Department of Electrical Engineering and Automation

Hilbert Space Projection Methods for Numerical Integration and State Estimation

Juha Sarmavuori

Aalto University

DOCTORAL THESES

Aalto University publication series
DOCTORAL THESES 174/2024

Hilbert Space Projection Methods for Numerical Integration and State Estimation

Juha Sarmavuori

A doctoral thesis completed for the degree of Doctor of Science (Technology) to be defended, with the permission of the Aalto University School of Electrical Engineering, at a public examination held at the lecture hall TU1 of the school on 27 September 2024 at 12.

Aalto University
School of Electrical Engineering
Department of Electrical Engineering and Automation
Sensor Informatics and Medical Technology

Supervising professor
Prof. Simo Särkkä, Aalto University, Finland

Thesis advisor
Prof. Simo Särkkä, Aalto University, Finland

Preliminary examiners
Prof. Jindřich Duník, University of West Bohemia, Czech Republic
Prof. Andreas Rupp, Lappeenranta-Lahti University of Technology, Finland

Opponent
Prof. Jindřich Duník, University of West Bohemia, Czech Republic

Aalto University publication series
DOCTORAL THESES 174/2024

ISBN 978-952-64-1984-8 (printed)
ISBN 978-952-64-1985-5 (pdf)
ISSN 1799-4934 (printed)
ISSN 1799-4942 (pdf)
http://urn.fi/URN:ISBN:978-952-64-1985-5

Unigrafia Oy
Helsinki 2024

Finland

NORDIC SWAN ECOLABEL
Printed matter
4041-0619

Aalto University

Abstract

Aalto University, P.O. Box 11000, FI-00076 Aalto www.aalto.fi

Author
Juha Sarmavuori

Name of the doctoral thesis
Hilbert Space Projection Methods for Numerical Integration and State Estimation

Publisher School of Electrical Engineering

Unit Department of Electrical Engineering and Automation

Series Aalto University publication series DOCTORAL THESES 174/2024

Field of research Automation, Systems and Control Engineering

Manuscript submitted 26 April 2024 **Date of the defence** 27 September 2024

Permission for public defence granted (date) 24 May 2024 **Language** English

☐ **Monograph** ☒ **Article thesis** ☐ **Essay thesis**

Abstract

The aim of this thesis is to develop Hilbert space methods for approximation of integrals appearing in filtering and smoothing of nonlinear state-space models. State-space models have many applications in real-world problems and have been studied extensively for almost a century. In filtering, the state is estimated at a given time instant based on measurements up to the given time instant. In smoothing, measurements after the given time instant are used as well. The used state-space models are stochastic and hence need to be estimated in probabilistic terms, which requires solving probability integrals. We consider two kinds of state-space models: discrete-time and continuous-discrete-time ones. In the latter case, the dynamics model is continuous time and the measurements are obtained in discrete time instants.

In linear state-space models with additive Gaussian noise, closed-form solutions are known for both filtering and smoothing problems. In a nonlinear case, we can use Gaussian approximations, which means that we approximate the probability distributions with Gaussian distributions. We study how to use Fourier–Hermite series for smoothing and filtering with Gaussian approximations. For computing terms of the Fourier–Hermite series, we develop a new method that uses partial differentials of a Weierstrass transform of a nonlinear function.

Even with the simplifying Gaussian approximation, in general, we cannot solve the resulting Gaussian integrals in closed form, but we need numerical approximations instead. We develop a new numerical integration method based on an approximation of a multiplication operator with a finite matrix, and it is not only applicable to Gaussian integrals but can be used for more general numerical integration. This new numerical integration method generalises Gaussian quadrature and has many similar properties, which are analysed using the theory of linear operators in Hilbert space. Specifically, we prove convergence for a large class of functions.

In the case of independent variables, it is possible to compute multidimensional integrals by product rule of unidimensional numerical integrals. With the new numerical integration method, we can generalise the product rule for non-independent variables. We apply this generalised product rule to filtering with arbitrary order moments.

Keywords Fourier–Hermite series, Numerical integration as a multiplication operator, stochastic filtering and smoothing, Hilbert space

ISBN (printed) 978-952-64-1984-8 **ISBN (pdf)** 978-952-64-1985-5

ISSN (printed) 1799-4934 **ISSN (pdf)** 1799-4942

Location of publisher Helsinki **Location of printing** Helsinki **Year** 2024

Pages 161 **urn** http://urn.fi/URN:ISBN:978-952-64-1985-5

Aalto-yliopisto

Tiivistelmä

Aalto-yliopisto, PL 11000, 00076 Aalto www.aalto.fi

Tekijä
Juha Sarmavuori

Väitöskirjan nimi
Hilbertin avaruuden projektiomenetelmät numeerisessa integroinnissa ja tilaestimoinnissa

Julkaisija Sähkötekniikan korkeakoulu

Yksikkö Sähkötekniikan ja automaation laitos

Sarja Aalto University publication series DOCTORAL THESES 174/2024

Tutkimusala Automaatio, systeemit ja säätötekniikka

Käsikirjoituksen pvm 26.04.2024 **Väitöspäivä** 27.09.2024

Väittelyluvan myöntämispäivä 24.05.2024 **Kieli** Englanti

☐ **Monografia** ☒ **Artikkeliväitöskirja** ☐ **Esseeväitöskirja**

Tiivistelmä

Tässä väitöstyössä on kehitetty Hilbertin avaruus -menetelmiä epälineaaristen tila-avaruusmallien suodatuksessa ja silotuksessa tarvittavien integraalien likimääräiseen laskentaan. Tila-avaruusmall voidaan soveltaa moniin reaalimaailman ongelmiin, ja niiden tutkimushistoria on miltei vuosisadan mittainen. Suodatuksessa tila tietyllä ajan hetkellä estimoidaan siihen hetkeen mennessä saatujen mittausten perusteella. Silotuksessa käytetään lisäksi myöhemmin saatuja mittauksia. Tässä työssä käsiteltävät tila-avaruusmallit ovat stokastisia, minkä vuoksi niitä voidaan estimoida ainoastaan todennäköisyyksin, mikä vuorostaan edellyttää todennäköisyysintegraalien ratkaisemista. Työsssä käsitellään kahdenlaisia tila-avaruusmalleja: diskreettiaikaisia ja jatkuva-diskreettiaikaisia. Jälkimmäisessä dynaaminen malli on jatkuva-aikainen, ja mittaustulokset ovat diskreettiaikaisia.

Lineaariselle tila-avaruusmallille, jossa kohina on additiivista ja normaalijakautunutta, suodatus ja silotusongelmille tunnetaan valmiit suljetun muodon ratkaisut. Epälineaarisessa tapauksessa voidaan käyttää Gaussisia approksimaatioita, mikä tarkoittaa todennäköisyysjakaumien approksimointia normaalijakaumilla. Tässä työssä tutkitaan, kuinka Fourier–Hermite-sarjaa voidaan käyttää suodatuksen ja silotuksen Gaussisissa approksimaatioissa. Lisäksi kehitetään menetelmä t sarjan termien laskentaan epälineaarisen funktion Weierstrass-muunnoksen osittaisderivaatoista.

Gaussisen approksimaation suomasta yksinkertaistuksesta huolimatta yleisessä tapauksessa vaadittavia Gaussisia integraaleja ei voida ratkaista suljetussa muodossa, vaan ratkaisuihin tarvitaan vielä lisäksi likimääräistä numeerista integrointia. Tässä työssä kehitetään myös uusi numeerinen integrointimenetelmä, joka perustuu kertolaskuoperaattorin likiarvoon äärellisenä matriisina, eikä rajoitu pelkästään Gaussisiin integraaleihin, vaan soveltuu myös muihin yleisempiin integraaleihin. Tämä uusi numeerinen integrointimenetelmä on yleistys Gaussin kvadratuurista ja sillä on monia samanlaisia ominaisuuksia, joita analysoidaan Hilbertin avaruuden lineaaristen operaattoreiden teorian avulla. Erityisesti uuden menetelmän suppeneminen todistetaan useille erilaisille funktioille.

Riippumattomille muuttujille voidaan laskea numeerinen integraali yksiulotteisten muuttujien numeeristen integrointisääntöjen tulona. Uudella numeerisen integroinnin menetelmällä tämä tulosääntö voidaan yleistää myös toisistaan riippuvien muuttujien tapauksessa. Uutta numeerista integrointimenetelmää sovelletaan suodatukseen käyttäen rajoittamattoman kertaluvun momentteja.

Avainsanat Fourier–Hermite-sarja, Numeerinen integrointi kertolaskuoperaattorilla, stokastinen suodatus ja silotus, Hilbertin avaruus

ISBN (painettu) 978-952-64-1984-8 **ISBN (pdf)** 978-952-64-1985-5

ISSN (painettu) 1799-4934 **ISSN (pdf)** 1799-4942

Julkaisupaikka Helsinki **Painopaikka** Helsinki **Vuosi** 2024

Sivumäärä 161 **urn** http://urn.fi/URN:ISBN:978-952-64-1985-5

Preface

The research work in this thesis has been carried out mainly in the Department of Electrical Engineering and Automation, Aalto University, during the years 2017–2024 while also working 80 % at Nokia. Before that, the work started already in the year 2009, in co-operation with Prof. Simo Särkkä. My doctoral research work at Aalto University has been financially supported by Academy of Finland.

My deepest gratitude goes to my supervisor, Prof. Simo Särkkä, whom I can also thank for almost three decades of good friendship. This whole research started from a lunch discussion during Simo's first filtering and smoothing course at Aalto University back in the year 2009. I am also grateful to other professors Olli Simula and Jouko Lampinen who acted as my supervisors before Simo started his professorship. I also learned a lot about stochastic processes at the lectures of late Prof. Esko Valkeila in the years 2009–2010.

During my part-time years at Aalto University, I have shared an office room on Mondays with numerous people whom I like to thank for friendship, at least: Prof. Roland Hostettler, Janne Myllärinen, Dr. Marco Soldati, Dr. Zheng Zhao, Ajinkya Gorad, Dr. Masay Murata, Casian Iacob, Xiaofeng Ma, Otto Kangasmaa, and Adrien Corenflos. I am especially thankful to Zheng for his friendship and interesting collaboration after his move to Sweden and Ajinkya for checking the language of this thesis and many lively conversations.

I also like to thank other colleagues who I have had the pleasure of getting to know at Aalto, especially Dr Matti Raitoharju for collaboration on a still unpublished paper, and Dr. Toni Karvonen for feedback and assistance during the research on Publication I. Other colleagues to mention are Dr Lauri Palva, Prof. Rui Gao, Prof. Arno Solin, Prof. Leo Kärkkäinen, Prof. Filip Tronarp, Dr Jouni Hartikainen, Dr Zenith Purisha, Dr Joel Jaskari, Dr Sakira Hassan, Fatemeh Yaghoobi, Dr Abubakar Yamin, Zaeed Khan, Dr Christos Merkatas, Prof. Muhammad Fuady Emzir, Prof. Ali Bahrami Rad, Prof. Ángel García-Fernández, Dr Kundan Kumar, Dr Hany Abdulsamad, Sahel Iqbal, Elhadi Gasmi, Mahdi Nasiri, Yvann Le Fay, Cristian

Galvis-Florez, Kimmo Suotsalo, Dr Ahmad Farooq, Dr Muhammad Iqbal, Hassan Razavi, Shreeram Murali, Dr Marcin Mińkowski, Hajiba Legrara, Dr Fahime Seyedheydari, Prof. Ivan Vujaklija, Dennis Yeung, Wendy Lam, Prof. Ilkka Laakso, and Dr Marko Mikkonen.

During this research, I have mostly been working at Nokia where I, of course, have far too many friends and colleagues to mention here. Thank you all for your great friendship over the years! A special thanks goes to my many line managers over the years Topi Rantalainen, Matti Lehtimäki, Dr. Matti Rintamäki, Henna Koskenniemi, and Jyri Suvanen who approved and allowed to continue the special arrangement of working one day per week at Aalto University. Already in the year 2005 before starting the research, my line manager at the time Teemu Himanen encouraged me to start the graduate studies and pointed out that at that time Nokia allowed two hours of work time per week to be used for graduate studies.

Of course, I am also very grateful to my family members: my mother Katri Karasma, late father Osmo Sarmavuori, and little brother Mikael Sarmavuori. My parents have always been very encouraging towards my education for which I am very grateful and my father's passing away during the end of this project was a great source of sorrow for me. Other two close relatives to thank are my odd parents who also regretfully passed away during this research: my aunt Marja-Liisa Kuittinen and her son, my cousin Jukka Kuittinen. Finally, last but not least, very special thanks go to my dear wife Terhi Tervo.

Espoo, August 26, 2024,

Juha Sarmavuori

Contents

List of Publications

This thesis consists of an overview and of the following publications which are referred to in the text by their Roman numerals.

I Juha Sarmavuori and Simo Särkkä. Numerical integration as a finite matrix approximation to multiplication operator. *Journal of Computational and Applied Mathematics*, 353: 283–291, June 2019.

II Juha Sarmavuori and Simo Särkkä. On the convergence of numerical integration as a finite matrix approximation to multiplication operator. *Calcolo*, 60:22, April 2023.

III Juha Sarmavuori and Simo Särkkä. Fourier–Hermite Kalman filter. *IEEE Transactions on Automatic Control*, 57(6): 1511–1515, June 2012.

IV Juha Sarmavuori and Simo Särkkä. Fourier–Hermite Rauch–Tung–Striebel smoother. In *Proceedings of EUSIPCO 2012*, Bucharest, Romania, 2109–2113, August 2012.

V Zheng Zhao and Juha Sarmavuori. Stochastic filtering with moment representation. Submitted to *SIAM Journal on Scientific Computing*, March 2023.

Author's Contribution

Publication I: "Numerical integration as a finite matrix approximation to multiplication operator"

Juha Sarmavuori wrote most of the article, and Simo Särkkä helped to generalise the method, and gave useful comments.

Publication II: "On the convergence of numerical integration as a finite matrix approximation to multiplication operator"

Juha Sarmavuori wrote most of the article. Simo Särkkä helped in writing and gave useful comments.

Publication III: "Fourier–Hermite Kalman filter"

Juha Sarmavuori wrote most of the article with many comments and revisions from Simo Särkkä, who also designed and implemented the numerical example.

Publication IV: "Fourier–Hermite Rauch–Tung–Striebel smoother"

Juha Sarmavuori wrote most of the article with many comments and revisions from Simo Särkkä, who also designed and implemented the numerical example.

Publication V: "Stochastic filtering with moment representation"

Zheng Zhao came up with the idea of the paper, did all the experiments, and wrote the initial draft. Juha Sarmavuori developed the moment quadrature methods and demonstrated them in Matlab. Juha Sarmavuori proved the convergence of the moment quadrature, and Zheng Zhao proved the convergence of the moment filter.

Abbreviations

FHKF	Fourier–Hermite Kalman filter
FHRTSS	Fourier–Hermite Rauch–Tung–Striebel smoother
KF	Kalman filter
MC	Monte Carlo
MF	Moment filter
ODE	Ordinary differential equation
PF	Particle filter
RTSS	Rauch–Tung–Striebel smoother
SDE	Stochastic differential equation
SSM	State-space model
TME	Taylor moment expansion

Symbols

$[a, b)$	Interval between real numbers a and b including a and excluding b
$[\boldsymbol{A}]_{i,j}$	Component i, j of a matrix $\boldsymbol{A}$
$\boldsymbol{A}^\top$	Matrix transpose
A^*	Adjoint operator, $\overline{\boldsymbol{A}}^\top$ for matrices
$\mathscr{D}(\mathsf{A})$	Domain of the operator A
$\mathsf{E}(t)$	Spectral family of an operator
$\boldsymbol{E}(t)$	Spectral family of a matrix
$\mathrm{E}[x]$	Expectation of a random variable x
$\mathrm{E}[x \mid y]$	Conditional expectation of x given y
$\mathrm{E}[\boldsymbol{x} \mid \boldsymbol{\mu}, \boldsymbol{\Sigma}]$	Expectation of $\boldsymbol{x} \sim N(\boldsymbol{x} \mid \boldsymbol{\mu}, \boldsymbol{\Sigma})$
$\boldsymbol{e}_i$	Unit vector at direction of orthogonal component i
f	Function, usually an outside function $f : \mathbb{R} \mapsto \mathbb{R}$
$\hat{f}$	Weierstrass transform of a function f
g	Function, usually an inside function $g : \Omega \mapsto \mathbb{R}$
$\mathscr{H}$	Hilbert space
H_k	Hermite polynomial of order k
$\mathrm{id}(\cdot)$	Identity function, $\mathrm{id}(\boldsymbol{x}) = \boldsymbol{x}$
$\mathrm{id}_k(\cdot)$	Scalar function that is component k of a vector argument of the function, $\mathrm{id}_k(\boldsymbol{x}) = [\boldsymbol{x}]_k$
$\boldsymbol{I}$	Identity matrix
I_k	Interval in dimension k in multidimensional space
$\inf$	Infimum

$\mathscr{L}^2_w$	Space of square-integrable functions with respect to the weight function w
ℓ^2	Space of square-summable sequences
$\mathsf{M}[g]$	Multiplication operator of the function g
$\boldsymbol{M}_\infty[g]$	Infinite matrix representation of the multiplication operator of the function g
$\boldsymbol{M}_n[g]$	Finite approximation of the infinite matrix representation of the multiplication operator of the function g
$\mathbb{N}$	Natural numbers
$\mathbb{N}_0$	Natural numbers starting from 0
$\mathbb{N}_0^d$	Set of multi-indices of dimension d
$N(\boldsymbol{x} \mid \boldsymbol{\mu}, \boldsymbol{\Sigma})$	Probability density of a Gaussian random variable $\boldsymbol{x}$ with a mean $\boldsymbol{\mu}$ and a covariance matrix $\boldsymbol{\Sigma}$
$P(A)$	Probability of an event A
$P(A \mid B)$	Probability of an event A conditional on an event B
$p(x)$	Probability density of a random variable x
$p(x \mid y)$	Probability density of a random variable x conditional on a random variable y
$\mathbb{R}$	Real numbers
$\mathbb{R}^n$	n-dimensional real space
sup	Supremum
$\xrightarrow{s}$	Strong convergence
$\xrightarrow{srs}$	Strong resolvent convergence
$\boldsymbol{u}_i$	Unit-length eigenvector i of a matrix
$[\boldsymbol{v}]_k$	Component k of a vector $\boldsymbol{v}$
$w(\cdot)$	Weight function for integral
$\boldsymbol{w}(t)$	Wiener process
$\xrightarrow{w}$	Weak convergence
x	Scalar variable
$\boldsymbol{x}$	Vector variable
$\boldsymbol{x} \sim p(\boldsymbol{x})$	Random variable $\boldsymbol{x}$ follows the probability density $p(\boldsymbol{x})$
$\boldsymbol{y}_{1:k}$	Values $\boldsymbol{y}_1, \boldsymbol{y}_2, \ldots, \boldsymbol{y}_k$
$\overline{z}$	Complex conjugate, $\overline{x + y\,i} = x - y\,i$

λ_i	Eigenvalue i of a matrix
$\sigma(\mathsf{A})$	Spectrum of the operator A
Ω	Domain of integration or sample space for probability measures
$\langle \cdot, \cdot \rangle$	Inner product
$\| \cdot \|$	Norm $\sqrt{\langle \cdot, \cdot \rangle}$
$\lvert \cdot \rvert$	Absolute value of a scalar or the order of a multi-index

1. Introduction

The topic of this thesis is nonlinear filtering and smoothing, and Hilbert space methods for approximation of the probability integrals in filtering and smoothing equations. In a linear case with additive Gaussian noise, closed-form solutions are known for the filtering and smoothing equations (Jazwinski, 2007; Särkkä and Svensson, 2023; Särkkä and Solin, 2019). In a nonlinear case, we often use Gaussian approximations which means that we approximate probability distributions with Gaussian ones. Even with this simplifying approximation, in general, the resulting integrals cannot be solved in closed form. Instead, we need numerical approximations. For this purpose, we develop a new numerical integration method in this thesis. The new method is based on the approximation of multiplication operators with finite matrices, and is not limited to Gaussian integrals.

Filtering and smoothing have applications in many fields, for example, communications (Rigatos, 2013), medicine (Hostettler et al., 2018), biological processes (Mansouri et al., 2014), brain imaging (Särkkä et al., 2012), finance (Date and Ponomareva, 2010), automatic control (Glad and Ljung, 2018), air quality (Agudelo et al., 2011), weather models (Sivagami et al., 2019), global temperature models (Dubinkina and Goosse, 2013), data analysis (Sitz et al., 2011), astrophysics (Giuliani et al., 2006), surveillance (Suliman et al., 2010), market analysis (Xie et al., 1997), machine learning (Yegenoglu et al., 2020), self-driving cars (Kianfar et al., 2012), positioning (Särkkä et al., 2007), maritime (Gorad et al., 2023), underwater target tracking (Kumar et al., 2022), direct current motors (Rigatos, 2009), combustion engines (Helm et al., 2012), speech recognition (Deshmukh, 2020), grammar (Kupiec, 1992), and washing machines (Nak et al., 2015), just to mention a few.

The terminology filtering and smoothing originates from Wiener (1950). Filtering means the estimation of the state at a given time based on observations up to the given time. Smoothing means the estimation of the state at a given time based on observations after the given time as well. Filtering is the method of choice for real-time applications where we want to know the state of the system as accurately as possible based

on the instantly available data for immediate control decisions. In offline time series analysis, smoothing gives a more accurate estimate because also future data, relative to the given time, is used for the estimate.

The original formulation of filtering and smoothing was restricted to stationary processes. Later, the concepts were extended to non-stationary stochastic processes with Markov property, which means that the probability for the state evolution depends only on the current state and not the state history, first for linear systems and linear observation model (Kalman, 1960; Bryson and Frazier, 1963; Rauch et al., 1965) and later for nonlinear ones too (Ho and Lee, 1964; Kushner, 1964; Zakai, 1969; Leondes et al., 1970). In the nonlinear case, a closed-form solution is known for special cases like discrete finite random variables (hidden Markov model) (Baum and Petrie, 1966) and Beneš system (Beneš, 1981).

For more general nonlinear cases, we must use approximations. The best-known and historically first approximations use Jacobian matrices for linearisation of the nonlinear functions (Smith et al., 1962; Bryson and Frazier, 1963). In the later developed statistical linearisation, the linearisation is selected as the optimal linearisation of the functions with respect to a given probability distribution (Gelb, 1974). The closed-form solution of the statistical linearisation is also not known in general. Several numerical methods approximate the statistical linearisation or can be interpreted to approximate it, such as the unscented transform (Julier et al., 2000) and numerical integration of a Gaussian approximation (Wu et al., 2006; Wang and Li, 2021).

In Publications III and IV, we generalise the statistical linearisation in Gaussian approximation up to arbitrary order with Fourier–Hermite series in the Gaussian Hilbert space of functions. The concept of Fourier series is generalised to an infinite matrix representation of multiplication operator, which leads to the new numerical integration method of Publications I and II that is used for non-Gaussian filtering in Publication V.

This thesis consists of five publications and an overview of the background for them. Chapter 2 presents the theoretical background of numerical integration as a matrix function, which is the topic of Publications I–II and used in Publication V. Part of this Hilbert space theory is also used in Publications III and IV. Chapter 3 presents the basics of state-space models (SSM), filtering and smoothing algorithms for state estimation, and the new contributions of Publications III–V to this field.

2. Numerical Integration as a Matrix Function

In this chapter, we show how to transform an integration problem into a linear algebra problem. Before the exact formulation of the actual integration problem to solve, we give an overview of the related operator theory and Gaussian quadrature literature. Then, we explain the relevant concepts of the operator theory such as convergence and spectral theorem.

2.1 Background

The mathematical background of numerical integration as a matrix function is in the operator theory in Hilbert spaces. There are many good textbooks on this topic starting already from the 1930s (Stone, 1932; Akhiezer and Glazman, 1993; Kato, 1995; Reed and Simon, 1981; Kreyszig, 1989; Segal and Kunze, 1978; Weidmann, 1980; Schmüdgen, 2012; Simon, 2015). The earliest two books (Stone, 1932; Akhiezer and Glazman, 1993) also cover how operators, including unbounded ones, have an infinite matrix representation in separable Hilbert space. Later textbooks often omit this topic due to its limitations for the development of more general theory. It is, however, an important bridge in approximating linear operators with finite matrices.

One important special case of the presented numerical integration method is Gaussian quadrature, one of the oldest classical methods for numerical integration. Some good textbooks on Gaussian quadrature are Davis and Rabinowitz (1984); Gautschi (2004); Golub and Meurant (2009). Although, the name could suggest that the Gaussian quadrature is somehow tied or limited to the Gaussian probability distribution used as an approximation in filtering and smoothing, that is not the case. There are several Gaussian quadrature rules for several weight functions. The special Gaussian quadrature for integrals over the Gaussian probability distribution is called Gauss–Hermite quadrature.

We can also recognise the finite matrix approximation of a multiplication operator already in the famous Golub and Welsch algorithm for computing

the Gaussian quadrature nodes and weights (Golub and Welsch, 1969). However, their derivation did not use infinite dimensional operator theory. Their source for using the Jacobi matrix was (Wilf, 1962, Chapter 2, Theorem 5), which is based on finite dimensional matrix algebra. The use of operator theory for the derivation of quadrature rules started later, and the focus has been on Gaussian quadratures on the unit circle of the complex plane (Gragg, 1993; Simon, 2009; Velázquez, 2008; Bultheel et al., 2015).

In the case of the real line, as in this thesis and the publications covered here, there have been some operator theoretic results for the Gaussian quadrature as well (Akhiezer, 1965; Simon, 2008), although the main focus in these works has been for the more theoretic moment problem. The operator theory based algebraic integration theory (Segal, 1965; Segal and Kunze, 1978) has been used mainly for extending the concept of integration to new types of more exotic measures instead of numerical integration of more common ordinary integrals.

2.2 Integration Problem Formulation

The following formulation is used as the starting point for numerical integration in Publications I and II. We are interested in computing an integral of the form

$$\int_\Omega f(g(\boldsymbol{x}))w(\boldsymbol{x})d\boldsymbol{x}, \tag{2.1}$$

where $\Omega \subset \mathbb{R}^d$, $g\colon \Omega \mapsto \mathbb{R}$ is the inside function, $f\colon \mathbb{R} \mapsto \mathbb{R}$ the outside function, and the weight function $w\colon \Omega \mapsto [0,\infty)$ is such that $\int_\Omega w(\boldsymbol{x})d\boldsymbol{x} = 1$. Function g is measurable with respect to weight function w and f is measurable with respect to measure $\mu(f) = \int_\Omega f(g(\boldsymbol{x}))w(\boldsymbol{x})d\boldsymbol{x}$. We define an inner product $\langle\cdot,\cdot\rangle$ based on the weight function w as

$$\langle\phi,\psi\rangle = \int_\Omega \overline{\phi(\boldsymbol{x})}\psi(\boldsymbol{x})w(\boldsymbol{x})d\boldsymbol{x}, \tag{2.2}$$

and a norm $\|\phi\| = \sqrt{\langle\phi,\phi\rangle}$. The function $\overline{\phi}$ is the complex conjugate of the function ϕ. The space of square-integrable functions with respect to the weight function w is

$$\mathscr{L}^2_w = \{\phi : \|\phi\| < \infty\}. \tag{2.3}$$

The space $\mathscr{L}^2_w$ is a complete inner product space, that is, a Hilbert space (Weidmann, 1980, Section 2.1, Examples 5 and 6).

We define a multiplication operator $\mathsf{M}[g]$ of the function g so that a function $\psi = \mathsf{M}[g]\phi$ is determined almost everywhere pointwise as

$$\psi(\boldsymbol{x}) = g(\boldsymbol{x})\phi(\boldsymbol{x}). \tag{2.4}$$

Almost everywhere means that equality may not hold for a point set with 0 measure over the integral weighted with w. With the definition of the

multiplication operator, we can rewrite integral (2.1) as

$$\int_\Omega f(g(\boldsymbol{x}))w(\boldsymbol{x})d\boldsymbol{x} = \langle 1, \mathsf{M}[f(g)]\,1\rangle = \langle 1, f(\mathsf{M}[g])\,1\rangle,$$

where the operator function of a multiplication operator $f(\mathsf{M}[g])$ is defined through the pointwise relation (2.4) as $(f(\mathsf{M}[g])\phi)(\boldsymbol{x}) = f(g(\boldsymbol{x}))\phi(\boldsymbol{x}) = (\mathsf{M}[f(g)]\phi)(\boldsymbol{x})$ (Schmüdgen, 2012, Example 5.3). This redefinition of the integral is not yet very useful in itself but we can approximate the operator $\mathsf{M}[g]$ with finite matrices $\boldsymbol{M}_n[g]$ so that

$$\int_\Omega f(g(\boldsymbol{x}))w(\boldsymbol{x})d\boldsymbol{x} = \langle 1, f(\mathsf{M}[g])\,1\rangle \approx \boldsymbol{e}_0^\top f(\boldsymbol{M}_n[g])\,\boldsymbol{e}_0 = [f(\boldsymbol{M}_n[g])]_{0,0}, \qquad (2.5)$$

where $f(\cdot)$ on the right side approximation is a matrix function (Higham, 2008; Golub and Meurant, 2009). We start the matrix element indexing from 0. The column vector $\boldsymbol{e}_i$ is a vector that has value 1 in component i and 0 on other components, that is, $[\boldsymbol{e}_i]_i = 1$, and $[\boldsymbol{e}_i]_j = 0$ for $i \neq j$.

For computing the matrix $\boldsymbol{M}_n[g]$, we select some functions $\phi_0, \phi_1, \ldots, \phi_n$ that are orthonormal, that is, $\langle\phi_i, \phi_i\rangle = 1$, $\langle\phi_i, \phi_j\rangle = 0$ for $i \neq j$, and $\phi_0 = 1$. The elements of the matrix $\boldsymbol{M}_n[g]$ are then

$$[\boldsymbol{M}_n[g]]_{i,j} = \langle\phi_i, \mathsf{M}[g]\phi_j\rangle = \int_\Omega \overline{\phi_i(\boldsymbol{x})}\, g(\boldsymbol{x})\phi_j(\boldsymbol{x})w(\boldsymbol{x})d\boldsymbol{x}. \qquad (2.6)$$

If the above integrals are easier to compute than (2.1), we can use them to compute the matrix function and its upper left element as in (2.5). There are multiple ways to compute the matrix function or its approximation (Higham, 2008; Golub and Meurant, 2009). The most general definition is by the eigenvalue decomposition.

The matrix $\boldsymbol{M}_n[g]$ in (2.6) is Hermitian, that is, $\boldsymbol{M}_n[g]^* = \boldsymbol{M}_n[g]$. A Hermitian $(n+1)\times(n+1)$ matrix has $n+1$ eigenvalues λ_i and orthonormal eigenvectors $\boldsymbol{u}_i$ for which $\boldsymbol{M}_n[g]\boldsymbol{u}_i = \lambda_i \boldsymbol{u}_i$. An eigenvalue decomposition of a Hermitian matrix is

$$\boldsymbol{U}\boldsymbol{\Lambda}\boldsymbol{U}^*, \qquad (2.7)$$

where $\boldsymbol{\Lambda}$ is a diagonal matrix with the eigenvalues on its diagonal

$$\boldsymbol{\Lambda} = \begin{bmatrix} \lambda_0 & 0 & \ldots & 0 \\ 0 & \lambda_1 & \ldots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \ldots & \lambda_n \end{bmatrix},$$

and $\boldsymbol{U}$ is a matrix with the eigenvectors as its columns

$$\boldsymbol{U} = \begin{bmatrix} \boldsymbol{u}_0 & \boldsymbol{u}_1 & \ldots & \boldsymbol{u}_n \end{bmatrix}.$$

Because the vectors $\boldsymbol{u}_i$ are orthonormal, matrix $\boldsymbol{U}$ is unitary, that is, $\boldsymbol{U}\boldsymbol{U}^* = \boldsymbol{I}$. By the eigenvalue decomposition, the matrix function is

$$f(\boldsymbol{M}_n[g]) = \boldsymbol{U} f(\boldsymbol{\Lambda}) \boldsymbol{U}^* = \boldsymbol{U} \begin{bmatrix} f(\lambda_0) & 0 & \dots & 0 \\ 0 & f(\lambda_1) & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & f(\lambda_n) \end{bmatrix} \boldsymbol{U}^*.$$

For integral (2.5), this gives an approximation

$$\int_\Omega f(g(\boldsymbol{x})) w(\boldsymbol{x}) d\boldsymbol{x} \approx [f(\boldsymbol{M}_n[g])]_{0,0} = \sum_{i=0}^{n} |[\boldsymbol{u}_i]_0|^2 f(\lambda_i).$$

Since the matrix functions approximate multiplication operators, we can also approximate products of functions, for example,

$$\int_\Omega f_1(g_1(\boldsymbol{x})) f_2(g_2(\boldsymbol{x})) w(\boldsymbol{x}) d\boldsymbol{x} \approx [f_1(\boldsymbol{M}_n[g_1]) f_2(\boldsymbol{M}_n[g_2])]_{0,0},$$

when $\int_\Omega w(\boldsymbol{x}) d\boldsymbol{x} = 1$. In the unidimensional special case, when $g(x) = x$ and the basis functions are polynomials, we get the Gaussian quadrature (Publication I).

2.3 Hilbert Space

We now briefly explain the concept of Hilbert space for the development of the matrix method of numerical integration in Publications I and II, and also for the approximation of the Fourier–Hermite series in Publications III and IV. A Hilbert space is a complete inner product space (Stone, 1932; Akhiezer and Glazman, 1993; Kato, 1995; Reed and Simon, 1981; Kreyszig, 1989; Segal and Kunze, 1978; Weidmann, 1980). We will focus on three kinds of Hilbert spaces: for the finite-dimensional space $\mathbb{R}^n$ (or more generally $\mathbb{C}^n$ if necessary), for the space of square-summable sequences ℓ^2, and for the space of square-integrable functions with weight function w, that is, $\mathscr{L}^2_w$. The most typical weight function in filtering and smoothing is the Gaussian probability distribution

$$w(\boldsymbol{x}) = N(\boldsymbol{x} \mid \boldsymbol{\mu}, \boldsymbol{\Sigma}) = \frac{\exp\left(-\frac{1}{2}(\boldsymbol{x}-\boldsymbol{\mu})^\top \boldsymbol{\Sigma}^{-1}(\boldsymbol{x}-\boldsymbol{\mu})\right)}{\sqrt{(2\pi)^n |\boldsymbol{\Sigma}|}}, \tag{2.8}$$

where $\boldsymbol{\mu}$ is the mean, $\boldsymbol{\Sigma}$ is the covariance matrix of the multidimensional Gaussian random variable $\boldsymbol{x} \in \mathbb{R}^n$, and $|\boldsymbol{\Sigma}|$ is the determinant of $\boldsymbol{\Sigma}$.

The inner product in $\mathscr{L}^2_w$ is (2.2), and (2.3) defines a complete space. Elements of ℓ^2 are sequences of scalars (v_i) for $i = 0, 1, \dots$ or infinite dimensional vectors such that the inner product is

$$\langle \boldsymbol{v}, \boldsymbol{u} \rangle = \sum_{k=0}^{\infty} \overline{[\boldsymbol{v}]_k} [\boldsymbol{u}]_k = \sum_{k=0}^{\infty} \overline{v_k}\, u_k = \boldsymbol{v}^* \boldsymbol{u},$$

the norm is $\|\boldsymbol{v}\| = \sqrt{\boldsymbol{v}^* \boldsymbol{v}}$ and the space $\ell^2 = \{\boldsymbol{v} : \boldsymbol{v}^* \boldsymbol{v} < \infty\}$ is complete.

The relation of the three Hilbert spaces is that when $\mathscr{L}^2_w$ is separable, it and ℓ^2 are isomorphic, and we can approximate vectors in ℓ^2 arbitrarily closely by vectors in $\mathbb{R}^n$ if we choose n large enough. A separable Hilbert space is a span of countable set of orthonormal basis vectors $\phi_0, \phi_1, \ldots$ such that $\|\phi_i\| = 1$ and $\langle \phi_i, \phi_j \rangle = 0$ for all $i, j = 0, 1, \ldots$ and $i \neq j$. For any function ψ in separable $\mathscr{L}^2_w$, we can compute the corresponding infinite vector $\boldsymbol{v} \in \ell^2$ through a relation $[\boldsymbol{v}]_k = \langle \phi_k, \psi \rangle$ for $k = 0, 1, \ldots$. Similarly, for any infinite vector $\boldsymbol{v} \in \ell^2$, we can compute the corresponding function

$$\psi = \sum_{k=0}^{\infty} [\boldsymbol{v}]_k \, \phi_k \in \mathscr{L}^2_w. \tag{2.9}$$

This representation of the function is called a generalised Fourier series (Herman, 2017). Coefficients $[\boldsymbol{v}]_k$ are called Fourier coefficients (Reed and Simon, 1981; Kreyszig, 1989).

The representation of a function as a sequence or an infinite vector is in certain sense simpler than as a function. A function $\psi \in \mathscr{L}^2_w$ is a map from an uncountable infinite set to a uncountable infinite set $\psi : \mathbb{R} \mapsto \mathbb{R}$. An isomorphic infinite vector $\boldsymbol{v}$ defines a map from a countable set to an uncountable set, that is, we can define function $\tilde{\psi} : \mathbb{N} \mapsto \mathbb{R}$ as $\tilde{\psi}(k) = [\boldsymbol{v}]_k$ for $k = 0, 1, \ldots$, and this simpler function retains the properties of addition and multiplication by scalar of the original function. When $\phi_0(\boldsymbol{x}) = 1$, its relation to integrals is

$$\int_\Omega \psi(\boldsymbol{x}) w(\boldsymbol{x}) \, d\boldsymbol{x} = \int_\Omega \overline{\phi_0(\boldsymbol{x})} \psi(\boldsymbol{x}) w(\boldsymbol{x}) \, d\boldsymbol{x} = \langle \phi_0, \psi \rangle = [\boldsymbol{v}]_0,$$

$$\int_\Omega \overline{\phi_k(\boldsymbol{x})} \psi(\boldsymbol{x}) w(\boldsymbol{x}) \, d\boldsymbol{x} = \langle \phi_k, \psi \rangle = [\boldsymbol{v}]_k,$$

$$\int_\Omega \left| \psi(\boldsymbol{x}) \right|^2 w(\boldsymbol{x}) \, d\boldsymbol{x} = \langle \psi, \psi \rangle = \boldsymbol{v}^* \boldsymbol{v} = \sum_{k=0}^{\infty} |[\boldsymbol{v}]_k|^2,$$

where the last equality is a generalisation of the Pythagorean theorem to infinite-dimensional spaces known as Parseval's theorem. We can also use Parseval's theorem to compute the integral of the product of two functions ϕ and ψ isomorphic to $\boldsymbol{u}$ and $\boldsymbol{v}$, respectively, as

$$\int_\Omega \phi(\boldsymbol{x}) \psi(\boldsymbol{x}) w(\boldsymbol{x}) \, d\boldsymbol{x} = \langle \overline{\phi}, \psi \rangle = \boldsymbol{u}^\top \boldsymbol{v} = \sum_{k=0}^{\infty} [\boldsymbol{u}]_k \, [\boldsymbol{v}]_k.$$

We can approximate the infinite vector arbitrarily closely by a finite vector that is just a truncation of the infinite vector. As a function, this approximation is $\tilde{\psi} : \{0, 1, \ldots, n\} \mapsto \mathbb{R}$ with $\tilde{\psi}(k) = [\boldsymbol{v}]_k$ for $k = 0, 1, \ldots n$, which we can store in computer memory as $n + 1$ floating point numbers. We can then approximate the Parseval relation arbitrarily closely in the finite-dimensional space for ϕ and ψ that are isomorphic to $\boldsymbol{u}$ and $\boldsymbol{v}$, respectively,

as

$$\int_\Omega \phi(\boldsymbol{x})\psi(\boldsymbol{x})w(\boldsymbol{x})\,d\boldsymbol{x} = \sum_{k=0}^{\infty}[\boldsymbol{u}]_k[\boldsymbol{v}]_k \approx \sum_{k=0}^{n}[\boldsymbol{u}]_k[\boldsymbol{v}]_k,$$

$$\int_\Omega \psi(\boldsymbol{x})^2 w(\boldsymbol{x})\,d\boldsymbol{x} = \sum_{k=0}^{\infty}[\boldsymbol{v}]_k^2 \approx \sum_{k=0}^{n}[\boldsymbol{v}]_k^2. \tag{2.10}$$

This truncated Parseval relation is used for computing covariance matrices in Publications III and IV.

To approximate higher powers than two, we can use three operators: a multiplication operator in $\mathscr{L}_w^2$, which is isomorphic to an operator in ℓ^2 that is an infinite matrix and its finite matrix approximation. We can use this matrix approximation, not only for powers of the function but also arbitrary outside functions $f(\cdot)$ in (2.5).

2.4 Linear Operators and Infinite Matrices in Hilbert Space

After the Hilbert space overview, we can now develop a more rigorous formulation of the numerical integration method of Section 2.2 in terms of linear operators in Hilbert spaces $\mathscr{L}_w^2$, ℓ^2, and $\mathbb{R}^n$. The earlier definition (2.4) of the multiplication operator gave only the mapping of the function ϕ to ψ as $\psi(\boldsymbol{x}) = g(\boldsymbol{x})\phi(\boldsymbol{x})$. It was not yet fully rigorously stated for which vectors ϕ this relation holds, that is, what is the domain of the multiplication operator. We define the multiplication operator $\mathsf{M}[g]$ maximally, that is, so that its domain is

$$\mathscr{D}(\mathsf{M}[g]) = \left\{\phi \in \mathscr{L}_w^2 : \|g\phi\| < \infty\right\}.$$

By such definition, the maximal multiplication operator is self-adjoint operator for any real function g that is measurable with respect to the weight function w (Weidmann, 1980, Chapter 5.1, Example 2). It needs a quite wide set of conditions, and among them, that

$$\int_\Omega |g(\boldsymbol{x})|\,w(\boldsymbol{x})\,d\boldsymbol{x} < \infty. \tag{2.11}$$

An infinite matrix is an important intermediate object between a multiplication operator and its finite matrix approximation. The infinite matrix representation is always possible for bounded operators. For unbounded operators, it is possible under favorable conditions for operators that are self-adjoint (Stone, 1932; Akhiezer and Glazman, 1993).

An operator T has an adjoint T^* if T is densely defined on Hilbert space $\mathscr{H}$. The domain of the adjoint operator is

$$\mathscr{D}(\mathsf{T}^*) = \{\phi : \text{there is an } \eta \in \mathscr{H} \text{ such that } \langle \mathsf{T}\psi, \phi\rangle = \langle \psi, \eta\rangle \text{ for all } \psi \in \mathscr{D}(\mathsf{T})\}.$$

The mapping of vectors by the adjoint operator is then given as $\mathsf{T}^* \phi = \eta$ (Stone, 1932; Akhiezer and Glazman, 1993; Kato, 1995; Reed and Simon, 1981; Kreyszig, 1989; Segal and Kunze, 1978; Weidmann, 1980; Schmüdgen, 2012; Simon, 2015). In the finite-dimensional Hilbert spaces, operators are finite matrices, and the adjoint is simply a matrix transpose and a complex conjugate, that is, a Hermitian transpose. An operator that is its own adjoint is called self-adjoint.

A self-adjoint operator A has an infinite matrix representation $\boldsymbol{A}_\infty$ if the Hilbert space is separable, that is, spanned by a countable set of orthonormal basis vectors $\phi_0, \phi_1, \ldots$, and

$$\|\mathsf{A}\phi_k\| < \infty$$

for all $k = 0, 1, \ldots$. The elements of the infinite matrix are given by

$$[\boldsymbol{A}_\infty]_{i,j} = \langle \phi_i, \mathsf{A}\phi_j \rangle$$

for $i, j = 0, 1, \ldots$ (Stone, 1932; Akhiezer and Glazman, 1993). The infinite matrix $\boldsymbol{A}_\infty$ is isomorphic to the operator A, which means that if arbitrary $\boldsymbol{v} \in \ell^2$ is isomorphic to $\phi \in \mathscr{L}^2_w$, then $\boldsymbol{A}_\infty \boldsymbol{v}$ is isomorphic to $\mathsf{A}\phi$. For a multiplication operator $\mathsf{M}[g]$, condition (2.4) is equal to

$$\int_\Omega \left| g(\boldsymbol{x}) \phi_k(\boldsymbol{x}) \right|^2 w(\boldsymbol{x}) \, d\boldsymbol{x} < \infty$$

for all $k = 0, 1, \ldots$. Under this condition, an infinite matrix $\boldsymbol{M}_\infty[g]$, with elements as in (2.6) for $i, j = 0, 1, \ldots$, is isomorphic to the multiplication operator $\mathsf{M}[g]$. Compared with the condition (2.11) alone we have now slightly stricter conditions for the function g.

We can approximate the infinite matrix by finite $n \times n$ matrices. For self-adjoint operators, the finite matrices are Hermitian. Just like a finite vector, a finite matrix can be stored in computer memory as a finite number of floating point numbers. These finite approximations converge to the infinite matrix, and we could hence approximate a multiplication operator arbitrarily closely by finite matrices if we had perfect real numbers instead of floating point numbers.

The infinite matrix representation of a multiplication operator further generalises generalised Fourier series (2.9) because it is possible to compute approximations of also higher powers than just 2, which was the limitation of the Parseval relation (2.10). In fact, we can compute approximations of any function, not just powers. For example, if $\phi_0 = 1$, we have the Fourier coefficients on the zeroth row of the infinite matrix and can write the generalised Fourier series (2.9) as

$$\psi = \sum_{k=0}^{\infty} [\boldsymbol{M}_\infty[\psi]]_{0,k} \phi_k \approx \sum_{k=0}^{n} [\boldsymbol{M}_\infty[\psi]]_{0,k} \phi_k = \sum_{k=0}^{n} [\boldsymbol{M}_n[\psi]]_{0,k} \phi_k .$$

Similarly,

$$f(\psi) = \sum_{k=0}^{\infty} [f(\boldsymbol{M}_\infty[\psi])]_{0,k}\phi_k \approx \sum_{k=0}^{n} [f(\boldsymbol{M}_\infty[\psi])]_{0,k}\phi_k \approx \sum_{k=0}^{n} [f(\boldsymbol{M}_n[\psi])]_{0,k}\phi_k.$$

2.5 Convergence in Hilbert Space

So far, we have presented approximations of objects in the infinite-dimensional $\mathscr{L}^2_w$ as objects in the finite-dimensional $\mathbb{R}^n$. We may then ask the natural question: Do those approximations really converge to the approximation target as $n \to \infty$. This topic was considered in Publication II. Different forms of convergence of vectors (and a big part of the convergence of operators) in Hilbert space are covered, for example, in Akhiezer and Glazman (1993); Kato (1995); Kreyszig (1989). We also use more specialised forms of operator convergence from (Reed and Simon, 1981; Weidmann, 1980; Simon, 2015).

Strong convergence of a sequence of vectors ϕ_n in a Hilbert space $\mathscr{H}$ means that there is a limit $\phi \in \mathscr{H}$ such that

$$\lim_{n\to\infty} \|\phi_n - \phi\| = 0,$$

that is, for any $\epsilon > 0$, there is N such that for $n > N$, $\|\phi_n - \phi\| < \epsilon$. For example, to approximate an infinite vector in ℓ^2, we can take its n first components and replace the rest with 0s. This convergence is strong when $n \to \infty$. This way, we can represent the infinite vector arbitrarily closely, to within the accuracy of the floating point numbers, in computer memory. The infinite vector is isomorphic to some function in $\mathscr{L}^2_w$, and the finite vectors also therefore give an arbitrarily close approximation of the function.

Weak convergence of a sequence of ϕ_n means that there is a limit $\phi \in \mathscr{H}$ such that for every $\psi \in \mathscr{H}$, we have

$$\lim_{n\to\infty} \langle \psi, \phi_n \rangle = \langle \psi, \phi \rangle.$$

As the terminology suggests, if a sequence of vectors converges strongly, it also converges weakly.

With operators, we have more different forms of convergence. Strong or weak convergence of a sequence of operators T_n to T means that for every $\phi \in \mathscr{H}$, we have strong or weak convergence of the sequence of vectors $\mathsf{T}_n\phi$ to $\mathsf{T}\phi$, respectively. We write strong convergence as $\mathsf{T}_n \xrightarrow{s} \mathsf{T}$ and weak convergence as $\mathsf{T}_n \xrightarrow{w} \mathsf{T}$. Weak convergence is enough for convergence of the numerical integral approximation in (2.5). However, the finite matrix approximations converge strongly to the infinite matrix, and strong convergence has the essential property that products and sums of operators

converge strongly for operators that converge strongly, that is,

$$\mathsf{A}_n \mathsf{B}_n \xrightarrow{s} \mathsf{A}\mathsf{B},$$

$$\mathsf{A}_n + \mathsf{B}_n \xrightarrow{s} \mathsf{A} + \mathsf{B},$$

if $\mathsf{A}_n \xrightarrow{s} \mathsf{A}$ and $\mathsf{B}_n \xrightarrow{s} \mathsf{B}$ (Kato 1995, Chapter III, Lemma 3.8 and Kreyszig 1989, Chapter 4.9, Problem 2). Strong convergence then follows for polynomials of an operator, for continuous functions by Stone–Weierstrass theorem, and by the spectral theorem also for Riemann–Stieltjes integrable functions, as shown in Publication II, Section 3.2.

Strong convergence applies only to operators whose domain is the whole Hilbert space $\mathscr{H}$ because the convergence of operators $\mathsf{T}_n \phi$ must hold for all $\phi \in \mathscr{H}$. This requirement rules out unbounded operators that cannot have the whole Hilbert space $\mathscr{H}$ as their domain. For multiplication operators, this would limit us to multiplication operators of bounded functions. For unbounded self-adjoint operators, we need strong resolvent convergence (Reed and Simon, 1981; Weidmann, 1980; Simon, 2015). The idea is to map the spectrum of the self-adjoint operator, which is the entire real line or part of it, into a bounded set on the complex plane.

A sequence of self-adjoint operators A_n converges in the strong resolvent sense to a self-adjoint operator A if $(\mathsf{A}_n - z)^{-1} \xrightarrow{s} (\mathsf{A} - z)^{-1}$ for all non-real numbers z. We write strong resolvent convergence as $\mathsf{A}_n \xrightarrow{srs} \mathsf{A}$. Strong convergence implies strong resolvent convergence, and therefore strong resolvent convergence is also called generalised strong convergence (Kato, 1995). With strong resolvent convergence, we can show convergence for a large class of functions as in Publication II, Section 3. With unbounded functions, convergence may also not hold, as is numerically demonstrated in Publication II, Section 4.

2.6 Spectral Theorem

The spectral theorem is one of the most central results of the operator theory (Stone, 1932; Akhiezer and Glazman, 1993; Kato, 1995; Reed and Simon, 1981; Kreyszig, 1989; Segal and Kunze, 1978; Weidmann, 1980; Schmüdgen, 2012; Simon, 2015). We also use it in our theory of numerical integration in Publication II on convergence proof for discontinuous Riemann–Stieltjes integrable functions, and in a generalisation of the product rule for non-independent variables of Publication V.

A self-adjoint operator has a spectral representation

$$\mathsf{A} = \int_{\sigma(\mathsf{A})} t \, d\mathsf{E}(t),$$

where $\sigma(\mathsf{A}) \subset \mathbb{R}$ is the spectrum of the operator and $\mathsf{E}(t)$ is a projection operator called the spectral family. For a Hermitian $n \times n$ matrix, the

spectrum is the set of eigenvalues $\{\lambda_i\}$ of the matrix, and since the Hilbert space is finite dimensional, the spectral family is an $n \times n$ matrix-valued function of t

$$\boldsymbol{E}(t) = \sum_{\lambda_i \le t} \boldsymbol{u}_i \boldsymbol{u}_i^*,$$

where $\boldsymbol{u}_i$ are the eigenvectors, and thus, the spectral representation is another way to express the eigenvalue decomposition (2.7).

For a multiplication operator $\mathsf{M}[g]$ of a real function g, the spectrum is the essential range of the function g, that is,

$$\sigma(\mathsf{M}[g]) = \left\{ y \in \mathbb{R} : \text{for all } \epsilon > 0, \int_{\{\boldsymbol{x} \in \Omega : |g(\boldsymbol{x}) - y| < \epsilon\}} w(\boldsymbol{x})\, d\boldsymbol{x} > 0 \right\}.$$

Usually, the essential range of a function is just the image of it. The spectral family of a multiplication operator is also a multiplication operator $\mathsf{E}(t) = \mathsf{M}[\chi_{\{\boldsymbol{x} \in \Omega : g(\boldsymbol{x}) \le t\}}]$, where χ is the characteristic function

$$\chi_{\{\boldsymbol{x} \in \Omega : g(\boldsymbol{x}) \le t\}}(\boldsymbol{x}) = \begin{cases} 0, & g(\boldsymbol{x}) > t, \\ 1, & g(\boldsymbol{x}) \le t, \end{cases}$$

(Weidmann, 1980, Chapter 7.2, Example 1).

We can define a function of self-adjoint operator A as

$$f(\mathsf{A}) = \int_{\sigma(\mathsf{A})} f(t)\, d\mathsf{E}(t)$$

for any function that is measurable with respect to the spectral family $\mathsf{E}(t)$. Strong resolvent convergence of self-adjoint operators $\mathsf{A}_n \xrightarrow{srs} \mathsf{A}$ also implies strong convergence of the corresponding spectral family operators $\mathsf{E}_n(t) \xrightarrow{s} \mathsf{E}(t)$. We can define the integral with respect to the spectral family as Riemann–Stieltjes integral in the strong operator topology and use that to prove convergence for Riemann–Stieltjes integrable functions, as shown in Publication II, Section 3.2.

2.7 Generalised Product Rule for Non-independent Variables

In Publication V, Section 2.1, the spectral theorem is used for computing multidimensional integrals in the following way. Let the multiplication operators for each variable x_i be $\mathsf{M}[\mathrm{id}_i]$, where function $\mathrm{id}_i(\boldsymbol{x}) = [\boldsymbol{x}]_i$. Let the spectral family of each operator $\mathsf{M}[\mathrm{id}_i]$ be $\mathsf{E}_i(t)$. The spectrum of each operator is $\sigma(\mathsf{M}[\mathrm{id}_i]) = [\inf x_i, \sup x_i]$. Then, for arbitrary function of these variables $f(\boldsymbol{x}) = f(x_1, x_2, \ldots, x_d)$, we can write the multiplication operator of the function as

$$\begin{aligned} \mathsf{M}[f] &= f(\mathsf{M}[\mathrm{id}_1], \mathsf{M}[\mathrm{id}_2], \ldots, \mathsf{M}[\mathrm{id}_d]) \\ &= \int_{I_d} \cdots \int_{I_2} \int_{I_1} f(t_1, t_2, \ldots, t_d)\, d\mathsf{E}_1(t_1)\, d\mathsf{E}_2(t_2) \cdots d\mathsf{E}_d(t_d), \end{aligned}$$

where $I_i = [\inf x_i, \sup x_i]$. This expression of the multiplication operator $\mathsf{M}[f]$ is valid because operators $\mathsf{M}[\mathrm{id}_i]$ commute with each other, that is, $\mathsf{M}[\mathrm{id}_i]\mathsf{M}[\mathrm{id}_j] = \mathsf{M}[\mathrm{id}_j]\mathsf{M}[\mathrm{id}_i]$ for all i, j. It means that also $\mathsf{E}_i(t_i)$ and $\mathsf{E}_j(t_j)$ commute, and therefore, the product spectral measure exists on the Borel algebra of $I_1 \times I_2 \times \ldots \times I_d$ by (Schmüdgen, 2012, Theorem 4.10).

For comparison, in general, the finite matrices $\boldsymbol{M}_n[\mathrm{id}_i]$ do not commute with $\boldsymbol{M}_n[\mathrm{id}_j]$ for $i \neq j$, and therefore, the definition of a matrix function $f(\boldsymbol{M}_n[\mathrm{id}_1], \boldsymbol{M}_n[\mathrm{id}]_2, \ldots, \boldsymbol{M}_n[\mathrm{id}_d])$ is tricky. For example, let us consider a second-order monomial $x_1 x_2 = x_2 x_1$. In general, for finite matrices, $\boldsymbol{M}_n[\mathrm{id}_1]\boldsymbol{M}_n[\mathrm{id}_2] \neq \boldsymbol{M}_n[\mathrm{id}_2]\boldsymbol{M}_n[\mathrm{id}_1]$. If we consider arbitrary order polynomial, we will face the same problem, that the matrix function will depend on the order of the matrices in the individual monomials of the polynomial. However, we can also make the definition for an arbitrary monomial so that the matrices are ordered by the corresponding subindex. For the example monomial $x_1 x_2 = x_2 x_1$, this would mean that the definition of the corresponding finite matrix function is always $\boldsymbol{M}_n[\mathrm{id}_1]\boldsymbol{M}_n[\mathrm{id}_2]$. This same ordering is preserved for polynomials as well. With this ordering, we can write the matrix function in terms of spectral integrals.

Let $\boldsymbol{E}_{i,n}(t_i)$ be the spectral family of $\boldsymbol{M}_n[\mathrm{id}_i]$. Then, for polynomials f, we can define

$$
\begin{aligned}
&f(\boldsymbol{M}_n[\mathrm{id}_1], \boldsymbol{M}_n[\mathrm{id}_2], \ldots, \boldsymbol{M}_n[\mathrm{id}_d]) \\
&= \int_{I_d} \ldots \int_{I_2} \int_{I_1} f(t_1, t_2, \ldots, t_d) d\boldsymbol{E}_{1,n}(t_1) d\boldsymbol{E}_{2,n}(t_2) \cdots d\boldsymbol{E}_{d,n}(t_d) \\
&= \sum_{i_d=0}^{n} \ldots \sum_{i_2=0}^{n} \sum_{i_1=0}^{n} f(\lambda_{n,1,i_1}, \lambda_{n,2,i_2}, \ldots, \lambda_{n,d,i_d}) \boldsymbol{P}_{n,1,i_1} \boldsymbol{P}_{n,2,i_2} \cdots \boldsymbol{P}_{n,d,i_d},
\end{aligned}
$$

where matrices $\boldsymbol{P}_{n,k,i_k} = \boldsymbol{u}_{n,k,i_k} \boldsymbol{u}^*_{n,k,i_k}$ are projection operators to the one-dimensional subspace spanned by the eigenvectors $\boldsymbol{u}_{n,k,i_k}$ of matrices $\boldsymbol{M}_n[\mathrm{id}_k]$, and λ_{n,k,i_k} are the eigenvalues of $\boldsymbol{M}_n[\mathrm{id}_k]$ for $i_k = 0, 1, \ldots, n$ and $k = 1, 2, \ldots, d$. Again, by the spectral integral, we can compute the integral of interest and approximate it as

$$
\begin{aligned}
\langle 1, \mathsf{M}[f]\, 1 \rangle &= \int_{I_d} \ldots \int_{I_2} \int_{I_1} f(x_1, x_2, \ldots, x_d) w(x_1, x_2, \ldots, x_d) d\boldsymbol{x} \\
&= \langle 1, f(\mathsf{M}[\mathrm{id}_1], \mathsf{M}[\mathrm{id}_2], \ldots, \mathsf{M}[\mathrm{id}_d])\, 1 \rangle \\
&= \left\langle 1, \int_{I_d} \ldots \int_{I_2} \int_{I_1} f(t_1, t_2, \ldots, t_d) d\mathsf{E}_1(t_1) d\mathsf{E}_2(t_2) \cdots d\mathsf{E}_d(t_d)\, 1 \right\rangle \\
&\approx \sum_{i_d=0}^{n} \ldots \sum_{i_2=0}^{n} \sum_{i_1=0}^{n} f(\lambda_{n,1,i_1}, \lambda_{n,2,i_2}, \ldots, \lambda_{n,d,i_d}) \\
&\quad \boldsymbol{e}_0^\top \boldsymbol{P}_{n,1,i_1} \boldsymbol{P}_{n,2,i_2} \cdots \boldsymbol{P}_{n,d,i_d} \boldsymbol{e}_0,
\end{aligned}
\tag{2.12}
$$

when the 0th orthonormal basis function $\phi_0 = 1$ is isomorphic to infinite vector $\boldsymbol{e}_0 \in \ell^2$.

In Publication V, Proposition 3, it was shown that (2.12) converges for bounded and continuous functions f when intervals $I_k = [\inf x_k, \sup x_k]$ are bounded and the basis functions $\phi_0, \phi_1, \ldots$ are polynomials. This convergence holds also for any other ordering of the matrices $\boldsymbol{P}_{n,k,i_k}$ in (2.12) and thus our ordering-dependent definition of a function of several matrices is justified. The proof of convergence could be extended for a larger family of functions, like non-continuous Riemann–Stieltjes integrable functions f and non-polynomial basis functions ϕ_k as in Publication II, Section 3.2 or unbounded functions of two variables similar to Publication II, Theorem 19.

The approximation of (2.12) is a generalisation of the product rule for weight functions that are factorable into a product of weight functions of one variable (Davis and Rabinowitz, 1984) which has also been used in filtering with Gaussian approximation (Ito and Xiong, 2000; Challa et al., 2000; Wu et al., 2006) and arbitrary order moments (Jia and Xin, 2019). For probability distributions, product rules are limited to distributions of independent variables (Ahlfeld et al., 2016).

For independent variables $x_1, x_2, \ldots, x_d$, we can factor the probability distribution as a product of probability distributions of the individual variables

$$p(x_1, x_2, \ldots, x_d) = p(x_1)\,p(x_2) \cdots p(x_d).$$

Here we adopt the usual convention that $p(x_1)$ stands for a function of x_1 that is the probability distribution of x_1 and $p(x_2)$ can be an entirely different function not just the same function with x_2 as its argument, that is, $p(x_1) = f_1(x_1)$ and $p(x_2) = f_2(x_2)$ where f_1 and f_2 are different functions. For dependent variables, the probability distribution is factored into a product of conditional probability distributions as

$$p(x_1, x_2, \ldots, x_d) = p(x_1 \mid x_2, \ldots, x_d)\,p(x_2 \mid x_3, \ldots, x_d) \cdots p(x_{d-1} \mid x_d)\,p(x_d),$$

which is identical for any permutation of the ordering of the dependent distributions like, for example, in case of two variables $p(x_1, x_2) = p(x_1 \mid x_2)\,p(x_2) = p(x_2 \mid x_1)\,p(x_1)$.

In order to use the normal product rule, we would need to define conditional numerical integration rules for the dependent variables. By using the spectral integral representation and approximating it with the finite matrices as in (2.12), we can avoid that. When the variables are independent, approximation (2.12) is the same as using the product rule of Gaussian quadratures for each independent variable.

3. Filtering and Smoothing

In this chapter, we show how to use the Hilbert space methods of the previous chapter in stochastic filtering and smoothing. We start with discrete-time filtering and smoothing after explaining their foundation on only two relatively simple principles: Bayes' theorem and Markov property. The contribution of using the Hilbert space method Fourier–Hermite series of Publications III and IV for nonlinear filtering and smoothing is explained.

Next, we move to the topic of continuous-discrete-time filtering and smoothing. There, we first give a short and only cursory explanation of the underlying continuous-time stochastic calculus.

3.1 Formulation of Filtering and Smoothing Problems

Filtering and smoothing mean estimating the state of a dynamic stochastic state-space model (SSM) based on noisy measurements of the states (Jazwinski, 2007; Särkkä and Svensson, 2023). For example, the dynamic discrete-time SSM can be

$$\boldsymbol{x}_k = \boldsymbol{f}_k(\boldsymbol{x}_{k-1}) + \boldsymbol{q}_{k-1},$$

where $\boldsymbol{q}_k$ are random numbers but $\boldsymbol{f}_k$ are deterministic known functions at each time step $k = 1, 2, \ldots$. The measurement model can be, for example,

$$\boldsymbol{y}_k = \boldsymbol{h}_k(\boldsymbol{x}_k) + \boldsymbol{r}_k,$$

where $\boldsymbol{r}_k$ are random, and $\boldsymbol{h}_k$ are deterministic. We then aim to estimate $\boldsymbol{x}_k$ based on the measurement values $\boldsymbol{y}_k$. The distributions of the random values are known. Because there is randomness involved, it is not generally even possible to know the exact value of $\boldsymbol{x}_k$. We can only know a probabilistic estimate of it.

In filtering, the state estimate at a given time instant is based on measurements only up to that instant and not after that. In terms of probabili-

ties, we try to estimate the probability distribution of $\boldsymbol{x}_k$ given measurements $\boldsymbol{y}_{1:k} = \boldsymbol{y}_1, \boldsymbol{y}_2, \ldots, \boldsymbol{y}_k$, that is,

$$p(\boldsymbol{x}_k \mid \boldsymbol{y}_{1:k}). \tag{3.1}$$

In smoothing, the state estimate at a given time is based on measurements up to the given time and after it, that is, we try to estimate

$$p(\boldsymbol{x}_k \mid \boldsymbol{y}_{1:T}) \tag{3.2}$$

for $k < T$. There is also a third form of state estimation: prediction, which means that we try to estimate $p(\boldsymbol{x}_{k+q} \mid \boldsymbol{y}_{1:k})$ for $q > 0$.

As a practical example of an SSM, we use the same model as in Publications III and IV, a two-dimensional pendulum model with the following discretised non-linear dynamics

$$\begin{aligned} x_{k,1} &= x_{k-1,1} + x_{k-1,2}\,\Delta t, \\ x_{k,2} &= x_{k-1,2} - g\,\sin(x_{k-1,1})\,\Delta t + q_{k-1}, \end{aligned} \tag{3.3}$$

where the state $\boldsymbol{x}_k \in \mathbb{R}^2$ consists of $x_{k,1}$, the angle of the pendulum, and $x_{k,2}$, its angular velocity at the time step k. Parameter $\Delta t = 1/1000$ is the sampling period of the discretised model, and $g = 9.81$ is the acceleration due to the gravity. Process noise is $q_{k-1} \sim N(0, Q)$ with variance $Q = \Delta t/100$. The measurement model is

$$y_k = h(x_{k,1}) + r_k, \qquad r_k \sim N(0, R), \tag{3.4}$$

where $R = 1/1000$ is the measurement noise variance, and h is a piecewise constant function

$$h(x) = \begin{cases} -1 & \text{if } x < -a/2 + b, \\ 0 & \text{if } -a/2 + b < x < a/2 + b, \\ 1 & \text{if } x > a/2 + b, \end{cases}$$

with constant values $a = 0.5$ and $b = 0.4$. The initial value $\boldsymbol{x}_0$ is a Gaussian around zero with a small error of variance of 0.01 in both components.

The process noise q_{k-1} in (3.3) adds a small non-deterministic unpredictable component to the dynamics through the angular velocity. This makes virtually any kind of path possible, but the most probable expected outcome is the same as for a deterministic pendulum. The dimension is only two, making this example the smallest possible multidimensional SSM, but in general, there is no upper limit on the dimension of the SSM. In general, also the measurements can be multidimensional.

3.2 Bayesian Background of Filtering and Smoothing

Next, we look at the elementary laws of probability that we need to derive the solutions for the discrete-time problems of the previous section and

the Markov properties that we must meet to make the problems tractable (Jazwinski, 2007; Särkkä and Svensson, 2023). The derivation of the discrete-time filtering and smoothing equations is based on the 18th century Bayes' theorem, that is, the relation between the joint probability and the conditional probability for discrete events A and B,

$$P(A \cap B) = P(A \mid B)P(B) = P(B \mid A)P(A),$$

from which we deduce

$$P(A \mid B) = \frac{P(B \mid A)P(A)}{P(B)},$$

when $P(B) > 0$ (De Finetti, 2017, Chapter 4). We use similar notation for probability densities of continuous random variables for which the Bayes' theorem is similar

$$p(x \mid y) = \frac{p(y \mid x)p(x)}{p(y)} = \frac{p(y \mid x)p(x)}{\int p(y \mid x)p(x)dx}, \tag{3.5}$$

when $p(y) > 0$.

Additionally, we use the Markov property that the state transition probability depends only on the previous state and the observation probability depends only on the current state, not on the longer history of the states. The state transition and the observation probability densities are

$$\begin{aligned} \boldsymbol{x}_k &\sim p(\boldsymbol{x}_k \mid \boldsymbol{x}_{k-1}, \boldsymbol{x}_{k-2}, \ldots, \boldsymbol{x}_0) = p(\boldsymbol{x}_k \mid \boldsymbol{x}_{k-1}), \\ \boldsymbol{y}_k &\sim p(\boldsymbol{y}_k \mid \boldsymbol{x}_k, \boldsymbol{x}_{k-1}, \ldots, \boldsymbol{x}_0) = p(\boldsymbol{y}_k \mid \boldsymbol{x}_k), \end{aligned} \tag{3.6}$$

where $\boldsymbol{x}_k \in \mathbb{R}^n$ is the state of the system at time step k and $\boldsymbol{y}_k \in \mathbb{R}^m$ is the observation also called the measurement at the time step k. We also have an initial state with a probability density that is not conditional on any previous state $\boldsymbol{x}_0 \sim p(\boldsymbol{x}_0)$. We notice that the Markov properties (3.6) are met in the pendulum example of (3.3) and (3.4).

3.3 Discrete-Time Filtering

Now, we can derive the general filtering equations using the Bayes' rule and the Markov property. In this section, we will also present the Kalman filter, the solution for the linear case.

By the Bayes' rule (3.5) and the Markov property (3.6), the recursive filtering equations for probability density of the state $\boldsymbol{x}_k$ at a time step k given the history of observations until the time step k are (Ho and Lee 1964; Särkkä and Svensson 2023, Equations (6.11)–(6.13))

$$\begin{aligned} p(\boldsymbol{x}_k \mid \boldsymbol{y}_{1:k-1}) &= \int p(\boldsymbol{x}_k \mid \boldsymbol{x}_{k-1}) p(\boldsymbol{x}_{k-1} \mid \boldsymbol{y}_{1:k-1}) d\boldsymbol{x}_{k-1}, \\ p(\boldsymbol{x}_k \mid \boldsymbol{y}_{1:k}) &= \frac{p(\boldsymbol{y}_k \mid \boldsymbol{x}_k) p(\boldsymbol{x}_k \mid \boldsymbol{y}_{1:k-1})}{\int p(\boldsymbol{y}_k \mid \boldsymbol{x}_k) p(\boldsymbol{x}_k \mid \boldsymbol{y}_{1:k-1}) d\boldsymbol{x}_k}. \end{aligned} \tag{3.7}$$

The first equation solves the prediction problem for one time step and is called the prediction step. Longer prediction is possible by repeating it. The second equation is called the update step. Thus, we see that the solution of the prediction problem is half of the solution of the filtering problem. These equations are not, in general, solvable in closed form. Kalman filter is a closed-form solution in the linear case. In the non-linear case, like our simple pendulum model in (3.3) and (3.4), different kinds of numerical approximations are used.

When the initial state is Gaussian, likewise, the state transition and observation models, which are also linear, the solution is given by Kalman filter (KF). The Gaussian distributions and linear SSM mean that (3.6) are given by

$$\begin{aligned} p(\boldsymbol{x}_0) &= N(\boldsymbol{x}_0 \mid \boldsymbol{m}_0, \boldsymbol{P}_0), \\ p(\boldsymbol{x}_k \mid \boldsymbol{x}_{k-1}) &= N(\boldsymbol{x}_k \mid \boldsymbol{A}_{k-1}\boldsymbol{x}_{k-1}, \boldsymbol{Q}_{k-1}), \\ p(\boldsymbol{y}_k \mid \boldsymbol{x}_k) &= N(\boldsymbol{y}_k \mid \boldsymbol{H}_k\boldsymbol{x}_k, \boldsymbol{R}_k), \end{aligned} \tag{3.8}$$

where N is the Gaussian distribution (2.8) and $\boldsymbol{Q}_{k-1}$, $\boldsymbol{R}_k$, and $\boldsymbol{P}_0$ are the covariance matrices of the process and measurement noises and the initial state, respectively, which are mutually independent of each other and between time steps. Dynamics and measurement functions are linear $\boldsymbol{f}_k(\boldsymbol{x}_{k-1}) = \boldsymbol{A}_{k-1}\boldsymbol{x}_{k-1}$ and $\boldsymbol{h}_k(\boldsymbol{x}_k) = \boldsymbol{H}_k\boldsymbol{x}_k$ respectively. In this simple linear Gaussian case, the general solution (3.7) is reduced to much simpler matrix algebra, starting from $\boldsymbol{m}_0$ and $\boldsymbol{P}_0$ and continuing recursively with (Särkkä and Svensson, 2023, Theorem 6.6)

$$\begin{aligned} \boldsymbol{m}_k^- &= \boldsymbol{A}_{k-1}\boldsymbol{m}_{k-1}, \\ \boldsymbol{P}_k^- &= \boldsymbol{A}_{k-1}\boldsymbol{P}_{k-1}\boldsymbol{A}_{k-1}^\top + \boldsymbol{Q}_{k-1}, \\ \boldsymbol{v}_k &= \boldsymbol{y}_k - \boldsymbol{H}_k\boldsymbol{m}_k^-, \\ \boldsymbol{S}_k &= \boldsymbol{H}_k\boldsymbol{P}_k^-\boldsymbol{H}_k^\top + \boldsymbol{R}_k, \\ \boldsymbol{K}_k &= \boldsymbol{P}_k^-\boldsymbol{H}_k^\top\boldsymbol{S}_k^{-1}, \\ \boldsymbol{m}_k &= \boldsymbol{m}_k^- + \boldsymbol{K}_k\boldsymbol{v}_k, \\ \boldsymbol{P}_k &= \boldsymbol{P}_k^- - \boldsymbol{K}_k\boldsymbol{S}_k\boldsymbol{K}_k^\top. \end{aligned} \tag{3.9}$$

The KF is also the best linear unbiased estimator for a larger family of probability distributions than just Gaussian distributions (Maryak et al., 2004). The system has to be linear and the random noise additive, so we cannot use it for our pendulum SSM of (3.3) and (3.4).

3.4 Fourier–Hermite Kalman Filter

In Publication III, the Fourier–Hermite series is used for an approximate solution of a non-linear filtering problem. The approach is still very similar to the Kalman filter. The probability distributions of the noise terms are Gaussian, but the dynamics and measurement functions can be non-linear. The solution probability densities and intermediate distributions in the computations are approximated as Gaussian distributions. This kind of Gaussian approximation has been used earlier in Ito and Xiong (2000); Maybeck (1982).

The non-linear SSM dynamics and measurement functions (3.6) are given by

$$\begin{aligned} p(\boldsymbol{x}_0) &= N(\boldsymbol{x}_0 \mid \boldsymbol{m}_0, \boldsymbol{P}_0), \\ p(\boldsymbol{x}_k \mid \boldsymbol{x}_{k-1}) &= N(\boldsymbol{x}_k \mid \boldsymbol{f}_k(\boldsymbol{x}_{k-1}), \boldsymbol{Q}_{k-1}), \\ p(\boldsymbol{y}_k \mid \boldsymbol{x}_k) &= N(\boldsymbol{y}_k \mid \boldsymbol{h}_k(\boldsymbol{x}_k), \boldsymbol{R}_k), \end{aligned} \tag{3.10}$$

where now $\boldsymbol{f}_k$ and $\boldsymbol{h}_k$ can be non-linear functions. With linear $\boldsymbol{f}_k$ and $\boldsymbol{h}_k$, we would get the exact solution as the KF (3.9). With the non-linear functions, we use Gaussian approximations of probability densities, and the result is very similar to KF (3.9) (Maybeck 1982, Section 12.7; Ito and Xiong 2000)

$$\begin{aligned} \boldsymbol{m}_k^- &= \mathrm{E}[\boldsymbol{f}_k(\boldsymbol{x}) \mid \boldsymbol{m}_{k-1}, \boldsymbol{P}_{k-1}], \\ \boldsymbol{P}_k^- &= \mathrm{E}[(\boldsymbol{f}_k(\boldsymbol{x}) - \boldsymbol{m}_k^-)(\boldsymbol{f}_k(\boldsymbol{x}) - \boldsymbol{m}_k^-)^T \mid \boldsymbol{m}_{k-1}, \boldsymbol{P}_{k-1}] + \boldsymbol{Q}_k, \\ \boldsymbol{\mu}_k &= \mathrm{E}[\boldsymbol{h}_k(\boldsymbol{x}) \mid \boldsymbol{m}_k^-, \boldsymbol{P}_k^-], \\ \boldsymbol{S}_k &= \mathrm{E}[(\boldsymbol{h}_k(\boldsymbol{x}) - \boldsymbol{\mu}_k)(\boldsymbol{h}_k(\boldsymbol{x}) - \boldsymbol{\mu}_k)^T \mid \boldsymbol{m}_k^-, \boldsymbol{P}_k^-] + \mathbf{R}_k, \\ \mathbf{K}_k &= \mathrm{E}[(\boldsymbol{x} - \mathbf{m}_k^-)(\mathbf{h}_k(\boldsymbol{x}) - \boldsymbol{\mu}_k)^T \mid \boldsymbol{m}_k^-, \boldsymbol{P}_k^-]\mathbf{S}_k^{-1}, \\ \mathbf{m}_k &= \mathbf{m}_k^- + \mathbf{K}_k(\mathbf{y}_k - \boldsymbol{\mu}_k), \\ \mathbf{P}_k &= \mathbf{P}_k^- - \mathbf{K}_k \mathbf{S}_k \mathbf{K}_k^T. \end{aligned} \tag{3.11}$$

In the above equations, we have several Gaussian weighted integrals of non-linear vector or matrix functions in the form of a Gaussian expectation

$$\mathrm{E}[\boldsymbol{g}(\boldsymbol{x}) \mid \boldsymbol{\mu}, \boldsymbol{\Sigma}] = \int_{\mathbb{R}^n} \boldsymbol{g}(\boldsymbol{x}) N(\boldsymbol{x} \mid \boldsymbol{\mu}, \boldsymbol{\Sigma}) \, d\boldsymbol{x}.$$

The extension is not an exact solution but an approximation that is based on a Fourier–Hermite expansion of the non-linear functions. We define an inner product of two scalar functions as

$$\langle f, g \rangle = \mathrm{E}[f(\boldsymbol{x}) g(\boldsymbol{x}) \mid \boldsymbol{\mu}, \boldsymbol{\Sigma}] = \int_{\mathbb{R}^n} f(\boldsymbol{x}) g(\boldsymbol{x}) N(\boldsymbol{x} \mid \boldsymbol{\mu}, \boldsymbol{\Sigma}) \, d\boldsymbol{x}. \tag{3.12}$$

This inner product naturally induces a norm $\|f\| = \sqrt{\langle f, f\rangle}$, and the functions that satisfy $\|f\| < \infty$ form a separable Hilbert space. The Hermite polynomials are (Särkkä and Svensson, 2023, Equation 8.18)

$$H_q(x) = (-1)^q \, e^{x^2/2} \frac{d^q}{dx^q} e^{-x^2/2} \tag{3.13}$$

for $q = 0, 1, \ldots$ and they are orthogonal to each other. By renormalising the Hermite polynomials, we can form an orthonormal basis $\{\phi_0 = 1, \phi_1, \ldots\}$ in this Hilbert space. As shown in Chapter 2.3, the choice of the basis defines an isomorphism between the Hilbert space of the functions and the Hilbert space of square-summable infinite sequences that is, for scalar functions f and g in the Hilbert space, there exist infinite square-summable sequences $(a_q)_{q=0}^{\infty}$ and $(b_q)_{q=0}^{\infty}$ such that $a_q = \langle f, \phi_q\rangle$, $b_q = \langle g, \phi_q\rangle$, and $\langle f, g\rangle = \sum_{q=0}^{\infty} a_q \, b_q$. We can use this isomorphism to compute approximations of expectations for Gaussian approximation of a non-linear filter by using

$$\begin{aligned} \mathrm{E}[f(\boldsymbol{x})] &= \langle f, \phi_0\rangle = a_0, \\ \mathrm{E}[f(\boldsymbol{x}) g(\boldsymbol{x})] &= \langle f, g\rangle = \sum_{q=0}^{\infty} a_q \, b_q \approx \sum_{q=0}^{Q} a_q \, b_q, \end{aligned} \tag{3.14}$$

where $\phi_0 = 1$ is the zeroth order Hermite polynomial and Q is some arbitrary finite order of truncation for the approximation.

The Fourier–Hermite series has the property that it resembles the Taylor series. In one dimension, the Fourier–Hermite coefficients can be computed using integration by parts as (Malliavin 1997, Theorem 2.5; Gelb 1974, Equation 6.4-12)

$$\mathrm{E}\left[H_q\left(\frac{x-\mu}{\sigma}\right) f(x) \,\middle|\, \mu, \sigma^2\right] = \mathrm{E}\left[\frac{d^q f(x)}{dx^q} \,\middle|\, \mu, \sigma^2\right] \sigma^q = \frac{d^q \, \mathrm{E}[f(x) \mid \mu, \sigma^2]}{d\mu^q} \sigma^q. \tag{3.15}$$

This means that, for example, the limit of the Fourier–Hermite series when $\sigma \to 0$ is the Taylor series for analytical functions

$$\begin{aligned} f(x) &= \sum_{q=0}^{\infty} \frac{1}{q!} \frac{d^q f(\mu)}{d\mu^q} (x-\mu)^q \\ &= \lim_{\sigma \to 0} \sum_{q=0}^{\infty} \frac{1}{q!} \mathrm{E}\left[\frac{d^q f(x)}{dx^q} \,\middle|\, \mu, \sigma^2\right] \sigma^q \, H_q\left(\frac{x-\mu}{\sigma}\right) \\ &= \lim_{\sigma \to 0} \sum_{q=0}^{\infty} \frac{1}{q!} \frac{d^q \, \mathrm{E}[f(x) \mid \mu, \sigma^2]}{d\mu^q} \sigma^q \, H_q\left(\frac{x-\mu}{\sigma}\right). \end{aligned}$$

The limits when $\sigma \to 0$ follow due to $\mathrm{E}[f(x) \mid \mu, 0] = f(\mu)$ and $\lim_{\sigma \to 0} \sigma^q \, H_q\left(\frac{x-\mu}{\sigma}\right) = (x-\mu)^q$ because the highest term in $H_q(x)$ is x^q.

In Publication III, Theorem 2, it was shown, by using the multidimensional version of relation (3.15), that we can compute the Fourier–Hermite

series coefficients by first computing the Weierstrass transformation (Zayed, 1996, Chapter 18) of the function

$$\hat{g}(\boldsymbol{\mu}, \boldsymbol{\Sigma}) = \int_{\mathbb{R}^n} g(\boldsymbol{x}) N(\boldsymbol{x} \mid \boldsymbol{\mu}, \boldsymbol{\Sigma}) d\boldsymbol{x}$$

and then the high order mixed partial differentials of the Weierstrass transform of the function with respect to μ_i like, for example, $\frac{\partial^3 \hat{g}(\boldsymbol{\mu}, \boldsymbol{\Sigma})}{\partial \mu_i \partial \mu_j, \partial \mu_k}$.

Thus, the approximate filtering solution of the non-linear filtering equations (3.10) requires a closed-form solution of two integrals

$$\begin{aligned}
\hat{\boldsymbol{f}}(\boldsymbol{\mu}, \boldsymbol{\Sigma}) &= \int_{\mathbb{R}^n} \boldsymbol{f}(\boldsymbol{x}) N(\boldsymbol{x} \mid \boldsymbol{\mu}, \boldsymbol{\Sigma}) d\boldsymbol{x}, \\
\hat{\boldsymbol{h}}(\boldsymbol{\mu}, \boldsymbol{\Sigma}) &= \int_{\mathbb{R}^n} \boldsymbol{h}(\boldsymbol{x}) N(\boldsymbol{x} \mid \boldsymbol{\mu}, \boldsymbol{\Sigma}) d\boldsymbol{x},
\end{aligned} \tag{3.16}$$

where the functions $\boldsymbol{f}$ and $\boldsymbol{h}$ are the dynamics and the measurement models in (3.10), respectively. Then, we just compute the derivatives of the desired order $q = 1, \ldots, N$ for $\boldsymbol{g} = \boldsymbol{f}, \boldsymbol{h}$

$$\hat{\boldsymbol{g}}^{(q)}_{i_1,\ldots,i_q}(\boldsymbol{\mu}, \boldsymbol{\Sigma}) = \frac{\partial^q}{\partial \mu_{i_1} \ldots \partial \mu_{i_q}} \hat{\boldsymbol{g}}(\boldsymbol{\mu}, \boldsymbol{\Sigma}),$$

and we can approximate (3.11) as

$$\begin{aligned}
\boldsymbol{m}_k^- &= \hat{\boldsymbol{f}}(\boldsymbol{m}_{k-1}, \boldsymbol{P}_{k-1}), \\
\boldsymbol{P}_k^- &= \boldsymbol{Q}_k + \sum_{q=1}^{N} \frac{1}{q!} \sum_{\substack{i_1,\ldots,i_q=1 \\ j_1,\ldots,j_q=1}}^{n} \hat{\boldsymbol{f}}^{(q)}_{i_1,\ldots,i_q}(\boldsymbol{m}, \boldsymbol{P}) \left(\prod_{r=1}^{q} [\boldsymbol{P}]_{i_r, j_r} \right) \left(\hat{\boldsymbol{f}}^{(q)}_{j_1,\ldots,j_q}(\boldsymbol{m}, \boldsymbol{P}) \right)^\top \Bigg|_{\substack{\boldsymbol{m}=\boldsymbol{m}_{k-1} \\ \boldsymbol{P}=\boldsymbol{P}_{k-1}}}, \\
\boldsymbol{\mu}_k &= \hat{\boldsymbol{h}}(\boldsymbol{m}_k^-, \boldsymbol{P}_k^-), \\
\boldsymbol{S}_k &= \boldsymbol{R}_k + \sum_{q=1}^{N} \frac{1}{q!} \sum_{\substack{i_1,\ldots,i_q=1 \\ j_1,\ldots,j_q=1}}^{n} \hat{\boldsymbol{h}}^{(q)}_{i_1,\ldots,i_q}(\boldsymbol{m}, \boldsymbol{P}) \left(\prod_{r=1}^{q} [\boldsymbol{P}]_{i_r, j_r} \right) \left(\hat{\boldsymbol{h}}^{(q)}_{j_1,\ldots,j_q}(\boldsymbol{m}, \boldsymbol{P}) \right)^\top \Bigg|_{\substack{\boldsymbol{m}=\boldsymbol{m}_k^- \\ \boldsymbol{P}=\boldsymbol{P}_k^-}}, \\
\boldsymbol{K}_k &= \boldsymbol{P}_k^- \boldsymbol{H} \boldsymbol{S}_k^{-1}, \text{ for } [\boldsymbol{H}]_{i,j} = \left[\hat{\boldsymbol{h}}_i^{(1)}(\boldsymbol{m}_k^-, \boldsymbol{P}_k^-) \right]_j, \\
\boldsymbol{m}_k &= \boldsymbol{m}_k^- + \boldsymbol{K}_k (\boldsymbol{y}_k - \boldsymbol{\mu}_k), \\
\boldsymbol{P}_k &= \boldsymbol{P}_k^- - \boldsymbol{K}_k \boldsymbol{S}_k \boldsymbol{K}_k^\top.
\end{aligned} \tag{3.17}$$

When $N \to \infty$, the matrices $\boldsymbol{P}_k^-$ and $\boldsymbol{S}_k$ are not only approximations of the corresponding matrices in (3.11) but exactly the same.

We can now apply the Fourier–Hermite Kalman filter (FHKF) to our non-linear SSM of the pendulum (3.3) and (3.4) as in Publication III. The results are presented in Figures 3.1–3.3. The horizontal axis represents time on the range $[0, 10]$, that is, discrete time steps $k = 0, 1, \ldots, 10\,000$. The vertical axis represents an angle in radians at a range of about $\left[-\frac{\pi}{2}, \frac{\pi}{2}\right]$. The

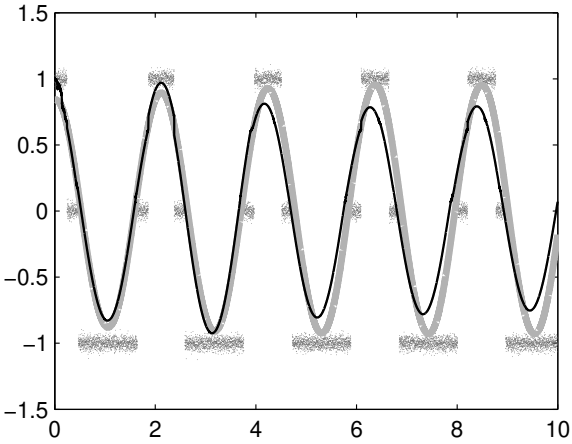

Figure 3.1. The first-order FHKF estimate of the discrete-time pendulum. The thick grey line is the true trajectory, the small grey dots are the measurements, and the thin black line is the estimate (Publication III) © 2011 IEEE.

thick grey line is the true trajectory, grey small dots are the measurements, and the thin black line is the estimated trajectory. The filter estimate is computed with the first, second, and third-order truncation of the Fourier–Hermite series and we see that the estimate is closer to the true trajectory as the order increases. The root-mean-square error is 0.16, 0.04, and 0.03 for the first, second, and third-order, respectively, so the improvement between the second and third-order filter is already quite small, which is visible in Figures 3.2–3.3 as well. We see that despite very inaccurate measurements, by using the knowledge of the dynamics, it is possible to reconstruct an accurate estimate of the true trajectory.

3.5 Discrete-Time Smoothing

Next, we look at the smoothing equations and present the approximate Gaussian solution in Publication IV. For the discrete-time smoothing, we use the same Bayes' rule (3.5) and the Markov property (3.6) as for the filtering equations (3.7), except future observations are also used instead of only observations until the timestep of the estimate (Särkkä and Svensson, 2023, Theorem 12.1)

$$
\begin{aligned}
p(\boldsymbol{x}_{k+1} \mid \boldsymbol{y}_{1:k}) &= \int p(\boldsymbol{x}_{k+1} \mid \boldsymbol{x}_k)\, p(\boldsymbol{x}_k \mid \boldsymbol{y}_{1:k})\, d\boldsymbol{x}_k, \\
p(\boldsymbol{x}_k \mid \boldsymbol{y}_{1:T}) &= p(\boldsymbol{x}_k \mid \boldsymbol{y}_{1:k}) \int \frac{p(\boldsymbol{x}_{k+1} \mid \boldsymbol{x}_k)\, p(\boldsymbol{x}_{k+1} \mid \boldsymbol{y}_{1:T})}{p(\boldsymbol{x}_{k+1} \mid \boldsymbol{y}_{1:k})}\, d\boldsymbol{x}_{k+1},
\end{aligned}
\tag{3.18}
$$

where we assume to have the filtering solution $p(\boldsymbol{x}_k \mid \boldsymbol{y}_{1:k})$ of (3.7) for time steps $1, 2, \ldots T$ and compute recursively backwards in time from steps T –

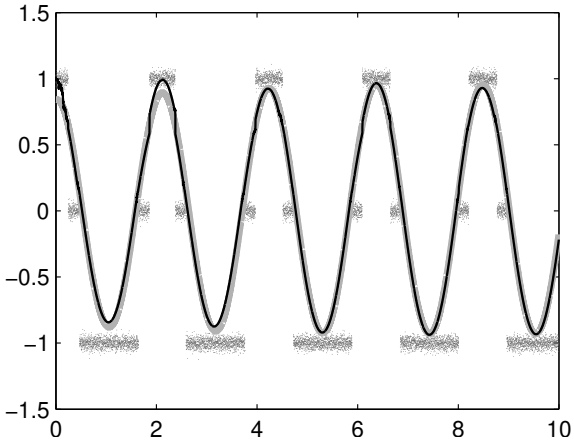

Figure 3.2. The second-order FHKF estimate of the discrete-time pendulum. The thick grey line is the true trajectory, the small grey dots are the measurements, and the thin black line is the estimate (Publication III) © 2011 IEEE.

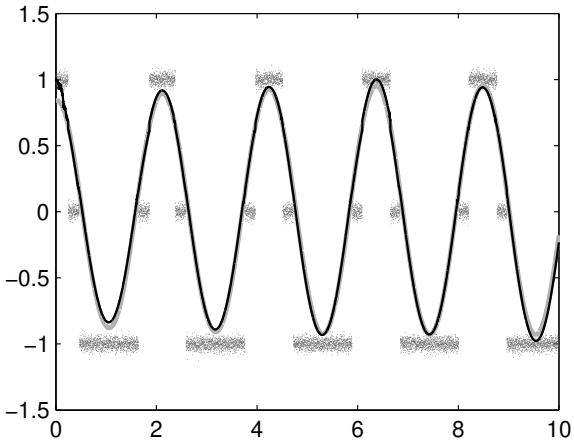

Figure 3.3. The third-order FHKF estimate of the discrete-time pendulum. The thick grey line is the true trajectory, the small grey dots are the measurements, and the thin black line is the estimate (Publication III) © 2011 IEEE.

$1, T-2, \ldots k+1$ based on the $p(\boldsymbol{x}_{k+1} \mid \boldsymbol{y}_{1:T})$. More precisely, these equations are for fixed-interval smoother. There are also fixed-point and fixed-lag smoothers, which have similar equations (Särkkä and Svensson, 2023).

With the linear Gaussian SSM (3.8), the solution to the smoother equations (3.18) is given by the Rauch–Tung–Striebel smoother (RTSS) (Rauch et al. 1965;Särkkä and Svensson 2023, Theorem 12.2)

$$\begin{aligned}
\boldsymbol{m}_{k+1}^{-} &= \boldsymbol{A}_k \boldsymbol{m}_k, \\
\boldsymbol{P}_{k+1}^{-} &= \boldsymbol{A}_k \boldsymbol{P}_k \boldsymbol{A}_k^\top + \boldsymbol{Q}_k, \\
\boldsymbol{G}_k &= \boldsymbol{P}_k \boldsymbol{A}_k^\top [\boldsymbol{P}_{k+1}^{-}]^{-1}, \\
\boldsymbol{m}_k^s &= \boldsymbol{m}_k + \boldsymbol{G}_k (\boldsymbol{m}_{k+1}^s - \boldsymbol{m}_{k+1}^{-}), \\
\boldsymbol{P}_k^s &= \boldsymbol{P}_k + \boldsymbol{G}_k (\boldsymbol{P}_{k+1}^s - \boldsymbol{P}_{k+1}^{-}) \boldsymbol{G}_k^\top,
\end{aligned} \tag{3.19}$$

where $\boldsymbol{m}_k$ and $\boldsymbol{P}_k$ are computed by the KF (3.9) and the recursion is running backwards starting from the last time step T with $\boldsymbol{m}_T^s = \boldsymbol{m}_T$ and $\boldsymbol{P}_T^s = \boldsymbol{P}_T$. Notice that also $\boldsymbol{m}_{k+1}^{-}$ and $\boldsymbol{P}_{k+1}^{-}$ are the same as in the KF(3.9).

In Publication IV, an approximate solution of smoothing is extended to non-linear SSM (3.10) using the relation between Fourier–Hermite series and the expectations (3.14) and it is called Fourier–Hermite Rauch–Tung–Striebel smoother (FHRTSS). The equations are similar to the RTSS (3.19) with a backwards recursion for $\boldsymbol{m}_k^s$ and $\boldsymbol{P}_k^s$ as

$$\begin{aligned}
\boldsymbol{G}_k &= \boldsymbol{P}_k \boldsymbol{F} [\boldsymbol{P}_{k+1}^{-}]^{-1}, \quad \text{for} \quad [\boldsymbol{F}]_{i,j} = \left[\hat{\boldsymbol{f}}_i^{(1)}(\boldsymbol{m}_k, \boldsymbol{P}_k) \right]_j, \\
\boldsymbol{m}_k^s &= \boldsymbol{m}_k + \boldsymbol{G}_k (\boldsymbol{m}_{k+1}^s - \boldsymbol{m}_{k+1}^{-}), \\
\boldsymbol{P}_k^s &= \boldsymbol{P}_k + \boldsymbol{G}_k (\boldsymbol{P}_{k+1}^s - \boldsymbol{P}_{k+1}^{-}) \boldsymbol{G}_k^\top,
\end{aligned}$$

where $\boldsymbol{m}_k$, $\boldsymbol{m}_{k+1}^{-}$, $\boldsymbol{P}_k$, and $\boldsymbol{P}_{k+1}^{-}$ are the same as for the FHKF (3.17), and $\hat{\boldsymbol{f}}$ is the Weierstrass transform (3.16) of the dynamics model function.

We demonstrate smoothing with the same pendulum SSM as before and also in Publication IV. The results are shown in Figures 3.4–3.6. The measurements are again in small grey dots. The true trajectory is in a blue solid line. The filter estimate of Section 3.4 is in a solid green line, and the smoother estimate in a red dashed line. Figures 3.4–3.6 show that the results are better for higher-order filters and smoothers, that is, the order of the truncation of the Fourier–Hermite series. The results are also always better for the smoother solution at any order, increasingly closer to timestep $k = 0$ and away from $k = T$. Remember that at $k = T$, the smoother and filter solutions are the same. This is even more clearly demonstrated in Table 3.1, which shows root-mean-square error for different order filters and smoothers.

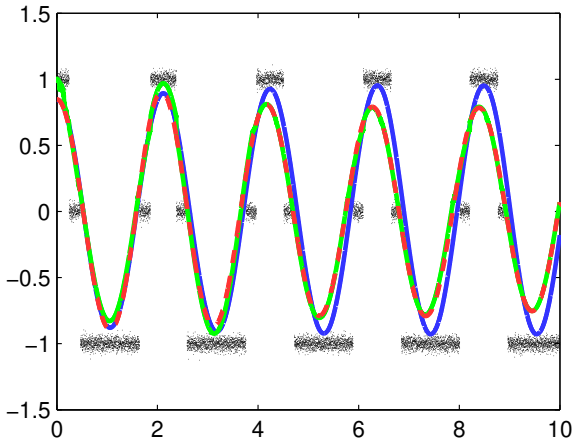

Figure 3.4. The first-order FHKF and FHRTSS for the pendulum. The thick blue line is the true trajectory, the small grey dots the measurements, the FHKF estimate the solid green line, and FHRTSS the red dashed line (Publication IV).

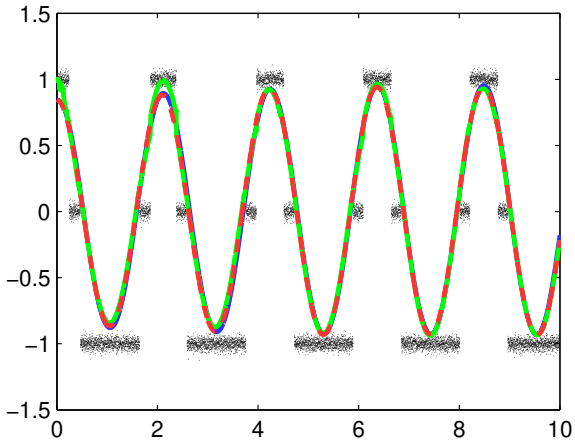

Figure 3.5. The second-order FHKF and FHRTSS for the pendulum. The thick blue line is the true trajectory, the small grey dots the measurements, the FHKF estimate the solid green line, and FHRTSS the red dashed line (Publication IV).

Table 3.1. The root-mean-square error for the FHKF and FHRTSS (Publication III).

Order	FHKF	FHRTSS
1	0.164	0.160
2	0.041	0.025
3	0.031	0.018

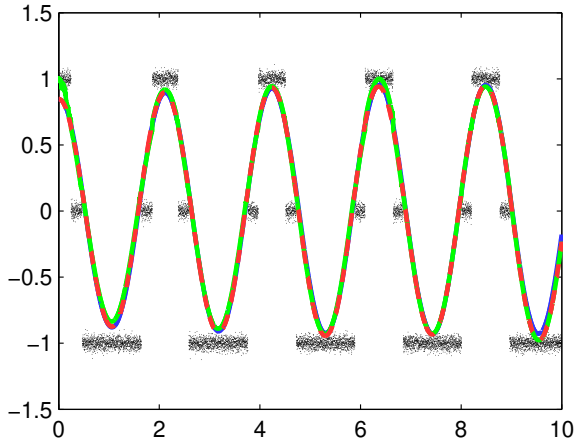

Figure 3.6. The third-order FHKF and FHRTSS for the pendulum. The thick blue line is the true trajectory, the small grey dots the measurements, the FHKF estimate the solid green line, and FHRTSS the red dashed line (Publication IV).

3.6 Itô Calculus for Continuous-Time State-Space Models

In the following sections, we change from discrete-time models to continuous-discrete-time models (Jazwinski, 2007; Särkkä and Solin, 2019), which are the topic of Publication V. Continuous-discrete-time means that the dynamics model is continuous and the measurements are obtained at discrete time instants. The continuous dynamics model is a stochastic differential equation (SDE) which consists of deterministic and stochastic parts and is, in fact, formulated in terms of stochastic Itô integrals rather than differentials (Øksendal, 2003; Jazwinski, 2007; Särkkä and Solin, 2019). If we would have a perfectly deterministic system, the ordinary differential equation (ODE) describing it would be

$$\frac{d\boldsymbol{x}(t)}{dt} = \boldsymbol{f}(\boldsymbol{x}(t), t), \tag{3.20}$$

where $t \in [0, \infty)$ is time, $\boldsymbol{x}(t) \in \mathbb{R}^n$ is the state of the system at time t, and $\boldsymbol{f} : \mathbb{R}^n \times [0, \infty) \mapsto \mathbb{R}^n$ describes the deterministic dynamics. We then add a stochastic component into the differential equation.

The idea is that the system is not perfectly deterministic but also stochastic. The deterministic part is perfectly predictable. The stochastic part is unpredictable in its realisation of the path but nevertheless does have a predictable, deterministic probability distribution. In practical applications, we usually can not identify the deterministic part of the system perfectly. We must make some degree of modelling choice about what parts of the real-world to try to embed into the mathematical model of it. The real-world part which is not present in the deterministic part of the mathematical model is then a possible outcome of the stochastic part if it

is rich enough to include a wide range of possibilities.

When we add the stochastic component to the differential equation, we face difficulty in defining the differential equation in terms of time derivatives because the stochastic component must be such that it is not differentiable. A differentiable function of time would be arguably too predictable. Loosely said, this follows because a differentiable function has a bounded variation (Apostol, 1974, Theorem 6.6), and a local martingale is a constant if it has a bounded variation (Kallenberg, 2002, Proposition 17.2). For an explanation of why we like the random continuous-time part to be a martingale, see (Kailath and Poor, 1998, Section II A–D).

In discrete-time dynamics, we could derive the equations using the end of the 18th century Bayes' rule, but in continuous-time, we have to use much more recent 1940s Itô integral theory and formulate the SDE in terms of Itô integrals. Example textbooks of stochastic calculus that cover the Itô integrals are (Kallenberg, 2002; Øksendal, 2003). For application of this to filtering and smoothing, see (Jazwinski, 2007; Särkkä and Solin, 2019).

First, we convert the differential equation (3.20) into an integral form, such that for any $t_0, t_1 \in [0,\infty)$, we have

$$\int_{t_0}^{t_1} d\boldsymbol{x}(t) = \boldsymbol{x}(t_1) - \boldsymbol{x}(t_0) = \int_{t_0}^{t_1} \boldsymbol{f}(\boldsymbol{x}(t), t)\, dt$$

or in a shorter notation

$$d\boldsymbol{x}(t) = \boldsymbol{f}(\boldsymbol{x}(t), t)\, dt.$$

When we add a stochastic component to the equation, we have

$$\int_{t_0}^{t_1} d\boldsymbol{x}(t) = \boldsymbol{x}(t_1) - \boldsymbol{x}(t_0) = \int_{t_0}^{t_1} \boldsymbol{f}(\boldsymbol{x}(t), t)\, dt + \int_{t_0}^{t_1} \boldsymbol{L}(\boldsymbol{x}(t), t)\, d\boldsymbol{w}(t)$$

or in a shorter form

$$d\boldsymbol{x}(t) = \boldsymbol{f}(\boldsymbol{x}(t), t)\, dt + \boldsymbol{L}(\boldsymbol{x}(t), t)\, d\boldsymbol{w}(t) \tag{3.21}$$

where $t, \boldsymbol{x}$, and $\boldsymbol{f}$ are as before and $\boldsymbol{L} : \mathbb{R}^n \times [0,\infty) \mapsto \mathbb{R}^n \times \mathbb{R}^s$ is a new deterministic matrix function, but $\boldsymbol{w}(t)$ is a new kind of entity, a non-deterministic function of time, an s-dimensional Wiener process. The integral with respect to $\boldsymbol{w}(t)$ is an Itô integral.

The Wiener process $\boldsymbol{w}(t)$ is defined by the following properties:

1. $\boldsymbol{w}(0) = \boldsymbol{0}$.

2. The increments $\Delta\boldsymbol{w}_k = \boldsymbol{w}(t_{k+1}) - \boldsymbol{w}(t_k)$ are zero mean Gaussian random variables with covariance $\boldsymbol{Q}\,\Delta t_k$, where $\boldsymbol{Q}$ is a Hermitian positive definite diffusion matrix and $\Delta t_k = t_{k+1} - t_k > 0$.

3. The increments $\Delta\boldsymbol{w}_k$ are independent of earlier values of $\boldsymbol{w}(t)$ for $t < t_k < t_{k+1}$.

The Itô integral is

$$\int_{t_0}^{t} \boldsymbol{L}(\boldsymbol{x}(\tau),\tau)\, d\boldsymbol{w}(\tau) = \lim_{N\to\infty} \sum_{k=0}^{N-1} \boldsymbol{L}(\boldsymbol{x}(t_k), t_k)\, \Delta\boldsymbol{w}_k,$$

where $t_0 < t_1 < \ldots < t_N = t$.

The Wiener process $\boldsymbol{w}(t)$ does have the properties that we wanted. It does have unpredictable paths. Any continuous path that starts from $\mathbf{0}$ is a possible realisation of it, although with probability 0. The probability distribution of the path is predictable, that is,

$$\boldsymbol{w}(t) \mid \boldsymbol{w}(t_0) \sim N(\boldsymbol{w}(t_0), (t-t_0)\boldsymbol{Q}) \tag{3.22}$$

for $t \geq t_0 \geq 0$. Although we noticed that the path of the Wiener process is very unpredictable, in stochastic process terminology, it is a predictable process like all continuous local martingales (Kallenberg, 2002, Proposition 25.16). The technical term predictable process refers to the predictability of the value of the process at any time instant t given the knowledge of the value before t. In continuous time, this before t can be arbitrarily short time before, and hence the continuous-time stochastic process is predictable if it is left-continuous (Kallenberg, 2002, Chapter 25). From (3.22), we see that, in this sense, the Wiener process is predictable.

3.7 Continuous-Discrete-Time Filtering

Now, we present the continuous-discrete-time SSM using the Itô calculus of the previous section. The dynamic model is continuous as in (3.21) with some initial distribution for the state $\boldsymbol{x}(0) \sim p(\boldsymbol{x}(0))$ that is independent of the Wiener process $\boldsymbol{w}(t)$ of (3.21). The m-dimensional measurement part is discrete-time and has the Markov property with respect to the state at time t_k

$$\boldsymbol{y}_k \sim p(\boldsymbol{y}_k \mid \boldsymbol{x}(t_k)).$$

In continuous-discrete-time filtering, we are interested in solving the continuous-time state trajectory $\boldsymbol{x}(t_k)$ at time steps $t_1, t_2, \ldots, t_k$ of discrete-time measurements $\boldsymbol{y}_1, \boldsymbol{y}_2, \ldots, \boldsymbol{y}_k$, that is,

$$p(\boldsymbol{x}(t_k) \mid \boldsymbol{y}_{1:k}). \tag{3.23}$$

A recursive solution is given as in (Särkkä and Solin, 2019, Algorithm 10.15). For each timestep $k = 1, 2, \ldots$, we first solve the prediction probability distribution

$$p(\boldsymbol{x}(t_k) \mid \boldsymbol{y}_{1:k-1}) = \int p(\boldsymbol{x}(t_k) \mid \boldsymbol{x}(t_{k-1}))\, p(\boldsymbol{x}(t_{k-1}) \mid \boldsymbol{y}_{1:k-1})\, d\boldsymbol{x}(t_{k-1}). \tag{3.24}$$

Second, we compute the posterior distribution by Bayes' rule

$$p(\boldsymbol{x}(t_k) \mid \boldsymbol{y}_{1:k}) = \frac{1}{Z_k} p(\boldsymbol{y}_k \mid \boldsymbol{x}(t_k)) p(\boldsymbol{x}(t_k) \mid \boldsymbol{y}_{1:k-1}), \tag{3.25}$$

where the normalisation constant Z_k is

$$Z_k = \int p(\boldsymbol{y}_k \mid \boldsymbol{x}(t_k)) p(\boldsymbol{x}(t_k) \mid \boldsymbol{y}_{1:k-1}) \, d\boldsymbol{x}(t_k). \tag{3.26}$$

In general, we can not solve the continuous-discrete-time filtering integrals (3.24)–(3.26) in closed form, as well as the transition probability density $p(\boldsymbol{x}(t_k) \mid \boldsymbol{x}(t_{k-1}))$ in (3.24). For practical numerical solutions, we need some kind of approximations. It is possible to use Gaussian approximations, that is, two first moments, the mean and the covariance as in Publications III and IV. The moments can be solved with Itô's formula (Jazwinski 2007; Øksendal 2003; Särkkä and Solin 2019) or, for example, with the Taylor moment expansion (TME) (Zhao et al., 2021; Zhao, 2021). Then the transition probabilities $p(\boldsymbol{x}(t_k) \mid \boldsymbol{x}(t_{k-1}))$ in (3.24) are also approximated as Gaussian. In Publication V, we the TME to compute arbitrary order moments and the product rule of the discrete approximations of the spectral measures (2.12).

It is also possible to compute transition probabilities $p(\boldsymbol{x}(t_k) \mid \boldsymbol{x}(t_{k-1}))$ in (3.24) using the Fokker–Planck–Kolmogorov partial differential equation (Jazwinski, 2007; Øksendal, 2003; Särkkä and Solin, 2019), also known as the Kolmogorov forward equation. There are also smoothing equations in continuous-discrete time (Jazwinski, 2007; Särkkä and Sarmavuori, 2013; Särkkä and Solin, 2019).

3.8 Higher Moments in Continuous-Discrete-Time Filtering

The Gaussian filter uses only the two lowest moments (mean and covariance) for the solution. In Publication V, an arbitrary number of moments is taken into account, and asymptotically, the filter would, in principle, provide the exact solution when the number of moments approaches infinity. The contribution is to use (2.12) to form an arbitrary order quadrature rule based on an arbitrary number of moments from an approximation of the filtering distribution $p(\boldsymbol{x}(t_k) \mid \boldsymbol{y}_{1:k})$ in (3.23). The moments are then updated accordingly using the quadrature rule. In principle, if we would have an infinite number of moments, they would uniquely define a continuous distribution. In practice, we use the quadrature rule based on a finite number of moments as an approximation of the filtering distribution $p(\boldsymbol{x}(t_k) \mid \boldsymbol{y}_{1:k})$.

We use the multi-index notation $\boldsymbol{n} \in \mathbb{N}_0^d$ with the usual definition of order $|\boldsymbol{n}| = \sum_{k=1}^{d} n_k$ (Reed and Simon 1981, Chapter IX.1;Dunkl and Xu 2014, Chapter 3.1). We define an arbitrary order moment of a d-dimensional

random variable $\boldsymbol{x}$ for a multi-index $\boldsymbol{n} \in \mathbb{N}_0^d$ as

$$m_{\boldsymbol{n}} = \mathrm{E}\left[\boldsymbol{x}^{\boldsymbol{n}}\right] = \mathrm{E}\left[\boldsymbol{x}_1^{n_1} \boldsymbol{x}_2^{n_2} \dots \boldsymbol{x}_d^{n_d}\right] = \int \boldsymbol{x}_1^{n_1} \boldsymbol{x}_2^{n_2} \dots \boldsymbol{x}_d^{n_d} \, p(\boldsymbol{x}) \, d\boldsymbol{x}.$$

In practical computations, we use a finite number of moments for some finite value N so that $|\boldsymbol{n}| \leq N$.

For filtering, we start with the initial moments and then compute, alternatingly, two different conditional moments for prediction and update steps. The prediction step conditional moments are

$$m_{\boldsymbol{n},k|k-1} = \mathrm{E}\left[\boldsymbol{x}(t_k)^{\boldsymbol{n}} \,\middle|\, \boldsymbol{y}_{1:k-1}\right] = \int \mathrm{E}\left[\boldsymbol{x}(t_k)^{\boldsymbol{n}} \,\middle|\, \boldsymbol{x}(t_{k-1})\right] \, p(\boldsymbol{x}(t_{k-1}) \mid \boldsymbol{y}_{1:k-1}) \, d\boldsymbol{x}(t_{k-1}),$$

where, for approximating $\mathrm{E}[\boldsymbol{x}(t_k)^{\boldsymbol{n}} \mid \boldsymbol{x}(t_{k-1})]$, we use Taylor moment expansion (TME) (Zhao et al., 2021; Zhao, 2021). The update conditional moments are computed by Bayes' rule

$$\begin{aligned} m_{\boldsymbol{n},k|k} &= \mathrm{E}\left[\boldsymbol{x}(t_k)^{\boldsymbol{n}} \,\middle|\, \boldsymbol{y}_{1:k}\right] \\ &= \int \boldsymbol{x}(t_k)^{\boldsymbol{n}} \, p(\boldsymbol{x}(t_k) \mid \boldsymbol{y}_{1:k}) \, d\boldsymbol{x}(t_k) \\ &= \frac{\int \boldsymbol{x}(t_k)^{\boldsymbol{n}} \, p(\boldsymbol{y}_k \mid \boldsymbol{x}(t_k)) \, p(\boldsymbol{x}(t_k) \mid \boldsymbol{y}_{1:k-1}) \, d\boldsymbol{x}(t_k)}{\int p(\boldsymbol{y}_k \mid \boldsymbol{x}(t_k)) \, p(\boldsymbol{x}(t_k) \mid \boldsymbol{y}_{1:k-1}) \, d\boldsymbol{x}(t_k)}. \end{aligned}$$

For approximating the integral of prediction moments $m_{\boldsymbol{n},k|k-1}$, we use the quadrature of (2.12) from a finite number of approximate initial moments on the first step $k = 1$ or from the previous update step approximations of $m_{\boldsymbol{n},k-1|k-1}$ on the later prediction steps $k > 1$. For the approximation of the update moments $m_{\boldsymbol{n},k|k}$, we use the quadrature from a finite number of approximate prediction moments $m_{\boldsymbol{n},k|k-1}$.

As an example, we use the same two-dimensional stochastic prey-predator model (Arató, 2003) as in Publication V

$$\begin{aligned} dx_1(t) &= x_1(t)(\alpha - \beta x_2(t)) \, dt + \sigma x_1(t) \, dw_1(t), \\ dx_2(t) &= x_2(t)(\zeta x_1(t) - \gamma) \, dt + \sigma x_2(t) \, dw_2(t), \end{aligned}$$

where $\alpha = \beta = \zeta = \gamma = 4$ and $\sigma = 0.1$. Variable $x_1(t)$ represents the number of preys and $x_2(t)$ the number of predators. The initial value is a mixed Gaussian random variable with small covariances $\boldsymbol{x}(0) \sim 0.5\, N([1\ 1]^\top, 10^{-3} I) + 0.5\, N([1\ 1]^\top, 2 \cdot 10^{-3} I)$. The measurements y_k come from a discrete Poisson distribution

$$p(y_k \mid x_1(t_k)) = \frac{\lambda_k^{y_k} e^{-\lambda_k}}{y_k!} \quad \text{for} \quad y_k \in \mathbb{N}_0,$$

where

$$\lambda_k = \frac{1}{1 + e^{1 - x_1(t_k)^3}}$$

and $t_k = 10^{-3} k$.

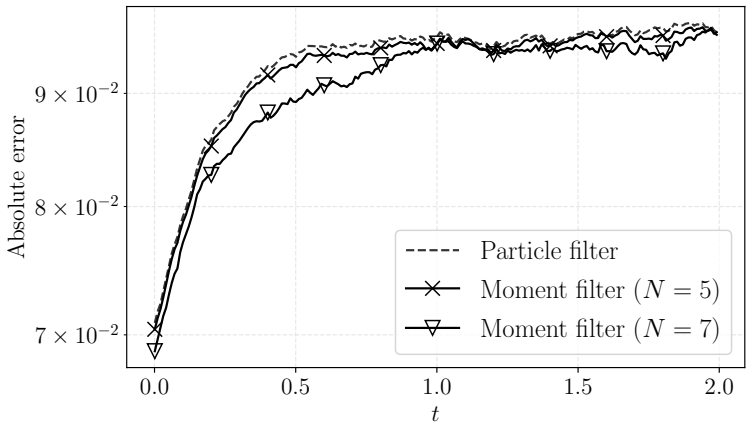

Figure 3.7. Estimation errors for the prey-predator SSM (Publication V).

We compare the performance of the moment filter (MF) with the bootstrap particle filter (PF) using 10,000 particles with stratified resampling (Särkkä and Svensson, 2023, Algorithm 11.6). For comparison quantity, we use the average of the absolute error of the filter estimate of $x_1(t_k)$ for $k = 1,\ldots,2000$ and 10,000 Monte Carlo runs. The results are shown in Figure 3.7 for the PF and MF with $N = 5$ and $N = 7$. Although the MF estimates are more accurate and it is computationally lighter, it has the problem of sometimes diverging. Out of the 10,000 MC runs, 909 and 4,746 runs diverged for $N = 5$ and $N = 7$, respectively.

4. Summary and Discussion

In this chapter, we present a summary of the publications included in this thesis and discussion about future research on these topics.

4.1 Summary of Publications

This section presents a summary of Publications I–V.

Publication I (Chapter 2)

This article shows how integrals can be numerically approximated as a function of a finite Hermitian matrix. It is shown that Gaussian quadrature is a special case of this new method. It is shown that the new method has similar properties to Gaussian quadrature, like the positivity of the weights and similar location properties of the nodes. In this article, all the proofs are applications of classical theorems of operator theory. For proof of convergence, a classical theorem is also used, but it is limited to bounded functions and does not cover all Riemann–Stieltjes integrable functions. The main significance is that similar matrix methods, that are effective for computing Gaussian quadrature, can be used for more general quadrature methods, and these more general quadratures share many properties of Gaussian quadratures.

Publication II (Chapter 2)

This article focuses on proofs of convergence for the new numerical integration method. In Publication I, the convergence theorem covered only bounded functions in limited form, which did not cover all Riemann–Stieltjes integrable functions. Here, the convergence proofs are extended to cover bounded Riemann–Stieltjes integrable functions and many cases of unbounded functions. Existing theorems from operator theory are used as far as possible, but several new theorems, lemmas, and inequalities have

to be developed in order to close some gaps in the pre-existing theory. The main significance is filling the missing pieces in the operator theory for establishing some theory of the convergence for unbounded functions. This has some significance, for example, for Gaussian filtering and smoothing, because Gaussian probability distribution has unbounded support and therefore many functions, like polynomials or even sublinear functions are unbounded.

Publication III (Section 3.4)

In this article, the Fourier–Hermite series is used for Gaussian approximation of nonlinear filtering. It is shown how the Fourier–Hermite coefficients can be computed as partial derivatives of a linear transformation (Weierstrass transformation) of the dynamics and measurement model functions.

Publication IV (Section 3.5)

The same methods, as in the previous article, are applied for smoothing. The performance of equal order filtering and smoothing are compared in the example problem.

Publication V (Section 3.8)

From the point of view of this thesis, the main significance of this article is the generalisation of product rule for dependent random variables by using the operator theoretic approach to numerical integration developed in Publications I and II. This novel method is combined with recently developed TME (Zhao et al., 2021; Zhao, 2021) for continuous-discrete-time filtering. Also, log-likelihoods of the filtering solution are used for parameter estimation.

4.2 Future Work

The most interesting future research topic seems like improving the generalised product rule introduced in Publication V. Instead of using the eigenvalue decomposition, more efficient results could be obtained by using Lanczos iteration (Golub and Meurant, 2009). With the Lanczos iteration, the idea is to use a smaller set of Ritz values and corresponding vectors instead of the eigenvalues and eigenvectors. Some sparse grid methods would improve the efficiency as well. It should be possible to develop a numerical approximation of spectral measures on the unit circle as an extension for quadratures on the unit circle (Gragg, 1993; Simon, 2009; Velázquez, 2008; Bultheel et al., 2015). These could be further developed

for multidimensional numerical integration rule as the product of the radial part and the unit circle parts. It could be investigated how to approximate the spectral measures for the multiplication operators with matrices that have the same eigenvectors. This would guarantee positive weights.

The filter and smoother solutions in Publications III and IV are limited in that they require the closed-form solution of the Weierstrass transform for the dynamics and measurement functions. Based on Publication I, a new method could be developed for numerical approximation of Fourier–Hermite series. This method could be used as part of a Gaussian mixture model (Murata and Hiramatsu, 2019; Raitoharju et al., 2019) or posterior linearisation (García-Fernández et al., 2015). This numerical integration method could also be used for Wiener chaos expansions (Lototsky, 2011).

The new numerical integration method of Publication I lacks some basic theoretical results like error estimates and convergence rates. These are important future research topics.

There are numerous applications of nonlinear SSMs for which the method of Publication I could be used. The new method has lots of degrees of freedom with respect to what basis functions to use and how to decompose the nonlinear function into inside and outside functions $h(\boldsymbol{x}) = f(g(\boldsymbol{x}))$. Different choices of these degrees of freedom could prove particularly useful depending on practical applications.

References

O. M. Agudelo, O. Barrero, V. Peter, and B. De Moor. Assimilation of ozone measurements in the air quality model AURORA by using the ensemble Kalman filter. In *2011 50th IEEE Conference on Decision and Control and European Control Conference*, pages 4430–4435. IEEE, 2011.

R. Ahlfeld, B. Belkouchi, and F. Montomoli. SAMBA: Sparse approximation of moment-based arbitrary polynomial chaos. *Journal of Computational Physics*, 320:1–16, 2016.

N. Akhiezer. *The Classical Moment Problem and Some Related Questions in Analysis*. University mathematical monographs. Oliver & Boyd, 1965.

N. Akhiezer and I. Glazman. *Theory of Linear Operators in Hilbert Space*. Dover Books on Mathematics. Dover Publications, New York, 1993.

T. Apostol. *Mathematical Analysis*. World student series edition. Addison-Wesley, Reading, Mass., second edition, 1974.

M. Arató. A famous nonlinear stochastic equation (Lotka–Volterra model with diffusion). *Mathematical and Computer Modelling*, 38(7):709–726, 2003.

L. E. Baum and T. Petrie. Statistical inference for probabilistic functions of finite state Markov chains. *The Annals of Mathematical Statistics*, 37(6):1554–1563, 1966.

V. E. Beneš. Exact finite-dimensional filters for certain diffusions with nonlinear drift. *Stochastics*, 5:65–92, 1981.

A. Bryson and M. Frazier. Smoothing for linear and nonlinear dynamic systems. In *Proceedings of the optimum system synthesis conference*, pages 353–364. DTIC Document Ohio, 1963.

A. Bultheel, M. J. Cantero, and R. Cruz-Barroso. Matrix methods for quadrature formulas on the unit circle. A survey. *Journal of Computational and Applied Mathematics*, 284:78–100, 2015.

S. Challa, Y. Bar-Shalom, and V. Krishnamurthy. Nonlinear filtering via generalized Edgeworth series and Gauss-Hermite quadrature. *IEEE Transactions on Signal Processing*, 48(6), June 2000.

P. Date and K. Ponomareva. Linear and non-linear filtering in mathematical finance: a review. *IMA Journal of Management Mathematics*, 22(3):195–211, 2010.

P. Davis and P. Rabinowitz. *Methods of Numerical Integration*. Computer Science and Applied Mathematics. Academic Press, San Diego, California, 2nd edition, 1984.

B. De Finetti. *Theory of Probability: A Critical Introductory Treatment*. Wiley series in probability and statistics. John Wiley & Sons, 2017.

A. Deshmukh. Comparison of hidden Markov model and recurrent neural network in automatic speech recognition. *European Journal of Engineering and Technology Research*, 5(8):958–965, 2020.

S. Dubinkina and H. Goosse. An assessment of particle filtering methods and nudging for climate state reconstructions. *Climate of the Past*, 9(3):1141–1152, 2013.

C. Dunkl and Y. Xu. *Orthogonal Polynomials of Several Variables*. Encyclopedia of Mathematics and its Applications. Cambridge University Press, Cambridge, second edition, 2014.

Á. F. García-Fernández, L. Svensson, M. R. Morelande, and S. Särkkä. Posterior linearization filter: Principles and implementation using sigma points. *IEEE Transactions on Signal Processing*, 63(20):5561–5573, 2015.

W. Gautschi. *Orthogonal Polynomials: Computation and Approximation*. Numerical mathematics and scientific computation. Oxford University Press, Oxford, 2004.

A. Gelb, editor. *Applied Optimal Estimation*. The MIT Press, 1974.

A. Giuliani, V. Cocco, S. Mereghetti, C. Pittori, and M. Tavani. The agile on-board Kalman filter. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 568(2): 692–699, 2006.

T. Glad and L. Ljung. *Control Theory*. CRC Press, 2018.

G. Golub and G. Meurant. *Matrices, Moments and Quadrature with Applications*. Princeton Series in Applied Mathematics. Princeton University Press, 2009.

G. H. Golub and J. H. Welsch. Calculation of Gauss quadrature rules. *Mathematics of Computation*, 23(106):221–230, 1969.

A. Gorad, S. Hassan, and S. Särkkä. Vessel bearing estimation using visible and thermal imaging. In R. Gade, M. Felsberg, and J.-K. Kämäräinen, editors, *Image Analysis*, pages 373–381, Cham, 2023. Springer Nature Switzerland.

W. B. Gragg. Positive definite Toeplitz matrices, the Arnoldi process for isometric operators, and Gaussian quadrature on the unit circle. *J. Comput. Appl. Math.*, 46(1–2):183–198, June 1993.

S. Helm, M. Kozek, and S. Jakubek. Combustion torque estimation and misfire detection for calibration of combustion engines by parametric Kalman filtering. *IEEE Transactions on Industrial Electronics*, 59(11):4326–4337, 2012.

R. Herman. *An Introduction to Fourier Analysis*. A Chapman & Hall Book. CRC Press, 2017.

N. Higham. *Functions of Matrices: Theory and Computation*. Society for Industrial and Applied Mathematics, 2008.

Y. C. Ho and R. C. K. Lee. A Bayesian approach to problems in stochastic estimation and control. *IEEE Transactions on Automatic Control*, 9:333–339, 1964.

R. Hostettler, T. Lumikari, L. Palva, T. Nieminen, and S. Särkkä. Motion artifact reduction in ambulatory electrocardiography using inertial measurement units and Kalman filtering. In *2018 21st International Conference on Information Fusion (FUSION)*, pages 1–8, 2018.

K. Ito and K. Xiong. Gaussian filters for nonlinear filtering problems. *IEEE Transactions on Automatic Control*, 45:910–927, 2000.

A. H. Jazwinski. *Stochastic Processes and Filtering Theory*. Dover Books on Electrical Engineering Series. Dover Publications, 2007.

B. Jia and M. Xin. Refined nonlinear Gaussian quadrature filter. In *2019 American Control Conference (ACC)*, pages 5366–5371, 2019.

S. J. Julier, J. K. Uhlmann, and H. F. Durrant-Whyte. A new method for the nonlinear transformation of means and covariances in filters and estimators. *IEEE Transactions on Automatic Control*, 45(3):477–482, March 2000.

T. Kailath and H. Poor. Detection of stochastic processes. *IEEE Transactions on Information Theory*, 44(6):2230–2231, 1998.

O. Kallenberg. *Foundations of Modern Probability*. Probability and Its Applications. Springer New York, 2002.

R. E. Kalman. A new approach to linear filtering and prediction problems. *Transactions of the ASME, Journal of Basic Engineering*, 82:35–45, March 1960.

T. Kato. *Perturbation Theory for Linear Operators*. Classics in Mathematics. Springer Berlin Heidelberg, Berlin, Heidelberg, second edition, 1995.

R. Kianfar, B. Augusto, A. Ebadighajari, U. Hakeem, J. Nilsson, A. Raza, R. S. Tabar, N. V. Irukulapati, C. Englund, P. Falcone, S. Papanastasiou, L. Svensson, and H. Wymeersch. Design and experimental validation of a cooperative driving system in the grand cooperative driving challenge. *IEEE Transactions on Intelligent Transportation Systems*, 13(3):994–1007, 2012.

E. Kreyszig. *Introductory Functional Analysis with Applications*. Wiley Classics Library. John Wiley & Sons, Inc., New York, 1989.

K. Kumar, S. Bhaumik, and S. Arulampalam. Tracking an underwater object with unknown sensor noise covariance using orthogonal polynomial filters. *Sensors*, 22(13), 2022.

J. Kupiec. Hidden Markov estimation for unrestricted stochastic context-free grammars. In *Acoustics, Speech, and Signal Processing, IEEE International Conference on*, volume 1, pages 177–180, Los Alamitos, CA, USA, 1992. IEEE Computer Society.

H. J. Kushner. On the differential equations satisfied by conditional probability densities of Markov processes, with applications. *J. SIAM Control Ser. A*, 2(1), 1964.

C. T. Leondes, J. B. Peller, and E. B. Stear. Nonlinear smoothing theory. *IEEE Transactions on Systems Science and Cybernetics*, 6(1), January 1970.

S. V. Lototsky. Chaos approach to nonlinear filtering. In D. Crisan and B. Rozovskii, editors, *The Oxford Handbook of Nonlinear Filtering*, pages 231–264. Oxford University Press, 2011.

P. Malliavin. *Stochastic Analysis*. Springer, 1997.

M. M. Mansouri, H. N. Nounou, M. N. Nounou, and A. A. Datta. State and parameter estimation for nonlinear biological phenomena modeled by S-systems. *Digital Signal Processing*, 28:1–17, 2014.

J. L. Maryak, J. C. Spall, and B. D. Heydon. Use of the Kalman filter for inference in state-space models with unknown noise distribution. *IEEE Transactions on Automatic Control*, 49(1):87–90, Jan. 2004.

P. Maybeck. *Stochastic Models, Estimation and Control, Volume 2*. Academic Press, 1982.

M. Murata and K. Hiramatsu. Non-Gaussian filter for continuous-discrete models. *IEEE Transactions on Automatic Control*, 64(12):5260–5264, 2019.

H. Nak, M. O. Gülbahce, M. Gokasan, and A. F. Ergene. Performance investigation of extended Kalman filter based observer for PMSM using in washing machine applications. In *2015 9th International Conference on Electrical and Electronics Engineering (ELECO)*, pages 618–623, 2015.

B. Øksendal. *Stochastic Differential Equations: An Introduction with Applications*. Springer, 6 edition, 2003.

M. Raitoharju, Á. F. García-Fernández, and S. Särkkä. Partitioned update binomial Gaussian mixture filter. In *2019 22th International Conference on Information Fusion (FUSION)*, pages 1–8, 2019.

H. E. Rauch, F. Tung, and C. T. Striebel. Maximum likelihood estimates of linear dynamic systems. *AIAA Journal*, 3(8):1445–1450, 1965.

M. Reed and B. Simon. *I: Functional Analysis*. Methods of Modern Mathematical Physics. Academic Press, New York, revised and enlarged edition, 1981.

G. G. Rigatos. Particle and Kalman filtering for state estimation and control of DC motors. *ISA Transactions*, 48(1):62–72, 2009.

G. G. Rigatos. Design of chaos-based communication system with use of the derivative-free nonlinear Kalman filter. In *2013 12th International Conference on Machine Learning and Applications*, volume 2, pages 8–14, 2013.

K. Schmüdgen. *Unbounded Self-Adjoint Operators on Hilbert Space*, volume 265 of *Graduate Texts in Mathematics*. Springer, Dordrecht, 2012.

I. Segal. Algebraic integration theory. *Bull. Amer. Math. Soc.*, 71:419–489, 05 1965.

I. Segal and R. Kunze. *Integrals and Operators*. Grundlehren der mathematischen Wissenschaften. Springer-Verlag, Berlin, second edition, 1978.

B. Simon. The christoffel-darboux kernel. In D. Mitrea and M. Mitrea, editors, *Perspectives in Partial Differential Equations, Harmonic Analysis and Applications: A Volume in Honor of Vladimir G. Maz'ya's 70th Birthday*, pages 295–335. American Mathematical Society, 2008.

B. Simon. *Orthogonal Polynomials on the Unit Circle: Part 1: Classical Theory*. Colloquium Publications Series. American Math. Soc., Providence, RI, 2009.

B. Simon. *Operator Theory: A Comprehensive Course in Analysis, Part 4*. American Mathematical Society, Providence, RI, 2015.

A. Sitz, U. Schwarz, and J. Kurths. The unscented Kalman filter, a powerful tool for data analysis. *International Journal of Bifurcation and Chaos*, 14, 11 2011.

Sivagami, A. Vaishali, R. Ramakrishnan, and A. Subasini. Weather prediction model using Savitzky-Golay and Kalman filters. *Procedia Computer Science*, 165:449–455, 2019.

G. Smith, S. Schmidt, and L. McGee. *Application of Statistical Filter Theory to the Optimal Estimation of Position and Velocity on Board a Circumlunar Vehicle*. NASA TR R-135. National Aeronautics and Space Administration, 1962.

M. Stone. *Linear Transformations in Hilbert Space and Their Applications to Analysis*. American Mathematical Society, Providence, RI, 1932.

C. Suliman, C. Cruceru, and F. Moldoveanu. Kalman filter based tracking in an video surveillance system. *Advances in Electrical and Computer Engineering*, 10:30–34, 2010.

S. Särkkä and J. Sarmavuori. Gaussian filtering and smoothing for continuous-discrete dynamic systems. *Signal Processing*, 93(2):500–510, 2013.

S. Särkkä and A. Solin. *Applied Stochastic Differential Equations*. Institute of Mathematical Statistics Textbooks. Cambridge University Press, 2019.

S. Särkkä and L. Svensson. *Bayesian Filtering and Smoothing*. Institute of Mathematical Statistics Textbooks. Cambridge University Press, second edition, 2023.

S. Särkkä, A. Vehtari, and J. Lampinen. Rao-Blackwellized particle filter for multiple target tracking. *Information Fusion*, 8(1):2–15, 2007.

S. Särkkä, A. Solin, A. Nummenmaa, A. Vehtari, T. Auranen, S. Vanni, and F.-H. Lin. Dynamic retrospective filtering of physiological noise in BOLD fMRI: DRIFTER. *NeuroImage*, 60(2):1517–1527, 2012.

L. Velázquez. Spectral methods for orthogonal rational functions. *Journal of Functional Analysis*, 254(4):954–986, 2008.

Z. Wang and Y. Li. Compressed positive quadrature filter. *IEEE Transactions on Automatic Control*, pages 1–1, 2021.

J. Weidmann. *Linear Operators in Hilbert Spaces*. Graduate Texts in Mathematics. Springer-Verlag, New York, 1980.

N. Wiener. *Extrapolation, Interpolation and Smoothing of Stationary Time Series with Engineering Applications*. John Wiley & Sons, Inc., New York, 1950.

H. Wilf. *Mathematics for the Physical Sciences*. Dover Books on Mathematics Series. Wiley, 1962.

Y. Wu, D. Hu, M. Wu, and X. Hu. A numerical-integration perspective on Gaussian filters. *IEEE Transactions on Signal Processing*, 54(8):2910–2921, 2006.

J. Xie, X. M. Song, M. Sirbu, and Q. Wang. Kalman filter estimation of new product diffusion models. *Journal of Marketing Research*, 34(3):378–393, 1997.

A. Yegenoglu, K. Krajsek, S. D. Pier, and M. Herty. Ensemble Kalman filter optimizing deep neural networks: An alternative approach to non-performing gradient descent. In G. Nicosia, V. Ojha, E. La Malfa, G. Jansen, V. Sciacca, P. Pardalos, G. Giuffrida, and R. Umeton, editors, *Machine Learning, Optimization, and Data Science*, pages 78–92, Cham, 2020. Springer International Publishing.

M. Zakai. On the optimal filtering of diffusion processes. *Zeit. Wahrsch.*, 11: 230–243, 1969.

A. Zayed. *Handbook of Function and Generalized Function Transformations*. Applied mathematics. Taylor & Francis, 1996.

Z. Zhao. *State-space deep Gaussian processes with applications*. PhD thesis, Aalto University, 2021.

Z. Zhao, T. Karvonen, R. Hostettler, and S. Särkkä. Taylor moment expansion for continuous-discrete Gaussian filtering. *IEEE Transactions on Automatic Control*, 66(9):4460–4467, 2021.

Aalto-DT 174/2024

ISBN 978-952-64-1984-8 (printed)
ISBN 978-952-64-1985-5 (pdf)
ISSN 1799-4934 (printed)
ISSN 1799-4942 (pdf)

Aalto University
School of Electrical Engineering
Department of Electrical Engineering and Automation
www.aalto.fi

BUSINESS + ECONOMY

ART + DESIGN + ARCHITECTURE

SCIENCE + TECHNOLOGY

CROSSOVER

DOCTORAL THESES