

PART I

PRELIMINARIES

COPYRIGHTED MATERIAL

1

INTRODUCTION

Estimation and tracking of dynamic systems has been the research focus of many a mathematician since the dawn of statistical mathematics. Many estimation methods have been developed over the past 50 years that allow statistical inference (estimation) for dynamic systems that are linear and Gaussian. In addition, at the cost of increased computational complexity, several methods have shown success in estimation when applied to nonlinear Gaussian systems. However, real-world dynamic systems, both linear and nonlinear, usually exhibit behavior that results in an excess of outliers, indicative of non-Gaussian behavior. The toolbox of standard Gaussian estimation methods have proven inadequate for these problems resulting in divergence of the estimation filters when applied to such real-world data.

With the advent of high-speed desktop computing, over the past decade the emphasis in mathematics has shifted to the study of dynamic systems that are non-Gaussian in nature. Much of the literature related to performing inference for non-Gaussian systems is highly mathematical in nature and is lacking in practical methodology that the average engineer can utilize without a lot of effort. In addition, several of the Gaussian methods related to estimation for nonlinear systems are presented ad hoc, without a cohesive derivation. Finally, there is a lack of continuity in the conceptual development to date between the Gaussian methods and their non-Gaussian counterparts.

In this book, we will endeavor to present a comprehensive study of the methods currently in use for statistical dynamic system estimation: linear and nonlinear, Gaussian

and non-Gaussian. Using a Bayesian framework, we will present a conceptually cohesive roadmap that starts at first principles and leads directly to derivations of many of the Gaussian estimation methods currently in use. We will then extend these concepts into the non-Gaussian estimation realm, where the theory leads directly to working Monte Carlo methods for estimation. Although the Bayesian approach leads to the estimation of statistical densities, in most cases we will develop point estimation methods that can be obtained through the evaluation of density-weighted integrals. Thus, this book is all about numerical methods for evaluating density-weighted integrals for both Gaussian and non-Gaussian densities.

For each estimation method that we discuss and derive, we present both pseudo-code and graphic block diagram that can be used as tools in developing a software-coded tracking toolbox. As an aid in understanding the methods presented, we also discuss what is required to develop simulations for several very specific real-world problems. These case-study problems will be addressed in great detail, with track estimation results presented for each. Since it is hard to compare tracking methods ad hoc, we also present multiple methods to evaluate the relative performance of the various tracking filters.

1.1 BAYESIAN INFERENCE

Inference methods consist of estimating the current values for a set of parameters based on a set of observations or measurements. The estimation procedure can follow one of two models. The first model assumes that the parameters to be estimated, usually unobservable, are nonrandom and constant during the observation window but the observations are noisy and thus have random components. The second model assumes that the parameters are random variables that have a prior probability and the observations are noisy as well. When the first model is used for parameter estimation, the procedure is called non-Baysian or Fisher estimation [1]. Parameter estimation using the second model is called Bayesian estimation.

Bayesian estimation is conceptually very simple. It begins with some initial prior belief, such as the statement “See that ship. It is about 1000 yards from shore and is moving approximately Northeast at about 10 knots.” Notice that the initial belief statement includes an indication that our initial guess of the position and velocity of the ship are uncertain or random and based on some prior probability distribution. Based on one’s initial belief, one can then make the prediction “Since the ship appears to be moving at a constant velocity, it will be over there in about 10 minutes.” This statement includes a mental model of the ship motion dynamics as well as some additional uncertainty. Suppose now, that one has a small portable radar on hand. The radar can be used to measure (observe) the line-of-sight range and range rate of the ship to within some measure of uncertainty. Given the right mathematical model, one that links the observations to the Cartesian coordinates of the ships position and velocity, a current radar measurement can be used to update the predicted ships state (position and velocity).

The above paragraph contains the essence of recursive Bayesian estimation:

1. begin with some prior belief statement,
2. use the prior belief and a dynamic model to make a prediction,
3. update the prediction using a set of observations and an observation model to obtain a posterior belief, and
4. declare the posterior belief our new prior belief and return to 2.

This concept was first formalized in a paper by the Reverend Thomas Bayes, read to the Royal Statistical Society in 1763 by Richard Price several years after Bayes’ death. An excellent review of the history and concepts associated with Bayesian statistical inference can be found in the paper by Stephen Brooks [2]. Brooks’ paper also has some interesting examples that contrast the Bayesian method with the so-called “Frequentist” method for statistical inference. Since this book is devoted completely to Bayesian methods, we will not address the frequentist approach further and refer the interested reader to Brooks’ paper.

1.2 BAYESIAN HIERARCHY OF ESTIMATION METHODS

As noted above, in this book we will present a cohesive derivation of a subset of modern tracking filters. Figure 1.1 shows the hierarchy of tracking filters that will

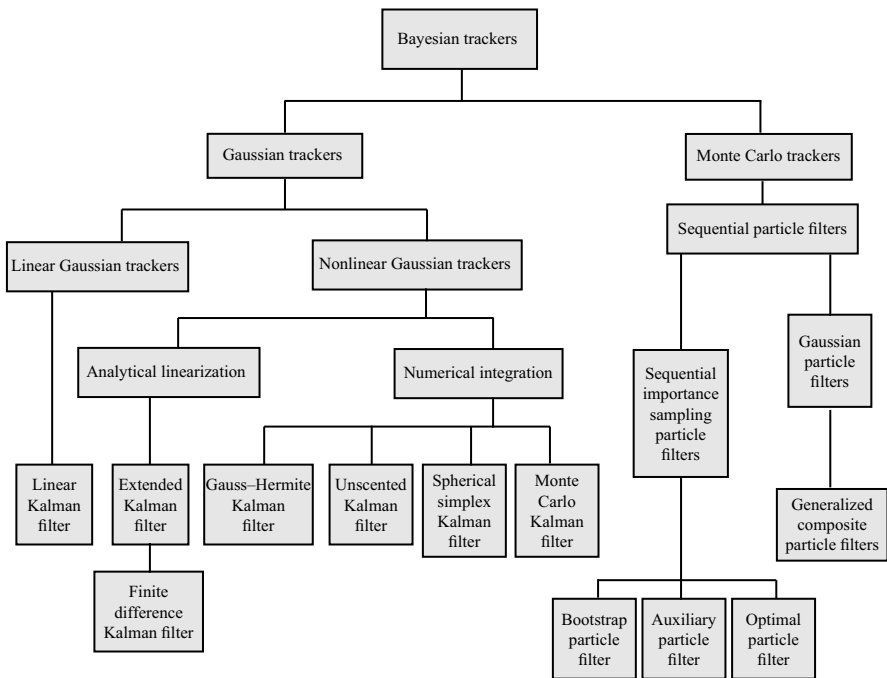


FIGURE 1.1 Hierarchy of Bayesian estimation tracker filters.

be addressed in this book. Along the left-hand side are all the Gaussian tracking filters and along the right-hand side are all of the Monte Carlo non-Gaussian filters. This figure will be our guide as we progress through our discussions on each tracking filter. We will use it to locate where we are in our developments. We may occasionally take a side trip into other interesting concepts, such as a discussion of performance measures, but for the most part we will stick to a systematic development from top to bottom and left to right. By the time we reach the bottom right, you the reader will have a comprehensive understanding of the interrelatedness of all of the Bayesian tracking filters.

1.3 SCOPE OF THIS TEXT

1.3.1 Objective

The objective of this book is to give the reader a firm understanding of Bayesian estimation methods and their interrelatedness. Starting with the first principles of Bayesian theory, we show how each tracking filter is derived from a slight modification to a previous filter. Such a development gives the reader a broader understanding of the hierarchy of Bayesian estimation and tracking. Following the discussions about each tracking filter, the filter is put into both pseudo-code and process flow block diagram form for ease in future recall and reference.

In his seminal book on filtering theory [3], originally published in 1970, Jazwinski stated that “The need for this book is twofold. First, although linear estimation theory is relatively well known, it is largely scattered in the journal literature and has not been collected in a single source. Second, available literature on the continuous nonlinear theory is quite esoteric and controversial, and thus inaccessible to engineers uninitiated in measure theory and stochastic differential equations.” A similar statement can be made about the current state of affairs in non-Gaussian Monte Carlo methods of estimation theory. Most of the published work is esoteric and inaccessible to engineers uninitiated in measure theory. The edited book of invited papers by Doucet et al. [4] is a prime example. This is an *excellent* book of invited papers, but is extremely esoteric in many of its stand-alone sections.

In this book, we will take Jazwinski’s approach and remove much of the esoteric measure theoretic-based mathematics that makes understanding difficult for the average engineer. Hopefully, we have not replaced it with equally esoteric alternative mathematics.

1.3.2 Chapter Overview and Prerequisites

This book is not an elementary book and is intended as a one semester graduate course or as a reference for anyone requiring or desiring a deeper understanding of estimation and tracking methods. Readers of this book should have a graduate level understanding of probability theory similar to that of the book by Papoulis [5]. The reader should also be familiar with matrix linear algebra and numerical methods

including finite differences. In an attempt to reduce the steep prerequisite requirements for the reader, we have included several review sections in the next chapter on some of these mathematical topics. Even though some readers may want to skip these sections, the material presented is integral to an understanding of what is developed in Parts II and III of this book.

Part I consists of this introduction followed by a chapter that presents an overview of some mathematical principles required for an understanding of the estimation methods that follow. The third chapter introduces the concepts of recursive Bayesian estimation for a dynamic system that can be modeled as a potentially unobservable discrete Markov process. The observations (measurements) are related to the system states through an observation model and the observations are considered to be discrete. Continuous estimation methods are generally not considered in this book. The last chapter of Part I is devoted to preliminary development of a case study that will be used as working examples throughout the book, the problem of tracking a ship through a distributed field of directional frequency analysis and recording (DIFAR) sonobuoys. Included for this case study will be demonstrations of methods for development of complete simulations of the system dynamics along with the generation of noisy observations.

Part II is devoted to the development and application of estimation methods for the Gaussian noise case. In Chapter 5, the general Bayesian estimation methods developed in Chapter 3 are rewritten in terms of Gaussian probability densities. Methods for specific Gaussian Kalman filters are derived and codified in Chapters 6 through 12, including the linear Kalman filter (LKF), extended Kalman filter (EKF), finite difference Kalman filter (FDKF), unscented Kalman filter (UKF), spherical simplex Kalman filter (SSKF), Gauss–Hermite Kalman filter (GHKF), and the Monte Carlo Kalman filter (MCKF). With the exception of the MCKF, four of latter five tracking filters can be lumped into the general category of sigma point Kalman filters where deterministic vector integration points are used in the evaluation of the Gaussian-weighted integrals needed to estimate the mean and covariance matrix of the state vector. In the MCKF, the continuous Gaussian distribution is replaced by a sampled distribution reducing the estimation integrals to sums leaving the nonlinear functions intact. It will be shown in Chapter 13 that the latter five Kalman filter methods can be summarized into a single estimation methodology requiring just a change in the number and location of the vector points used and their associated weights.

An important aspect of estimation, usually ignored in most books on estimation, is the quantification of performance measures associated with the estimation methods. In Chapter 14 this topic is addressed, with sections on methods for computing and plotting error ellipses based on the estimated covariance matrices for use in real-system environments, as well as methods for computing and plotting root mean squared (RMS) errors and their Cramer–Rao lower bounds (CRLB) for use in Monte Carlo simulation environments. The final section of this chapter is devoted to application of these estimation methods to the DIFAR buoy tracking case study and includes a comparison of performance results as a function of decreasing input signal-to-noise ratio (SNR).

Estimation methods for use primarily with non-Gaussian probability densities is the topic addressed in Part III. For the MCKF introduced in Chapter 12 of Part II, the Gaussian density is approximated by a set of discrete Monte Carlo samples, reducing the mean and covariance estimation integrals to weighted sums, usually referred to as sample mean and sample covariance, respectively. For Gaussian densities, the sample weight is always $1/N$, where N is the number of samples used. Non-Gaussian densities present two problems: first, it is usually very difficult to generate a set of Monte Carlo samples directly from the density. A second problem arises if the first or second moment does not exist for the density, with the Cauchy density as a prime example. To address the sampling problem, in Chapter 15 Monte Carlo methods are introduced and the concept of importance sampling developed that leads to estimation methods called particle filters, where the particles are the Monte Carlo sample points. Several problems arise when implementing these particle filters and potential enhancements are considered that correct for these problems. For importance sampling, weighting for each sample is calculated as the ratio of the non-Gaussian density to the importance density at the sample point. Under certain assumptions, the weights can be calculated recursively, giving rise to the sequential importance sampling (SIS) class of particle filters, the topic of Chapter 16. In Chapter 17, the case where the weights are recalculated every filter iteration step is addressed, leading to the Gaussian class of combination particle filters. Performance results for all of the particle filter track estimation methods applied to the DIFAR case study are presented as the conclusion of Chapter 17.

Several recently published books provide additional insight into the topics presented in this book. For Gaussian Kalman filters of Part II, books by Bar Shalom et al. [6] and Candy [7] are good companion books. For non-Gaussian filtering methods of Part III, books by Doucet et al. [4] and Ristic et al. [8] are excellent reference books.

1.4 MODELING AND SIMULATION WITH MATLAB[®]

It is important to the learning process that the reader be given concrete examples of application of estimation methods to a set of complex problems. This will be accomplished in this book through the use of simulations using MATLAB[®]. We present a set of four case studies that provide an increase in complexity from the first to the last. Each case study will include an outline of how to set up a simulation that models both the dynamics and observations of the system under study. We then show how to create a set of randomly generated observational data using a Monte Carlo methodology. This simulated observational data can then be used to exercise each tracking filter, producing sets of track data that can be compared across multiple track filters.

The first case study examines the problem of tracking a ship as it moves through a distributed field of DIFAR buoys. A DIFAR buoy uses the broadband noise signal radiated from the ship as an input and produces noisy observations of the bearing to the ship as an output. As we will show in Chapter 4, the probability density of the bearing estimates at the DIFAR buoy output is dependent on the SNR of the input

signal. The density will be Gaussian for high SNR but will transition to a uniform distribution as the SNR falls. The purpose of this case study will be to examine what happens to the filter tracking performance for each track estimation method as the observation noise transitions from Gaussian to non-Gaussian.

This DIFAR case study will be the primary tool used throughout this book to illustrate each track estimation filter in turn. In Chapter 4, we show how to set up a simulation of the DIFAR buoy processing so as to produce simulated SNR dependent observation sets. Using these sets of bearing observations, in subsequent chapters we exercise each tracking algorithm to produce Monte Carlo sets of track estimates, allowing us to see the impact of the Gaussian to non-Gaussian observation noise transition on each tracking method.

In Part IV of this book, we present three additional case studies that illustrate the use of many of the tracking filters developed in Parts II and III. In Chapter 18, we address the important problem of tracking a maneuvering object in three dimension space. In this chapter, we introduce a new approach that uses a constant spherical velocity model vice the more traditional constant Cartesian velocity model. We show how this spherical model shows improved performance for tracking a maneuvering object using most of the Gaussian tracking filters.

The third case study, found in Chapter 19, considers the rather complex problem of tracking the dynamics of a falling bomb through the use of video frames of multiple tracking points on both the plane dropping the bomb and the bomb itself. This is a particular example of a complex process called photogrammetry, in which the geometric and dynamic properties of an object are inferred from successive photographic image frames. Thus, this case study consists of a very complex nonlinear multidimensional observational process as well as a nonlinear multidimensional dynamic model. In addition, both the dynamic and observational models are of high dimension, a particularly taxing problem for tracking filters. This will illustrate the effects of the so-called “curse” of dimensionality, showing that it is computationally impractical to utilize all tracking filters.

The final case study, the topic of Chapter 20, improves on the use of photogrammetric methods in estimation by showing how a separate estimator can be used for fusing data from additional sensors, such as multiple cameras, translational accelerometers, and angular rate gyroscopes. When used independently, each data source has its unique strengths and weaknesses. When several different sensors are used jointly in an estimator, the resulting solution is usually more accurate and reliable. The resulting analysis shows that estimator aided sensor fusion can recover meaningful results from flight tests that would otherwise have been considered failures.

REFERENCES

1. Fisher R. *Statistical Methods and Scientific Inference*, Revised Edition. Macmillan Pub. Co.; 1973.
2. Brooks SP. Bayesian computation: a statistical revolution. *Phil. Trans. R. Soc. Lond. A* 2003;361:2681–2697.

3. Jazwinski AH. *Stochastic Processes and Filtering Theory*. Academic Press (1970), recently republished in paperback by Dover Publications; 2007.
4. Doucet A, de Freitas JFG, Gordon NJ, editors. *Sequential Monte Carlo Methods in Practice*. New York, NY: Springer-Verlag; 2001.
5. Papoulis A. *Probability, Random Variables, and Stochastic Processes*, 4th ed. McGraw-Hill; 2002.
6. Bar Shalom Y, Li XR, Kirubarajan T. *Estimation with Application to Tracking and Navigation: Theory, Algorithms and Software*. Wiley; 2001.
7. Candy JV. *Bayesian Signal Processing: Classical, Modern, and Particle Filtering Methods*. Hoboken, NJ: Wiley; 2009.
8. Ristic B, Arulampalam S, Gordon N. *Beyond the Kalman Filter: Particle Filters for Tracking Applications*. Boston, MA: Artech House; 2004.