State Estimation for Robotics

A key aspect of robotics today is estimating the state (e.g., position and orientation) of a robot, based on noisy sensor data. This book targets students and practitioners of robotics by presenting not only classical state estimation methods (e.g., the Kalman filter) but also important modern topics such as batch estimation, Bayes filter, sigma-point and particle filters, robust estimation for outlier rejection, and continuous-time trajectory estimation and its connection to Gaussian-process regression. Since most robots operate in a three-dimensional world, common sensor models (e.g., camera, laser rangefinder) are provided, followed by practical advice on how to carry out state estimation for rotational state variables. The book covers robotic applications such as point-cloud alignment, pose-graph relaxation, bundle adjustment, and simultaneous localization and mapping.

Highlights of this expanded second edition include a new chapter on variational inference, a new section on inertial navigation, more introductory material on probability, and a primer on matrix calculus.

TIMOTHY D. BARFOOT is Professor at the University of Toronto Institute for Aerospace Studies. He has been conducting research in the area of navigation of mobile robotics for over 20 years, in both industry and academia, for applications including space exploration, mining, military and transportation. He is a Fellow of the IEEE Robotics and Automation Society.

State Estimation for Robotics

Second Edition

TIMOTHY D. BARFOOT
University of Toronto

CAMBRIDGE
UNIVERSITY PRESS

CAMBRIDGE
UNIVERSITY PRESS

Shaftesbury Road, Cambridge CB2 8EA, United Kingdom

One Liberty Plaza, 20th Floor, New York, NY 10006, USA

477 Williamstown Road, Port Melbourne, VIC 3207, Australia

314–321, 3rd Floor, Plot 3, Splendor Forum, Jasola District Centre, New Delhi – 110025, India

103 Penang Road, #05-06/07, Visioncrest Commercial, Singapore 238467

Cambridge University Press is part of Cambridge University Press & Assessment, a department of the University of Cambridge.

We share the University's mission to contribute to society through the pursuit of education, learning and research at the highest international levels of excellence.

www.cambridge.org
Information on this title: www.cambridge.org/9781009299893
DOI: 10.1017/9781009299909

First published 2017

Second Edition 2024

A catalogue record for this publication is available from the British Library

A Cataloging-in-Publication data record for this book is available from the Library of Congress

ISBN 978-1-009-29989-3 Hardback

Cambridge University Press & Assessment
978-1-009-29989-3 — State Estimation for Robotics
Timothy D. Barfoot
Frontmatter
More Information

Contents

Preface for the First Edition

My interest in state estimation stems from the field of mobile robotics, particularly for space exploration. Within mobile robotics, there has been an explosion of research referred to as *probabilistic robotics*. With computing resources becoming very inexpensive, and the advent of rich new sensing technologies, such as digital cameras and laser rangefinders, robotics has been at the forefront of developing exciting new ideas in the area of state estimation.

In particular, this field was probably the first to find practical applications of the so-called Bayes filter, a much more general technique than the famous Kalman filter. In just the last few years, mobile robotics has even started going beyond the Bayes filter to batch, nonlinear optimization-based techniques, with very promising results. Because my primary area of interest is navigation of robots in outdoor environments, I have often been faced with vehicles operating in three dimensions. Accordingly, I have attempted to provide a detailed look at how to approach state estimation in three dimensions. In particular, I show how to treat rotations and poses in a simple and practical way using matrix Lie groups.

Introductio Geographica by Petrus Apianus (1495–1552), a German mathematician, astronomer, and cartographer. Much of three-dimensional state estimation has to do with *triangulation* and/or *trilateration*; we measure some angles and lengths and infer the others through trigonometry.

The reader should have a background in undergraduate linear algebra and calculus, but otherwise, this book is fairly stand-alone (Appendix A is new in the second edition and serves as a primer/reminder on matrix algebra and calculus). I hope readers of these pages will find something useful; I know I learned a great deal while creating them.

I have provided some historical notes in the margins throughout the book, mostly in the form of biographical sketches of some of the researchers after whom various concepts and techniques are named; I primarily used Wikipedia as the source for this information. Also, the first part of Chapter 7 (up to the alternate rotation parameterizations), which introduces three-dimensional geometry, is based heavily on notes originally produced by Chris Damaren at the University of Toronto Institute for Aerospace Studies.

This book would not have been possible without the collaborations of many fantastic graduate students along the way. Paul Furgale's PhD thesis extended my understanding of matrix Lie groups significantly by introducing me to their use for describing poses; this led us on an interesting journey into the details of transformation matrices and how to use them effectively in estimation problems. Paul's later work led me to become interested in continuous-time estimation. Chi Hay Tong's PhD thesis introduced me to the use of Gaussian processes in estimation theory, and he helped immensely in working out the details of the continuous-time methods presented herein; my knowledge in this area was further improved through collaborations with Simo Särkkä from Aalto University while on sabbatical at the University of Oxford. Additionally, I learned a great deal by working with Sean Anderson, Patrick Carle, Hang Dong, Andrew Lambert, Keith Leung, Colin McManus, and Braden Stenning; each of their projects added to my understanding of state estimation. Colin, in particular, encouraged me several times to turn my notes from my graduate course on state estimation into this book.

I am indebted to Gabriele D'Eleuterio, who set me on the path of studying rotations and reference frames in the context of dynamics; many of the tools he showed me transferred effortlessly to state estimation. He also taught me the importance of clean, unambiguous notation.

Finally, thanks to all those who read and pointed out errors in the drafts of this book, particularly Marc Gallant and Shu-Hua Tsao, who found many typos, and James Forbes, who volunteered to read and provide comments.

Preface for the Second Edition

It has been just over seven years since the first edition of this book was released. I have been delighted with the reception, with many colleagues and students providing useful feedback, comments, and errata over the years. Since publication of the first edition, I have kept a working copy on my personal webpage and attempted to correct any minor problems as they came in. Thank you very much to all those who took the time to give me feedback; please keep it coming for this second edition.

I am also excited that the first edition has been translated into simplified Chinese and has become extremely popular with Chinese readers; 感谢读者对于本书的支持 (thank you for reading the book!). Thanks very much to 高翔 (Gao, Xiang) and 谢晓佳 (Xie, Xiaojia) for their hard work on producing the translation.

The second edition brings about 160 pages of new material. Highlights of the new additions are as follows (chapter numbers refer to the new edition):

- **Chapter 2, Primer on Probability Theory:** expanded to cover several new topics including cumulative distributions, quantifying the difference between probability density functions, and randomly sampling from a probability density function
- **Chapter 3, Linear-Gaussian Estimation:** expanded to include computing the posterior covariance in the Cholesky and Rauch–Tung–Striebel smoothers, and a short section on recursive continuous-time smoothing and filtering
- **Chapter 4, Nonlinear Non-Gaussian Estimation:** added a new section giving some details for sliding-window filters
- **Chapter 5, Handling Nonidealities in Estimation:** expanded the scope and renamed this chapter to include information on what properties a good estimator should have, and a new section on adaptive covariance estimation
- **Chapter 6, Variational Inference:** a new chapter that frames estimation as finding a Gaussian approximation that is closest to the full Bayesian posterior in terms of the Kullback–Leibler divergence; also enables parameter learning from a common data-likelihood objective
- **Chapter 8, Matrix Lie Groups:** expanded to include sections on Riemannian optimization, computing the statistics of compounded and differenced poses with correlations, a discussion of symmetry, invariance, and equivariance
- **Chapter 9, Pose Estimation Problems:** added a large new section on inertial navigation from a matrix Lie group perspective including IMU pre-integration for batch estimation

- **Chapter 11, Continuous-Time Estimation:** rewrote this chapter from scratch to be consistent with the Simultaneous Trajectory Estimation and Mapping framework that my research group uses regularly
- **Appendix A, Matrix Primer:** a new appendix on linear algebra and matrix calculus that can serve as a primer and reference
- **Appendix B, Rotation and Pose Extras:** some extra derivations for rotations and poses including eigen/Jordan decomposition of rotation and pose matrices
- **Appendix C, Miscellaneous Extras:** a collection of useful results including the Fisher information matrix for a multivariate Gaussian, a derivation of Stein's lemma, converting continuous-time models to discrete time, connection to invariant EKF
- **Exercises:** several new exercises added plus solutions to almost all exercises in Appendix D

In addition to those who I thanked in the first edition, the following people were instrumental in this new version. First and foremost I would like to thank Professor James Forbes from McGill University who has been a wonderful collaborator over the years. He gave me several great suggestions for this second edition and provided invaluable advice on the invariant EKF and inertial navigation sections, in particular. The new chapter on variational inference was also a collaboration with James Forbes, and my student David Yoon and I thank them both for their help. Thanks also to Professor Gabriele D'Eleuterio; we worked together on the eigen/Jordan decomposition of rotation and pose matrices. Charles Cossette and Keenan Burnett helped find typos/issues in some new sections; thank you! A big thank you to my postdoc Dr. Johann Laconte who implemented the methods in the new inertial navigation section to make sure they worked as written and provided feedback to improve the readability. Many thanks to Lauren Cowles, my publisher at Cambridge, who encouraged me to put this second edition together.

Acronyms and Abbreviations

BA	bundle adjustment	199
BCH	Baker–Campbell–Hausdorff	272
BLUE	best linear unbiased estimate	73
CDF	cumulative distribution function	9
CRLB	Cramér–Rao lower bound	16
DARCES	data-aligned rigidity-constrained exhaustive search	170
EKF	extended Kalman filter	74
ELBO	evidence lower bound	185
EM	expectation minimization	204
ESGVI	exactly sparse Gaussian variational inference	191
FIM	Fisher information matrix	16
GN	Gauss–Newton	202
GP	Gaussian process	35
GPS	Global Positioning System	3
GVI	Gaussian variational inference	182
HMM	hidden Markov model	207
ICP	iterative closest point	347
IEKF	iterated extended Kalman filter	109
IMU	inertial measurement unit	4
IRLS	iteratively reweighted least squares	174
ISPKF	iterated sigmapoint Kalman filter	126
KF	Kalman filter	40
KL	Kullback–Leibler	13
LDU	lower-diagonal-upper	31
LG	linear-Gaussian	41
LOTUS	law of the unconscious statistician	11
LTI	linear time-invariant	85
LTV	linear time-varying	79
MAP	maximum a posteriori	4
ML	maximum likelihood	140
MMSE	minimum mean-squared error	73
NASA	National Aeronautics and Space Administration	3
NEES	normalized estimation error squared	160
NGD	natural gradient descent	183
NIS	normalized innovation squared	161
NLNG	nonlinear, non-Gaussian	97
PDF	probability density function	9

RAE	range-azimuth-elevation	249
RANSAC	random sample consensus	172
RTS	Rauch–Tung–Striebel	58
SDE	stochastic differential equation	79
SLAM	simultaneous localization and mapping	93
SMW	Sherman–Morrison–Woodbury	31
SP	sigmapoint	114
SPKF	sigmapoint Kalman filter	122
STEAM	simultaneous trajectory estimation and mapping	423
SVD	singular-value decomposition	449
SWF	sliding-window filter	144
UDL	upper-diagonal-lower	31
UKF	unscented Kalman filter (also called SPKF)	122

Notation

General Notation

a	This font is used for quantities that are real scalars
$\mathbf{a}$	This font is used for quantities that are real column vectors
$\mathbf{A}$	This font is used for quantities that are real matrices
$\boldsymbol{A}$	This font is used for time-invariant system quantities
$p(\mathbf{a})$	The probability density of $\mathbf{a}$
$p(\mathbf{a}\|\mathbf{b})$	The probability density of $\mathbf{a}$ given $\mathbf{b}$
$\mathcal{N}(\mathbf{a}, \mathbf{B})$	Gaussian probability density with mean $\mathbf{a}$ and covariance $\mathbf{B}$
$\mathcal{GP}(\boldsymbol{\mu}(t), \boldsymbol{\mathcal{K}}(t, t'))$	Gaussian process with mean function, $\boldsymbol{\mu}(t)$, and covariance function, $\boldsymbol{\mathcal{K}}(t, t')$
$\mathcal{O}$	Observability matrix
$(\cdot)_k$	The value of a quantity at timestep k
$(\cdot)_{k_1:k_2}$	The set of values of a quantity from timestep k_1 to timestep k_2, inclusive
$\underrightarrow{\mathcal{F}}_a$	A vectrix representing a reference frame in three dimensions
$\underrightarrow{a}$	A vector quantity in three dimensions
$(\cdot)^\times$	The cross-product operator, which produces a skew-symmetric matrix from a 3×1 column
$\mathbf{1}$	The identity matrix
$\mathbf{0}$	The zero matrix
$\mathbb{R}^{M \times N}$	The vector space of real $M \times N$ matrices
$\hat{(\cdot)}$	A posterior (estimated) quantity
$\check{(\cdot)}$	A prior quantity

Matrix-Lie-Group Notation

$SO(3)$	The special orthogonal group, a matrix Lie group used to represent rotations
$\mathfrak{so}(3)$	The Lie algebra associated with $SO(3)$
$SE(3)$	The special Euclidean group, a matrix Lie group used to represent poses
$\mathfrak{se}(3)$	The Lie algebra associated with $SE(3)$
$(\cdot)^\wedge$	An operator associated with the Lie algebra for rotations and poses
$(\cdot)^\curlywedge$	An operator associated with the adjoint of an element from the Lie algebra for poses
$\mathrm{Ad}(\cdot)$	An operator producing the adjoint of an element from the Lie group for rotations and poses
$\mathrm{ad}(\cdot)$	An operator producing the adjoint of an element from the Lie algebra for rotations and poses
$\mathbf{C}_{ba}$	A 3×3 rotation matrix (member of $SO(3)$) that takes points expressed in $\underrightarrow{\mathcal{F}}_a$ and re-expresses them in $\underrightarrow{\mathcal{F}}_b$, which is rotated with respect to $\underrightarrow{\mathcal{F}}_a$
$\mathbf{T}_{ba}$	A 4×4 transformation matrix (member of $SE(3)$) that takes points expressed in $\underrightarrow{\mathcal{F}}_a$ and re-expresses them in $\underrightarrow{\mathcal{F}}_b$, which is rotated/translated with respect to $\underrightarrow{\mathcal{F}}_a$
$\boldsymbol{\mathcal{T}}_{ba}$	A 6×6 adjoint of a transformation matrix (member of $\mathrm{Ad}(SE(3))$)