Preface

This work is aimed at mathematics students in the area of stochastic dynamical systems and at engineering graduate students in signal processing and control systems. First-year graduate-level students with some background in systems theory and probability theory can tackle much of this material, at least once the techniques of Chapter 2 are mastered (with reference to the Appendices and some tutorial help). Even so, most of this work is new and would benefit more advanced graduate students. Familiarity with the language of the general theory of random processes and measure-theoretic probability will be a help to the reader. Well-known results such as the Kalman filter and Wonham filter, and also H^2, H^∞ control, emerge as special cases. The motivation is from advanced signal processing applications in engineering and science, particularly in situations where signal models are only partially known and are in noisy environments. The focus is on optimal processing, but with a counterpoint theme in suboptimal, adaptive processing to achieve a compromise between performance and computational effort.

The central theme of the book is the exploitation, in novel ways, of the so-called reference probability methods for optimal estimation and control. These methods supersede, for us at least, the more familiar innovation and martingale representation methods of earlier decades. They render the theory behind the very general and powerful estimation and control results accessible to the first-year graduate student. We claim that these reference probability methods are powerful and, perhaps, comprehensive in the context of discrete-time stochastic systems; furthermore, they turn out to be relevant for systems control. It is in the nature of mathematics that these methods were first developed for the technically more demanding area of continuous time stochastic systems, starting with the theorems of Cameron and Martin (1944), and Girsanov (1960). The reference probability approach to optimal filtering was introduced in continuous-time in Duncan (1967), Mortensen (1966) and Zakai (1969). This material tends to be viewed as inaccessible to graduate students in engineering. However, apart from contributions in Boel (1976), Brémaud and van Schuppen (1976), di Masi and Runggaldier (1982), Segall (1976b), Kumar and Varaiya (1986b) and Campillo and le Gland (1989), there has been little work on discrete-time filtering and control using the measure change approach.

An important feature of this book is the systematic introduction of new, equivalent probability measures. Under the new measure the variables of the observation process, and at times the state process, are independent, and the computations are greatly simplified, being no more difficult than processing for linear models. An inverse change of measure returns the variables to the "real world" where the state influences the observations. Our methods also apply in continuous time, giving simpler proofs of known theorems together with new results. However, we have chosen to concentrate on models whose state is a noisily observed Markov chain. We thus avoid much of the delicate mathematics associated with continuous-time diffusion processes.

The signal models discussed in this text are, for the main part, in discrete time and, in the first instance, with states and measurements in a discrete set. We proceed from discrete time to continuous time, from linear models to nonlinear ones, from completely known models to partially known models, from one-dimensional signal processing to two-dimensional processing, from white noise environments to colored noise environments, and from general formulations to specific applications.

Our emphasis is on recent results, but at times we cannot resist the temptation to provide "slicker" derivations of known theorems.

This work arose from a conversation two of the authors had at a conference twenty years ago. We talked about achieving adaptive filter stability and performance enhancement using martingale theory. We would have been incredulous then at what we have recently achieved and organized as this book. Optimal filtering and closed-loop control objectives have been attained for quite general nonlinear signal models in noisy environments. The optimal algorithms are simply stated. They are derived in a systematic manner with a minimal number of steps in the proofs.

Of course, twenty years ago we would have been absolutely amazed at the power of supercomputers and, indeed, desktop computers today, and so would not have dreamt that *optimal* processing could actually be implemented in applications except for the simplest examples. It is still true that our simply formulated optimal algorithms can be formidable to implement, but there are enough applications areas where it is possible to proceed effectively from the foundations laid here, in spite of the dreaded curse of dimensionality.

Our work starts with discrete-time signal models and with states and measurements belonging to a discrete set. We first apply the change-of-measure technique so that the observations under a probability measure are independent and uniformly distributed. We then achieve our optimization objectives, and, in a final step, translate these results back to the real world. Perhaps at first glance, the work looks too mathematical for the engineers of today, but all the results have engineering motivation, and our pedagogical style should allow an engineer to build the mathematical tools without first taking numerous mathematics courses in probability theory and stochastic systems. The advanced mathematics student may find later chapters immediately accessible and see earlier chapters as special cases. However, we believe many of the key insights are right there in the first technical chapter. For us, these first results were the key to most of what follows, but it must be admitted that only by tackling the harder, more general problems did we develop proofs which we now use to derive the first results.

Actually, it was just two years ago that we got together to work on hidden Markov model (HMM) signal processing. One of us (JBM) had just developed exciting application studies for such models in biological signal processing. It turns out that ionic channel currents in neuron cell membranes can now be observed using Nobel prize winning apparatus measuring femto (10^{-15}) amps. The noise is white and Gaussian but dominates the signals. By assuming that the signals are finite-state Markov chains, and adaptively estimating transition probability and finite state values, much information can be obtained about neural synapses and the synaptic response to various new drug formulations. We believed that the on-line biological signal processing techniques which we developed could be applied to communication systems involving fading channels, such as mobile radio communications.

The key question for us, two years ago, was how could we do all this signal processing, with uncertain models in noisy environments, *optimally*? Then, if this task was too formidable for implementation, how could we achieve a reasonable compromise between computational effort and performance? We believed that the martingale approach would be rewarding, and it was, but it was serendipitous to find just how powerful were the reference probability methods for discrete-time stochastic systems. This book has emerged somewhat as a surprise.

In our earlier HMM studies, work with Ph.D. student Vikram Krishnamurthy and postdoctoral student Dr. Lige Xia set the pace for adaptive HMM signal processing. Next, work with Ph.D. student Hailiang Yang helped translate some continuous-time domain filtering insights to discrete time. The work of some of our next generation of Ph.D. students, including Iain Collings, features quite significantly in our final manuscript. Also, discussions with Matt James, Alain Bensoussan, and John Baras have been very beneficial in the development of the book. We wish to acknowledge to seminal thinking of Martin Clarke in the area of nonlinear filtering and his influence on our work. Special thanks go to René Boel for his review of the first version of the book and to N. Krylov for supplying corrections to the first printing.

The support of the Cooperative Research Centre for Robust and Adaptive Systems, the Boeing Commercial Airplane Company, and the NSERC Grant A7964 are gratefully acknowledged. We acknowledge the typing support of Shelley Hey, and Marita Rendina, and LaTeX programming support of James Ashton.

http://www.springer.com/978-0-387-94364-0

Hidden Markov Models
Estimation and Control
Elliott, R.J.; Aggoun, L.; Moore, J.B.
1995, XIV, 382 p., Hardcover
ISBN: 978-0-387-94364-0