Beyond the Kalman Filter

Branko Ristic, Sanjeev Arulampalam & Neil Gordon
Artech House, Norwood, MA, USA
2004, 299 pages, Hardcover
ISBN 1-58053-631-x

This new book, cleverly entitled ***Beyond the Kalman Filter,*** is by far the best book on filtering published in many years: It is clearly written, well organized, authoritative, and contains a wealth of useful information. It explains a state-of-the-art class of practical nonlinear filtering algorithms, called *"particle filters,"* using examples and theory accessible to normal engineers. The book contains many examples of practical nonlinear filters worked out in detail. One of the best features of the book is the large variety of plots of filter estimation accuracy for different types of filters, including the extended Kalman filter (EKF), the unscented Kalman filter (UKF), particle filters (PFs), as well as a bound on the theoretical optimal performance. It is refreshing to see such thorough quantitative work. The main focus of this book is particle filtering, which is a new class of nonlinear filter which often results in vastly superior performance compared with the Kalman filter. The Kalman filter is the most widely used and most successful algorithm for real world aerospace and other applications. It is used in many radars, passive infrared systems, sonar, and other systems, as well as for multi-sensor data fusion. Thousands of papers have been written about Kalman filters. Often the Kalman filter gives excellent estimation accuracy – but not always. There are many real-world applications in which Kalman filter accuracy is poor relative to theoretical expectations, owing to nonlinearity in the measurement equations or the dynamical model of the physical system. Almost all real-world problems are nonlinear. Only textbooks and papers create the impression that filtering problems are linear.

Particle filters use real-time on-line Monte Carlo simulations to model the decrease of uncertainty due to sensor measurements and the increase of uncertainty as time evolves in the physical system of interest. This is a very direct approach to modelling uncertainty, facilitated by the availability of fast, low-cost modern computers with large memories. Computer speed and memory have increased by roughly eight orders of magnitude per unit cost since Kalman published his famous paper in 1960. This is what allows us to run particle filters in real-time on a single PC for significant real-world applications. Each particle represents one possible random path of the state vector of the system. The probability density of the state vector conditioned on the measurements is represented by these particles, similar to a histogram. The expected value of the state vector is computed by adding up all the particles, with appropriate weights (some particles are more important than others). One can code a pretty good particle filter in a page or two of MATLAB. One does not need to understand the finer points of stochastic calculus or the Fokker-Planck equation, or any fancy numerical methods for solving partial differential equations. On the other hand, the best particle filters use sophisticated (but easy to code) sampling methods (e.g., Hastings-Metropolis), rather than brute force simple Monte Carlo sampling. A crucial detail in particle filters is called "importance sampling," which means putting the particles in the most important regions in state space. This is accomplished by using some simple Bayesian probability theory.

The key issue in nonlinear filters of any kind is the curse of dimensionality. This is the phrase coined by Richard Bellman over forty years ago to describe the exponential growth of computational complexity as a function of dimension of the state vector of the system. The computational complexity of the Kalman filter grows as the cube of the dimension, but for general nonlinear problems using filters that achieve optimal accuracy, the computational complexity grows more rapidly with dimension. It has been asserted (but not in this book) that particle filters avoid the curse of dimensionality. But, contrary to popular opinion, particle filters generally do not avoid the curse of dimensionality. The authors should be commended for not making such an assertion. Moreover, they avoid hype and buzzwords, and use good solid quantitative analysis and credible simulations to compare the performance of particle filters with other competing algorithms.

Another outstanding feature of this book is the extensive use of a new theoretical bound on optimal estimation accuracy. Many plots show quantitative comparisons between Kalman filters, unscented Kalman filters, particle filters, and theoretical bounds on accuracy. For parameter estimation problems, engineers have used the Cramer-Rao bound (CRB) for many decades. However, the CRB only applies to filtering problems with zero process noise. This book uses a theoretical bound that is a generalization of the CRB, which was published a few years ago. This generalized CRB gives surprisingly tight bounds on optimal estimation accuracy for many practical problems with non-zero process noise.

A word of caution about comparisons with the extended Kalman filter:

> *Hard-boiled engineers know that there is no such thing as "the" EKF, just as there is no such thing as "the" UKF or "the" PF. But rather, there are hundreds of different types of EKF, UKFs, and PFs, which use various engineering tricks (e.g., second order terms, iterations), tuning of process noise, different coordinate systems, various factorizations of the covariance matrix, etc.*

Therefore, when someone says that "the" PF is much better than "the" EKF, the first question to ask is: *Exactly which of the hundreds of possible EKFs are we talking about?* Similarly, there are hundreds of different PFs using bells and whistles invented to reduce computational complexity. The second question to ask about such comparisons is: *Are the Monte Carlo simulation results correct?* You should never trust a Monte Carlo simulation without some method to verify its correctness. I am appalled by engineers who blindly believe the results of Monte Carlo simulations. Years ago, it was standard to demand a simple back-of-the-envelope calculation that explains a Monte Carlo simulation result, but, unfortunately, this good practice has diminished with time. There are many examples of engineers who have come a cropper as a result of blindly believing Monte Carlo simulations.

One of my favorite stories happened only a few years ago on a big expensive high-tech missile system, which shall remain nameless. The Monte Carlo simulation predicted that the probability of correct target association between Sensor A and Sensor B was not good enough. A long expensive study was undertaken to fix this problem. The focus of the study was a detailed cost/performance/schedule/risk system trade-off of five or ten methods to improve the accuracy of Sensor A. Nobody questioned the correctness of the Monte Carlo simulation results. During the study, I compared the Monte Carlo simulation results with a simple back-of-the-envelope formula, and it was obvious that the Monte Carlo simulation was performing much worse than it should compared with theory. Theory predicted better than 99% probability of correct target association, whereas the simulation produced roughly 50%. Not a minor discrepancy! Needless to say, this announcement was not greeted with joy by the authors of the Monte Carlo simulation. In fact, they were in denial. After many months, the simulation code was scrutinized by an engineer who had not coded the simulation and who was not in the same company or state of the Union or state of mind. Sure enough, there was a bug in the Monte Carlo simulation! After fixing the bug, the Monte Carlo simulation predicted good target association probabilities, in excellent agreement with theory. This saved millions of dollars in needless system improvements. Much more can be said about this subject, but suffice it to say that this book on particle filters builds confidence in the correctness of its results by the comparison of multiple filters, as well as the theoretical bound on estimation accuracy.

This book can be read by any engineer who understands Kalman filtering. The prerequisites are modest: elementary probability theory, calculus and linear algebra. ***Beyond the Kalman Filter*** is also ideal for graduate or advanced undergraduate courses. Such a course could be a mixture of new interesting theory and fun numerical experiments.

– Reviewed by Fred Daum
317 Astor Street
Carlisle, MA 01741