

Probabilistic modeling for sensor fusion with inertial measurements

Manon Kok

Cover illustration: Norm of the magnetic field estimated from experimental data using Gaussian process regression as described in Paper F. The experimental setup is shown in Figure 6 in this paper. The cover illustration is an adapted version of Figure 7c.

Linköping studies in science and technology. Dissertations.
No. 1814

Probabilistic modeling for sensor fusion with inertial measurements

Manon Kok

manko@isy.liu.se
www.control.isy.liu.se
Division of Automatic Control
Department of Electrical Engineering
Linköping University
SE-581 83 Linköping
Sweden

ISBN 978-91-7685-621-5

ISSN 0345-7524

Copyright © 2016 Manon Kok

Printed by LiU-Tryck, Linköping, Sweden 2016

To everyone who reads this

Abstract

In recent years, inertial sensors have undergone major developments. The quality of their measurements has improved while their cost has decreased, leading to an increase in availability. They can be found in stand-alone sensor units, so-called inertial measurement units, but are nowadays also present in for instance any modern smartphone, in Wii controllers and in virtual reality headsets.

The term inertial sensor refers to the combination of accelerometers and gyroscopes. These measure the external specific force and the angular velocity, respectively. Integration of their measurements provides information about the sensor's position and orientation. However, the position and orientation estimates obtained by simple integration suffer from drift and are therefore only accurate on a short time scale. In order to improve these estimates, we combine the inertial sensors with additional sensors and models. To combine these different sources of information, also called sensor fusion, we make use of probabilistic models to take the uncertainty of the different sources of information into account. The first contribution of this thesis is a tutorial paper that describes the signal processing foundations underlying position and orientation estimation using inertial sensors.

In a second contribution, we use data from multiple inertial sensors placed on the human body to estimate the body's pose. A biomechanical model encodes the knowledge about how the different body segments are connected to each other. We also show how the structure inherent to this problem can be exploited. This opens up for processing long data sets and for solving the problem in a distributed manner.

Inertial sensors can also be combined with time of arrival measurements from an ultrawideband (UWB) system. We focus both on calibration of the UWB setup and on sensor fusion of the inertial and UWB measurements. The UWB measurements are modeled by a tailored heavy-tailed asymmetric distribution. This distribution naturally handles the possibility of measurement *delays* due to multipath and non-line-of-sight conditions while not allowing for the possibility of measurements arriving *early*, i.e. traveling faster than the speed of light.

Finally, inertial sensors can be combined with magnetometers. We derive an algorithm that can calibrate a magnetometer for the presence of metallic objects attached to the sensor. Furthermore, the presence of metallic objects in the environment can be exploited by using them as a source of position information. We present a method to build maps of the indoor magnetic field and experimentally show that if a map of the magnetic field is available, accurate position estimates can be obtained by combining inertial and magnetometer measurements.

Populärvetenskaplig sammanfattning

När regissören Seth MacFarlane animerade teddybjörnen Ted i den storsäljande filmen med samma namn, lånade han ut inte bara sin röst utan också sin kropp till Ted. Genom att montera en mängd sensorer på kroppen, kan man lagra rörelsemönster digitalt, och sedan spela upp dem i exempelvis en animerad teddybjörn. Tekniken har använts inte bara i en stor mängd filmer, utan också av spelindustrin för att utveckla verklighetstroga avатарer, i medicinsk rehabilitering, och för att analysera och optimera rörelsemönster inom elit-idrott.

Denna avhandling behandlar en rad forskningsproblem kring denna typ av sensorer, av exakt samma modell som användes i filmen Ted. Sensorerna som används är en kombination av så kallade *tröghetssensorer* (eng. *inertial sensors*) sammansatta i små enheter. Varje enhet mäter acceleration inklusive tyngdaccelerationen med accelerometer och rotationshastigheter med gyroskop. Dessa sensorer kan tillsammans ge information om enhetens orientering och position. Här används även andra sensortyper och annan information, såsom matematiska modeller. Eftersom dessa modeller är en förenkling av verkligheten och sensormätningar aldrig är exakta, vill vi kombinera olika informationskällor, och ange hur mycket vi kan lita på varje källa. Detta kallas *sensorfusion* och kan göras med *probabilistiska modeller* som kan representera osäkerhet.

En sådan modell som används för att skatta kroppens rörelser är en biomekanisk modell som beskriver kroppens olika delar och hur dessa kan röra sig. I vår modell är dessa kroppsdelar sammankopplade. Vi antar alltså att personen inte förlorar kroppsdelar under experimenten. Denna typ av information kan användas för att animera teddybjörnen Ted eller för att skapa avатарer i dataspel. Om vi även vill att de ska interagera, till exempel hålla hand, behöver vi veta var de är. För att åstadkomma detta kan vi lägga till positionsmätningar.

En typ av sensor som ofta kombineras med tröghetssensorer är magnetometrar. Dessa mäter magnetfältet och man kan likna den vid en kompass som tillhandahåller information om sensorns orientering. I denna avhandling används magnetometern även för att bestämma sensorns position. Magnetometern mäter om det finns magnetiskt material i till exempel möbler eller i byggnaden. Denna information kan man använda för att avgöra var i byggnaden sensorn befinner sig.

Utvecklingen av tröghetssensorer har gått snabbt de senaste åren. Kvaliteten på mätningarna har ökat samtidigt som kostnaden har minskat, vilket har lett till en ökad tillgänglighet. Idag finns de exempelvis i mobiltelefoner, handkontroller till Wii tv-spel och i virtual reality headsets. Allt detta öppnar upp möjligheter för flera spännande tillämpningar inom detta intressanta forskningsområde. Denna avhandling visar att bra information om orientering och position kan fås genom att kombinera olika sorters mätningar och modeller. Kanske kan det leda till att vi kan animera teddybjörnar i våra egna vardagsrum om några år!

Acknowledgments

“The only reason for time is so that everything doesn’t happen at once.”

– Often attributed to Albert Einstein

The quote above symbolizes one of the things that I sometimes forget: I often try to do everything at once. My supervisor, Thomas Schön, has often had to remind me that this is a bad plan and to tell me that I’m “not allowed” to do more things. The reason that I am trying to do all these things, however, is that I enjoy my work so much. The Automatic Control Group at Linköping University is a very inspiring environment and working in the group has been a truly wonderful experience. I am very grateful to Thomas Schön, Fredrik Gustafsson and Svante Gunnarsson for giving me the opportunity to join the group. I am also grateful for the financial support of MC Impulse, a European Commission, FP7 research project and of CADICS, a Linnaeus Center funded by the Swedish Research Council (VR).

Writing this thesis would not have been as easy without the \LaTeX template developed and maintained by Henrik Tidefelt and Gustaf Hendeby. Gustaf, your help with \LaTeX and any other work-related questions, even late in the evenings and during the weekends is very much appreciated. The thesis would not have been in such good shape without the help of my supervisors Thomas Schön, Jeroen Hol and Fredrik Gustafsson. I am also very thankful to Johan Dahlin and Hanna Nyqvist for proofreading my kappa and to Michael Roth, Erik Hedberg, Zoran Sjanic and Jonas Linder for being available to comment on parts of the thesis, even though I asked them only last-minute. The Swedish popular scientific abstract would have neither been as popular scientific nor as Swedish without the great help of Ylva Jung, Fredrik Gustafsson and Thomas Schön. Thanks a lot to you all! Thanks also to Michael Lorenz for sending me brain fuel in the form of Mozartkugeln to give me the energy to write a lot in the past weeks!

I would not have enjoyed these five years as much if I would not have had such a great supervisor. Thomas, your enthusiasm about your work has a great positive influence on the people around you. I am very grateful for all our meetings and all the time that you’ve always made available for me, even when you were busy or when I had to ask a bit too much of your time. You taught me a lot about how to do research, how to write papers, structure my work, etc, etc. Furthermore, the fact that you are so organized has helped structure my work and brought some order in the chaos that I can sometimes create. I am in the lucky situation to also have two co-supervisors, Jeroen Hol and Fredrik Gustafsson. Jeroen, thanks for all the nice technical discussions we have had and for always being welcoming when I was visiting Xsens! Fredrik, thanks for always being available when I needed your help!

I am grateful for having had great collaborations over the past years with Johan Dahlin, Kjartan Halvorsen, Anders Hansson, Alf Isaksson, Daniel Jönsson, Sina Khoshfetrat Pakazad, Joel Kronander, Lennart Ljung, Henk Luinge, Fredrik Olsson, Johan Sjöberg, Arno Solin, Andreas Svensson, Simo Särkkä, David Törnqvist, Jonas Unger, Adrian Wills and Niklas Wahlström. Thanks specifically to Simo and Arno for being so kind as to welcome me for a 2.5 month pre-doc in the Bayesian Methodology Group at Aalto University. Thanks also to Niklas for great

collaborations on our magnetic field papers. I am also very thankful for having gotten the opportunity to supervise the master thesis project by Michael Lorenz, together with Thomas Seel and Philipp Müller from TU Berlin. I hope to have the opportunity for further collaborations with you also in the future!

I consider my colleagues in the group not so much as colleagues but more as my friends. I want to thank all of you for creating such an amazing work environment! Specifically, I would like to thank our head Svante Gunnarsson for making sure that there is always a good atmosphere and Ninna Stensgård for always being there for help with administrative tasks.

I have shared some amazing time with my colleagues over the years, both at work and outside of work. It has always been a great experience to go to conferences. Thanks everyone for also being up for so many fun activities around the conferences such as going on safari in Kruger park, wine tasting around Cape Town, whale watching in Vancouver, diving in Malaysia, and eating lots of dumplings in Beijing! Thanks also for all the nice times we shared in Linköping. Thanks to Sina Khoshfetrat Pakazad for always being the one to arrange fun things to do during weekends and evenings, thanks to Zoran Sjanic for always arranging for drinking beer on Wednesday's and to Marek Syldatk for making our corridor more lively. Thanks Johan Dahlin for being a great friend and for always being there to answer all of my questions. Thanks also to Zoran Sjanic, Ylva Jung, Jonas Linder, Niklas Wahlström and Erik Hedberg for being there for me when I needed someone to talk to. Last but not least, I would like to thank Hanna Nyqvist, Clas Veibäck, Michael Roth, George Mathai, Gustaf Hendeby, Emre Özkan, Daniel Petersson, Emina Alickovic, Mahdieh Sadabadi, André Carvalho Bittencourt, Martin Lindfors and all the other people from the Automatic Control Group for a wonderful time in the past years!

I would also like to thank my friends from outside our group. Our friendship is an important reason that Linköping feels like home! Furthermore, I would like to thank my family and friends back in the Netherlands for always being welcoming when I go back home and for making sure that I always have a great time when I'm visiting.

Unfortunately, my time as a PhD student is coming to an end. I have learned so much and met so many great people over the past years ... thank you all for this! I'm looking forward to seeing what the future brings!

Linköping, December 2016

Manon Kok

Contents

Notation	xvii
I Background	
1 Introduction	3
1.1 Background	3
1.2 Additional sensors and models	9
1.2.1 Magnetometers	9
1.2.2 Ultrawideband	10
1.2.3 Biomechanical models	10
1.3 Main contributions	11
1.4 Outline	12
2 Probabilistic models	19
2.1 Models for position and orientation estimation	21
2.2 Maps of the magnetic field	24
2.3 Visualizing the resulting model structures	25
3 Inference	29
3.1 Building maps of the magnetic field	29
3.2 Estimating position and orientation	30
3.3 Estimating calibration parameters	33
4 Conclusions and future work	35
4.1 Position and orientation estimation using inertial sensors	35
4.2 Inertial sensor motion capture	36
4.3 Combining UWB with inertial sensors	38
4.4 Magnetometer calibration	38
4.5 Mapping and localization using magnetic fields	39
4.6 Concluding remarks	40
Bibliography	41

II Publications

A	Using inertial sensors for position and orientation estimation	49
1	Introduction	51
1.1	Background and motivation	52
1.2	Using inertial sensors for pose estimation	55
1.3	Tutorial content and its outline	57
2	Inertial sensors	58
2.1	Coordinate frames	59
2.2	Angular velocity	60
2.3	Specific force	60
2.4	Sensor errors	62
3	Probabilistic models	65
3.1	Introduction	66
3.2	Parametrizing orientation	69
3.3	Probabilistic orientation modeling	75
3.4	Measurement models	77
3.5	Choosing the state and modeling its dynamics	82
3.6	Models for the prior	83
3.7	Resulting probabilistic models	85
4	Estimating position and orientation	88
4.1	Smoothing in an optimization framework	88
4.2	Filtering estimation in an optimization framework	95
4.3	Extended Kalman filtering	97
4.4	Evaluation based on experimental and simulated data	103
4.5	Extending to pose estimation	119
5	Calibration	122
5.1	Maximum a posteriori calibration	123
5.2	Maximum likelihood calibration	124
5.3	Orientation estimation with an unknown gyroscope bias	126
5.4	Identifiability	127
6	Concluding remarks	128
A	Orientation parametrizations	130
A.1	Quaternion algebra	130
A.2	Conversions between different parametrizations	132
B	Pose estimation	133
B.1	Smoothing in an optimization framework	133
B.2	Filtering in an optimization framework	133
B.3	EKF with quaternion states	134
B.4	EKF with orientation deviation states	134
C	Gyroscope bias estimation	135
C.1	Smoothing in an optimization framework	135
C.2	Filtering in an optimization framework	135
C.3	EKF with quaternion states	135
C.4	EKF with orientation deviation states	136
	Bibliography	137

B	An optimization-based approach to motion capture using inertial sensors	145
1	Introduction	147
2	Problem formulation	149
3	Biomechanical model	150
4	Dynamic and sensor models	153
4.1	Dynamic model	153
4.2	Sensor model	154
5	Resulting algorithm	155
6	Experiments	155
7	Conclusions and future work	159
	Bibliography	161
C	A scalable and distributed solution to the inertial motion capture problem	163
1	Introduction	165
2	Problem formulation	168
3	Model	169
3.1	Dynamics of the state $x_t^{S_i}$	170
3.2	Placement of the sensors on the body segments	170
3.3	Biomechanical constraints	171
4	Problem reformulation enabling structure exploitation	171
4.1	Reordering based on time	172
4.2	Reordering based on sensors and body segments	173
5	Tree structure in coupled problems and message passing	175
6	Scalable and distributed solutions using message passing	177
7	Results and discussion	179
8	Conclusions and future work	180
	Bibliography	182
D	Indoor positioning using ultrawideband and inertial measurements	185
1	Introduction	187
2	Related work	189
3	Problem formulation	191
4	Sensor models	192
4.1	Modeling the ultrawideband measurements	193
4.2	Modeling the inertial measurements	194
5	Ultrawideband calibration	195
5.1	Initial estimate: step I	197
5.2	Initial estimate: step II - multilateration	198
5.3	Resulting calibration algorithm	199
6	Sensor fusion	199
7	Experimental results	201
7.1	Experimental validation of the asymmetric noise distribution	201
7.2	Calibration	202
7.3	Pose estimation	205

8	Conclusions and future work	210
	Bibliography	211
E	Magnetometer calibration using inertial sensors	215
1	Introduction	217
2	Related work	219
3	Problem formulation	220
4	Magnetometer measurement model	221
5	Calibration algorithm	224
	5.1 Optimization algorithm	225
	5.2 Evaluation of the cost function	225
	5.3 The parameter vector θ	226
6	Finding good initial estimates	227
	6.1 Ellipsoid fitting	227
	6.2 Determine misalignment of the inertial and magnetometer sensor axes	229
7	Experimental results	230
	7.1 Experimental setup	230
	7.2 Calibration results	230
	7.3 Heading estimation	233
8	Simulated heading accuracy	236
9	Conclusions	238
	Bibliography	240
F	Modeling and interpolation of the ambient magnetic field by Gaussian processes	243
1	Introduction	245
2	Related work	248
3	The ambient magnetic field	250
4	Modeling the magnetic field using Gaussian process priors	251
	4.1 Gaussian process regression	252
	4.2 Interpolation of magnetic fields	254
	4.3 Separate modeling of the magnetic field components	254
	4.4 Modeling the magnetic field as the gradient of a scalar potential	256
5	Efficient GP modeling of the magnetic field	257
	5.1 Reduced-rank GP modeling	257
	5.2 Batch estimation	259
	5.3 Sequential estimation	261
	5.4 Spatio-temporal modeling	261
6	Experiments	263
	6.1 Simulated experiment	263
	6.2 Empirical proof-of-concept data	264
	6.3 Mapping the magnetic field in a building	268
	6.4 Online mapping	270
7	Discussion	273

8	Conclusion	275
	Bibliography	277
G	MEMS-based inertial navigation based on a magnetic field map	283
1	Introduction	285
2	Models	286
	2.1 Dynamical model	287
	2.2 Magnetometer measurement model	289
	2.3 Some additional words about the magnetic field model	290
3	Computing the estimate	290
	3.1 RBPF-MAP	291
4	Experimental results	293
	4.1 Experimental setup	293
	4.2 Results	294
5	Conclusions and future work	294
	Bibliography	296

Notation

ABBREVIATIONS

Abbreviation	Meaning
BFGS	Broyden-Fletcher-Goldfarb-Shanno
EKF	Extended Kalman filter
GP	Gaussian process
GPS	Global positioning system
IMU	Inertial measurement unit
KF	Kalman filter
LS	Least squares
MAP	Maximum a posteriori
MEKF	Multiplicative extended Kalman filter
MEMS	Micro-machined electromechanical system
MHE	Moving horizon estimation
ML	Maximum likelihood
NLOS	Non-line-of-sight
NLS	Nonlinear least squares
PDF	Probability density function
PDR	Pedestrian dead-reckoning
PF	Particle filter
PF-MAP	Maximum a posteriori estimate for the particle filter
RMS	Root mean square
RMSE	Root mean square error
RTS	Rauch-Tung-Striebel
RBPF	Rao-Blackwellized particle filter
RBPF-MAP	Maximum a posteriori estimate for the Rao-Blackwellized particle filter
SLAM	Simultaneous localization and mapping
SQP	Sequential quadratic programming
TOA	Time of arrival
TDOA	Time difference of arrival
UWB	Ultrawideband
VR	Virtual reality

SYMBOLS AND OPERATORS

Notation	Meaning
x_t	Vector x at time t
$x_{1:N}$	Vector x from time $t = 1$ to $t = N$
\hat{x}	Estimate of x
x^u	Vector x expressed in the u -frame
R^{uv}	Rotation matrix from the v -frame to the u -frame
\odot	Quaternion multiplication
q^L	Left quaternion multiplication of the quaternion q
q^R	Right quaternion multiplication of the quaternion q
q_v	Vector part of the quaternion q
\mathbb{R}	Set of real numbers
$SO(3)$	Special orthogonal group in three dimensions
$\det A$	Determinant of the matrix A
$\text{Tr } A$	Trace of the matrix A
A^T	Transpose of the matrix A
\times	Cross product
$[a \times]$	Cross product matrix of the vector a
\otimes	Kronecker product
A^{-1}	Inverse of the matrix A
A^\dagger	Pseudo-inverse of the matrix A
$\mathcal{N}(\mu, \sigma^2)$	Gaussian distribution with mean μ and covariance σ^2
$\text{Cauchy}(\mu, \gamma)$	Cauchy distribution with location parameter μ and scale parameter γ
$\mathcal{U}(a, b)$	Uniform distribution on the interval $[a, b]$
$\mathcal{GP}(\mu, k)$	Gaussian process with mean μ and covariance function k
$p(\cdot)$	Probability density function
$p(a b)$	Conditional probability of a given b
$p(a, b)$	Joint probability of a and b
\sim	Is distributed according to
\mathbb{E}	Expected value
cov	Covariance
\mathcal{I}_n	Identity matrix of size $n \times n$
$0_{m \times n}$	Zero matrix of size $m \times n$
\triangleq	Defined as
\emptyset	Empty set
\in	Is a member of
$A \subseteq B$	A is a subset of or is included in B
$\arg \max$	Maximizing argument
$\arg \min$	Minimizing argument
$\ a\ _2$	Two-norm of the vector a

Part I

Background

1

Introduction


In this thesis, we consider the problem of estimating position and orientation using inertial sensors. In Section 1.1, we give some example applications and introduce what inertial sensors are and what their measurements look like. We will also discuss why inertial sensors typically need to be combined with additional sensors or models to obtain accurate position and orientation information. Examples of additional sensors and models used in this thesis are given in Section 1.2. In Sections 1.3 and 1.4, we will introduce the contributions of the thesis and give an outline of the rest of the thesis.

1.1 Background

Sensors can be used to provide information about the position and orientation of a person or an object. For instance, it is possible to place sensors on a human body to see how the person moves. This information can be useful for rehabilitation or for improving sports performance. An example can be seen in Figure 1.1a where Olympic and world champion speed skating Ireen Wüst wears sensors on her body that give information about her posture while ice skating. One can imagine that she can use this information to analyze which angles her knees and hips should have to skate as fast as possible and if her posture changes when she gets more tired. It is also possible to use the information about how a person moves for motion capture in movies and games, as illustrated in Figure 1.1b, where the actor Seth MacFarlane wears sensors on his body that measure his movements to animate the bear Ted. Sensors can also be placed in or on objects, for example cars, to provide information about their position and orientation as illustrated in Figure 1.1c. This information is for instance useful for self-driving cars. There is a wide range of other examples that one can think of, such as using sensors to



(a) Left: Olympic and world champion speed skating Ireen Wüst wearing sensors on her body. Right: graphical representation of the estimated orientation and position of her body segments.




(b) Actor Seth MacFarlane wearing sensors on his body to capture his motion and animate the bear Ted.



(c) Sensors can be used to provide information about the position of the cars in a challenge on cooperative and autonomous driving.

Figure 1.1: Example applications of using sensors to obtain information about the position and orientation of cars and of the various body segments of a person. Courtesy of Xsens Technologies.



(a) Gyroscope measurements $y_{\omega,t}$ in the x - (blue), y - (green) and z -axis (red) of the sensor. (b) Accelerometer measurements $y_{a,t}$ in the x - (blue), y - (green) and z -axis (red) of the sensor.

Figure 1.2: Inertial measurements collected with a smartphone.

obtain information about the position and orientation of robots, unmanned aerial vehicles, trains and people.

The sensors placed on the people and in the cars in Figure 1.1 are *inertial sensors*. The term inertial sensor is used to refer to the combination of accelerometers and gyroscopes. A gyroscope measures the rate of change of the orientation of the sensor, called the *angular velocity*. The gyroscopes that we consider have three axes, implying that they measure the angular velocity in three directions. This is illustrated in Figure 1.2a, which shows gyroscope measurements collected with a Sony Xperia Z5 Compact smartphone using the app described in Hendeby et al. (2014). For the first 10 seconds, the smartphone was lying stationary on a table. Afterwards, the gyroscope was rotated back and forth around its x -, y - and z -axis. An accelerometer measures both the earth's gravity and the acceleration of the sensor. The accelerometers that we consider also have three axes as illustrated in Figure 1.2b. During the first 10 seconds, the smartphone was again lying stationary on a table. The accelerometer measurements can be seen to be around zero in the x - and y - axis, while the z -axis measures a value of around 10 m/s^2 which is due to the earth's gravity. When rotating the smartphone, the accelerometer measures the gravity in different axes. After around 37 seconds, the smartphone was shaken, resulting in a significant acceleration that is measured in addition to the earth's gravity.

Over recent years, inertial sensors have undergone major developments. They have become smaller, lighter and cheaper while providing more accurate measurements. Because of this, they are nowadays available in a large number of devices such as smartphones, Wii controllers and virtual reality (VR) headsets, as shown in Figure 1.3. They are also present in dedicated devices called inertial measurement units (IMUs). The sensor devices placed on the persons and in the cars in Figure 1.1 are IMUs.

Gyroscopes can be used to provide information about the orientation of the sensor, by adding up the changes in orientation over time. This process is called *integration* of the signal. Accelerometers can be used to provide information both about the position and about the orientation of the sensor. If the sensor is not accelerated, the accelerometer measurements can be used to provide information



(a) Left bottom: an Xsens MTx IMU (Xsens Technologies B.V., 2016). Left top: a Trivisio Colibri Wireless IMU (Trivisio Prototyping GmbH, 2016). Right: a Samsung Galaxy S4 mini smartphone.



(b) A Samsung gear VR.¹



(c) A Wii controller containing an accelerometer and a Motion-Plus expansion device containing a gyroscope.²

Figure 1.3: Examples of devices containing inertial sensors.

¹ 'Samsung Gear VR' available at flic.kr/photos/pestoverde/15247458515 under CC BY 2.0 (<http://creativecommons.org/licenses/by/2.0>).

² 'WiiMote with MotionPlus' by Asmodai available at https://commons.wikimedia.org/wiki/File:WiiMote_with_MotionPlus.JPG under CC BY SA (<https://creativecommons.org/licenses/by-sa/3.0/>).





Figure 1.4: Schematic illustration of dead-reckoning, where accelerometer and gyroscope measurements are integrated to position and orientation.

about the orientation of the sensor, because they measure the direction of the earth's gravity with respect to the axes of the sensor. If the sensor is accelerated, the measurements provide information about the change in velocity, which in turn provides information about the change in position. Hence, to obtain position information from the acceleration of the sensor, the signal needs to be integrated twice. To be able to distinguish between the acceleration of the sensor and the earth's gravity, the orientation needs to be known so that the gravity component can be subtracted from the measurements. Because of this, when using inertial sensors, the estimation of the sensor's position is inextricably linked to the estimation of its orientation. The process of integrating the inertial sensor measurements to obtain position and orientation information is often called *dead-reckoning*. This process is summarized in Figure 1.4.


In practice, the position and orientation estimates obtained using dead-reckoning are only accurate for a short time. The reason is that the gyroscope and accelerometer measurements are both biased and noisy, as illustrated in Figure 1.5, where we zoom in on the first 10 seconds of the data shown in Figure 1.2. Because of this, the integration steps from angular velocity to rotation and from acceleration to position introduce *integration drift*. The integration drift in orientation for simulated gyroscope data is illustrated in Figure 1.6. This simulated data has the same bias as the gyroscope measurements in Figure 1.5a, and the same spread in the noise. Because of the constant bias, the orientation error grows linearly with time. The different lines in Figure 1.6 represent the orientation error for different realizations of this noise. The variation in the orientation error for different noise realizations increases over time. The integration drift is more severe for position, which relies both on double integration of the acceleration and on accurate orientation estimates to subtract the earth's gravity.

Because the process of dead-reckoning only gives accurate position and orientation information on a short time scale, inertial sensors are typically combined with additional sensors or additional models. In this thesis, we consider two separate problems related to position and orientation estimation using inertial sensors. The first is concerned only with orientation estimation. The three-dimensional orientation can be described in terms of the roll, pitch and yaw or heading angles. The combination of the roll and pitch angles is often also called inclination. In a



(a) Gyroscope measurements $y_{\omega,t}$ in the x - (blue), y - (green) and z -axis (red) the sensor.

(b) Accelerometer measurements $y_{a,t}$ in the x - (blue), y - (green) and z -axis (red) of the sensor.



(c) Histogram of the gyroscope measurements in the x -axis.

(d) Histogram of the accelerometer measurements in the x -axis.

Figure 1.5: The first 10 seconds of the gyroscope and accelerometer measurements shown in Figure 1.2, during which the smartphone is lying stationary on a table (a,b) and the histograms of one of the axes of the gyroscope and of the accelerometer (c,d).




Figure 1.6: Integration of simulated one-dimensional gyroscope measurements to orientation for 50 different noise realizations having the same characteristics as in Figure 1.5a.

second problem, we consider the combined estimation of position and orientation, which is often also called *pose estimation*. In this case, we are interested both in the three-dimensional orientation and in the three-dimensional position.

1.2 Additional sensors and models


In this section, we will discuss a number of additional sensors and additional models that are used in this thesis to complement the inertial sensors.

1.2.1 Magnetometers

A magnetometer measures the strength and the direction of the magnetic field. The magnetic field consists of contributions both from the local earth magnetic field and from the field due to the presence of magnetic material. The magnitude and the direction of the earth magnetic field depend on the location on the earth. The horizontal component points to the earth magnetic north. The properties of the earth magnetic field are accurately known from geophysical studies, see e.g. National Centers for Environmental Information (2016).

In combination with inertial sensors, magnetometers typically serve the purpose of a compass and are used to provide information about the sensor's heading. This relies on the assumption that the magnetic field is at least locally constant and that it points in the direction of a local magnetic north. There are two reasons why this assumption is frequently violated in practice. Firstly, the sensor can be mounted such that it is rigidly attached to magnetic material. This is for instance the case when the magnetometer is integrated in a smartphone or when it is placed in a car. Secondly, objects containing magnetic material can be present in the vicinity of the sensor, specifically in indoor environments. For instance, there is typically a large amount of magnetic material present in the structures of buildings and in the furniture present in the building.

If the magnetic material is rigidly attached to the sensor, the magnetometer can be calibrated for the presence of this material. Afterwards, the measurements can be used for heading estimation as if the material was not present. The presence of magnetic material in the vicinity of the sensor, however, can not be calibrated for and is typically considered an undesired disturbance. An alternative view is that the presence of magnetic material in indoor environments can be exploited by using it as a source of position information, see e.g. Angermann et al. (2012); Frassl et al. (2013); Solin et al. (2016). This can be done by building a *map* of the magnetic field. Both information about the strength and about the direction of the field can be included in the map. An example of an indoor magnetic field map is shown in Figure 1.7a. It is built from data collected using the mobile robot shown in Figure 1.7b. After the map has been constructed, magnetometer measurements can be compared to it in order to obtain information about possible sensor locations, see e.g. Solin et al. (2016) and Paper G.



(a) Map of the magnitude of the indoor magnetic field.



(b) Mobile robot.


Figure 1.7: Left: Magnitude of an indoor magnetic field estimated using the method presented in Paper F. Right: Mobile robot that was used to collect data.

1.2.2 Ultrawideband

Time of arrival (TOA) measurements from an ultrawideband (UWB) system can be used to provide information about the position of the sensor. UWB is a radio technology which uses a very large frequency band. An example of a UWB system consisting of a number of stationary UWB receivers and a number of small, mobile transmitters is depicted in Figure 1.8a. Each UWB transmitter sends out a UWB pulse as illustrated in Figure 1.8b. The pulse travels with the speed of light towards the receivers, which each measure when the pulse arrives. Combining the measurements from different receivers, it is possible to obtain an estimate of the position of the transmitter. Note that the time when the pulses arrive needs to be measured with very high accuracy. For instance, if the transmitter is 10 meters away from the receiver, it will take the pulse only approximately 33 nanoseconds to reach the receiver.


1.2.3 Biomechanical models

In the examples shown in Figures 1.1a and 1.1b, multiple IMUs are placed on the human body to estimate its movements. More specifically, the IMUs are placed on a large number of body segments and the position and the orientation of each body segment is estimated. This is schematically illustrated in Figure 1.9a. The two body segments can be thought of as the upper and the lower leg, each having an IMU attached to it. The sensors are attached as rigidly as possible to the body segments. This is illustrated in Figure 1.9b, which shows a suit containing 17 IMUs. The suit is meant to be a tight fit such that the sensors move as little as possible with respect to the body. For this application, knowledge about how the human body can move is available to complement the inertial measurements. For instance, the different body segments are known to be connected to each other. This can be captured in biomechanical models.



(a) Hardware used in a UWB setup. More specifically, a UWB receiver and a small, battery-powered UWB transmitter. Courtesy of Xsens Technologies. (b) A UWB setup consisting of a number of stationary receivers obtaining TOA measurements of signal pulses originating from a mobile transmitter.

Figure 1.8: Illustration of the TOA measurements and the hardware used in a UWB setup.



(a) Schematic illustration of two connected body segments (purple and green), each with a sensor (orange) attached to it.



(b) Suit containing 17 IMUs placed on the human body. Courtesy of Xsens Technologies.

Figure 1.9: Illustration of using IMUs placed on the human body to estimate its movements.

1.3 Main contributions

In this thesis, inertial sensors are combined with additional sensors and additional models for position and orientation estimation. Examples of sensors and models that can be used for this were discussed in Section 1.2. The choice of these examples was highly inspired by the contributions of this thesis. In short, these contributions are:

- A tutorial paper describing the signal processing foundations, i.e. the algorithms and models, underlying position and orientation estimation using inertial sensors [Paper A].
- An approach to estimate the pose of the human body using inertial sensors placed on the body, as illustrated in Figure 1.9 [Paper B]. We also present a method that allows us to solve this problem for large data sets. The same approach can be used to distribute the computations needed to solve the problem over the sensors on the body [Paper C].
- An approach to combine inertial measurements with TOA measurements from a UWB system for indoor positioning. We provide solutions to the pose estimation problem using inertial and UWB measurements, and to the calibration of the UWB setup shown in Figure 1.8 [Paper D].
- We have developed a magnetometer calibration algorithm which uses inertial sensors to calibrate the magnetometer for the presence of magnetic disturbances attached to the sensor. It also calibrates for magnetometer sensor errors and for misalignment between the magnetometer and the inertial sensor axes [Paper E].
- An approach to build maps of the indoor magnetic field, taking into account the well-known physical properties of the magnetic field [Paper F]. An example of a magnetic field map obtained using this method is illustrated in Figure 1.7. We also show that the magnetic field can be used as a source of position information for an experiment where we generate a known magnetic field [Paper G].

1.4 Outline

The thesis consists of two parts. In Part II, seven papers are presented. The contributions of these papers were discussed in Section 1.3. Below we provide a summary of each paper in Part II together with a discussion of the background and of the author's contributions. A background to these papers is provided in Part I. In this introductory chapter, we have briefly introduced the problem at hand, the sensors and models involved and the contributions of the thesis. To combine these different sources of information, also called *sensor fusion*, we make use of *probabilistic models* to take into account each source of information and its accuracy. In Chapters 2 and 3 we discuss the subjects of probabilistic models and inference using these models. Having introduced these topics we revisit the contributions of the thesis in Chapter 4 and discuss them in more technical detail, followed by a discussion of some directions for future work.

Paper A: Using inertial sensors for position and orientation estimation

Paper A is an edited version of

M. Kok, J. D. Hol, and T. B. Schön. Using inertial sensors for position and orientation estimation. Technical Report LiTH-ISY-R-3093, Department of Electrical Engineering, Linköping University, Linköping, Sweden, December 2016a.

Summary: In recent years, micro-machined electromechanical system (MEMS) inertial sensors (3D accelerometers and 3D gyroscopes) have become widely available due to their small size and low cost. Inertial sensor measurements are obtained at high sampling rates and can be integrated to obtain position and orientation (pose) estimates. These pose estimates are accurate on a short time scale, but suffer from integration drift over longer time scales. To overcome this issue, inertial sensors are typically combined with additional sensors and models. In this tutorial we focus on the signal processing aspects of pose estimation using inertial sensors, discussing different modeling choices and a selected number of important algorithms. These algorithms are meant to provide the reader with a starting point to implement their own pose estimation algorithm. The algorithms include optimization-based smoothing and filtering as well as computationally cheaper extended Kalman filter implementations.

Background and contributions: A couple of years ago, Prof. Thomas Schön came up with the idea of writing a tutorial paper on pose estimation using inertial sensors. Towards the end of the PhD of the author of this thesis, the plans for writing this paper became more concrete since it is a nice way of rounding up the work we have done together in the past years. The paper has been written together with Dr. Jeroen Hol.

Paper B: An optimization-based approach to motion capture using inertial sensors

Paper B is an edited version of

M. Kok, J. D. Hol, and T. B. Schön. An optimization-based approach to human body motion capture using inertial sensors. In *Proceedings of the 19th World Congress of the International Federation of Automatic Control*, pages 79–85, Cape Town, South Africa, August 2014.

Summary: In inertial human motion capture, a multitude of body segments are equipped with inertial measurement units, consisting of 3D accelerometers, 3D gyroscopes and 3D magnetometers. Relative position and orientation estimates can be obtained using the inertial data together with a biomechanical model. In this work we present an optimization-based solution to magnetometer-free inertial motion capture. It allows for natural inclusion of biomechanical constraints, for handling of nonlinearities and for using all data in obtaining an estimate. As a proof-of-concept we apply our algorithm to a lower body configuration, illustrating that the estimates are drift-free and match the joint angles from an optical reference system.

Background and contributions: The co-authors Dr. Jeroen Hol and Prof. Thomas Schön came up with the idea of solving the human body motion capture

problem as an optimization problem. The implementation of the optimization algorithm has been done using a framework developed by Xsens Technologies. With this framework, it is possible to define the optimization problem at a high level. The author of this thesis has been involved in developing and implementing the algorithm, in the data collection and has written a major part of the paper.

Paper C: A scalable and distributed solution to the inertial motion capture problem

Paper C is an edited version of

M. Kok, S. Khoshfetrat Pakazad, T. B. Schön, A. Hansson, and J. D. Hol. A scalable and distributed solution to the inertial motion capture problem. In *Proceedings of the 19th International Conference on Information Fusion*, pages 1348–1355, Heidelberg, Germany, July 2016b.

Summary: In inertial motion capture, a multitude of body segments are equipped with inertial sensors, consisting of 3D accelerometers and 3D gyroscopes. Using an optimization-based approach to solve the motion capture problem allows for natural inclusion of biomechanical constraints and for modeling the connection of the body segments at the joint locations. The computational complexity of solving this problem grows both with the length of the data set and with the number of sensors and body segments considered. In this work, we present a scalable and distributed solution to this problem using tailored message passing, capable of exploiting the structure that is inherent in the problem. As a proof-of-concept we apply our algorithm to data from a lower body configuration.

Background and contributions: This work solves the inertial motion capture problem from Paper B using the message passing algorithm developed by Khoshfetrat Pakazad et al. (2016). After the author of this thesis presented the inertial motion capture problem during an internal group meeting, Dr. Sina Khoshfetrat Pakazad suggested that the structure of the motion capture problem can be exploited using the message passing algorithm. The implementation and the writing of the paper has been done together with Dr. Sina Khoshfetrat Pakazad.

Paper D: Indoor positioning using ultrawideband and inertial measurements

Paper D is an edited version of

M. Kok, J. D. Hol, and T. B. Schön. Indoor positioning using ultrawideband and inertial measurements. *IEEE Transactions on Vehicular Technology*, 64(4):1293–1303, 2015b.

Summary: In this work we present an approach to combine measurements from inertial sensors (accelerometers and gyroscopes) with time of arrival measurements from an ultrawideband system for indoor positioning. Our algorithm

uses a tightly-coupled sensor fusion approach, where we formulate the problem as a maximum a posteriori problem that is solved using an optimization approach. It is shown to lead to accurate 6D position and orientation estimates when compared to reference data from an independent optical tracking system. To be able to obtain position information from the ultrawideband measurements, it is imperative that accurate estimates of the ultrawideband receivers' positions and their clock offsets are available. Hence, we also present an easy-to-use algorithm to calibrate the ultrawideband system using a maximum likelihood formulation. Throughout this work, the ultrawideband measurements are modeled by a tailored heavy-tailed asymmetric distribution to account for measurement outliers. The heavy-tailed asymmetric distribution works well on experimental data, as shown by analyzing the position estimates obtained using the ultrawideband measurements via a novel multilateration approach.

Background and contributions: The co-authors of this paper, Dr. Jeroen Hol and Prof. Thomas Schön, have been working on the subject of indoor positioning using ultrawideband and inertial measurements, resulting in Hol et al. (2009, 2010) and in the results presented in Hol (2011). The author of this thesis has substantially extended and adapted the previously presented algorithms for sensor fusion, calibration and multilateration. The paper has been written together with Dr. Jeroen Hol.

Paper E: Magnetometer calibration using inertial sensors

Paper E is an edited version of

M. Kok and T. B. Schön. Magnetometer calibration using inertial sensors. *IEEE Sensors Journal*, 16(14):5679 – 5689, 2016.

Earlier versions of this work were presented in:

M. Kok and T. B. Schön. Maximum likelihood calibration of a magnetometer using inertial sensors. In *Proceedings of the 19th World Congress of the International Federation of Automatic Control*, pages 92–97, Cape Town, South Africa, August 2014,

M. Kok, J. D. Hol, T. B. Schön, F. Gustafsson, and H. Luinge. Calibration of a magnetometer in combination with inertial sensors. In *Proceedings of the 15th International Conference on Information Fusion*, pages 787–793, Singapore, July 2012.

Summary: In this work we present a practical algorithm for calibrating a magnetometer for the presence of magnetic disturbances and for magnetometer sensor errors. To allow for combining the magnetometer measurements with inertial measurements for orientation estimation, the algorithm also corrects for misalignment between the magnetometer and the inertial sensor axes. The calibration algorithm is formulated as the solution to a maximum likelihood problem and the computations are performed offline. The algorithm is shown to give good

results using data from two different commercially available sensor units. Using the calibrated magnetometer measurements in combination with the inertial sensors to determine the sensor's orientation is shown to lead to significantly improved heading estimates.

Background and contributions: Before the author of this thesis started her work as a PhD student at Linköping University, she worked at Xsens Technologies. During this time she studied the topic of magnetometer calibration. Hence, the magnetometer calibration problem provided a good starting point for research during her PhD. A first paper on this subject has therefore been co-authored by Dr. Jeroen Hol and Dr. Henk Luinge from Xsens Technologies. Later work has mainly been done in cooperation with Prof. Thomas Schön. Dr. Henk Luinge and Laurens Slot from Xsens Technologies and Dr. Gustaf Hendeby from Linköping University have been so kind as to help in collecting the data sets presented in the paper. The author of this thesis has implemented the calibration algorithm and has written a major part of the paper.

Paper F: Modeling and interpolation of the ambient magnetic field by Gaussian Processes

Paper F is an edited version of

A. Solin, M. Kok, N. Wahlström, T. B. Schön, and S. Särkkä. Modeling and interpolation of the ambient magnetic field by Gaussian processes. *ArXiv e-prints*, September 2015. arXiv:1509.04634.

Summary: Anomalies in the ambient magnetic field can be used as features in indoor positioning and navigation. By using Maxwell's equations, we derive and present a Bayesian non-parametric probabilistic modeling approach for interpolation and extrapolation of the magnetic field. We model the magnetic field components jointly by imposing a Gaussian process (GP) prior on the latent scalar potential of the magnetic field. By rewriting the GP model in terms of a Hilbert space representation, we circumvent the computational pitfalls associated with GP modeling and provide a computationally efficient and physically justified modeling tool for the ambient magnetic field. The model allows for sequential updating of the estimate and time-dependent changes in the magnetic field. The model is shown to work well in practice in different applications: we demonstrate mapping of the magnetic field both with an inexpensive Raspberry Pi powered robot and on foot using a standard smartphone.

Background and contributions: This paper has largely been written during the author's PreDoc visit to the Bayesian Methodology Group at Aalto University in January – March 2015. It combines the approaches from Wahlström et al. (2013) and Solin and Särkkä (2014) and builds on the common interest of the authors in localization using magnetic fields as a source of position information. The map of the indoor magnetic field obtained using the method presented in this paper, has been used in Solin et al. (2016) for localization. In the future we hope to find time to combine these ideas into a working simultaneous localization and mapping (SLAM) solution. The work on implementation and writing of the paper

has been split more or less equally between Dr. Arno Solin and the author of this thesis.

Paper G: MEMS-based inertial navigation based on a magnetic field map

Paper G is an edited version of

M. Kok, N. Wahlström, T. B. Schön, and F. Gustafsson. MEMS-based inertial navigation based on a magnetic field map. In *Proceedings of the 38th International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, pages 6466–6470, Vancouver, Canada, May 2013.

Summary: This paper presents an approach for 6D pose estimation where MEMS inertial measurements are complemented with magnetometer measurements assuming that a model (map) of the magnetic field is known. The resulting estimation problem is solved using a Rao-Blackwellized particle filter. In our experimental study the magnetic field is generated by a magnetic coil giving rise to a magnetic field that we can model using analytical expressions. The experimental results show that accurate position estimates can be obtained in the vicinity of the coil, where the magnetic field is strong.

Background and contributions: The idea of looking into pose estimation using magnetometers as a source of position information was started through discussions with Dr. Slawomir Grzonka during the CADICS “Learning World Models” workshop in 2010 in Linköping. The experiments used in the paper were performed while the author of this thesis was working at Xsens Technologies. During this time, a first implementation of the pose estimation algorithm was made, using an extended Kalman filter. During the author’s time at Linköping University, the work has been extended with an implementation using a Rao-Blackwellized particle filter. The author of this thesis wrote a major part of this paper. This paper was the start of our work towards SLAM using magnetic measurements.

Publications of related interest, but not included in this thesis

F. Olsson, M. Kok, K. Halvorsen, and T. B. Schön. Accelerometer calibration using sensor fusion with a gyroscope. In *Proceedings of the IEEE Workshop on Statistical Signal Processing*, pages 660–664, Palma de Mallorca, Spain, June 2016.

M. Kok, J. Dahlin, T. B. Schön, and A. Wills. Newton-based maximum likelihood estimation in nonlinear state space models. In *Proceedings of the 17th IFAC Symposium on System Identification*, pages 398–403, Beijing, China, October 2015a.

A. Svensson, T. B. Schön, and M. Kok. Nonlinear state space smoothing using the conditional particle filter. In *Proceedings of the 17th IFAC Symposium on System Identification*, pages 975–980, Beijing, China, October 2015.

A. J. Isaksson, J. Sjöberg, D. Törnqvist, L. Ljung, and M. Kok. Using horizon estimation and nonlinear optimization for grey-box identification. *Journal of Process Control*, 30:69–79, June 2015.

J. Kronander, J. Dahlin, D. Jönsson, M. Kok, T. B. Schön, and J. Unger. Real-time video based lighting using GPU raytracing. In *Proceedings of the 2014 European Signal Processing Conference (EUSIPCO)*, pages 1627–1631, Lisbon, Portugal, September 2014.

M. Kok. *Probabilistic modeling for positioning applications using inertial sensors*. Licentiate’s thesis no. 1656, Linköping University, Linköping, Sweden, June 2014.

N. Wahlström, M. Kok, T. B. Schön, and F. Gustafsson. Modeling magnetic fields using Gaussian processes. In *Proceedings of the 38th International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, pages 3522 – 3526, Vancouver, Canada, May 2013.

2

Probabilistic models

As discussed in Chapter 1, our interest lies in position and orientation estimation using inertial sensors. For general estimation problems, two key questions need to be answered to set up a description of the problem:


What are we interested in? And which information is available?

For the inertial motion capture problem illustrated in Figure 1.9 for instance, we are interested in estimating the relative position and orientation of each of the body segments. The information that is available are the inertial measurements from each of the 17 IMUs. Furthermore, knowledge is available from biomechanical models. For instance, the body segments are known to be connected to each other.

Our answers to these two key questions will guide us when we *model* the relation between the quantities that we are interested in and the information that is available. It is important to realize that models are *simplifications* of reality, which implies that they are never completely true. Since our sensors are not perfect (see Figure 1.5) and since our models are not perfect descriptions of reality, we typically want to combine multiple sources of information. This is illustrated in Example 2.1.


Example 2.1: Estimating orientation using inertial measurements

As described in Chapter 1, the gyroscope measures the angular velocity of the sensor and integration of the measurements provides information about the sensor's orientation. Modeling the accelerometer measurements as measuring only the gravity, its measurements can be used to estimate the inclination of the sensor. In practice, however, the measurements are biased and contain noise, as illustrated in Figure 1.5. We simulate noisy accelerometer and gyroscope measurements, assuming that the sensor is lying still. Note that compared to the data in Figure 1.5,



(a) Inclination estimates using only gyroscope measurements.

(b) Inclination estimates using only accelerometer measurements.



(c) Inclination estimates from combining accelerometer and gyroscope measurements.

Figure 2.1: Estimated inclination using integration of simulated gyroscope measurements (a), by using accelerometer measurements, assuming that the sensor is stationary (b) and by combining the measurements (c). The roll is depicted in black, the pitch in grey.

we have assumed that the measurements do not contain any bias. Furthermore, the noise levels are chosen slightly differently for illustrational purposes.

The inclination estimates obtained by integration of the gyroscope data are shown in Figure 2.1a. Instead of staying around 0° , they drift over time. The inclination estimated from the simulated accelerometer measurements is shown in Figure 2.1b. As can be seen, the orientation estimates are centered around 0° . However, they are quite noisy. We would like to combine the accelerometer and gyroscope measurements to estimate the inclination such that our estimates look as smooth as the ones using the gyroscope data but at the same time do not exhibit any integration drift. An example of our desired outcome is shown in Figure 2.1c.

To effectively combine multiple sources of information, it is beneficial to take the *uncertainty* of the different sources into account. For instance, to obtain Figure 2.1c, we explicitly made use of the knowledge of the noise levels of the (simulated) measurements. This is an important reason for why we are interested in using *probabilistic* models.

We express our models in terms of mathematical relations. For this, we denote all the quantities that we are interested in the *states* x_t or the *parameters* θ . The

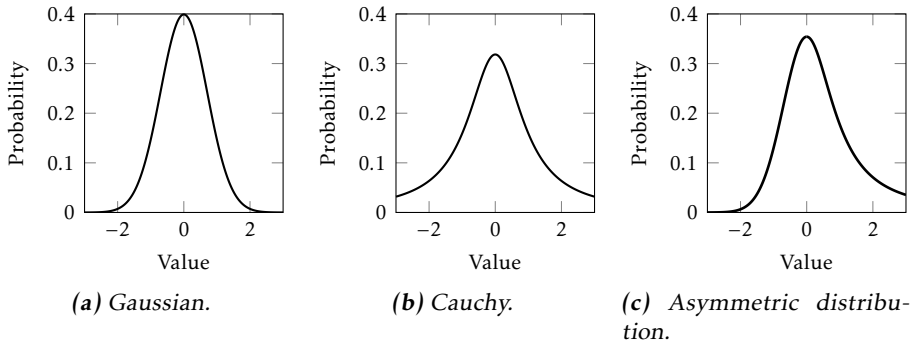


Figure 2.2: A number of probability density functions.

subscript t on x implies that we assume that x changes over time and has value x_t at time t . We model the states to be in discrete time from time $t = 1$ to $t = N$. The set of states at all time steps is denoted $x_{1:N}$. In Example 2.1, the state x_t consists of the inclination of the sensor. The parameters θ do not have a subscript t . With this we explicitly indicate that they are *constant*. We will encounter examples of parameters θ in Chapter 3. We denote the measurements at time t by y_t and the set of all measurements from $t = 1, \dots, N$ by $y_{1:N}$. In Example 2.1, the measurements y_t consist of both the gyroscope and the accelerometer measurements.

To take the uncertainty of the states $x_{1:N}$ and the measurements $y_{1:N}$ into account, we represent both the states and the measurements as *random variables* distributed according to some *probability distribution*. Examples of probability distributions that we encounter throughout this thesis are given in Figure 2.2. The Gaussian distribution shown in Figure 2.2a has a *mean* of zero and a *covariance* of one. This implies that the variable is most likely to have a value around 0. In fact, there is a 68% chance that the random variable is between -1 and $+1$ and a 99.7% chance that it is within -3 and $+3$. A general Gaussian distribution with mean μ and covariance Σ is denoted $\mathcal{N}(\mu, \Sigma)$.

For the Gaussian distribution in Figure 2.2a, the probability of the variable to have a value smaller than -3 or larger than $+3$ is very small. The Cauchy distribution shown in Figure 2.2b on the other hand, assigns a larger probability to values deviating more from zero. A distribution that models the probability of large positive values to be higher than the probability of large negative values is shown in Figure 2.2c.

2.1 Models for position and orientation estimation

In this section we discuss a number of probabilistic models to illustrate the types of models that we use for position and orientation estimation in the papers in Part II. We start with an example of a *dynamic model* in Example 2.2. Dynamic

models are used to describe the relation between the state x_{t+1} and x_t as

$$x_{t+1} | x_t \sim p(x_{t+1} | x_t), \quad (2.1)$$

where $p(\cdot)$ denotes a probability density function. The dynamic model describes the *conditional distribution* of the state x_{t+1} given the state x_t .

Example 2.2: Dynamic model

For almost all applications, we have some knowledge about the *dynamics*. For instance, when estimating the position of a person, it is very unlikely, if not impossible, for the person to be in Linköping, Sweden at a specific time instance, and in Amsterdam, the Netherlands half an hour later. In other words, *conditioned* on the fact that we know that the person is in Linköping at time t , we know something about where the person can be at time $t + 1$.

Since inertial sensors measure the acceleration and the angular velocity of the sensor, they can be used to provide information about the *change* in position and orientation from time t to time $t + 1$. This can be used in a dynamic model. The inertial measurements are both noisy and biased as illustrated in Figure 1.5. Comparing the histograms in Figures 1.5c and 1.5d to the distributions in Figure 2.2, it can be seen that the inertial sensor measurement noise is quite Gaussian with a non-zero mean value (bias) and a covariance that is significantly smaller than one. The presence of Gaussian noise and of a sensor bias can be represented by the probabilistic dynamic model (2.1).

The model discussed in Example 2.2 is used in Papers A – E and Paper G. In some applications, additional knowledge is available about the relation between different parts of the state vector x_t . This can explicitly be modeled in terms of the conditional distribution

$$x_t^a | x_t^b \sim p(x_t^a | x_t^b), \quad (2.2)$$

where x_t^a and x_t^b are subsets of the states x_t . Two examples related to Papers B and C are discussed in Examples 2.3 and 2.4.

Example 2.3: Sensors placed on body segments

To estimate the pose of the human body, sensors can be placed on different body segments, as discussed in Section 1.2.3. It is not possible to place the sensors directly on the bone. Instead, they are placed on the skin and because of the presence of soft tissue, they will move slightly with respect to the bone. It is difficult to model this movement exactly. Instead, we assume that the position and orientation of the sensors on the body segments are constant up to some Gaussian noise.

Example 2.4: Connection of body segments at joints

When sensors are placed on a human body, it is possible to make use of the knowledge that the body segments are connected to each other at the joints. This assumption is actually *exactly* true. Hence, we would like to model this as a *deterministic constraint* instead of using a probability distribution.

Finally, we can model the knowledge provided by the sensor measurements about the states. This can be represented as

$$y_t | x_t \sim p(y_t | x_t), \quad (2.3)$$

i.e. in terms of the conditional distribution of the measurements y_t given the state x_t . Examples 2.5 and 2.6 discuss the inclusion of UWB and magnetometer measurements. UWB measurements are used in Paper D, while magnetometers are used in Papers E – G.

Example 2.5: Ultrawideband measurements

In Section 1.2.2, we discussed the use of TOA measurements from a UWB system in combination with inertial sensors. In practice, a small number of pulses sent by the transmitter to the receivers can be *delayed*. This can be because the pulse did not take the shortest path to the receiver, but instead traveled via for instance the floor or a wall in the building. This is called *multipath*. It can also be because the pulse had to travel through some material other than air to reach the receiver. This is called *non-line-of-sight (NLOS)* and causes a delayed pulse since the speed of light in material is lower than the speed of light in air. The presence of a small number of delayed measurements can be modeled by assuming that the TOA measurements y_t given the state x_t are distributed according to an asymmetric distribution such as the one shown in Figure 2.2c. This distribution allows for measurements to be delayed while not allowing for the possibility of measurements arriving *earlier*, i.e. traveling faster than the speed of light.

Example 2.6: Magnetometer measurements

Magnetometers measure the local magnetic field. This field consists of contributions both from the local earth magnetic field and from the magnetic field due to magnetic material such as metallic structures of buildings and furniture. Because of this, especially in indoor environments, it can vary significantly over different locations in the building. Let us define a function $f(p_t^n)$ that gives the magnetic field at each position p_t^n . The magnetometer measurements $y_{m,t}$ can then be modeled as

$$y_{m,t} = R_t^{\text{bn}} f(p_t^n) + e_{m,t}, \quad (2.4)$$

where $e_{m,t}$ is Gaussian measurement noise. The *rotation matrix* R_t^{bn} rotates the magnetic field from the coordinate frame in which the sensor is localized to the coordinate frame in which the sensor obtains its measurements. Note that we

use slightly different notation here compared to (2.3). A subscript m is added to the measurements y_t to explicitly indicate that we consider magnetometer measurements. Furthermore, the state x_t in this case consists of both the position of the sensor p_t^n and the orientation R_t^{bn} .

When the magnetic field is used for heading information, it is typically assumed that the magnetic field is constant, i.e. that $f(p_t^n)$ is a constant three-dimensional vector. Because of this, local variations of the magnetic field are considered undesired disturbances. On the other hand, it is also possible to make use of the changes in the magnetic field to provide position and orientation information. For this we would like to know the function $f(p_t^n)$. In practice, it is typically hard to obtain $f(p_t^n)$ because a large number of magnetic field sources contribute to the magnetic field, severely complicating the modeling process. However, it is possible to estimate the function $f(p_t^n)$ by *learning a map of the magnetic field*. This can be done by collecting *training data*, which can be used to *predict* the magnetic field at previously unknown locations.

The models discussed in Examples 2.2 – 2.6 can be combined and used for position and orientation estimation, which is the topic of Section 3.2. In Section 2.2 we will first discuss a method to build maps of the magnetic field.

2.2 Maps of the magnetic field

In Example 2.6, we introduced the problem of building maps of the magnetic field. An example of a map of the magnetic field is shown in Figure 1.7. The map is obtained by interpolation and extrapolation of magnetic field measurements at different locations, collected by a small robot. Hence, based on a number of measurements, so-called training data, we learn the local magnetic field. This allows us to predict the magnetic field in previously unobserved locations. In Paper F, we build these maps by assuming that the magnetic field can be modeled as a *Gaussian process (GP)*. GPs are defined by Rasmussen and Williams (2006) as:

Definition 2.7. “A GP is a collection of random variables, any finite number of which have a joint Gaussian distribution.”

Consider the slightly more general notation as compared to Example 2.6 and model the measurements y_t as

$$\begin{aligned} y_t &= f(x_t) + e_t, \\ f(x) &\sim \mathcal{GP}(\mu(x), k(x, x')), \end{aligned} \tag{2.5}$$

where $e_t \sim \mathcal{N}(0, \sigma_n^2)$ and $\mathcal{GP}(\mu(x), k(x, x'))$ denotes a GP with mean $\mu(x)$ and covariance $k(x, x')$. Hence, the magnetic field at different locations x_t is jointly

Gaussian distributed as

$$\begin{pmatrix} f(x_1) \\ \vdots \\ f(x_N) \end{pmatrix} \sim \mathcal{N}(\mu, K), \quad (2.6a)$$

with

$$\mu = \begin{pmatrix} \mu(x_1) \\ \vdots \\ \mu(x_N) \end{pmatrix}, \quad K = \begin{pmatrix} K(x_1, x_1) & \cdots & K(x_1, x_N) \\ \vdots & \ddots & \vdots \\ K(x_N, x_1) & \cdots & K(x_N, x_N) \end{pmatrix}. \quad (2.6b)$$

Using GPs, it is possible to incorporate prior information about the physical properties of the magnetic field. This is illustrated in Example 2.8.

Example 2.8: Encoding prior knowledge

The choice of covariance function $k(x, x')$ can encode prior knowledge about the function $f(x)$. Two well-known covariance functions are the squared exponential covariance function $k_{\text{SE}}(x, x')$ and the exponential covariance function $k_{\text{E}}(x, x')$, defined as

$$k_{\text{SE}}(x, x') = \sigma_f^2 \exp\left(-\frac{\|x-x'\|_2^2}{2\ell^2}\right), \quad (2.7a)$$

$$k_{\text{E}}(x, x') = \sigma_f^2 \exp\left(-\frac{\|x-x'\|_2}{\ell}\right). \quad (2.7b)$$

They model the mutual dependence of $f(x)$ on $f(x')$ in terms of the *hyperparameters* σ_f and ℓ . Figures 2.3a – 2.3c show samples drawn from a GP prior using a squared exponential covariance function with different hyperparameters. As can be seen, the parameter σ_f determines the magnitude and the parameter ℓ the length scale. The hyperparameters σ and ℓ therefore influence the shapes of the functions. They can be learned from data. As shown in Figure 2.3d, the choice of covariance function can incorporate prior knowledge for instance about the smoothness of the function.

In Example 2.8, we assumed that x and y are one-dimensional. For the case of modeling the magnetic field, however, the magnetometer measurements $y_{m,t}$ and the position p_t^n are both three-dimensional vectors. This opens up for additional modeling choices, for instance on how these three components are related. Physical knowledge of magnetic fields is available through Maxwell's equations, see also Griffiths (1999); Jackson (1999). We incorporate this into the GP prior in Wahlström et al. (2013) and in Paper F.

2.3 Visualizing the resulting model structures

In this chapter, we have introduced models for two different estimation problems. The first is estimation of position and orientation, for which we have discussed

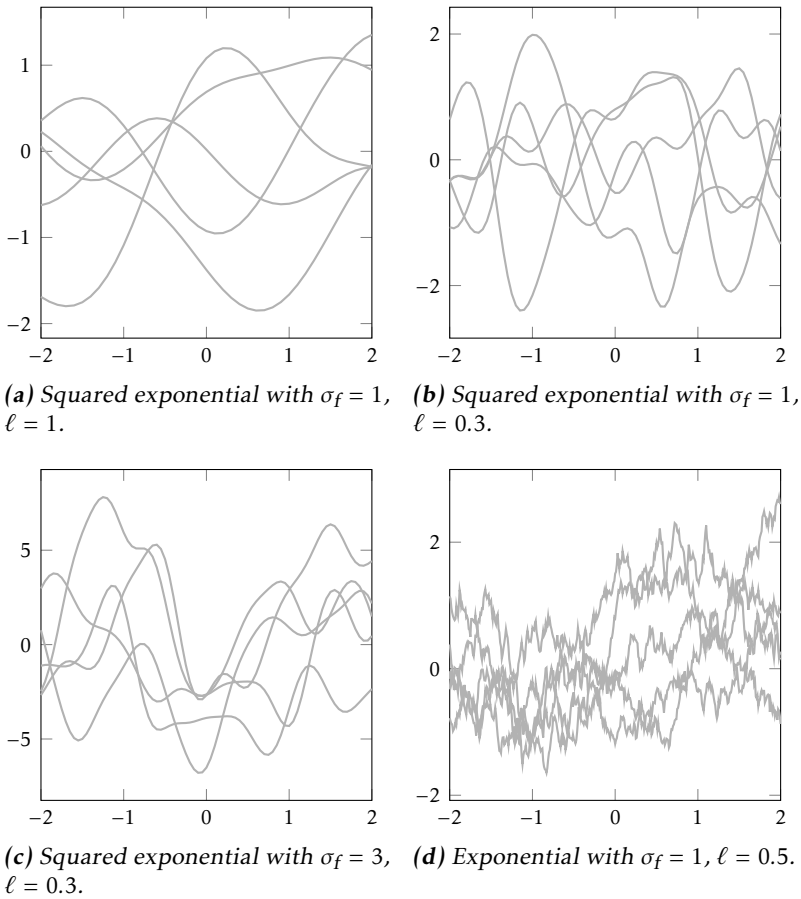


Figure 2.3: Samples from two different GP priors for a number of different hyperparameters.

several models in Section 2.1. The second is estimation of the magnetic field using GPs discussed in Section 2.2. In this section we discuss the structure of these two resulting models and visualize them in terms of their corresponding *graphical models* (Bishop, 2006).

Combining (2.1) and (2.3), the structure of our model for position and orientation estimation is graphically illustrated in Figure 2.4. The state x_{t+1} can be seen to depend on x_t and result in measurements y_{t+1} . Note that x_{t+1} is *conditionally independent* of $x_{1:t-1}$ given the state x_t . This implies that if the sensor's current position, velocity and acceleration are known, it is possible to predict the position and velocity at the next time instance. It is not necessary to know where the sensor has been or how fast it has traveled before arriving in this state. This property of the model is called the *Markov property*. The algorithms used to compute posi-




Figure 2.4: Structure of pose estimation problem where x_{t+1} depends on x_t according to (2.1) and results in measurements y_{t+1} according to (2.3).




Figure 2.5: Structure of the GP model where noisy measurements y_t of the magnetic field $f(x_t)$ at locations x_t are available.

tion and orientation estimates which will be introduced in Chapter 3 rely on this property.

In the GP model (2.5), noisy measurements y_t of the magnetic field $f(x_t)$ at different locations x_t are available. Having a closer look at the covariance matrix in (2.6), it can be seen that the magnetic field at each location x_t depends on the magnetic field at all other locations, i.e. that all components $K(x, x')$ are non-zero. This is graphically illustrated in Figure 2.5. The fact that $f(x_{1:N})$ are all connected to each other, results in a high computational complexity to build a map of the magnetic field if N is large. Because of this, approaches that use GP models to learn function values $f(x)$ often approximate the model in Figure 2.5, see e.g. Quiñonero-Candela and Rasmussen (2005) for an early survey and Chapter 4 of Bijl (2016) for a more recent overview. The simplest approximation would be to simply discard some of the measurements. An alternative approach is used in Paper F.

3

Inference

In Chapter 2, we focused on modeling the quantities that we are interested in and the information that is available. In this chapter, we will focus on the question:

How can we use the available models and measurements to infer knowledge about the quantities we are interested in?

We will use the models on the forms introduced in Chapter 2 to obtain information about the sensor's position and orientation and about the magnetic field. More formally, our aim is to *infer* information about the states $x_{1:N}$ and the parameters θ using the available models and the measurements $y_{1:N}$. For this, we make extensive use of the basic relations of probabilities, (see e.g. Gut (1995); Bishop (2006)),

$$p(a) = \int p(a, b)db, \quad (3.1a)$$

$$p(a, b) = p(a | b)p(b). \quad (3.1b)$$

where $p(a | b)$ denotes the *conditional probability* of a given b and $p(a, b)$ denotes the *joint probability* of a and b .

3.1 Building maps of the magnetic field

Given a data set with measurements $y = \{y_t\}_{t=1}^N$ at locations $x = \{x_t\}_{t=1}^N$, it is possible to infer knowledge about the magnetic field $f(x_*)$ at some new location x_* using *GP regression*. In the remainder, we will use the short-hand notation f_* to denote $f(x_*)$. Using the fact that y and f_* are jointly Gaussian as

$$p(f_*, y | x, x_*) = \mathcal{N} \left(\begin{pmatrix} y \\ f_* \end{pmatrix}; \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} K(x, x) + \sigma_n^2 & K(x, x_*) \\ K(x_*, x) & K(x_*, x_*) \end{pmatrix} \right), \quad (3.2)$$

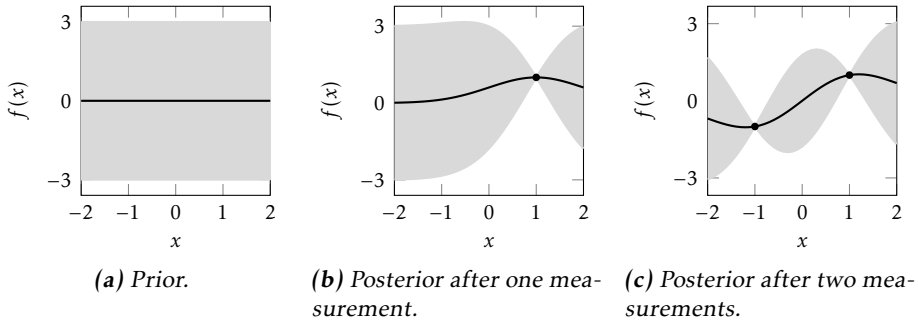


Figure 3.1: Illustration of GP regression where the mean is represented by a black line and the grey area represents the uncertainty.

the conditional distribution $p(f_* | x, y, x_*)$ can be computed. For Gaussian distributions, this conditional distribution can be computed exactly using (3.1b), see e.g. Schön and Lindsten (2011). It results in the Gaussian distribution

$$p(f_* | x, y, x_*) = \mathcal{N}(\mathbb{E}[f_*], \text{cov}[f_*]), \quad (3.3)$$

with

$$\mathbb{E}[f_*] = K(x_*, x) \left(K(x, x) + \sigma_n^2 \mathcal{I}_N \right)^{-1} y, \quad (3.4a)$$

$$\text{cov}[f_*] = K(x_*, x_*) - K(x_*, x) \left(K(x, x) + \sigma_n^2 \mathcal{I}_N \right)^{-1} K(x, x_*). \quad (3.4b)$$

An example of GP regression can be found in Figure 3.1. Before any measurements are observed, the mean of the GP prior is zero and the uncertainty is the same for each position x as illustrated in Figure 3.1a. Observing a measurement y_t at position x_t provides knowledge about the function value $f(x_t)$. This is illustrated in Figures 3.1b and 3.1c. Note that for illustrational purposes we assume that the measurements are noiseless, i.e. $\sigma_n = 0$. Since the function values $f(x)$ at the different locations x depend on each other, the measurements also provide information about the function values at surrounding locations.

3.2 Estimating position and orientation

In this section, we are concerned with inferring knowledge about the states $x_{1:N}$ (containing the position and orientation) from the measurements $y_{1:N}$. This can be expressed in terms of the *joint smoothing distribution*

$$p(x_{1:N} | y_{1:N}). \quad (3.5)$$

In the pose estimation problem, we are typically interested in obtaining *point estimates*, denoted by $\hat{x}_{1:N}$. One way of obtaining these is by solving

$$\begin{aligned}\hat{x}_{1:N} &= \arg \max_{x_{1:N}} p(x_{1:N} | y_{1:N}) \\ &= \arg \max_{x_{1:N}} \frac{p(x_{1:N}, y_{1:N})}{p(y_{1:N})} \\ &= \arg \max_{x_{1:N}} p(x_{1:N}, y_{1:N}),\end{aligned}\tag{3.6}$$

where $\arg \max$ denotes the maximizing argument. The first equality sign in (3.6) follows from (3.1) and the second from the fact that $p(y_{1:N})$ in the denominator does not depend on $x_{1:N}$. The estimate $\hat{x}_{1:N}$ is called the *maximum a posteriori* (MAP) estimate. Hence, the MAP estimate $\hat{x}_{1:N}$ is the pose for which the probability of the states given the measurements is maximized.

In (3.6) we assume that all measurements $y_{1:N}$ are used to obtain the posterior distribution of $x_{1:N}$. Although it makes sense to use all available information to obtain the best estimates, a downside of smoothing is that one needs to wait until all measurements are collected before the pose can be computed. Because of this, in many applications, we are also interested in *filtering*, where the aim is to compute the *filtering distribution*

$$p(x_t | y_{1:t}).\tag{3.7}$$

The MAP estimate can in this case be computed as

$$\hat{x}_t = \arg \max_{x_t} p(x_t | y_{1:t}).\tag{3.8}$$

Using the relations (3.1) in combination with the Markov property discussed in Section 2.3, the full probabilistic model $p(x_{1:N}, y_{1:N})$ in (3.6) can be decomposed as

$$p(x_{1:N}, y_{1:N}) = \underbrace{p(x_1)}_{\text{Prior}} \underbrace{\prod_{t=2}^N p(x_t | x_{t-1})}_{\text{Dynamics}} \underbrace{\prod_{t=1}^N p(y_t | x_t)}_{\text{Measurements}}.\tag{3.9a}$$

Here, $p(x_1)$ is a *prior distribution* over x_1 . The *dynamics* is modeled in terms of $p(x_{t+1} | x_t)$. The distribution $p(y_t | x_t)$ models the information given by the measurements about the state. Note that we have encountered explicit examples of these kinds of models in Chapter 2. Similarly, the filtering distribution can be decomposed as

$$p(x_t | y_{1:t}) = \frac{\overbrace{p(y_t | x_t)}^{\text{Measurements}} \overbrace{p(x_t | y_{1:t-1})}^{\text{Prediction}}}{p(y_t | y_{1:t-1})},\tag{3.9b}$$

where

$$p(x_t | y_{1:t-1}) = \int \underbrace{p(x_t | x_{t-1})}_{\text{Dynamics}} \underbrace{p(x_{t-1} | y_{1:t-1})}_{\text{Filtering distribution}} dx_{t-1}.$$

If our models in Section 2.1 would be linear and Gaussian, closed form expressions for the MAP estimates of the position and orientation of the sensor could be obtained. For filtering, this results in the well-known *Kalman filter* (Kalman, 1960). For smoothing, this results in the linear smoothing equations, see e.g. Särkkä (2013). Estimation of position and orientation, however, is inherently a nonlinear problem due to the nonlinear nature of orientations. In Examples 3.1 – 3.3 we give a number of examples illustrating how the MAP estimates (3.6) and (3.8) can be obtained for this nonlinear case. The choice of which algorithm to use is highly application-dependent. It depends on the properties of the problem, as well as on the computational resources and the desired accuracy.

Example 3.1: Optimization-based smoothing and filtering

The smoothing and filtering problems can be solved using an *optimization-based approach* (Nocedal and Wright, 2006; Boyd and Vandenberghe, 2004). In this approach, the solutions $\hat{x}_{1:N}$ to (3.6) or \hat{x}_t to (3.8) can be found by studying the *shape* of the smoothing or filtering distributions as a function of the states $x_{1:N}$ or x_t . This can be characterized in terms of their *slope* and *curvature*. If n_x is the size of the state vector x_t , the filtering problem solves N problems of size n_x . The smoothing problem instead solves one problem of size $n_x N$. For example, for orientation estimation the size of x_t is 3. Estimating a smoothing solution of the pose of the lower body for 37 seconds as in Paper B, the state vector has 40 284 elements instead. The framework of optimization naturally allows for including equality and inequality constraints.

Example 3.2: Extended Kalman filters

An alternative approach to solving the filtering problem (3.8) is to use an extended Kalman filter (EKF) (Särkkä, 2013; Gustafsson, 2012). EKFs make a linear approximation of the nonlinear models and use the relations (3.1) to approximate the filtering distribution. They are computationally less expensive than the optimization-based methods from Example 3.1. However, they can not straightforwardly include equality and inequality constraints.

Example 3.3: Particle filters

Particle filters (PFs) approximate the distribution (3.7) using a number of samples and their associated weights. For an introduction to particle filtering, see e.g. Doucet and Johansen (2011). PFs are specifically useful in cases when the models are severely nonlinear and a Gaussian assumption on the state x_t is a poor description. An example when this is the case is localization of the robot in the map in Figure 1.7.

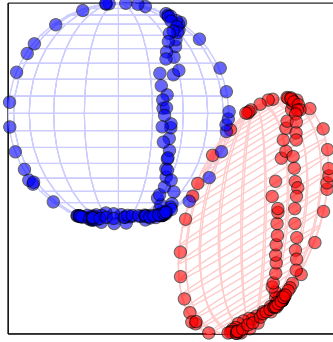


Figure 3.2: Example calibration results with an ellipsoid of magnetometer data before calibration (red) and a sphere of data after calibration (blue).

3.3 Estimating calibration parameters

In Section 3.2, we focused on estimating $x_{1:N}$ from the measurements $y_{1:N}$. We did not consider the presence of unknown parameters θ . Parameter estimation is typically of concern when calibrating the sensors. One approach to estimate constant parameters θ is to straightforwardly include them as additional unknowns in the smoothing optimization problem (3.6), resulting in MAP estimates of the parameters. They can also be included as slowly time-varying parameters in the filtering problem (3.8).

An alternative way of estimating the parameters θ is to solve the *maximum likelihood (ML)* problem defined as

$$\hat{\theta} = \arg \max_{\theta \in \Theta} \mathcal{L}(\theta; y_{1:N}), \quad (3.10)$$

where $\mathcal{L}(\theta; y_{1:N})$ is referred to as the likelihood function, see e.g. Ljung (1999). It is defined as $\mathcal{L}(\theta; y_{1:N}) \triangleq p_{\theta}(Y_{1:N} = y_{1:N})$, where $Y_{1:N}$ are random variables and $y_{1:N}$ are particular realizations of these variables. Hence, $\mathcal{L}(\theta; y_{1:N})$ is a *deterministic* function of a deterministic unknown parameter vector θ . Here, θ is an n_{θ} -dimensional vector which can be limited to a subset Θ of $\mathbb{R}^{n_{\theta}}$. An example of a calibration problem that we encounter in this thesis is calibration of a magnetometer for the presence of magnetic material attached to the sensor, as illustrated in Example 3.4.

Example 3.4: Magnetometer calibration

In Example 2.6, we discussed the use of magnetometers to provide heading information. In that case, the magnetic field at the different locations is assumed to be constant. We will denote this constant magnetic field as m^n . Furthermore, we extend the measurement model (2.4) with calibration parameters D and o as

$$y_{m,t} = DR_t^{\text{bn}} m^n + o + e_{m,t}, \quad (3.11)$$

where $D \in \mathbb{R}^{3 \times 3}$ and $o \in \mathbb{R}^3$. If the sensor would be properly calibrated, rotation of the sensor would lead to a *sphere* of magnetometer data. When magnetic material

is rigidly attached to the sensor, the magnetometer data obtained by rotating the sensor lies on an *ellipsoid* instead. An example based on experimental data is shown in Figure 3.2. Estimating D and o can be interpreted as fitting the red ellipsoid of data to a sphere. Magnetometer calibration is the topic of Paper E.

4

Conclusions and future work

In Chapters 1 – 3 of Part I of this thesis, a background was given to the seven papers that will be presented in Part II. In Chapter 1, an introduction to the sensors and models that we use was given and the contributions of the papers were summarized. Chapter 2 focused on describing the concept of a probabilistic model, illustrated with a few examples of models used in Part II. In Chapter 3, we subsequently discussed how to infer knowledge from these models and the sensor measurements. In the present chapter, we revisit the contributions of this thesis and summarize them in more technical detail. We also discuss possible directions for future work.

4.1 Position and orientation estimation using inertial sensors

Contributions The tutorial in Paper A describes the topic of position and orientation estimation using inertial sensors. Different modeling choices for the dynamics, the measurements and the priors are presented. We introduce smoothing and filtering algorithms solved as optimization problems as well as computationally attractive EKF implementations. The estimates from the different algorithms are compared. Furthermore, some general characteristics are discussed both for the problem of orientation estimation and for the problem of combined position and orientation estimation. Finally, the topic of sensor calibration is discussed and illustrated in terms of the estimation of an unknown gyroscope bias.

Future work In Paper A we discuss four different algorithms for pose estimation. An interesting direction for future work would be to consider also other popular algorithms such as the complementary filter, see e.g. Mahony et al. (2008); Bald-

win et al. (2007), and the filter presented by Madgwick et al. (2011). Furthermore, it would be interesting to have a closer look at using spherical distributions to represent the orientation. These distributions explicitly model the orientation to lie on a manifold. In recent years, a number of approaches have been proposed to estimate the orientation using these kinds of distributions. For instance, in Kurz et al. (2013); Gilitschenski et al. (2016); Glover and Kaelbling (2013), algorithms are presented to estimate orientation by using Bingham distributions.

4.2 Inertial sensor motion capture

Contributions Paper B presents an approach for inertial sensor motion capture to estimate the pose of the human body. We assume that the sensors are more or less rigidly attached to the body and that the body segments are connected to each other. Solving the problem as a smoothing optimization problem allows us to straightforwardly incorporate the connection between the body segments as equality constraints into the problem. The algorithm is applied to experimental data with promising results.

As discussed in Example 3.1, the number of states in the smoothing formulation of the motion capture problem is large. It grows both with time and with the number of sensors and body segments that are considered. However, as illustrated in Figure 2.4, the states x_{t+1} in the motion capture problem are conditionally independent of the states $x_{1:t-1}$ given the states x_t . This structure can be exploited using a technique called message passing as presented in Paper C. This allows for solving the smoothing problem for large data sets. A similar structure can be seen in the human body where the left foot is connected to the left lower leg (and not to any other body segments), the left lower leg is connected to the left upper leg, and so forth. Using this structure, message passing can also be used to solve the problem in a distributed manner using the sensors on the body.

Future work The topic of inertial sensor motion capture is an interesting and large field of study in itself. Because of this, there are many interesting directions of future work for Papers B and C, a few of which will be highlighted here. First of all, Paper B only presents a proof-of-concept, showing that inertial sensor motion capture can indeed be solved using an optimization approach. The algorithm is applied to data from a lower-body configuration with promising results. However, an extensive analysis of the quality of the estimates has not been performed and is an important direction for future work. We also see clear possibilities for future work in the following directions:

- **Relative pose estimation without using magnetometers** Solving the inertial motion capture problem using an optimization-based approach allows us to straightforwardly incorporate equality constraints to model the connections of the different body segments. By incorporating these constraints, the sensors' relative position and orientation become observable as long as the subject is not standing completely still (Hol, 2011). Because of this, it is

not necessary to make use of magnetometer data to provide heading information. The use of magnetometers is known to be problematic specifically in motion capture applications because the magnetic field at the different sensor locations is typically different, see e.g. Ligorio and Sabatini (2016); Seel et al. (2014); El-Gohary and McNames (2015) and the references in Paper B. The claim in Hol (2011) that the relative position and orientation become observable if the subject is not completely standing still, is mainly based on experimental results and on physical insight. It would be an interesting direction for future work to undertake a more rigorous observability study and to derive the necessary and sufficient conditions under which the relative pose is indeed observable.

- **Estimating calibration parameters** In Paper B, we assume that the position and orientation of the sensors on the body segments are known from pre-calibration. It would be an interesting direction of future work to add these calibration parameters as unknowns in the optimization approach. We expect that these parameters will only be identifiable under “sufficient” motion of the body. It would be interesting to derive conditions under which the calibration parameters can be identified from the available data and to study the quality of the parameter estimates as a function of the motion of the body.
- **Online pose estimation** The algorithm derived in Paper B obtains smoothing estimates of the pose. To allow for online estimation, it would be interesting to consider filtering or sliding window approaches. A first step in this direction has been taken in the master thesis project by Lorenz (2016). To run the filtering or sliding window approach in real-time, an efficient implementation is necessary. It would be interesting to see if the message passing algorithm presented in Paper C will be useful for this.
- **Activity recognition** In Paper B we have focused on pose estimation using inertial sensors placed on the human body. A separate but related topic is that of activity recognition using sensors placed on the human body, see e.g. Bulling et al. (2014); Hardegger et al. (2016); Reiss et al. (2010). This field focuses on recognizing the activities that the subject performs. It would be interesting to combine the motion capture approach from Paper B with activity recognition.

A number of these directions of future work have already been addressed in Taetz et al. (2016) and Miezal et al. (2016). In Miezal et al. (2016), an optimization-based solution to the inertial motion capture problem is presented that is inspired by the approach in Paper B. A sliding window approach is used to allow for online estimation. The approach is compared to two different EKF -based approaches, particularly in terms of performance in the presence of calibration errors and dependence on magnetometer usage. The optimization-based approach was shown to result in more accurate pose estimates than the EKF approaches. In Taetz et al. (2016), an optimization-based approach using a sliding window of data was used

to estimate both the human body pose and the calibration parameters, i.e. the position and orientation of the sensors on the body segments. To this end, additional biomechanical models were included in the optimization approach.

4.3 Combining UWB with inertial sensors

Contributions Paper D presents our approach to combine inertial measurements with TOA measurements from a UWB system for indoor positioning. We consider both pose estimation using inertial and UWB measurements, and calibration of the UWB setup as shown in Figure 1.8. By using the asymmetric distribution illustrated in Figure 2.2c, we explicitly model the possibility of delayed measurements due to multipath or NLOS while not allowing for measurements arriving *earlier*, i.e. traveling faster than the speed of light. It is experimentally shown that taking the possibility of delayed measurements into account leads to significantly improved position estimates.

Future work In Paper D we use a tailored asymmetric heavy-tailed distribution to model the possibility of delayed measurements. Related studies have also modeled this presence of delayed measurements using skew- t distributions (Nurminen et al., 2015; Müller et al., 2016) and Gaussian mixture models (Müller et al., 2014). It would be an interesting topic for future work to study how these approaches compare to each other. Another interesting direction of future work is to combine measurements from multiple UWB transmitters and multiple IMUs for human body motion capture. This would be a combination of the work presented in Paper B and that presented in Paper D and would allow us to estimate the absolute pose of the human body.

4.4 Magnetometer calibration

Contributions Paper E presents a magnetometer calibration algorithm which uses inertial sensors to calibrate the magnetometer for the presence of magnetic disturbances attached to the sensor. The algorithm also calibrates for magnetometer sensor errors and for misalignment between the magnetometer and the inertial sensor axes. The calibration algorithm is based on an ML formulation and is shown to give good results using data from two different commercially available IMUs. Using the calibrated magnetometer measurements in combination with the inertial sensors to determine the sensor's orientation is shown to lead to significantly improved heading estimates.

Future work In Paper E we show that our magnetometer calibration algorithm leads to significantly improved heading estimates based on measurements from two different commercially available IMUs. An interesting line of future work is to apply the magnetometer calibration algorithm to inertial and magnetometer

measurements from a smartphone. Smartphones typically use their own magnetometer calibration algorithm, thereby complicating the testing of other calibration algorithms. However, as of Android API level 18 (Jelly Bean MR2), it is possible to log uncalibrated magnetometer data. Hence, it is possible to apply our calibration algorithm to measurements from a smartphone.

The calibration algorithm is formulated as a batch, offline, method. It would be interesting to extend it to an online approach. Using this approach, it might be possible to automatically recalibrate the sensor once it enters a different magnetic environment.

4.5 Mapping and localization using magnetic fields

Contributions Paper F presents an approach to build maps of the indoor magnetic field. Physical knowledge is incorporated by modeling the magnetic field as the gradient of a scalar potential. The magnetic field map is built using GP regression. Since the magnetometers typically sample at 50 or 100 Hz, the amount of data that can be used for building the map of the magnetic field grows quickly over time. As illustrated in Figure 2.5, the GP model assumes that the magnetic field at each location depends on the magnetic field at all other locations. Because of this, GP regression becomes intractable for large amounts of data. To circumvent this issue, we use a computationally efficient implementation using the approach introduced in Solin and Särkkä (2014). Using this approach in combination with the sequential approach introduced in Särkkä et al. (2013) allows for online updating of the magnetic field estimate. It also opens up the possibility to focus on the spatio-temporal problem in which the magnetic field can change over time, for instance due to furniture being moved around.

In Paper G we use the magnetic field as a source of position information. We make use of a magnetic coil which generates a magnetic field that we can model using analytical expressions. Combining magnetometer and inertial measurements, it is possible to estimate the pose of the sensor. Based on experimental results we show that accurate pose estimates can be obtained in the vicinity of the coil, where the magnetic field is strong.

Future work In Paper F we discussed a method to build maps of the magnetic field. Paper G focused on localization in a known map. A natural direction of future work is to consider the problem of simultaneous localization and mapping (SLAM). A first step in this direction has been set by Solin et al. (2016), where a smartphone is localized in a map which is built using the approach presented in Paper F. Using such a SLAM approach, it would be very interesting to analyze what localization accuracy we can achieve. Furthermore, it would be interesting to analyze if this accuracy varies significantly over different buildings because of differences in the amount of magnetic material that is present.

4.6 Concluding remarks

In general, these are interesting times to work with inertial sensors. The quality of their measurements has improved while their cost has decreased, leading to an increase in availability. The fact that they have become so widely available opens up for many exciting possibilities. For instance, a lab for master students has been developed at Linköping University in which students implement their own orientation estimation algorithm using data from a smartphone (Hendeby et al., 2014). Other sensors have also undergone significant developments. For instance, very small devices have been developed which can both act as UWB receiver and transmitter (DecaWave, 2016; BeSpoon, 2016). Most experiments in this thesis have been done with standalone IMUs. In Paper A, however, some experiments have been included using measurements collected with a smartphone. It will be very interesting to see how the measurements from smartphones and from small UWB devices can be used and what the quality of the resulting estimates will be. It will also be interesting to see what new and exciting applications for position and orientation estimation using inertial sensors will open up in the future.

Bibliography

- M. Angermann, M. Frassl, M. Doniec, B. J. Julian, and P. Robertson. Characterization of the indoor magnetic field for applications in localization and mapping. In *Proceedings of the International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, pages 1–9, Sydney, Australia, November 2012.
- G. Baldwin, R. Mahony, J. Trumpf, T. Hamel, and T. Cheviron. Complementary filter design on the special Euclidean group $SE(3)$. In *Proceedings of the IEEE European Control Conference (ECC)*, pages 3763–3770, Kos, Greece, July 2007.
- BeSpoon. <http://spoonphone.com/en/>, 2016. Accessed on November 30.
- H. Bijl. *Gaussian process regression techniques – With applications to wind turbines*. PhD thesis, Delft University of Technology, Delft, the Netherlands, 2016.
- C. M. Bishop. *Pattern recognition and machine learning*. Springer Science + Business Media, 2006.
- S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- A. Bulling, U. Blanke, and B. Schiele. A tutorial on human activity recognition using body-worn inertial sensors. *ACM Computing Surveys (CSUR)*, 46(3):33, 2014.
- DecaWave. <http://www.decawave.com/>, 2016. Accessed on November 30.
- A. Doucet and A. M. Johansen. A tutorial on particle filtering and smoothing: Fifteen years later. In *The Oxford Handbook of Nonlinear Filtering*. Oxford University Press, 2011.
- M. El-Gohary and J. McNames. Human joint angle estimation with inertial sensors and validation with a robot arm. *IEEE Transactions on Biomedical Engineering*, 62(7):1759–1767, 2015.
- M. Frassl, M. Angermann, M. Lichtenstern, P. Robertson, B. J. Julian, and M. Doniec. Magnetic maps of indoor environments for precise localization

- of legged and non-legged locomotion. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 913–920, Tokyo, Japan, November 2013.
- I. Gilitschenski, G. Kurz, S. J. Julier, and U. D. Hanebeck. Unscented orientation estimation based on the Bingham distribution. *IEEE Transactions on Automatic Control*, 61(1):172–177, 2016.
- J. Glover and L. P. Kaelbling. Tracking 3-D rotations with the quaternion Bingham filter. *Computer Science and Artificial Intelligence Laboratory Technical Report*, 2013.
- D. J. Griffiths. *Introduction to electrodynamics*, volume 3. Prentice Hall, New Jersey, 1999.
- F. Gustafsson. *Statistical Sensor Fusion*. Studentlitteratur, 2012.
- A. Gut. *An intermediate course in probability*. Springer-Verlag, New York, 1995.
- M. Hardegger, D. Roggen, A. Calatroni, and G. Tröster. S-SMART: A unified bayesian framework for simultaneous semantic mapping, activity recognition, and tracking. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 7(3):34, 2016.
- G. Hendeby, F. Gustafsson, and N. Wahlström. Teaching sensor fusion and Kalman filtering using a smartphone. In *Proceedings of the 19th World Congress of the International Federation of Automatic Control*, pages 10586–10591, Cape Town, South Africa, August 2014.
- J. D. Hol. *Sensor fusion and calibration of inertial sensors, vision, ultra-wideband and GPS*. Dissertation no. 1368, Linköping University, Linköping, Sweden, June 2011.
- J. D. Hol, F. Dijkstra, H. Luinge, and T. B. Schön. Tightly coupled UWB/IMU pose estimation. In *Proceedings of the IEEE International Conference on Ultra-Wideband (ICUWB)*, pages 688–692, Vancouver, Canada, September 2009.
- J. D. Hol, T. B. Schön, and F. Gustafsson. Ultra-wideband calibration for indoor positioning. In *Proceedings of the IEEE International Conference on Ultra-Wideband (ICUWB)*, pages 1–4, Nanjing, China, September 2010.
- A. J. Isaksson, J. Sjöberg, D. Törnqvist, L. Ljung, and M. Kok. Using horizon estimation and nonlinear optimization for grey-box identification. *Journal of Process Control*, 30:69–79, June 2015.
- J. D. Jackson. *Classical Electrodynamics*. Wiley, New York, 3rd edition, 1999.
- R. E. Kalman. A new approach to linear filtering and prediction problems. *Journal of basic engineering*, 82(1):35–45, 1960.

- S. Khoshfetrat Pakazad, A. Hansson, M. S. Andersen, and I. Nielsen. Distributed primal-dual interior-point methods for solving tree-structured coupled convex problems using message-passing. *Optimization Methods and Software*, pages 1–35, 2016.
- M. Kok. *Probabilistic modeling for positioning applications using inertial sensors*. Licentiate’s thesis no. 1656, Linköping University, Linköping, Sweden, June 2014.
- M. Kok and T. B. Schön. Maximum likelihood calibration of a magnetometer using inertial sensors. In *Proceedings of the 19th World Congress of the International Federation of Automatic Control*, pages 92–97, Cape Town, South Africa, August 2014.
- M. Kok and T. B. Schön. Magnetometer calibration using inertial sensors. *IEEE Sensors Journal*, 16(14):5679 – 5689, 2016.
- M. Kok, J. D. Hol, T. B. Schön, F. Gustafsson, and H. Luinge. Calibration of a magnetometer in combination with inertial sensors. In *Proceedings of the 15th International Conference on Information Fusion*, pages 787–793, Singapore, July 2012.
- M. Kok, N. Wahlström, T. B. Schön, and F. Gustafsson. MEMS-based inertial navigation based on a magnetic field map. In *Proceedings of the 38th International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, pages 6466–6470, Vancouver, Canada, May 2013.
- M. Kok, J. D. Hol, and T. B. Schön. An optimization-based approach to human body motion capture using inertial sensors. In *Proceedings of the 19th World Congress of the International Federation of Automatic Control*, pages 79–85, Cape Town, South Africa, August 2014.
- M. Kok, J. Dahlin, T. B. Schön, and A. Wills. Newton-based maximum likelihood estimation in nonlinear state space models. In *Proceedings of the 17th IFAC Symposium on System Identification*, pages 398–403, Beijing, China, October 2015a.
- M. Kok, J. D. Hol, and T. B. Schön. Indoor positioning using ultrawideband and inertial measurements. *IEEE Transactions on Vehicular Technology*, 64(4):1293–1303, 2015b.
- M. Kok, J. D. Hol, and T. B. Schön. Using inertial sensors for position and orientation estimation. Technical Report LiTH-ISY-R-3093, Department of Electrical Engineering, Linköping University, Linköping, Sweden, December 2016a.
- M. Kok, S. Khoshfetrat Pakazad, T. B. Schön, A. Hansson, and J. D. Hol. A scalable and distributed solution to the inertial motion capture problem. In *Proceedings of the 19th International Conference on Information Fusion*, pages 1348–1355, Heidelberg, Germany, July 2016b.

- J. Kronander, J. Dahlin, D. Jönsson, M. Kok, T. B. Schön, and J. Unger. Real-time video based lighting using GPU raytracing. In *Proceedings of the 2014 European Signal Processing Conference (EUSIPCO)*, pages 1627–1631, Lisbon, Portugal, September 2014.
- G. Kurz, I. Gilitschenski, S. Julier, and U. D. Hanebeck. Recursive estimation of orientation based on the Bingham distribution. In *Proceedings of the 16th International Conference on Information Fusion*, pages 1487–1494, Istanbul, Turkey, July 2013.
- G. Ligorio and A. M. Sabatini. Dealing with magnetic disturbances in human motion capture: A survey of techniques. *Micromachines*, 7(3):43, 2016.
- L. Ljung. *System Identification, Theory for the User*. Prentice Hall PTR, 2nd edition, 1999.
- M. Lorenz. Towards realtime estimation of human motion using inertial sensors without magnetometers – an optimization-based filtering approach. Master thesis, TU Berlin, Germany, 2016.
- S. O. H. Madgwick, A. J. L. Harrison, and R. Vaidyanathan. Estimation of IMU and MARG orientation using a gradient descent algorithm. In *IEEE International Conference on Rehabilitation Robotics*, pages 1–7, Zürich, Switzerland, Jun – Jul 2011.
- R. Mahony, T. Hamel, and J.-M. Pflimlin. Nonlinear complementary filters on the special orthogonal group. *IEEE Transactions on Automatic Control*, 53(5):1203–1218, 2008.
- M. Miezal, B. Taetz, and G. Bleser. On inertial body tracking in the presence of model calibration errors. *Sensors*, 16(7):1132, 2016.
- P. Müller, H. Wymeersch, and R. Piché. UWB positioning with generalized gaussian mixture filters. *IEEE Transactions on Mobile Computing*, 13(10):2406–2414, 2014.
- P. Müller, J. A. del Peral-Rosado, R. Piché, and G. Seco-Granados. Statistical trilateration with skew-t distributed errors in LTE networks. *IEEE Transactions on Wireless Communications*, 15(10):7114–7127, 2016.
- National Centers for Environmental Information. <https://www.ngdc.noaa.gov/geomag/geomag.shtml>, 2016. Accessed on October 10.
- J. Nocedal and S. J. Wright. *Numerical Optimization*. Springer Series in Operations Research, 2nd edition, 2006.
- H. Nurminen, T. Ardeshiri, R. Piché, and F. Gustafsson. Robust inference for state-space models with skewed measurement noise. *IEEE Signal Processing Letters*, 22(11):1898–1902, 2015.

- F. Olsson, M. Kok, K. Halvorsen, and T. B. Schön. Accelerometer calibration using sensor fusion with a gyroscope. In *Proceedings of the IEEE Workshop on Statistical Signal Processing*, pages 660–664, Palma de Mallorca, Spain, June 2016.
- J. Quiñonero-Candela and C. E. Rasmussen. A unifying view of sparse approximate Gaussian process regression. *Journal of Machine Learning Research*, 6: 1939–1959, 2005.
- C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- A. Reiss, G. Hendeby, G. Bleser, and D. Stricker. Activity recognition using biomechanical model based pose estimation. In *Proceedings of the 5th European Conference on Smart Sensing and Context*, pages 42–55, Passau, Germany, November 2010.
- S. Särkkä. *Bayesian Filtering and Smoothing*. Cambridge University Press, 2013.
- S. Särkkä, A. Solin, and J. Hartikainen. Spatiotemporal learning via infinite-dimensional Bayesian filtering and smoothing. *IEEE Signal Processing Magazine*, 30(4):51–61, 2013.
- T. B. Schön and F. Lindsten. Manipulating the multivariate Gaussian density. *Technical Report, Division of Automatic Control, Linköping University, Sweden*, Jan 2011.
- T. Seel, J. Raisch, and T. Schauer. IMU-based joint angle measurement for gait analysis. *Sensors*, 14(4):6891–6909, 2014.
- A. Solin and S. Särkkä. Hilbert space methods for reduced-rank Gaussian process regression. *ArXiv e-prints*, January 2014. arXiv:1401.5508.
- A. Solin, M. Kok, N. Wahlström, T. B. Schön, and S. Särkkä. Modeling and interpolation of the ambient magnetic field by Gaussian processes. *ArXiv e-prints*, September 2015. arXiv:1509.04634.
- A. Solin, S. Särkkä, J. Kannala, and E. Rahtu. Terrain navigation in the magnetic landscape: Particle filtering for indoor positioning. In *Proceedings of the European Navigation Conference*, Helsinki, Finland, May–June 2016.
- A. Svensson, T. B. Schön, and M. Kok. Nonlinear state space smoothing using the conditional particle filter. In *Proceedings of the 17th IFAC Symposium on System Identification*, pages 975–980, Beijing, China, October 2015.
- B. Taetz, G. Bleser, and M. Miezal. Towards self-calibrating inertial body motion capture. In *19th International Conference on Information Fusion*, pages 1751–1759, Heidelberg, Germany, July 2016.
- Trivisio Prototyping GmbH. <http://www.trivisio.com>, 2016. Accessed on November 14.

N. Wahlström, M. Kok, T. B. Schön, and F. Gustafsson. Modeling magnetic fields using Gaussian processes. In *Proceedings of the 38th International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, pages 3522 – 3526, Vancouver, Canada, May 2013.

Xsens Technologies B.V. <http://www.xsens.com>, 2016. Accessed on October 4.

Part II

Publications

Papers

The articles associated with this thesis have been removed for copyright reasons. For more details about these see:

<http://urn.kb.se/resolve?urn=urn:nbn:se:liu:diva-133083>

PhD Dissertations
Division of Automatic Control
Linköping University

- M. Millnert:** Identification and control of systems subject to abrupt changes. Thesis No. 82, 1982. ISBN 91-7372-542-0.
- A. J. M. van Overbeek:** On-line structure selection for the identification of multivariable systems. Thesis No. 86, 1982. ISBN 91-7372-586-2.
- B. Bengtsson:** On some control problems for queues. Thesis No. 87, 1982. ISBN 91-7372-593-5.
- S. Ljung:** Fast algorithms for integral equations and least squares identification problems. Thesis No. 93, 1983. ISBN 91-7372-641-9.
- H. Jonson:** A Newton method for solving non-linear optimal control problems with general constraints. Thesis No. 104, 1983. ISBN 91-7372-718-0.
- E. Trulsson:** Adaptive control based on explicit criterion minimization. Thesis No. 106, 1983. ISBN 91-7372-728-8.
- K. Nordström:** Uncertainty, robustness and sensitivity reduction in the design of single input control systems. Thesis No. 162, 1987. ISBN 91-7870-170-8.
- B. Wahlberg:** On the identification and approximation of linear systems. Thesis No. 163, 1987. ISBN 91-7870-175-9.
- S. Gunnarsson:** Frequency domain aspects of modeling and control in adaptive systems. Thesis No. 194, 1988. ISBN 91-7870-380-8.
- A. Isaksson:** On system identification in one and two dimensions with signal processing applications. Thesis No. 196, 1988. ISBN 91-7870-383-2.
- M. Viberg:** Subspace fitting concepts in sensor array processing. Thesis No. 217, 1989. ISBN 91-7870-529-0.
- K. Forsman:** Constructive commutative algebra in nonlinear control theory. Thesis No. 261, 1991. ISBN 91-7870-827-3.
- F. Gustafsson:** Estimation of discrete parameters in linear systems. Thesis No. 271, 1992. ISBN 91-7870-876-1.
- P. Nagy:** Tools for knowledge-based signal processing with applications to system identification. Thesis No. 280, 1992. ISBN 91-7870-962-8.
- T. Svensson:** Mathematical tools and software for analysis and design of nonlinear control systems. Thesis No. 285, 1992. ISBN 91-7870-989-X.
- S. Andersson:** On dimension reduction in sensor array signal processing. Thesis No. 290, 1992. ISBN 91-7871-015-4.
- H. Hjalmarsson:** Aspects on incomplete modeling in system identification. Thesis No. 298, 1993. ISBN 91-7871-070-7.
- I. Klein:** Automatic synthesis of sequential control schemes. Thesis No. 305, 1993. ISBN 91-7871-090-1.
- J.-E. Strömberg:** A mode switching modelling philosophy. Thesis No. 353, 1994. ISBN 91-7871-430-3.
- K. Wang Chen:** Transformation and symbolic calculations in filtering and control. Thesis No. 361, 1994. ISBN 91-7871-467-2.
- T. McKelvey:** Identification of state-space models from time and frequency data. Thesis No. 380, 1995. ISBN 91-7871-531-8.
- J. Sjöberg:** Non-linear system identification with neural networks. Thesis No. 381, 1995. ISBN 91-7871-534-2.
- R. Germundsson:** Symbolic systems – theory, computation and applications. Thesis No. 389, 1995. ISBN 91-7871-578-4.
- P. Pucar:** Modeling and segmentation using multiple models. Thesis No. 405, 1995. ISBN 91-7871-627-6.

H. Fortell: Algebraic approaches to normal forms and zero dynamics. Thesis No. 407, 1995. ISBN 91-7871-629-2.

A. Helmersson: Methods for robust gain scheduling. Thesis No. 406, 1995. ISBN 91-7871-628-4.

P. Lindskog: Methods, algorithms and tools for system identification based on prior knowledge. Thesis No. 436, 1996. ISBN 91-7871-424-8.

J. Gunnarsson: Symbolic methods and tools for discrete event dynamic systems. Thesis No. 477, 1997. ISBN 91-7871-917-8.

M. Jirstrand: Constructive methods for inequality constraints in control. Thesis No. 527, 1998. ISBN 91-7219-187-2.

U. Forssell: Closed-loop identification: Methods, theory, and applications. Thesis No. 566, 1999. ISBN 91-7219-432-4.

A. Stenman: Model on demand: Algorithms, analysis and applications. Thesis No. 571, 1999. ISBN 91-7219-450-2.

N. Bergman: Recursive Bayesian estimation: Navigation and tracking applications. Thesis No. 579, 1999. ISBN 91-7219-473-1.

K. Edström: Switched bond graphs: Simulation and analysis. Thesis No. 586, 1999. ISBN 91-7219-493-6.

M. Larsson: Behavioral and structural model based approaches to discrete diagnosis. Thesis No. 608, 1999. ISBN 91-7219-615-5.

E. Gunnarsson: Power control in cellular radio systems: Analysis, design and estimation. Thesis No. 623, 2000. ISBN 91-7219-689-0.

V. Einarsson: Model checking methods for mode switching systems. Thesis No. 652, 2000. ISBN 91-7219-836-2.

M. Norrlöf: Iterative learning control: Analysis, design, and experiments. Thesis No. 653, 2000. ISBN 91-7219-837-0.

F. Tjärnström: Variance expressions and model reduction in system identification. Thesis No. 730, 2002. ISBN 91-7373-253-2.

J. Löfberg: Minimax approaches to robust model predictive control. Thesis No. 812, 2003. ISBN 91-7373-622-8.

J. Roll: Local and piecewise affine approaches to system identification. Thesis No. 802, 2003. ISBN 91-7373-608-2.

J. Elbornsson: Analysis, estimation and compensation of mismatch effects in A/D converters. Thesis No. 811, 2003. ISBN 91-7373-621-X.

O. Härkegård: Backstepping and control allocation with applications to flight control. Thesis No. 820, 2003. ISBN 91-7373-647-3.

R. Wallin: Optimization algorithms for system analysis and identification. Thesis No. 919, 2004. ISBN 91-85297-19-4.

D. Lindgren: Projection methods for classification and identification. Thesis No. 915, 2005. ISBN 91-85297-06-2.

R. Karlsson: Particle Filtering for Positioning and Tracking Applications. Thesis No. 924, 2005. ISBN 91-85297-34-8.

J. Jansson: Collision Avoidance Theory with Applications to Automotive Collision Mitigation. Thesis No. 950, 2005. ISBN 91-85299-45-6.

E. Geijer Lundin: Uplink Load in CDMA Cellular Radio Systems. Thesis No. 977, 2005. ISBN 91-85457-49-3.

M. Enqvist: Linear Models of Nonlinear Systems. Thesis No. 985, 2005. ISBN 91-85457-64-7.

T. B. Schön: Estimation of Nonlinear Dynamic Systems — Theory and Applications. Thesis No. 998, 2006. ISBN 91-85497-03-7.

I. Lind: Regressor and Structure Selection — Uses of ANOVA in System Identification. Thesis No. 1012, 2006. ISBN 91-85523-98-4.

J. Gillberg: Frequency Domain Identification of Continuous-Time Systems Reconstruction and Robustness. Thesis No. 1031, 2006. ISBN 91-85523-34-8.

M. Gerdin: Identification and Estimation for Models Described by Differential-Algebraic Equations. Thesis No. 1046, 2006. ISBN 91-85643-87-4.

C. Grönwall: Ground Object Recognition using Laser Radar Data – Geometric Fitting, Performance Analysis, and Applications. Thesis No. 1055, 2006. ISBN 91-85643-53-X.

A. Eidehall: Tracking and threat assessment for automotive collision avoidance. Thesis No. 1066, 2007. ISBN 91-85643-10-6.

F. Eng: Non-Uniform Sampling in Statistical Signal Processing. Thesis No. 1082, 2007. ISBN 978-91-85715-49-7.

E. Wernholt: Multivariable Frequency-Domain Identification of Industrial Robots. Thesis No. 1138, 2007. ISBN 978-91-85895-72-4.

D. Axehill: Integer Quadratic Programming for Control and Communication. Thesis No. 1158, 2008. ISBN 978-91-85523-03-0.

G. Hendeby: Performance and Implementation Aspects of Nonlinear Filtering. Thesis No. 1161, 2008. ISBN 978-91-7393-979-9.

J. Sjöberg: Optimal Control and Model Reduction of Nonlinear DAE Models. Thesis No. 1166, 2008. ISBN 978-91-7393-964-5.

D. Törnqvist: Estimation and Detection with Applications to Navigation. Thesis No. 1216, 2008. ISBN 978-91-7393-785-6.

P-J. Nordlund: Efficient Estimation and Detection Methods for Airborne Applications. Thesis No. 1231, 2008. ISBN 978-91-7393-720-7.

H. Tidfelt: Differential-algebraic equations and matrix-valued singular perturbation. Thesis No. 1292, 2009. ISBN 978-91-7393-479-4.

H. Ohlsson: Regularization for Sparseness and Smoothness — Applications in System Identification and Signal Processing. Thesis No. 1351, 2010. ISBN 978-91-7393-287-5.

S. Moberg: Modeling and Control of Flexible Manipulators. Thesis No. 1349, 2010. ISBN 978-91-7393-289-9.

J. Wallén: Estimation-based iterative learning control. Thesis No. 1358, 2011. ISBN 978-91-7393-255-4.

J. D. Hol: Sensor Fusion and Calibration of Inertial Sensors, Vision, Ultra-Wideband and GPS. Thesis No. 1368, 2011. ISBN 978-91-7393-197-7.

D. Ankelhed: On the Design of Low Order H-infinity Controllers. Thesis No. 1371, 2011. ISBN 978-91-7393-157-1.

C. Lundquist: Sensor Fusion for Automotive Applications. Thesis No. 1409, 2011. ISBN 978-91-7393-023-9.

P. Skoglar: Tracking and Planning for Surveillance Applications. Thesis No. 1432, 2012. ISBN 978-91-7519-941-2.

K. Granström: Extended target tracking using PHD filters. Thesis No. 1476, 2012. ISBN 978-91-7519-796-8.

C. Lyzell: Structural Reformulations in System Identification. Thesis No. 1475, 2012. ISBN 978-91-7519-800-2.

J. Callmer: Autonomous Localization in Unknown Environments. Thesis No. 1520, 2013. ISBN 978-91-7519-620-6.

D. Petersson: A Nonlinear Optimization Approach to H2-Optimal Modeling and Control. Thesis No. 1528, 2013. ISBN 978-91-7519-567-4.

Z. Sjanic: Navigation and Mapping for Aerial Vehicles Based on Inertial and Imaging Sensors. Thesis No. 1533, 2013. ISBN 978-91-7519-553-7.

F. Lindsten: Particle Filters and Markov Chains for Learning of Dynamical Systems. Thesis No. 1530, 2013. ISBN 978-91-7519-559-9.

P. Axelsson: Sensor Fusion and Control Applied to Industrial Manipulators. Thesis No. 1585, 2014. ISBN 978-91-7519-368-7.

A. Carvalho Bittencourt: Modeling and Diagnosis of Friction and Wear in Industrial Robots. Thesis No. 1617, 2014. ISBN 978-91-7519-251-2.

M. Skoglund: Inertial Navigation and Mapping for Autonomous Vehicles. Thesis No. 1623, 2014. ISBN 978-91-7519-233-8.

S. Khoshfetrat Pakazad: Divide and Conquer: Distributed Optimization and Robustness Analysis. Thesis No. 1676, 2015. ISBN 978-91-7519-050-1.

T. Ardehshiri: Analytical Approximations for Bayesian Inference. Thesis No. 1710, 2015. ISBN 978-91-7685-930-8.

N. Wahlström: Modeling of Magnetic Fields and Extended Objects for Localization Applications. Thesis No. 1723, 2015. ISBN 978-91-7685-903-2.

J. Dahlin: Accelerating Monte Carlo methods for Bayesian inference in dynamical models. Thesis No. 1754, 2016. ISBN 978-91-7685-797-7.