation.

Because of the diverse areas in which this material applies, the book is not directed toward any specific application area. Nevertheless, it does include *some* applications-oriented examples, but these are designed to be easily understandable and not to require specialized knowledge. These simple examples are drawn from areas such as control engineering, transportation, music, chemistry, economics, medicine, and artificial intelligence. Further, the book includes an extensive list of references, many of which provide detail on specific applications.

TRUTH IN ADVERTISING

There are perfect books and there are finished books. As such, allow me to mention a few items and topics the book does *not* include. First, as mentioned above, this book does not include detailed discussions of many “serious” applications, although there are countless real-world uses of the methods here. Rather, the focus is on generally useful principles for a broad audience without devoting significant attention to the inevitable application-specific details of a serious implementation. Related to this, this book bucks the trend (particularly among statistics books) of including analyses of numerous real-world data sets. The aim is to encourage the reader to grasp basic principles and not be distracted by the nuances and scale of many such data sets. In an educational context, an instructor can introduce such data as desired or have students pursue one or more of the many listed references (and associated Web sites) that analyze such data.

A related distinction between this work and many others on algorithms is that this book is not filled with a large number of Monte Carlo simulation (numerical) studies. Certainly, the book includes *some* such studies, but the overall philosophy is to emphasize approaches that have a solid theoretical foundation and to present this foundation together with a *limited* number of simulation studies that augment the theory. The rationale for such an approach is that a strong reliance on numerical studies alone can provide misleading indications about the effectiveness (or lack thereof) of an approach. All simulation studies are problem specific, with results that may not transfer to other problems. Further, it is difficult to provide a fully objective numerical comparison when the algorithms must be “tuned” to the application. The temptation exists to spend more time tuning favored algorithms than tuning other algorithms.

Although this book presents the theoretical basis of the algorithms, it does not include lengthy proofs of such results. Rather, the reader is directed to the literature for detailed derivations and proofs. (Many of the proofs in the literature are quite intricate and use mathematics beyond the level of this book.) Nevertheless, results are presented rigorously, with the aim of having the reader gain an appreciation for the rationale and limits of various algorithms.

There is no standard list of topics to be covered in a book on stochastic search and optimization. Some experienced researchers and practitioners may find their favorite algorithms excluded here. Please accept my apologies for such omissions. *Some* of the specific technical topics not covered—which could well qualify as stochastic search and optimization—are stochastic programming, dynamic programming, Markov decision processes, and many of the optimization metaheuristics (such as tabu search and ant colony optimization). Some of these are omitted in the interest of providing balanced coverage of different general approaches (gradient-based versus nongradient-based, single path versus population-based, etc.) while keeping the book between only two covers. Others are omitted because there is little evidence of greater general efficiency or ease of use when compared to some of the basic methods that *are* discussed in the book. Although some may challenge the choice of topics, it is incontrovertible that the topics included are important, are relevant to the book’s theme, and are widely used in industrial and other practice.

Finally, this book does not review, endorse, or seriously discuss available commercial software in stochastic search and optimization. (I am referring to specialized software—e.g., a commercial package implementing a genetic algorithm—rather than generic packages such as MATLAB or Microsoft EXCEL.) Many “serious” implementations of the methods here rely on such specialized software. An Internet search or perusal of the advertisements appearing in many technical magazines or related publications will provide a number of possible vendors.