Book Review

Bayesian Filtering and Smoothing

Simo Särkkä

Cambridge University Press, Cambridge, UK
ISBN: 978-1-107-61928-9, Paperback, 2013, 252 pages

Reviewed by Lennart Svensson

Bayesian Filtering and Smoothing by Professor Simo Särkkä is an excellent contribution to the filtering and smoothing literature. It is very well written, accurate, and it covers essentially all the key techniques in the field in an elegant and pedagogical manner. Apart from being an outstanding book, another advantage is that a soft copy (a PDF version) is freely available online from the homepage of the author.

The filtering and smoothing area is important, thanks to its wide applicability ranging from positioning, navigation, and control to brain imaging and audio signal processing. The field was long dominated by the Kalman filter and the extended Kalman filter, but has received many important contributions over the past 20 years, most of which are nicely explained in this book. Different types of sigma point techniques, such as unscented, Gauss-Hermite, and cubature filters, take up a central position. The book also covers particle filtering solutions as well as an introduction to Bayesian inference in general. However, considering that the book is relatively short it is only natural that it cannot cover everything and readers who are primarily interested in multiple model filtering methods or solutions to the data association problems for target tracking should look for a different book.

An excellent feature of this book is how clearly it explains the relations between the many filtering and smoothing solutions, by first outlining a general solution and later indicating how well-known filtering and smoothing algorithms try to approximate that. Doing so helps the reader to view many of the existing algorithms as instantiations of a single general solution, which is an essential perspective in order to understand and get a manageable overview of the field.

The chapters that deal with sigma point techniques for filtering and smoothing are my favorite parts of this book, where the author clarifies both the motivations and the technical details behind existing algorithms. Presenting a clear summary of the sigma point techniques may sound like a small achievement, but considering that some early papers gave rise to significant misunderstanding of the theory, see, e.g., [1], [2], it is refreshing to read a text that does not leave any room for misconceptions. The chapters are further enhanced by the fact that the same arguments and exposition are used first in the development of Gaussian filtering methods and later reused to obtain the corresponding smoothing algorithms. A technical remark is that the author emphasizes the moment matching perspective on Gaussian filtering, whereas the statistical linear regression motivation, see, e.g., [3], [4], is not mentioned.

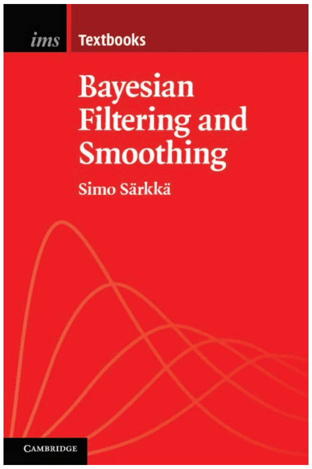

Another important tool for nonlinear filtering is the family of particle filters, also known as sequential Monte Carlo or sequential importance resampling methods [5]. The book dedicates one chapter to particle filtering and another chapter to particle smoothing, in which the background theory and motivations are nicely presented along with many of the key algorithms. These chapters are sufficiently detailed in order to enable the reader to use particle filters and smoothers to solve a wide range of problems, but they are less complete than the corresponding chapters on Gaussian methods and do not attempt to cover all existing algorithms. For instance, even though the book covers algorithms such as the Rao-Blackwellised particle filters, the reader is instead referred to the literature in order to understand the details regarding the auxiliary particle filter.

Perhaps one of the most obvious proofs that I really like this book is that I have selected it to be the main literature in a Masters course named "Sensor fusion and nonlinear filtering" that I teach. The course is still fairly new and I have not yet received much feedback on the literature, though many students were grateful that they could download a soft copy for free.

Reading the book requires basic knowledge in statistics, calculus, and linear algebra and I would recommend the book to Master students, PhD students and to most people working in the field. Anyone who has been using extended Kalman filters in the past and is curious to investigate more modern alternatives should find this book a rich source of inspiration, in particular thanks to its brilliant overview on Gaussian filters. Some may find the text a bit too brief, but succinct descriptions certainly also have their advantages and I personally find its style very appealing; especially since it is still complete in all its arguments and details.

When reviewing a book on such a mature field as nonlinear filtering it is inevitable to compare it with existing literature. Some of my favorite books on filtering are [6], [7], and [8] that all have their pros and cons. All of these three books are actually in some ways more suitable for a practitioner than the *Bayesian Filtering and Smoothing* book; [7] contains many chapters that are dedicated to various applications, whereas both [6] and [8] provide detailed descriptions regarding, e.g., how to tune filters and check their consistency. Perhaps my main reservation regarding Professor Särkkä's new book is that it lacks a chapter on motion and measurement models, which are essential components in most filters and nicely covered in both [6] and [8]. On the other hand, the topics that are covered in the *Bayesian Filtering and Smoothing* book are covered very nicely and I regard the chapters that deal with Gaussian filtering and smoothing as the most complete and accessible summary of those techniques that I have seen.

It is clear to me that most people who work with nonlinear filtering would benefit from reading this book and should have it in their bookshelf (or on a hard drive) for future references. I will personally continue using it in my courses on this topic, simply because I think it contains a brilliantly elegant and illuminating description of the field.

REFERENCES

1. Julier, S., and Uhlmann, J. Unscented filtering and nonlinear estimation. *Proceedings of the IEEE*, Vol. 92, 3 (2004), 401–422.
2. Gustafsson, F., and Hendeby, G. Some relations between extended and unscented Kalman filters. *IEEE Transactions on Signal Processing*, Vol. 60, 2 (2012), 545–555.
3. Arasaratnam, I., and Haykin, S. Cubature Kalman filters. *IEEE Transactions on Automatic Control*, Vol. 54, 6 (June 2009), 1254–1269.
4. García-Fernández, A. F., Svensson, L., and Morelande, M. Iterated statistical linear regression for Bayesian updates. In *Proceedings of the 17th International Conference on Information Fusion (FUSION)*, Salamanca, Spain, July 2014.
5. Doucet, A., Godsill, S. J., and Andrieu, C. On sequential Monte Carlo sampling methods for Bayesian filtering. *Statistics and Computing*, Vol. 10, 3 (2000), 197–208.
6. Bar-Shalom, Y., Li, X. R., and Kirubarajan, T. *Estimation with Applications to Tracking and Navigation: Theory, Algorithms, and Software*. New York: Wiley, 2001.
7. Ristic, B., Arulampalam, S., and Gordon, N. *Beyond the Kalman Filter: Particle Filters for Tracking Applications*. Norwood, MA: Artech House, 2004.
8. Gustafsson, F. *Statistical Sensor Fusion*. Lund, Sweden: Studentlitteratur, 2010.