

Preface

Engineering is in many ways an exercise in managing uncertainty or its alternate manifestation, risk. Uncertainty arises because real problems, when looked at with enough detail, have too many variables to track and physics that are too complicated to describe succinctly. We end up making simplifications and assumptions, and, sometimes, we just ignore phenomena altogether. This leaves a gap in our models, something that must be accounted for one way or another. In this book, we will use probability and statistics.

Probabilities and statistics are ultimately numbers that describe outcomes. They provide a piece of information about a process whose cause we probably do not understand, a description in the absence of an explanation. Probability often gets us out of a hopeless situation by giving us a place to start. Moreover, with a model in hand, even a probabilistic one, techniques from linear algebra, optimization, and least squares become available to us. Finally, they give us metrics which we can use to make decisions.

Now for a few words of caution. Statistics provide a snapshot but not the whole picture. No model can. Another thing to keep in mind is that we need to make assumptions about the foundational probabilities upon which to build the rest of the model. This is essentially a leap of faith, and there is no way to avoid it. We can gather data, but this data has uncertainty built into it as well.

In this text, we will try to show how probability can be used to model uncertainty in control and estimation problems. In the process, we will learn about probability theory, stochastic processes, estimation, and stochastic control strategies. Our chief objective is to provide insight. We do not aim to be the most mathematically rigorous in our presentation, though we have a great affinity for the math. The material that you will learn here is both wonderfully practical and rich in research opportunities. It has historical connections to Newton, Gauss, Wiener, Einstein, Kalman, and many of other great names in physics, mathematics, and control theory.

Book Content

The book can be considered to be an exposition in three parts: probability theory and stochastic processes; estimation theory; and stochastic optimal control. However, these divisions are integrated due to their intimate connections, though the derivations of certain concepts are concentrated in one place. In the following, each chapter is described and how it is related to other material found in other chapters.

Probability Theory: Chapter 1

In this chapter the rudiments of probability theory are introduced. These concepts are used throughout the book and therefore are essential, although elementary, concepts. For example, Bayes' rule, which forms the basis of statistical estimation theory, is easily developed.

Random Variables and Stochastic Processes: Chapter 2

In this chapter the concept of a random variable is introduced and its probabilistic characterization is described. The notions of probability distribution function and probability density function are developed as well as their use in the calculation of expected value with respect to a random variable. Furthermore, the concept of conditional expectation and its special case of conditional probability are explained. Very important is the extension of the notion of a random variable to be indexed by another variable, such as time, to produce a stochastic sequence or stochastic process. It is here that the discrete-time linear system with additive Gaussian noise (Gauss–Markov system) is introduced that will play an important role in the structure of linear estimators.

Conditional Expectations and Discrete-Time Kalman Filtering: Chapter 3

Although the concept of conditional probability and expectation is introduced in Chapters 1 and 2, the theory and application of conditional expectation to dynamic filtering are described in this chapter. Since the development of an estimator for linear systems with additive Gaussian process and measurement noise is an important example of conditional expectation, the discrete-time conditional mean estimator, called the “Kalman filter,” is derived and illustrated.

Least Squares, the Orthogonal Projection Lemma, and Discrete-Time Kalman Filtering: Chapter 4

In this chapter classical least squares is shown to be related to the discrete-time Kalman filter through the orthogonal projection lemma, which is a necessary and sufficient condition for a quadratic function to be at a minimum. In fact, for linear systems with additive but non-Gaussian noise, the best linear filter is derived by direct application of the orthogonal projection lemma and when restricted to Gaussian noise reduces to the Kalman filter.

Stochastic Processes and Stochastic Calculus: Chapter 5

In the previous chapters the statistical characteristics of stochastic sequences are described. Although the stochastic process was defined in Chapter 2, it is in this chapter that stochastic processes are characterized by their own calculus. This calculus is needed in developing the model for estimation problems found in Chapter 6 and in the development of the dynamic programming algorithm needed for the solution of continuous-time optimal control problems

found in Chapter 9. The chapter opens with a demonstration of the convergence of a discrete-time random walk to a continuous-time Brownian motion process, illustrating the special character and difficulties of stochastic processes.

Continuous-Time, Gauss–Markov Systems: Continuous-Time Kalman Filter, Stationarity, Power Spectral Density, and the Wiener Filter: Chapter 6

This chapter builds upon the mathematical foundation laid out in the previous chapter and provides the continuous-time versions of the Gauss–Markov theory laid out in Chapters 3 and 4. We begin this chapter by deriving the continuous-time Kalman filter by application of the orthogonal projection lemma. In this chapter we also introduce stationarity, ergodicity, and the power spectral density, useful concepts in engineering applications. This chapter culminates in an introduction of a foundational result in estimation theory: the Wiener filter. Although the Wiener filter is derived by spectral methods, it is shown to be equivalent to the stationarity form of the time-domain Kalman filter.

The Extended Kalman Filter: Chapter 7


The solution of the estimation problem for nonlinear system requires the construction of the conditional probability density function. Based on the conditional probability density function, state estimates, such as the conditional mean estimates, are not implementable for real-time application. Therefore, approximate filters are presented, called the *extended Kalman filter*.

A Selection of Results from Estimation Theory: Chapter 8

Special, but important, extensions to the basic Kalman filter are developed such as measurement noise that is time correlated and the smoothing problem that estimates the state at intermediate times using data over an entire time interval. An especially important extension is that of the estimator derived in Chapter 10, which is an optimal linear filter with respect to a particular cost criterion but is not a conditional mean estimator. Within this context certain classical theorems are interpreted.

Stochastic Control and the Linear Quadratic Gaussian Control Problem: Chapter 9

In this chapter the stochastic control problem is formulated for both the discrete-time and continuous-time problems with full information of the state and with noisy partial measurement information structure. The solution is obtained using a dynamic programming methodology where in the continuous-time derivation the stochastic calculus of Chapter 5 is critical. The dynamic programming algorithms are demonstrated on problems formulated with a linear system with additive Gaussian noise and the expected value of a quadratic function of the state and control, as the cost criterion, the so-called *linear quadratic Gaussian*



control problem. The robustness properties interpreted by the classical control criterion are given.

Linear Exponential Gaussian Control and Estimation: Chapter 10

In this chapter an important extension of the linear quadratic Gaussian control problem, called the *linear exponential Gaussian control*, is described. Here again the system assumes a linear system with additive Gaussian noise, but the cost criterion is the expectation of the exponential of a quadratic function of the state and control. The solution is obtained for both discrete-time and continuous-time formulations. Also, the linear exponential Gaussian estimator is derived, but its characteristics and properties are presented in Chapter 8. It is shown that the resulting controller is equivalent to that obtained from the H_∞ control syntheses of deterministic robust control theory.

