

Journal of Advances in Information Fusion

A semi-annual archival publication of the International Society of Information Fusion

Regular Papers	Page
Stochastic Filtering Using Periodic Cost Functions	123
<i>Eyal Nitzan, Ben-Gurion University of the Negev, Israel</i>	
<i>Tirza Routtenberg, Ben-Gurion University of the Negev, Israel</i>	
<i>Joseph Tabrikian, Ben-Gurion University of the Negev, Israel</i>	
Methods for Deterministic Approximation of Circular Densities	138
<i>Gerhard Kurz, Karlsruhe Institute of Technology (KIT), Germany</i>	
<i>Igor Gilitschenski, Swiss Federal Institute of Technology Zurich, Switzerland</i>	
<i>Roland Y. Siegwart, Swiss Federal Institute of Technology Zurich, Switzerland</i>	
<i>Uwe D. Hanebeck, Karlsruhe Institute of Technology (KIT), Germany</i>	
Multitarget Tracking with the von Mises-Fisher Filter and Probabilistic Data Association	157
<i>Ivan Marković, University of Zagreb, Croatia</i>	
<i>Mario Bukal, University of Zagreb, Croatia</i>	
<i>Josip Česić, University of Zagreb, Croatia</i>	
<i>Ivan Petrović, University of Zagreb, Croatia</i>	
Density Estimation on the Rotation Group Using Diffusive Wavelets	173
<i>Nicolas Le Bihan, CNRS/Gipsa-Lab, France</i>	
<i>Julien Flamant, Université de Lille, France</i>	
<i>Jonathan H. Manton, The University of Melbourne, Australia</i>	
Uncertainty Propagation of Correlated Quaternion and Euclidean States Using the Gauss-Bingham Density	186
<i>Jacob E. Darling, Missouri University of Science and Technology, USA</i>	
<i>Kyle J. Demars, Missouri University of Science and Technology, USA</i>	
Multivariate Angular Filtering Using Fourier Series	206
<i>Florian Pfaff, Karlsruhe Institute of Technology (KIT), Germany</i>	
<i>Gerhard Kurz, Swiss Federal Institute of Technology Zurich, Switzerland</i>	
<i>Uwe D. Hanebeck, Swiss Federal Institute of Technology Zurich, Switzerland</i>	
Particle Filtering with Observations in a Manifold: A Proof of Convergence and Two Worked Examples	227
<i>Salem Said, Université de Bordeaux, France</i>	
<i>Jonathan H. Manton, The University of Melbourne, Australia</i>	
Direct Position Determination for TDOA-based Single Sensor Localization	250
<i>Christian Steffes, Fraunhofer Institute for Communication, Germany</i>	
<i>Marc Oispuu, Fraunhofer Institute for Communication, Germany</i>	

From the
Editor-In-Chief
Guest Editorial:
Foreword to the
Special Issue on
Estimation
Involving
Directional
Quantities



INTERNATIONAL SOCIETY OF INFORMATION FUSION

The International Society of Information Fusion (ISIF) is the premier professional society and global information resource for multidisciplinary approaches for theoretical and applied INFORMATION FUSION technologies. Technical areas of interest include target tracking, detection theory, applications for information fusion methods, image fusion, fusion systems architectures and management issues, classification, learning, data mining, Bayesian and reasoning methods.

JOURNAL OF ADVANCES IN INFORMATION FUSION: DECEMBER 2016

Editor-In-Chief	Uwe D. Hanebeck	Karlsruhe Institute of Technology (KIT), Germany; +49-721-608-43909; uwe.hanebeck@ieee.org
Associate	Stefano Coraluppi	Systems & Technology Research, USA; +1 781-305-4055; stefano.coraluppi@ieee.org
Administrative Editor	David W. Krout	University of Washington, USA; +1 206-616-2589; dkrou@apl.washington.edu
Associate	Ruixin Niu	Virginia Commonwealth University, Richmond, Virginia, USA; +1 804-828-0030; rniu@vcu.edu
Associate	Marcus Baum	Karlsruhe Institute of Technology (KIT), Germany; +49-721-608-46797; marcus.baum@kit.edu

EDITORS FOR TECHNICAL AREAS

Tracking	Stefano Coraluppi	Systems & Technology Research, USA; +1 781-305-4055; stefano.coraluppi@ieee.org
Associate	Paolo Braca	NATO Science & Technology Organization, Centre for Maritime Research and Experimentation, Italy; +39 0187 527 461; paolo.braca@cmre.nato.int
Detection	Pramod Varshney	Syracuse University, Syracuse, New York, USA; +1 315-443-1060; varshney@syr.edu
Fusion Applications	Ben Slocumb	Numerica Corporation; Loveland, Colorado, USA; +1 970-461-2000; bjslocumb@numerica.us
Associate	Ramona Georgescu	United Technologies Research Center, East Hartford, Connecticut, USA; 860-610-7890; georgera@utrc.utc.com
Image Fusion	Lex Toet	TNO, Soesterberg, 3769de, Netherlands; +31 346356237; lex.toet@tno.nl
Associate	Ting Yuan	Mercedes Benz R&D North America, USA; +1 669-224-0443; dr.ting.yuan@ieee.org
Fusion Architectures and Management Issues	Chee Chong	BAE Systems, Los Altos, California, USA; +1 650-210-8822; chee.chong@baesystems.com
Classification, Learning, Bayesian and Other Reasoning Methods	Nageswara S. V. Rao	Oak Ridge National Laboratory, USA; +1 865-574-7517;
	Claude Jauffret	Université de Toulon, La Garde, France; +33 (0) 4 94 14 24 14; jauffret@univ-tln.fr
Associate	Jean Dezert	ONERA, Chatillon, 92320, France; +33 146734990; jean.dezert@onera.fr

Manuscripts are submitted at <http://jaif.msubmit.net>. If in doubt about the proper editorial area of a contribution, submit it under the unknown area.

INTERNATIONAL SOCIETY OF INFORMATION FUSION

Jean Dezert, *President*

Lyudmila Mihaylova, *President-elect*

Stefano Coraluppi, *Secretary*

Chee Chong, *Treasurer*

Dale Blair, *Vice President Publications*

David W. Krout, *Vice President Communications*

Lance Kaplan, *Vice President Conferences*

Anne-Laure Jousselme, *Vice President Membership*

Garfield Mellema, *Vice President Working Groups*

Uwe Hanebeck, *JAIF EIC*

Roy Streit, *Perspectives EIC*

Journal of Advances in Information Fusion (ISSN 1557-6418) is published semi-annually by the International Society of Information Fusion. The responsibility for the contents rests upon the authors and not upon ISIF, the Society, or its members. ISIF is a California Nonprofit Public Benefit Corporation at P.O. Box 4631, Mountain View, California 94040. **Copyright and Reprint Permissions:** Abstracting is permitted with credit to the source. For all other copying, reprint, or republication permissions, contact the Administrative Editor. Copyright© 2016 ISIF, Inc.

From the Editor-in-Chief:

December 2016



Gerhard Kurz



Igor Gilitschenski

Guest Editorial: Foreword to the Special Issue on Estimation Involving Directional Quantities

Estimation problems that involve directional quantities naturally arise in applications that range from signal processing, robotics, and aerospace to bioinformatics and geosciences. Directional approaches differ from traditional estimation methods intended for real vector spaces because they consider the underlying manifold structure of directional problems, e.g., the unit circle, the hypertorus, or the group of rigid body motions.

The question of how to apply methods from directional statistics in the context of estimation and filtering problems has recently gained significant interest as directional problems are abundant in a plethora of applications and classical estimation methods are being pushed to their boundaries. As a result, we believe that this field of research is highly relevant for the estimation and information fusion community.

We are proud to present this special issue of the *Journal of Advances in Information Fusion*. It comprises eight papers in total, some of which are extended versions of papers originally presented in the special session on directional estimation at the International Conference on Information Fusion in 2014 and 2015.

We start the special issue with the paper “Stochastic Filtering Using Periodic Cost Functions” by Eyal Nitzan, Tirza Routtenberg, and Joseph Tabrikian. This paper proposes two novel approaches for estimation of circular states, a sample-based method and a Fourier-based method, which can be used to minimize a predefined cost function.

The second paper “Methods for Deterministic Approximation of Circular Densities” by Gerhard Kurz, Igor Gilitschenski, Roland Y. Siegwart, and Uwe D. Hanebeck deals with deterministic sample-based approximations of circular densities and presents novel algorithms based on superposition of multiple sample sets as well as an algorithm that relies on a binary tree.

Afterwards, we turn our attention to the sphere with the paper “Multitarget Tracking with the von Mises–Fisher Filter and Probabilistic Data Association” by Ivan Markovic, Mario Bukal, Josip Cestic, and Ivan Petrovic. Here, the authors apply the idea of the joint

probabilistic data association filter (JPDAF) to spherical multitarget tracking problems and combine it with von Mises-Fisher-based filtering.

On the topic of the rotation group $SO(3)$, the special issue includes the paper “Density Estimation on the Rotation Group using Diffusive Wavelets” by Nicolas Le Bihan, Julien Flamant, and Jonathan H. Manton. Their paper is devoted to density estimation using two novel methods, the first being based on characteristic functions and the second relying on wavelets using the heat kernel.

Then, the paper “Uncertainty Propagation of Correlated Quaternion and Euclidean States using the Gauss-Bingham Density” by Jacob E. Darling and Kyle J. DeMars considers the problem of uncertainty propagation on the group of rigid body motions. For this purpose, a generalization of the Bingham distribution is presented that can consider Euclidean vectors along with quaternions while taking their correlation into account.

Another relevant manifold in directional estimation is the hypertorus, which is investigated in the paper “Multivariate Angular Filtering Using Fourier Series” by Florian Pfaff, Gerhard Kurz, and Uwe D. Hanebeck. The paper presents novel algorithms for multimodal filtering on this manifold based on a multivariate Fourier series representation either of the probability density or its square root.

More generally, Salam Said and Jonathan H. Manton do not restrict themselves to a particular manifold in their paper “Particle filtering with observations in a manifold: A proof of convergence and two worked examples,” but rather present some very universal results. The paper includes a proof of convergence as well as examples involving $SO(3)$ and the unit sphere.

We conclude the special issue with the paper “Direct Position Determination for TDOA-based Single Sensor Localization” by Christian Steffes and Marc Oispuu. This paper deals with the application of localization using TDOA sensors, a practical problem that involves directional quantities. Four different localization methods are proposed and compared with the Cramér-Rao Lower Bound.

At this point, we would like to thank the Journal of Advances in Information Fusion (JAIF) Editorial Board, the editor in chief Uwe D. Hanebeck, as well as all involved authors and reviewers for making this special issue possible.

Gerhard Kurz
Karlsruhe Institute for Technology
Igor Gilitschenski
Swiss Federal Institute of Technology in Zurich
Guest Associate Editors

Stochastic Filtering Using Periodic Cost Functions

EYAL NITZAN
TIRZA ROUTTENBERG
JOSEPH TABRIKIAN

Stochastic filters attempt to estimate an unobservable state of a stochastic dynamical system from a set of noisy measurements. In this paper, we consider circular stochastic filtering and develop two dynamic methods for estimation of circular states, named sample-based stochastic filtering via root-finding (SB-SFRF) and Fourier-based stochastic filtering via root-finding (FB-SFRF). The proposed SB-SFRF and FB-SFRF methods attempt to dynamically minimize Bayes periodic risks by using Fourier series representation of their corresponding cost functions. The performance of the proposed methods is evaluated in the problem of direction-of-arrival (DOA) tracking.

Manuscript received November 7, 2015; revised March 4, 2016; released for publication May 18, 2016. Refereeing of this contribution was handled by Dr. Gerhard Kurz.

Autors' address: Department of Electrical and Computer Engineering, Ben-Gurion University of the Negev, Beer-Sheva 84105, Israel (E-mail: eyalni@ee.bgu.ac.il, tirzar@bgu.ac.il, joseph@bgu.ac.il).

This research was partially supported by THE ISRAEL SCIENCE FOUNDATION (grant No. 1160/15).

1557-6418/16/\$17.00 © 2016 JAIF

1. INTRODUCTION

In stochastic filtering problems, it is required to estimate the state of a dynamic system using a sequence of noisy measurements. Bayesian stochastic filtering is a commonly-used estimation technique that based on a state space model, recursively updates the posterior probability density function (pdf) of the current state given current and previous measurements. Under a chosen risk, optimal estimators can be obtained from the posterior pdf at each time step.

The mean-squared-error (MSE) risk is widely used for performance evaluation in stochastic filtering problems. Due to the dynamic nature of these problems, computation/approximation of the minimum MSE (MMSE) estimator is performed recursively by using stochastic filters and recursive computation of the posterior pdf. For linear dynamic systems with Gaussian noise, the well-known Kalman filter [17] provides, at each time step, a closed-form expression for the MMSE estimator. However, for the general nonlinear and/or non-Gaussian case there is no optimal filtering method that provides an analytic expression for the MMSE estimator or for the MMSE performance. Suboptimal filtering methods include the extended Kalman filter (e.g. [1], [12], [16], [56]), the unscented Kalman filter (e.g. [15], [51]), as well as discrete (approximate grid-based) filters and particle filters (e.g. [2], [10], [13]).

In many stochastic filtering problems, the unknown state has a circular nature, for example, phase, frequency, and direction-of-arrival (DOA) (see e.g. [26], [43], [55], [59]). We denote these problems as circular stochastic filtering problems. In this case, at each time step, we are interested in the modulo- 2π estimation errors and not in the plain error values. In fact, the plain error values may be absurd for estimation of circular states, especially if the unknown state is close to the edges of the circular domain [31], [45], [49]. Thus, the MSE risk and the MMSE estimator are inappropriate for circular stochastic filtering and alternative periodic risks that are based on 2π -periodic cost functions, should be used [5], [37], [38], [40], [49]. As a result, recursively-computed estimators under these periodic risks should be derived.

Several circular estimation methods, also known as directional estimation methods, have been proposed for obtaining estimators under periodic risks in static estimation problems. In [57] and [58], infinite-dimensional equations for optimal estimation under periodic risks are derived by using infinite Fourier series. However, the solution to these infinite-dimensional equations is not presented. In [47], the parameter estimation via root-finding (PERF) method is proposed. The PERF method expresses 2π -periodic cost functions via their Fourier series and derives corresponding optimal Bayes estimators for these cost functions, by using a polynomial root-finding algorithm. This method is computationally

manageable and avoids a grid search for the optimal estimator. A new approach for estimating the mean direction of a circular random variable is presented in [35], based on the minimum squared arc length criterion.

There are two main approaches in the literature for circular stochastic filtering that utilize the circular nature of the state. In [3], [20]–[23], [33], circular filters attempt to estimate the current posterior pdf under the assumption that it belongs to specific distributions on the circle. The second approach is representing the posterior pdf via a Fourier series, without assuming any specific distribution, and predicting/updating its Fourier coefficients. In [7], this approach was utilized for nonlinear filtering on linear domains under a general state and measurement model. Circular filters that utilize this approach are proposed in [57], [58] for a specific state model and in [44] for a more general state model. The circular filter in [44] maintains a valid approximation for the posterior pdf by efficiently predicting and updating the Fourier coefficients of its square root and normalizing it accordingly.

In general stochastic filtering problems and in particular in circular problems, the posterior pdf, from which the state is estimated, is usually computed or approximated. In circular stochastic filtering problems, the circular mean of the posterior pdf, which describes the pdf location on the circle, is usually used as an estimator of the circular state (see e.g. [20], [21], [44], [53]). However, the posterior circular mean is not the optimal estimator under a general periodic risk [47].

In this paper, we consider discrete-time circular stochastic filtering problems. We propose two methods for circular stochastic filtering via root-finding (SFRF): the sample-based SFRF (SB-SFRF) and the Fourier-based SFRF (FB-SFRF). The proposed methods enable the implementation of PERF method from [47], which is suitable for off-line estimation with batch data, for filtering problems, where the data is processed sequentially. The SB-SFRF and FB-SFRF methods derive estimators under a general periodic Bayes risk in circular stochastic filtering problems. The two methods are based on representation of the corresponding periodic cost function by a Fourier series and then, implementation of a root-finding algorithm. The SB-SFRF method approximates the current posterior pdf with a finite sum of weighted Dirac components while the FB-SFRF method approximates the posterior pdf with a finite Fourier series. We examine the following periodic cost functions: 1) squared-periodic-error (SPE) (see e.g. [22], [48]); 2) absolute-periodic-error (APE) (see e.g. [32, pp. 19–20], [36]); and 3) cyclic-error (CE) (see e.g. [3], [39], [50], [57]). The performance of the proposed SB-SFRF and FB-SFRF methods is demonstrated in the problem of DOA tracking.

The remainder of the paper is organized as follows. In Section 2, we formulate the circular stochastic filtering model and review the properties of common periodic risks. In Section 3, the SB-SFRF and FB-SFRF

methods are derived. The proposed methods are evaluated via simulations for DOA tracking problem in Section 4. Finally, our conclusions appear in Section 5.

In the sequel, we denote vectors and matrices by boldface lowercase and uppercase letters, respectively. The m th element of the vector \mathbf{b} is denoted by b_m and $j \triangleq \sqrt{-1}$. The notations $(\cdot)^T$ and $(\cdot)^*$ denote the transpose and complex conjugate operators, respectively. The notation $\angle \cdot$ stands for the phase of a complex scalar, which is assumed to be restricted to the interval $[-\pi, \pi)$. The modulo- 2π operator, which maps $\rho \in \mathbb{R}$ to $[-\pi, \pi)$, is denoted as $[\rho]_{2\pi} \triangleq \rho - 2\pi \lfloor \frac{\rho}{2\pi} + \frac{1}{2} \rfloor$, where $\lfloor \cdot \rfloor$ is the floor operator. The operators of expectation and conditional expectation given an event Z , are denoted as $E[\cdot]$ and $E[\cdot | Z]$, respectively.

2. CIRCULAR STOCHASTIC FILTERING

Consider the following nonlinear discrete-time state space model

$$\begin{cases} \theta_n = a_n(\theta_{n-1}, \mathbf{w}_n) \\ \mathbf{x}_n = \mathbf{h}_n(\theta_n, \boldsymbol{\nu}_n) \end{cases}, \quad n = 1, 2, \dots, \quad (1)$$

where for any $n = 1, 2, \dots$

- $\theta_n \in \Omega_\theta \triangleq [-\pi, \pi)$ —circular state for which we are interested in the modulo- 2π estimation error.
- $\theta_0 \in \Omega_\theta$ —initial state with known a priori pdf f_{θ_0} .
- $\mathbf{x}_n \in \mathbb{C}^L$ —measurement vector.
- $\{\mathbf{w}_n\}$ —sequence of mutually independent $P \times 1$ noise vectors with known pdfs, $\{f_{\mathbf{w}_n}\}$, that are independent of past and present states.
- $\{\boldsymbol{\nu}_n\}$ —sequence of mutually independent $Q \times 1$ complex noise vectors with known pdfs, $\{f_{\boldsymbol{\nu}_n}\}$, that are independent of past and present states and the state noise.
- $a_n : \Omega_\theta \times \mathbb{R}^P \rightarrow \Omega_\theta$ —state transition function.
- $\mathbf{h}_n : \Omega_\theta \times \mathbb{C}^Q \rightarrow \mathbb{C}^L$ —measurement function.

The conditional pdfs $f_{\theta_n|\theta_{n-1}}$ and $f_{\mathbf{x}_n|\theta_n}$ can be obtained from (1) and the pdfs of \mathbf{w}_n and $\boldsymbol{\nu}_n$, respectively. The filtering goal is to estimate the circular state θ_n at each time step $n = 1, 2, \dots$, based on $\mathbf{x}^{(n)} \triangleq [\mathbf{x}_1^T, \dots, \mathbf{x}_n^T]^T \in \Omega_{\mathbf{x}}^{(n)}$, which is the augmented measurement vector containing all the measurements up to time step n , where $\Omega_{\mathbf{x}}^{(n)}$ is the n th step measurement space. An estimator of θ_n , based on $\mathbf{x}^{(n)}$, is denoted by $\hat{\theta}_n : \Omega_{\mathbf{x}}^{(n)} \rightarrow \Omega_\theta$. The posterior pdf of θ_n given $\mathbf{x}^{(n)}$ and the predicted pdf of θ_n given $\mathbf{x}^{(n-1)}$ are denoted by $f_{\theta_n|\mathbf{x}^{(n)}}$ and $f_{\theta_n|\mathbf{x}^{(n-1)}}$, respectively, $\forall n = 1, 2, \dots$, where $f_{\theta_1|\mathbf{x}^{(0)}} \triangleq f_{\theta_1}$ denotes the a priori pdf of $\theta_1 \in \Omega_\theta$. In addition, we define $f_{\theta_0|\mathbf{x}^{(0)}} \triangleq f_{\theta_0}$.

In the Bayesian framework, optimal estimators are obtained via minimization of Bayes risks. In circular stochastic filtering problems, the appropriate Bayes risk at time step n is based on 2π -periodic cost function,

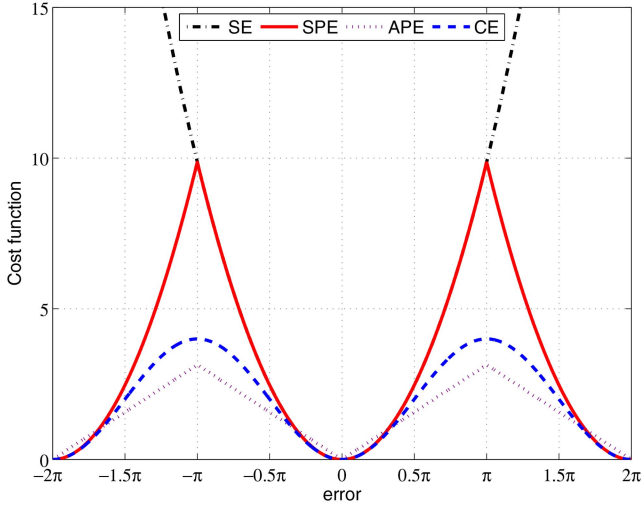


Fig. 1. SE and periodic cost functions: SPE, APE, and CE.

$C(\hat{\theta}_n - \theta_n)$, with respect to (w.r.t.) the estimation error, $\hat{\theta}_n - \theta_n$. The corresponding risk is the mean of the cost function, $E[C(\hat{\theta}_n - \theta_n)]$. It should be noted that restriction of the state estimator to the region $[-\pi, \pi]$ does not prevent the resulting estimation error from taking values in the region $(-2\pi, 2\pi)$. Thus, the inappropriate nature of the conventional risks, such as the MSE, cannot be resolved. For example, consider the case of close parameter and its estimate on the circle, in which $\theta_n = -\pi + \delta_1$ and $\hat{\theta}_n = \pi - \delta_2$, $0 < \delta_i \ll 1$, $i = 1, 2$. In this case, direct computation of the error results in $\hat{\theta}_n - \theta_n = 2\pi - (\delta_1 + \delta_2)$, which is a large error. However, computation of the periodic error results in $[\hat{\theta}_n - \theta_n]_{2\pi} = -(\delta_1 + \delta_2)$, which is a small error.

In the following, we describe three examples for periodic cost functions: SPE, APE, and CE. These periodic cost functions and the nonperiodic squared-error (SE) cost function are presented in Fig. 1 versus the estimation error. It can be seen that as the absolute value of the estimation error grows from π to 2π , the SE cost function increases, since it does not take the circular nature of the error into account, while the periodic cost functions decrease.

- SPE: The SPE cost function at the n th time step is defined as

$$C(\hat{\theta}_n - \theta_n) = \text{SPE}(\hat{\theta}_n - \theta_n) \triangleq ([\hat{\theta}_n - \theta_n]_{2\pi})^2. \quad (2)$$

Given a measurement vector, $\mathbf{x}^{(n)}$, the minimum mean SPE (MMSPE) estimator is given by

$$\hat{\theta}_{n,\text{MMSPE}} = \arg \min_{\hat{\theta}_n \in \Omega_\theta} E[(\hat{\theta}_n - \theta_n)_{2\pi}^2 | \mathbf{x}^{(n)}]. \quad (3)$$

- APE: The APE cost function at the n th time step is defined as

$$C(\hat{\theta}_n - \theta_n) = \text{APE}(\hat{\theta}_n - \theta_n) \triangleq |[\hat{\theta}_n - \theta_n]_{2\pi}|. \quad (4)$$

Given a measurement vector, $\mathbf{x}^{(n)}$, the minimum mean APE (MMAPE) estimator is given by

$$\hat{\theta}_{n,\text{MMAPE}} = \arg \min_{\hat{\theta}_n \in \Omega_\theta} E[|[\hat{\theta}_n - \theta_n]_{2\pi}| | \mathbf{x}^{(n)}]. \quad (5)$$

- CE: The CE cost function at the n th time step is defined as

$$C(\hat{\theta}_n - \theta_n) = \text{CE}(\hat{\theta}_n - \theta_n) \triangleq 2 - 2 \cos(\hat{\theta}_n - \theta_n). \quad (6)$$

In [57], [58], it is shown that given a measurement vector, $\mathbf{x}^{(n)}$, the minimum mean CE (MMCE) estimator is given by

$$\hat{\theta}_{n,\text{MMCE}}(\mathbf{x}^{(n)}) = \begin{cases} \angle E[e^{j\theta_n} | \mathbf{x}^{(n)}], & E[e^{j\theta_n} | \mathbf{x}^{(n)}] \neq 0 \\ 0, & \text{otherwise} \end{cases}, \quad (7)$$

which is the posterior circular mean at time step n .

In general, optimal estimators under periodic risks cannot be analytically derived except for a few special cases, such as the estimator from (7), which is optimal under the mean CE (MCE) risk. Therefore, a grid-search method is used for their derivation. A grid-search method involves the computation of conditional expectation, as in (3) and (5), for any point on the grid. The disadvantages of a grid-search method are: (a) its accuracy depends on the chosen grid; and (b) using a dense grid can be computationally prohibitive. In [47], a method called PERF was proposed for obtaining optimal estimators under arbitrary periodic risks. This method was derived for batch data. In this paper, we extend this method to circular stochastic filtering in which the estimators can be computed recursively at each time step, based on results from the previous steps and the new measurement. In the following section, we first review the PERF method. Then, we describe the SB-SFRF and FB-SFRF methods that utilize the PERF approach and obtain estimators under a general periodic risk in a dynamic setting.

3. SB-SFRF AND FB-SFRF METHODS

3.1. Review—PERF method

In circular stochastic filtering, our goal is minimization of an arbitrary periodic risk at each time step. First, we discuss periodic risks, whose corresponding cost function can be represented by a finite Fourier series. We refer to the case of infinite Fourier series in Subsection 3.4.

Let $C : \Omega_\varepsilon \rightarrow \mathbb{R}$ be a real, 2π -periodic, and even cost function, where $\Omega_\varepsilon \triangleq (-2\pi, 2\pi)$, that can be expressed as the following finite Fourier series [61]:

$$C(\hat{\theta}_n - \theta_n) = \sum_{k=-K}^K c_k e^{jk(\hat{\theta}_n - \theta_n)}, \quad (8)$$

for any $n = 1, 2, \dots$, where c_k , $k = -K, \dots, K$, are the corresponding Fourier coefficients of C . Since C is even, its Fourier coefficients satisfy $c_k = c_{-k}$. Given $\mathbf{x}^{(n)}$, the estimator that minimizes the Bayes risk, $E[C(\hat{\theta}_n - \theta_n)]$, is obtained by minimization of the conditional expectation

$$Q(e^{j\hat{\theta}_n}) \triangleq E[C(\hat{\theta}_n - \theta_n) | \mathbf{x}^{(n)}]. \quad (9)$$

By substituting (8) in (9) and using the linearity of the expectation operator, one obtains

$$\begin{aligned} Q(e^{j\hat{\theta}_n}) &= E \left[\sum_{k=-K}^K c_k e^{jk(\hat{\theta}_n - \theta_n)} | \mathbf{x}^{(n)} \right] \\ &= \sum_{k=-K}^K c_k m_{-k}^{(n)}(\mathbf{x}^{(n)}) e^{jk\hat{\theta}_n}, \end{aligned} \quad (10)$$

where

$$m_k^{(n)}(\mathbf{x}) \triangleq E[e^{jk\theta_n} | \mathbf{x}] \quad (11)$$

is the k th trigonometric moment of $f_{\theta_n|\mathbf{x}}$ (see e.g. [14, p. 26], [32, pp. 28–29]), where \mathbf{x} is a random vector. The term in (11) can be interpreted as a sample of the conditional characteristic function of θ_n given \mathbf{x} at an integer k . Since $f_{\theta_n|\mathbf{x}^{(n)}}$ is a real pdf, it can be verified that $m_0^{(n)}(\mathbf{x}^{(n)}) = 1$ and $m_k^{(n)}(\mathbf{x}^{(n)}) = (m_{-k}^{(n)}(\mathbf{x}^{(n)}))^*$, $k = 1, \dots, K$.

In order to obtain the estimator that minimizes $Q(e^{j\hat{\theta}_n})$ from (10), we first find stationary points of $Q(e^{j\hat{\theta}_n})$ and then find the minimum point from the set of stationary points. Since $Q(e^{j\hat{\theta}_n})$ is a real, smooth, and 2π -periodic function w.r.t. $\hat{\theta}_n$, it can be verified by using Rolle's theorem (see e.g. [54, p. 132]) that it has at least two stationary points in Ω_θ . The stationary points are obtained by equating the derivative of $Q(e^{j\hat{\theta}_n})$ w.r.t. $\hat{\theta}_n$ to zero, which yields

$$\sum_{k=-K}^K kc_k m_{-k}^{(n)}(\mathbf{x}^{(n)}) e^{jk\bar{\theta}_n} = 0, \quad (12)$$

where $\bar{\theta}_n$ is a stationary point of $Q(e^{j\hat{\theta}_n})$. In general, (12) cannot be analytically solved. Finding the stationary points and consequently the minimum point of $Q(e^{j\hat{\theta}_n})$ can be obtained by a grid-search method, whose drawbacks are discussed in Section 2. An alternative approach is utilizing the PERF method [47].

In the PERF method, the term $e^{j\hat{\theta}_n}$ in (12) is replaced by a general complex scalar $z \in \mathbb{C}$, resulting in

$$\sum_{k=-K}^K kc_k m_{-k}^{(n)}(\mathbf{x}^{(n)}) z^k = 0. \quad (13)$$

Then, a polynomial root-finding algorithm is applied on (13) and the $2K$ roots of (13), $\bar{z}_1, \dots, \bar{z}_{2K}$, are obtained. Finally, the optimal estimator is obtained by computing

$$\hat{\theta}_{n,\text{opt}} = \angle \bar{z}_{\text{opt}} = \arg \min_{\hat{\theta}_n \in \{\bar{\theta}_{n,1}, \dots, \bar{\theta}_{n,2K}\}} Q(e^{j\hat{\theta}_n}), \quad (14)$$

TABLE I
PERF method

Initialization:

- Choose a real, 2π -periodic, and even cost function, C , with Fourier series order K .
- Compute $\{c_k\}_{k=0}^K$, the Fourier coefficients of the periodic cost function C .

Algorithm stages:

- Compute $\{m_k^{(n)}(\mathbf{x}^{(n)})\}_{k=1}^K$, the trigonometric moments of $f_{\theta_n|\mathbf{x}^{(n)}}$, as defined in (11).
- Find the roots of (13), $\bar{z}_1, \dots, \bar{z}_{2K}$, and compute their corresponding phases, $\bar{\theta}_{n,1}, \dots, \bar{\theta}_{n,2K}$, respectively.
- Find $\hat{\theta}_{n,\text{opt}}$ using (10) and (14).

where $\bar{\theta}_{n,1} = \angle \bar{z}_1, \dots, \bar{\theta}_{n,2K} = \angle \bar{z}_{2K}$. It should be noted that as opposed to a grid-search method, which involves the computation of conditional expectation for any point on the grid (as mentioned in Section 2), the PERF method involves the computation of only K conditional expectations for derivation of the trigonometric moments in (11). The PERF method is summarized in Table I.

It can be seen that in order to use the PERF method in a circular stochastic filtering problem, the trigonometric moments, $\{m_k^{(n)}(\mathbf{x}^{(n)})\}_{k=1}^K$, from (11) should be computed at each time step. The SB-SFRF and FB-SFRF methods, derived in the following, enable the implementation of PERF method in a dynamic setting.

3.2. SB-SFRF method

In this subsection, we describe SB-SFRF method that can be implemented by using any sample-based filter. A sample-based filter, such as discrete filter and particle filter [2], approximates the n th step posterior pdf, $f_{\theta_n|\mathbf{x}^{(n)}}$, with a finite sum of weighted Dirac components, $f_{n,\text{SB}}^{(S)}: \Omega_\theta \rightarrow \mathbb{R}$, given by

$$f_{n,\text{SB}}^{(S)}(y) \triangleq \sum_{s=1}^S \omega_{n,s} \delta(y - \beta_{n,s}), \quad (15)$$

where y is the argument of $f_{n,\text{SB}}^{(S)}$, S is the number of samples, δ is the Dirac delta function, $\beta_{n,1}, \dots, \beta_{n,S} \in \Omega_\theta$ are the Dirac positions, and $\omega_{n,1}, \dots, \omega_{n,S}$ are nonnegative weights. The values of S , $\{\beta_{n,s}\}_{s=1}^S$, and $\{\omega_{n,s}\}_{s=1}^S$ depend on the chosen sample-based filter. In order to apply the PERF method, the first K trigonometric moments $m_1^{(n)}(\mathbf{x}^{(n)}), \dots, m_K^{(n)}(\mathbf{x}^{(n)})$ should be computed at each time step. By substituting the pdf approximation from (15) in (11), the approximations of the trigonometric moments, $\{m_k^{(n)}(\mathbf{x}^{(n)})\}_{k=1}^K$, are

$$m_{k,\text{SB}}^{(n,S)} = \sum_{s=1}^S \omega_{n,s} e^{jk\beta_{n,s}}, \quad \forall k = 1, \dots, K. \quad (16)$$

Thus, by substituting (16) in (10) and (13), the PERF method can be applied to obtain an approximation for the optimal estimator. We denote the resulting estimator

TABLE II
SB-SFRF method

Initialization:

- Choose a real, 2π -periodic, and even cost function, C , with Fourier series order K .
- Compute $\{c_k\}_{k=0}^K$, the Fourier coefficients of the periodic cost function C .

Algorithm stages for the n th time step:

- Compute the Dirac positions, $\beta_{n,1}, \dots, \beta_{n,S} \in \Omega_\theta$, and the weights, $\omega_{n,1}, \dots, \omega_{n,S}$, of the approximated posterior pdf, $f_{n,SB}^{(S)}$, from (15) by using a sample-based filter.
- Compute the approximated trigonometric moments, $m_{k,SB}^{(n,S)}$, $k = 1, \dots, K$, according to (16).
- Find the roots of (13), $\bar{z}_1, \dots, \bar{z}_{2K}$, and compute their corresponding phases, $\bar{\theta}_{n,1}, \dots, \bar{\theta}_{n,2K}$, respectively.
- Find $\hat{\theta}_{n,SB}^{(S)}$ using (10) and (14).

as $\hat{\theta}_{n,SB}^{(S)}$. The complete SB-SFRF method at time step n is summarized in Table II.

3.3. FB-SFRF method

In this subsection, we describe the FB-SFRF method that is implemented by using a Fourier-based circular filter, denoted as Fourier filter, which is proposed in [44]. First, we define

$$f_{\theta_0}^{(p)}(\alpha_0) \triangleq \sum_{l=-\infty}^{\infty} f_{\theta_0}(\alpha_0 + 2\pi l) \quad (17)$$

and for any $n = 1, 2, \dots$,

$$f_{\theta_n|\theta_{n-1}}^{(p)}(\alpha_n | \alpha_{n-1}) \triangleq \sum_{l=-\infty}^{\infty} \sum_{m=-\infty}^{\infty} f_{\theta_n|\theta_{n-1}}(\alpha_n + 2\pi l | \alpha_{n-1} + 2\pi m), \quad (18)$$

$$f_{\mathbf{x}_n|\theta_n}^{(p)}(\beta_n | \alpha_n) \triangleq \sum_{l=-\infty}^{\infty} f_{\mathbf{x}_n|\theta_n}(\beta_n | \alpha_n + 2\pi l), \quad (19)$$

$$f_{\theta_n|\mathbf{x}^{(n)}}^{(p)}(\alpha_n | \beta^{(n)}) \triangleq \sum_{l=-\infty}^{\infty} f_{\theta_n|\mathbf{x}^{(n)}}(\alpha_n + 2\pi l | \beta^{(n)}), \quad (20)$$

$$f_{\theta_n|\mathbf{x}^{(n-1)}}^{(p)}(\alpha_n | \beta^{(n-1)}) \triangleq \sum_{l=-\infty}^{\infty} f_{\theta_n|\mathbf{x}^{(n-1)}}(\alpha_n + 2\pi l | \beta^{(n-1)}), \quad (21)$$

which are the 2π -periodic extensions of the pdfs f_{θ_0} , $f_{\theta_n|\theta_{n-1}}$, $f_{\mathbf{x}_n|\theta_n}$, $f_{\theta_n|\mathbf{x}^{(n)}}$, and $f_{\theta_n|\mathbf{x}^{(n-1)}}$, respectively, w.r.t. $\{\theta_n\}_{n \geq 0}$.

Since $f_{\theta_0}^{(p)}$ is a 2π -periodic function w.r.t. θ_0 and since $f_{\mathbf{x}_n|\theta_n}^{(p)}$, $f_{\theta_n|\mathbf{x}^{(n)}}$, and $f_{\theta_n|\mathbf{x}^{(n-1)}}$ are 2π -periodic functions w.r.t. θ_n , it is assumed that they can be represented via Fourier series with Fourier coefficients $\{\eta_l^{(0)}\}_{l \in \mathbb{Z}}$, $\{d_l^{(n)}(\mathbf{x}_n)\}_{l \in \mathbb{Z}}$, $\{\eta_l^{(n|n)}(\mathbf{x}^{(n)})\}_{l \in \mathbb{Z}}$, and $\{\eta_l^{(n|n-1)}(\mathbf{x}^{(n-1)})\}_{l \in \mathbb{Z}}$, respectively. That is,

$$f_{\theta_0}^{(p)}(\alpha_0) = \sum_{l=-\infty}^{\infty} \eta_l^{(0)} e^{jl\alpha_0}, \quad (22)$$

$$f_{\mathbf{x}_n|\theta_n}^{(p)}(\beta_n | \alpha_n) = \sum_{l=-\infty}^{\infty} d_l^{(n)}(\beta_n) e^{jl\alpha_n}, \quad (23)$$

$$f_{\theta_n|\mathbf{x}^{(n)}}^{(p)}(\alpha_n | \beta^{(n)}) = \sum_{l=-\infty}^{\infty} \eta_l^{(n|n)}(\beta^{(n)}) e^{jl\alpha_n}, \quad (24)$$

and

$$f_{\theta_n|\mathbf{x}^{(n-1)}}^{(p)}(\alpha_n | \beta^{(n-1)}) = \sum_{l=-\infty}^{\infty} \eta_l^{(n|n-1)}(\beta^{(n-1)}) e^{jl\alpha_n}. \quad (25)$$

Similarly, the function $f_{\theta_n|\theta_{n-1}}^{(p)}$ is a 2π -periodic function w.r.t. both θ_n and θ_{n-1} , and therefore, it is assumed that it can be represented via a two-dimensional Fourier series with Fourier coefficients $\{\phi_{l,m}^{(n)}\}_{l,m \in \mathbb{Z}}$, i.e.

$$f_{\theta_n|\theta_{n-1}}^{(p)}(\alpha_n | \alpha_{n-1}) = \sum_{l=-\infty}^{\infty} \sum_{m=-\infty}^{\infty} \phi_{l,m}^{(n)} e^{jl\alpha_n} e^{jm\alpha_{n-1}}. \quad (26)$$

Based on the Fourier-based filters proposed in [7], [44], [57], [58], we derive prediction and update stages, which are applied on the Fourier coefficients of $f_{\theta_n|\mathbf{x}^{(n)}}$ and $f_{\theta_n|\mathbf{x}^{(n-1)}}$.

• Prediction:

According to Chapman-Kolmogorov equation and by using the Markovian nature of the state model in (1), the predicted pdf $f_{\theta_n|\mathbf{x}^{(n-1)}}$ is given by (see e.g. [2, Eq. (3)])

$$f_{\theta_n|\mathbf{x}^{(n-1)}}(\alpha_n | \beta^{(n-1)}) = \int_{\Omega_\theta} f_{\theta_n|\theta_{n-1}}(\alpha_n | \alpha_{n-1}) \times f_{\theta_{n-1}|\mathbf{x}^{(n-1)}}(\alpha_{n-1} | \beta^{(n-1)}) d\alpha_{n-1}, \quad \forall \alpha_n \in \Omega_\theta. \quad (27)$$

For $\theta_n, \theta_{n-1} \in \Omega_\theta$, the pdfs in the left hand side (l.h.s.) and right hand side (r.h.s.) of (27) are equal to their periodic extensions. Thus, (27) can be rewritten as

$$f_{\theta_n|\mathbf{x}^{(n-1)}}^{(p)}(\alpha_n | \beta^{(n-1)}) = \int_{\Omega_\theta} f_{\theta_n|\theta_{n-1}}^{(p)}(\alpha_n | \alpha_{n-1}) \times f_{\theta_{n-1}|\mathbf{x}^{(n-1)}}^{(p)}(\alpha_{n-1} | \beta^{(n-1)}) d\alpha_{n-1}, \quad \forall \alpha_n \in \Omega_\theta. \quad (28)$$

By substituting the corresponding Fourier series from (22), (24), and (26) in the r.h.s. of (28), we obtain the Fourier series representation of $f_{\theta_n|\mathbf{x}^{(n-1)}}$ from (25), whose l th Fourier coefficient is given by [7]

$$\eta_l^{(n|n-1)}(\mathbf{x}^{(n-1)}) = 2\pi \sum_{m=-\infty}^{\infty} \eta_{-m}^{(n-1|n-1)}(\mathbf{x}^{(n-1)}) \phi_{l,m}^{(n)}, \quad (29)$$

$\forall l \in \mathbb{Z}$, $n = 1, 2, \dots$, where $\eta_l^{(0|0)}(\mathbf{x}^{(0)}) \triangleq \eta_l^{(0)}$.

• Update:

According to Bayes' rule, the posterior pdf $f_{\theta_n|\mathbf{x}^{(n)}}$ is given by (see e.g. [2, Eq. (4)])

$$f_{\theta_n|\mathbf{x}^{(n)}}(\alpha_n | \beta^{(n)}) = \frac{f_{\mathbf{x}_n|\theta_n}(\beta_n | \alpha_n) f_{\theta_n|\mathbf{x}^{(n-1)}}(\alpha_n | \beta^{(n-1)})}{\int_{\Omega_\theta} f_{\mathbf{x}_n|\theta_n}(\beta_n | \alpha_n) f_{\theta_n|\mathbf{x}^{(n-1)}}(\alpha_n | \beta^{(n-1)}) d\alpha_n}, \quad (30)$$

$\forall \alpha_n \in \Omega_\theta$. For $\theta_n \in \Omega_\theta$, the pdfs in the l.h.s. and r.h.s. of (30) are equal to their periodic extensions. Thus, (30) can be rewritten as

$$f_{\theta_n|\mathbf{x}^{(n)}}^{(p)}(\alpha_n | \beta^{(n)}) = \frac{f_{\mathbf{x}_n|\theta_n}^{(p)}(\beta_n | \alpha_n) f_{\theta_n|\mathbf{x}^{(n-1)}}^{(p)}(\alpha_n | \beta^{(n-1)})}{\int_{\Omega_\theta} f_{\mathbf{x}_n|\theta_n}^{(p)}(\beta_n | \alpha_n) f_{\theta_n|\mathbf{x}^{(n-1)}}^{(p)}(\alpha_n | \beta^{(n-1)}) d\alpha_n}, \quad (31)$$

$\forall \alpha_n \in \Omega_\theta$. By substituting the corresponding Fourier series from (23) and (25) in the r.h.s. of (31), we obtain the Fourier series representation of $f_{\theta_n|\mathbf{x}^{(n)}}^{(p)}$ from (24), whose l th Fourier coefficient is given by [7]

$$\eta_l^{(n|n)}(\mathbf{x}^{(n)}) = \frac{\gamma_l^{(n)}(\mathbf{x}^{(n)})}{2\pi\gamma_0^{(n)}(\mathbf{x}^{(n)})}, \quad (32)$$

$\forall l \in \mathbb{Z}$, $n = 1, 2, \dots$, where

$$\gamma_l^{(n)}(\mathbf{x}^{(n)}) \triangleq \sum_{m=-\infty}^{\infty} \eta_m^{(n|n-1)}(\mathbf{x}^{(n-1)}) d_{l-m}^{(n)}(\mathbf{x}_n).$$

Since $f_{\theta_n|\mathbf{x}^{(n)}}$ from (30) is equal to $f_{\theta_n|\mathbf{x}^{(n)}}^{(p)}$ from (31) for $\theta_n \in \Omega_\theta$, the Fourier series of $f_{\theta_n|\mathbf{x}^{(n)}}^{(p)}$ can be used to represent $f_{\theta_n|\mathbf{x}^{(n)}}$. In practice, $f_{\theta_n|\mathbf{x}^{(n)}}$ is approximated by a finite Fourier series, $f_{n,\text{FB}}^{(D)} : \Omega_\theta \rightarrow \mathbb{R}$, given by

$$f_{n,\text{FB}}^{(D)}(y) \triangleq \sum_{l=-D}^D \eta_{l,\text{FB}}^{(n,D)} e^{jly}, \quad (33)$$

where y is the argument of $f_{n,\text{FB}}^{(D)}$ and D is the chosen Fourier series order. The series order, D , is determined by taking into account the trade-off between estimation quality and required rate of convergence. The approximation accuracy, i.e. the distance between the actual pdf and its Fourier series approximation, can be measured for example by using Hellinger metric [7] or Kullback-Leibler divergence [44]. At time step n , $\eta_{l,\text{FB}}^{(n,D)}$ approximates $\eta_l^{(n|n)}(\mathbf{x}^{(n)})$, $\forall l = -D, \dots, D$, and for $|l| > D$, $\eta_l^{(n|n)}(\mathbf{x}^{(n)})$ is approximated by zero. The Fourier series approximation of $f_{\theta_n|\mathbf{x}^{(n)}}$ in (33) represents a pdf and thus, should be nonnegative and integrate to 1. In order to maintain the approximation in (33) as a valid pdf, the approach in [44] that approximates $\sqrt{f_{\theta_n|\mathbf{x}^{(n)}}$ by a finite Fourier series, is applied. It should be noted that it is assumed in [44] that $f_{\theta_n|\theta_{n-1}}^{(p)}(\alpha_n | \alpha_{n-1}) = g(\alpha_n - \alpha_{n-1})$, $\forall \alpha_n, \alpha_{n-1} \in \Omega_\theta$, for some function g , which simplifies the prediction stage. An explanation is added in the appendix for implementing the prediction stage, under the approach of [44], with a general $f_{\theta_n|\theta_{n-1}}^{(p)}$.

By substituting the pdf approximation from (33) in (11) and since $\int_{\Omega_\theta} e^{jl\alpha} d\alpha = 2\pi\delta_{\text{kr}}(l)$, where δ_{kr} is the Kronecker delta function, the approximations of the trigonometric moments, $\{m_k^{(n)}(\mathbf{x}^{(n)})\}_{k=1}^K$, are

$$m_{k,\text{FB}}^{(n,D)} = 2\pi\eta_{-k,\text{FB}}^{(n,D)}, \quad \forall k = 1, \dots, K. \quad (34)$$

TABLE III
FB-SFRF method

Initialization:

- Choose a real, 2π -periodic, and even cost function, C , with Fourier series order K .
- Compute $\{c_k\}_{k=0}^K$, the Fourier coefficients of the periodic cost function C .

Algorithm stages for the n th time step:

- Compute the Fourier coefficients, $\eta_{l,\text{FB}}^{(n,D)}$, $l = 1, \dots, D$, of the approximated posterior pdf, $f_{n,\text{FB}}^{(D)}$, from (33) by using the Fourier-based circular filter from [44].
- Compute the approximated trigonometric moments, $m_{k,\text{FB}}^{(n,D)}$, $k = 1, \dots, K$, according to (34).
- Find the roots of (13), $\tilde{z}_1, \dots, \tilde{z}_{2K}$, and compute their corresponding phases, $\tilde{\theta}_{n,1}, \dots, \tilde{\theta}_{n,2K}$, respectively.
- Find $\hat{\theta}_{n,\text{FB}}^{(D)}$ using (10) and (14).

In case $K > D$, then $m_{k,\text{FB}}^{(n,D)} = 0$, $\forall k = D + 1, \dots, K$. Thus, by substituting (34) in (10) and (13), the PERF method can be applied to obtain an approximation for the optimal estimator. We denote the resulting estimator as $\hat{\theta}_{n,\text{FB}}^{(D)}$. The complete FB-SFRF method at time step n is summarized in Table III.

3.4. SB-SFRF and FB-SFRF methods with a general periodic cost function

Consider a 2π -periodic cost function, C , with a convergent infinite Fourier series representation. It is shown in [47] that under some regularity conditions, PERF method, applied on a truncated Fourier series representation of C with $K < \infty$, converges to the corresponding optimal estimator in the limit $K \rightarrow \infty$. For such cost functions the SB-SFRF and FB-SFRF methods are applied on a truncated Fourier series.

In the following, we describe three examples for periodic risks, the mean SPE (MSPE), mean APE (MAPE), and MCE risks and discuss the implementation of SB-SFRF and FB-SFRF methods under these risks. In order to apply SB-SFRF and FB-SFRF methods with the MSPE, MAPE, and MCE risks, their corresponding Fourier coefficients should be computed. The Fourier series order of the CE cost function is $K = 1$ and SB-SFRF and FB-SFRF methods can be directly applied. As opposed to the CE cost function, the Fourier series of the SPE and APE cost functions are infinite and therefore, SB-SFRF and FB-SFRF methods are applied on their truncated Fourier series representations. It is shown in [47] that under mild conditions PERF method converges to the optimal estimators under the MSPE and MAPE risks in the limit $K \rightarrow \infty$, which justifies applying SB-SFRF and FB-SFRF methods under these risks. The Fourier series, Fourier coefficients, and explicit form of (13) for the SPE, APE, and CE cost functions, are presented in Table IV.

It is shown in [47] that the optimal estimator under first-order Fourier series approximation of the SPE and the APE cost functions, i.e. for choosing $K = 1$, is given

TABLE IV
The Fourier series of the cost functions

Cost function	Squared periodic error (SPE)	Absolute periodic error (APE)	Cyclic error (CE)
Fourier series of $C(\varepsilon)$	$\frac{\pi^2}{3} + 2 \sum_{k \in \mathbb{Z}, k \neq 0} \frac{(-1)^k}{k^2} e^{jk\varepsilon}$	$\frac{\pi}{2} - 2 \sum_{k \in \mathbb{Z}} \frac{1}{\pi(2k+1)^2} e^{i(2k+1)\varepsilon}$	$2 - e^{j\varepsilon} - e^{-j\varepsilon}$
Fourier coefficients c_k	$\begin{cases} \frac{\pi^2}{3}, & k = 0 \\ \frac{2(-1)^k}{k^2}, & k \neq 0 \end{cases}$	$\begin{cases} \frac{\pi}{2}, & k = 0 \\ 0, & k \neq 0 \text{ and even} \\ -\frac{2}{\pi k^2}, & k \text{ odd} \end{cases}$	$\begin{cases} 2, & k = 0 \\ -1, & k = \pm 1 \\ 0, & \text{otherwise} \end{cases}$
Explicit form of (13) at n th time step	$\sum_{k=-K, k \neq 0}^K \frac{(-1)^k m_{-k}^{(n)}(\mathbf{x}^{(n)})}{k} z^k = 0$	$\sum_{k=-\lfloor (K+1)/2 \rfloor}^{\lfloor (K-1)/2 \rfloor} \frac{m_{-2k-1}^{(n)}(\mathbf{x}^{(n)})}{2k+1} z^{2k} = 0$	$m_{-1}^{(n)}(\mathbf{x}^{(n)})z + m_1^{(n)}(\mathbf{x}^{(n)})z^{-1} = 0$

by the posterior circular mean estimator, which is the MMCE estimator from (7). In addition, it is shown in [47] and [30] that the posterior circular mean estimator is optimal under the MSPE and MAPE risks, in case the posterior pdf is unimodal and even (as a function supported on the circle). In the general case, where the posterior pdf is not necessarily unimodal and even, choosing $K > 1$ for the SPE and APE approximations and applying SB-SFRF or FB-SFRF methods can improve the performance, comparing to the posterior circular mean estimator, under the MSPE and MAPE risks. We denote SB-SFRF and FB-SFRF methods under the MSPE risk as SB-SFRF-MSPE and FB-SFRF-MSPE, respectively. Similarly, SB-SFRF and FB-SFRF methods under the MAPE risk are denoted by SB-SFRF-MAPE and FB-SFRF-MAPE, respectively.

Remarks:

- 1) *Periodic cost functions:* In the conventional Bayesian framework, the SE and absolute-error (AE) cost functions are commonly used for performance evaluation. One of the differences between these cost functions is that the SE increases faster than the AE and thus, the SE is more sensitive to outliers [25, p. 51]. The SPE and APE are the natural periodic equivalents of the SE and AE, respectively. They are obtained by periodically extending their conventional counterparts. The CE can be viewed as a smooth first order approximation, in terms of Fourier series, of the SPE and APE [47]. In the small error region, the APE and SPE coincide with the AE and SE, respectively. Similar to the SPE, in the small error region, the CE coincides with the SE.
- 2) *Computational complexity:* The additional run-time complexity induced from using SB-SFRF and FB-SFRF methods with $K > 1$ series order is dominated by the polynomial root-finding applied on (13). Since the order of the polynomial in (13) is limited by $2K$, the additional run-time at each time step is of the order $O(K^3)$ [47]. The prediction and update stages of SB-SFRF and FB-SFRF methods

are mainly affected by the chosen filter and chosen number of samples/Fourier coefficients that are used for approximating the posterior pdf. For FB-SFRF method, the asymptotic run-time complexity of both the prediction and update stages of the Fourier filter is $O(D \log D)$ [44], where D is the chosen Fourier series order of the approximated posterior pdf in (33). For SB-SFRF method, the asymptotic run-time complexity depends on the complexity of the chosen sample-based filter.

- 3) *Choosing the value of K :* In general, for periodic cost functions with infinite Fourier series, the value of K can be determined by taking into account the trade-off between the additional computational complexity, induced from using SB-SFRF and FB-SFRF methods, and the accuracy of approximating the cost function Fourier series with a finite K . In addition, the performance improvement with a larger K depends on the posterior pdf approximation accuracy. The periodic cost functions approximation error can be assessed by using Parseval's formula [61, pp. 12–13].

4. EXAMPLE—DOA TRACKING

In this section, we consider the problem of single source DOA tracking by using a uniform circular array (UCA). At time step n , the measurement at the l th sensor is modeled as (see e.g. [43])

$$x_{n,l} = \xi e^{j\zeta \cos(\theta_n - (2\pi l/L))} + \nu_{n,l}, \quad l = 1, \dots, L. \quad (35)$$

where ξ is the signal complex amplitude, which is assumed to be known, $\zeta = 2\pi r/\lambda$, where r is the UCA radius and λ is the signal wavelength, θ_n is the signal DOA, and $\{\nu_n\}$ is an independent identically distributed (i.i.d.) complex circularly symmetric zero mean Gaussian noise vector sequence with known covariance matrix $\sigma^2 \mathbf{I}_L$, where \mathbf{I}_L is the identity matrix of size L . The DOA state model is given by

$$\theta_n = [\theta_{n-1} + w_n]_{2\pi}, \quad (36)$$

where $\{w_n\}$ is an i.i.d. noise in which each element is distributed according to a mixture of two von Mises distributions [14], [32] with known circular means μ_1, μ_2 and concentration parameters κ_1, κ_2 , i.e.

$$f_{w_n}(v) = \begin{cases} \epsilon \frac{e^{\kappa_1 \cos(v-\mu_1)}}{2\pi I_0(\kappa_1)} + (1-\epsilon) \frac{e^{\kappa_2 \cos(v-\mu_2)}}{2\pi I_0(\kappa_2)}, & v \in [-\pi, \pi) \\ 0, & \text{otherwise} \end{cases}, \quad (37)$$

where I_m is the modified Bessel function of order m and $0 \leq \epsilon \leq 1$ is a parameter that determines the weights between the two von Mises distributions. The 2π -periodic extension of f_{w_n} w.r.t. w_n is given by

$$f_{w_n}^{(p)}(v) = \epsilon \frac{e^{\kappa_1 \cos(v-\mu_1)}}{2\pi I_0(\kappa_1)} + (1-\epsilon) \frac{e^{\kappa_2 \cos(v-\mu_2)}}{2\pi I_0(\kappa_2)}, \quad \forall v \in \mathbb{R}.$$

The von Mises distribution is one of the most popular distributions for modeling random parameters with circular nature and is analogous to the Gaussian distribution on the real axis (see e.g. [11], [14], [32], [60]). Many noncircular distributions can be approximated to any desired degree of approximation in terms of Kullback-Leibler divergence, using a finite mixture of Gaussian distributions [19], [27], [46], [52]. Similarly, it is claimed in [32, p. 90] that some circular distributions are fitted well by mixtures of von Mises distributions. In [9] and [34], mixtures of von Mises distributions are used for modeling multimodal distributions on the circle. Multimodal state noise is considered e.g. in [6], [18], [28] for modeling abrupt changes in the state. It is assumed that the sequences $\{w_n\}$ and $\{\nu_n\}$ are statistically independent as well as independent of past and present states. In addition, it is assumed that the prior distribution of θ_0 is uniform, i.e. $\theta_0 \sim U(-\pi, \pi)$. Under this model, the Fourier coefficients of $f_{\theta_0}^{(p)}$ and $f_{\theta_n|\theta_{n-1}}^{(p)}$ are given by

$$\eta_l^{(0)} = \begin{cases} \frac{1}{2\pi}, & l = 0 \\ 0, & \text{otherwise} \end{cases} \quad (38)$$

and

$$\phi_{l,m}^{(n)} = \left(\epsilon \frac{I_{|l|}(\kappa_1)}{2\pi I_0(\kappa_1)} e^{-jl\mu_1} + (1-\epsilon) \frac{I_{|l|}(\kappa_2)}{2\pi I_0(\kappa_2)} e^{-jl\mu_2} \right) \delta_{\text{kr}}(l+m), \quad (39)$$

$\forall l, m \in \mathbb{Z}$, $n = 1, 2, \dots$, respectively. In this case, $f_{\theta_n|\theta_{n-1}}^{(p)}(\alpha_n | \alpha_{n-1}) = f_{w_n}^{(p)}(\alpha_n - \alpha_{n-1})$, $\forall \alpha_n, \alpha_{n-1} \in \Omega_\theta$, which simplifies the prediction stage. As proposed in [44], the Fourier coefficients of $f_{x_n|\theta_n}$ are approximated using the fast Fourier transform [42].

For this problem, the SB-SFRF and FB-SFRF methods are implemented under the MSPE and MAPE risks. For computation of the approximated trigonometric moments in SB-SFRF method, the particle filter from [2]

with the state transition pdf as importance function, is used. The root-finding step is employed by the function ‘roots’ of Matlab.

The SB-SFRF method is implemented with $S = 500$ samples (particles in this case) and the FB-SFRF method is implemented with $D = 40$ Fourier coefficients. In addition, we assume $L = 4$, $\zeta = 10$, $\sigma^2 = 1$, $\mu_1 = 0.95\pi$, $\kappa_1 = 20$, $\mu_2 = 0$, $\kappa_2 = 10$, $\epsilon = 0.5$, and $\xi = (1/\sqrt{2}) + j(1/\sqrt{2})$. The MSPEs and MAPEs of the considered methods are evaluated using 10,000 Monte-Carlo trials. In the following, the posterior mean estimator obtained by particle filter is denoted as Particle-PM. The posterior circular mean estimators obtained by particle and Fourier filters are denoted as Particle-CM and Fourier-CM, respectively. The proposed SB-SFRF-MSPE and SB-SFRF-MAPE methods implemented with particle filter are denoted as Particle-SB-SFRF-MSPE and Particle-SB-SFRF-MAPE, respectively.

In Figs. 2–3, the MSPEs of particle and Fourier estimators are presented versus the time step n , where the corresponding SB-SFRF and FB-SFRF methods are evaluated with $K = 2, 12$. It should be noted that for $K = 1$, the proposed Particle-SB-SFRF-MSPE and FB-SFRF-MSPE methods coincide with the Particle-CM and Fourier-CM, respectively. It can be seen that in both cases, the proposed SFRF-MSPE methods with $K = 12$ result in lower MSPEs than with $K = 2$ and that SFRF-MSPE methods with $K = 2, 12$ have lower MSPEs than the MSPEs of the posterior mean and posterior circular mean estimators.

Figs. 4–5 show the MSPEs of Particle-SB-SFRF-MSPE and FB-SFRF-MSPE methods as a function of the series order, K , averaged over all time steps $n = 1, \dots, 30$. It can be seen that in both cases, the MSPE decreases as K increases. The non-monotonic decrease can be explained by the fact that the error is evaluated w.r.t. the actual periodic cost function, while the minimization is w.r.t. its truncated approximation, which suffers from inaccuracies due to, for example, Gibbs phenomenon [61]. It can be seen that for Particle-SB-SFRF-MSPE and FB-SFRF-MSPE, the MSPEs with $K = 20$ are lower than the MSPEs with $K = 1$ by approximately 16.4% and 15.6%, respectively. For both methods, the MSPEs with $K \geq 5$ are very close to the MSPEs with $K = 20$. Therefore, the choice $K = 5$ seems appropriate in this case. The SPE cost function and its Fourier series approximations are depicted in Fig. 6, with series orders $K = 1, 2, 12$.

In Figs. 7–8 the MAPEs of particle and Fourier estimators are presented versus the time step n , where the corresponding SB-SFRF and FB-SFRF methods are evaluated with $K = 3, 23$. Similar to the MSPE case, for $K = 1$, the proposed Particle-SB-SFRF-MAPE and FB-SFRF-MAPE methods coincide with the Particle-CM and Fourier-CM, respectively. It can be seen that in both cases, SFRF-MAPE methods with $K = 23$ result in lower MAPEs than with $K = 3$ and that SFRF-MAPE

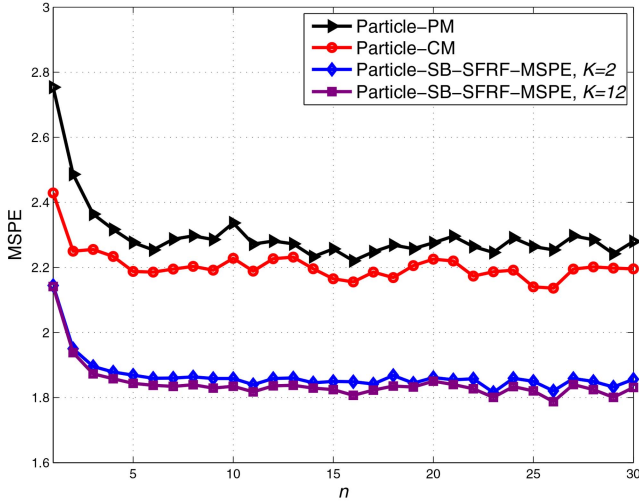


Fig. 2. The MSPEs of Particle-PM, Particle-CM, and Particle-SB-SFRF-MSPE method $K = 2, 12$, with $S = 500$ samples, versus the time step n .




Fig. 4. The MSPE of Particle-SB-SFRF-MSPE method with $S = 500$ samples versus K , averaged over time steps $n = 1, \dots, 30$.




Fig. 3. The MSPEs of Fourier-CM and FB-SFRF-MSPE method $K = 2, 12$, with $D = 40$ Fourier coefficients, versus the time step n .

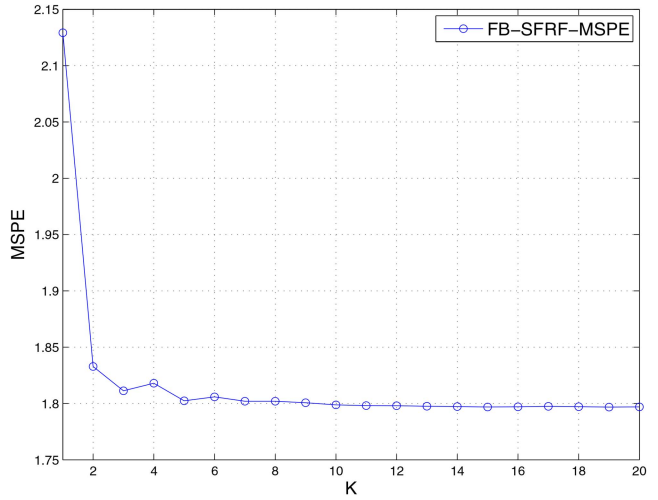


Fig. 5. The MSPE of FB-SFRF-MSPE method with $D = 40$ Fourier coefficients versus K , averaged over time steps $n = 1, \dots, 30$.

methods with $K = 3, 23$ have lower MAPEs than the MAPEs of the posterior circular mean estimators.

Figs. 9–10 show the MAPEs of Particle-SB-SFRF-MAPE and FB-SFRF-MAPE methods as a function of the series order, K , averaged over all time steps $n = 1, \dots, 30$. It can be seen that in both cases, the MAPE decreases as K increases. For Particle-SB-SFRF-MAPE and FB-SFRF-MAPE, the MAPEs with $K = 39$ are lower than the MAPEs with $K = 1$ by approximately 3.5% and 4.2%, respectively. For both methods, the MAPEs with $K \geq 9$ are very close to the MAPEs with $K = 39$. Therefore, the choice $K = 9$ seems appropriate in this case. It can be seen that the improvement in MAPE is smaller for both methods in comparison to the improvement in MSPE. The reason for this phenomenon may be that the SPE increases faster than the APE and thus, a low value of K is highly penalized by the SPE.




Fig. 6. Fourier series approximations of SPE with $K = 1, 2, 12$.

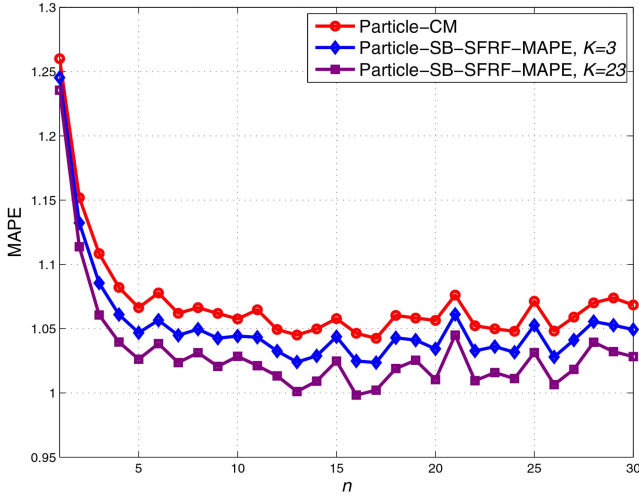


Fig. 7. The MAPEs of Particle-CM and Particle-SB-SFRF-MAPE method $K = 3, 23$, with $S = 500$ samples, versus the time step n .




Fig. 9. The MAPE of Particle-SB-SFRF-MAPE method with $S = 500$ samples versus K , averaged over time steps $n = 1, \dots, 30$.

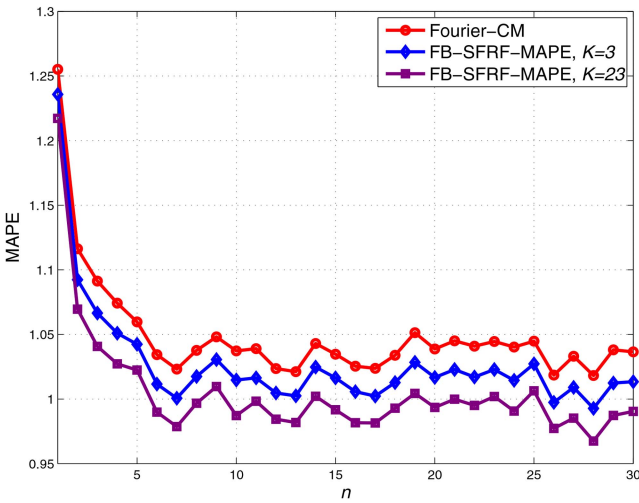


Fig. 8. The MAPEs of Fourier-CM and FB-SFRF-MAPE method $K = 3, 23$, with $D = 40$ Fourier coefficients, versus the time step n .

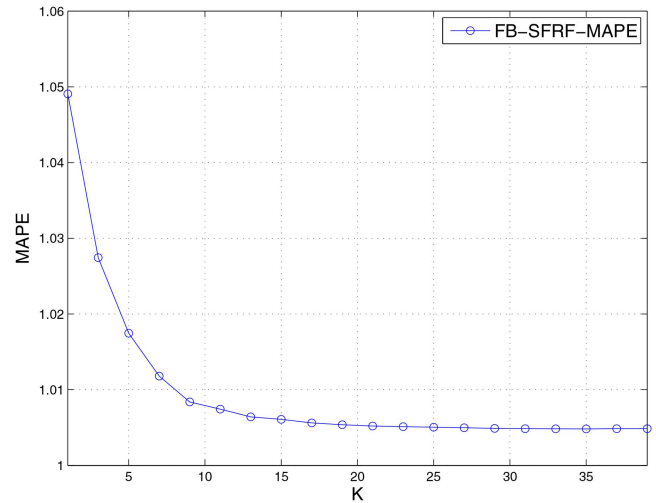


Fig. 10. The MAPE of FB-SFRF-MAPE method with $D = 40$ Fourier coefficients versus K , averaged over time steps $n = 1, \dots, 30$.

In Figs. 11–12, the MSPEs of particle and Fourier estimators are presented versus $\text{SNR} \triangleq |\xi|^2/\sigma^2$, averaged over all time steps $n = 1, \dots, 30$, where the corresponding SB-SFRF and FB-SFRF methods are evaluated with $K = 2, 12$. It can be seen that for both methods, SFRF-MSPE methods with $K = 12$ result in lower MSPEs than SFRF-MSPE methods with $K = 2$ and the corresponding posterior circular mean estimators. In addition, it can be seen that the difference between $K = 2$ and $K = 12$ is small.

As a comparison between SB-SFRF and FB-SFRF methods, we examine their performance in terms of MSPE for a similar number of samples or Fourier coefficients for the posterior pdf approximation, i.e. $S = F = 2D + 1$. In order to broaden the comparison, we implemented a discrete filter [24], used in [44], which is based on wrapped Dirac distribution. In the following, the SB-SFRF-MSPE method implemented with discrete filter is denoted as Discrete-SB-SFRF-MSPE. The

MSPEs of the considered SFRF-MSPE methods with $K = 2$, are presented in Fig. 13 versus the time step n . It can be seen that for $S = F = 9$, FB-SFRF method significantly outperforms SB-SFRF method with particle and discrete filters in all time steps. For $S = F = 201$, FB-SFRF method outperforms SB-SFRF method with particle filter in all time steps and the difference between FB-SFRF method and SB-SFRF method with discrete filter is small for $n \geq 7$. In general, in the considered scenario, the FB-SFRF method seems favourable over SB-SFRF method with both discrete and particle filters, especially in the case, in which only few samples/Fourier coefficients are available. The reason for FB-SFRF advantage may be the discrete representation of the posterior pdf in the frequency domain rather than parameter domain.

Finally, we examine the performance of SB-SFRF and FB-SFRF methods for unimodal state noise. Similar to [11], the a priori pdf of θ_0 is assumed to be a

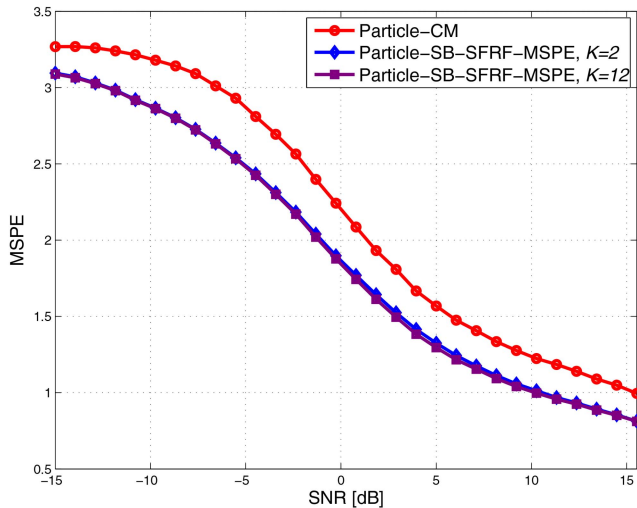


Fig. 11. The MSPEs of Particle-CM and Particle-SB-SFRF-MSPE method $K = 2, 12$, with $S = 500$ samples, versus SNR, averaged over time steps $n = 1, \dots, 30$.




Fig. 12. The MSPEs of Fourier-CM and FB-SFRF-MSPE method $K = 2, 12$, with $D = 40$ Fourier coefficients, versus SNR, averaged over time steps $n = 1, \dots, 30$.

mixture of two von Mises distributions, as in (37), with known circular means $\mu_{0,1} = 0.21$, $\mu_{0,2} = 1.91$, concentration parameters $\kappa_{0,1} = 9$, $\kappa_{0,2} = 39$, and weighting parameter $\epsilon_0 = 0.3$. The state-space model is assumed to be as in (35) and (36), except that each state noise element is assumed to be von Mises distributed with circular mean $\mu = 0$ and concentration parameter $\kappa = 3$. In this scenario, the discrete filter is also used for implementation of SB-SFRF method. In the following, the posterior mean and posterior *circular* mean estimators obtained by discrete filter are denoted as Discrete-PM and Discrete-CM, respectively. In Figs. 14–16 the MSPEs of particle, discrete, and Fourier estimators, are presented versus the time step n , where the corresponding SB-SFRF and FB-SFRF methods are evaluated with $K = 2, 12$. It can be seen that in all cases, SFRF-MSPE methods with $K = 12$ result in lower MSPEs than with

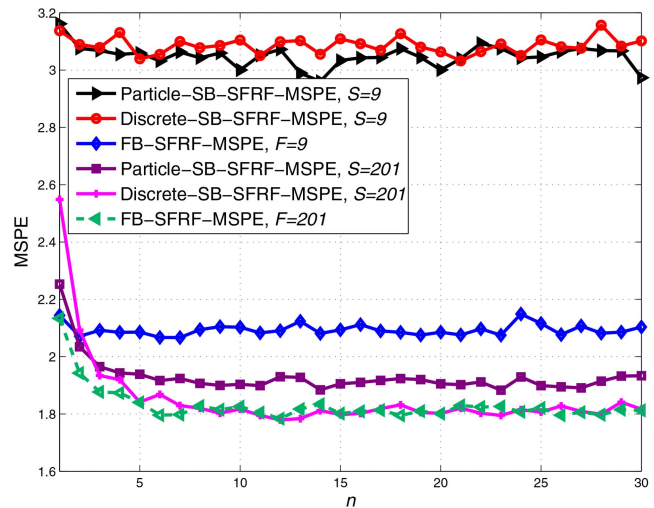


Fig. 13. The MSPEs of Particle-SB-SFRF-MSPE, Discrete-SB-SFRF-MSPE, and FB-SFRF-MSPE methods with $K = 2$, $S = F = 9$ and $S = F = 201$ samples/Fourier coefficients, versus the time step n .

$K = 2$ and that SFRF-MSPE methods with $K = 2, 12$ have lower MSPEs than the MSPEs of the posterior mean and posterior circular mean estimators.

5. CONCLUSION

In this paper, we propose two methods, SB-SFRF and FB-SFRF, for derivation of estimators under general periodic Bayes risks in circular stochastic filtering problems. Both methods utilize the PERF method [47], which is based on Fourier series representation of an arbitrary periodic cost function and polynomial root-finding. The proposed methods are not based on a grid search whose accuracy depends on the chosen grid density. The SFRF methods use approximated trigonometric moments. If the accurate trigonometric moments are available, the SFRF methods coincide with the optimal Bayes solution. Three examples of periodic risks are considered, the MSPE, MAPE, and MCE risks. It is shown that under the MCE risk there exist a tractable estimator. However, under the MSPE and MAPE risks the corresponding estimators cannot be derived analytically for the general case. The SB-SFRF and the FB-SFRF methods can be applied in circular stochastic filtering problems under these risks or any other periodic risk. The superiority of SB-SFRF and FB-SFRF methods w.r.t. state estimation using posterior circular mean is demonstrated in the problem of DOA tracking. A topic for future research is derivation of a method for estimation of a mixed state vector containing both circular and noncircular elements. Another topic is investigating the connection between the considered circular stochastic filtering problems and emerging new approaches, such as sequential deep learning [4], learning of dynamical systems [29], and novel methods for Bayesian filtering with discrete states [41].

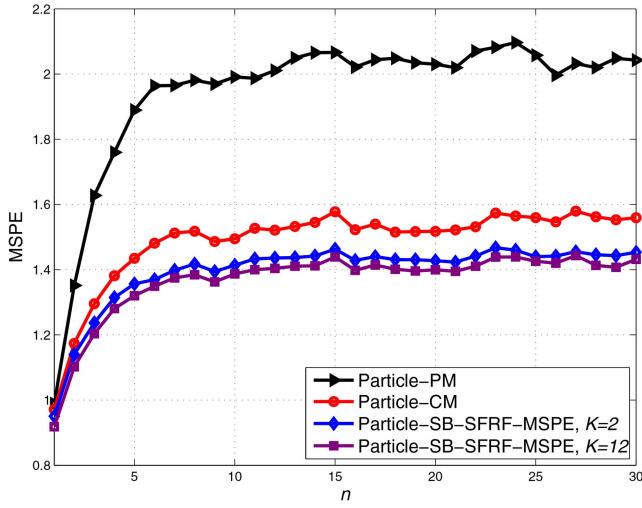


Fig. 14. Unimodal state noise scenario: the MSPEs of Particle-PM, Particle-CM, and Particle-SB-SFRF-MSPE method $K = 2, 12$, with $S = 500$ samples, versus the time step n .

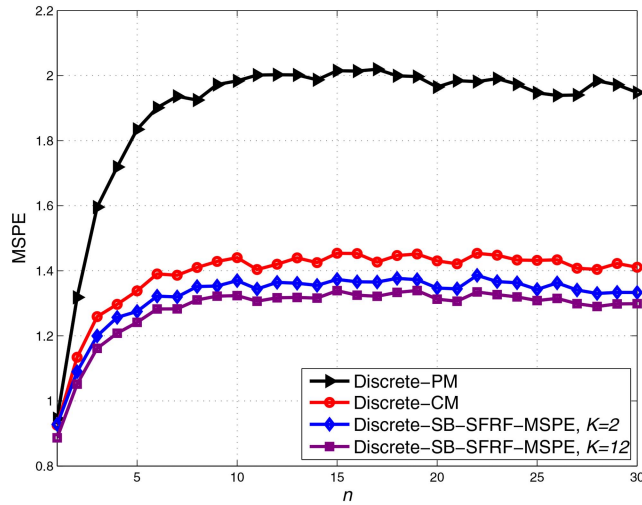


Fig. 15. Unimodal state noise scenario: the MSPEs of Discrete-PM, Discrete-CM, and Discrete-SB-SFRF-MSPE method $K = 2, 12$, with $S = 500$ samples, versus the time step n .

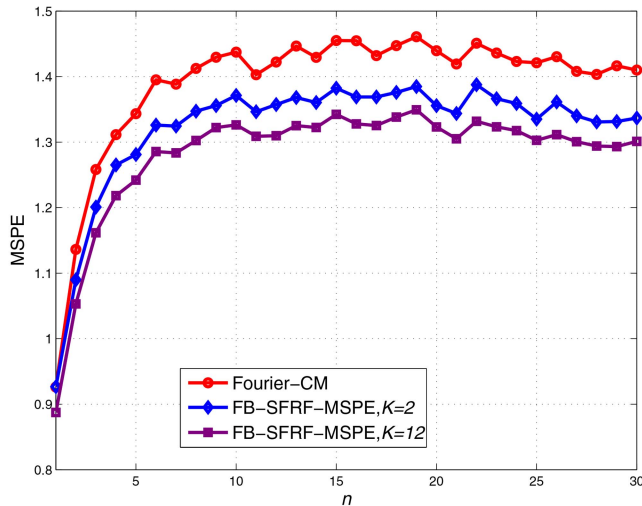


Fig. 16. Unimodal state noise scenario: the MSPEs of Fourier-CM and FB-SFRF-MSPE method $K = 2, 12$, with $D = 40$ Fourier coefficients, versus the time step n .

APPENDIX GENERALIZATION OF THE PREDICTION STAGE UNDER THE APPROACH OF [44] TO AN ARBITRARY STRUCTURE OF $f_{\theta_n|\theta_{n-1}}^{(p)}$

In this appendix, we describe the implementation of the prediction stage, under the approach of [44], with a general $f_{\theta_n|\theta_{n-1}}^{(p)}$, which is not necessarily equal to a function of $\theta_n - \theta_{n-1}$. At each time step n , it is assumed that the function $\sqrt{f_{\theta_n|\theta_{n-1}}^{(p)}}$ can be represented via a two-dimensional Fourier series with Fourier coefficients $\{\phi_{l,m}^{(n,\text{sqr})}\}_{l,m \in \mathbb{Z}}$, i.e.

$$\sqrt{f_{\theta_n|\theta_{n-1}}^{(p)}(\alpha_n | \alpha_{n-1})} = \sum_{l=-\infty}^{\infty} \sum_{m=-\infty}^{\infty} \phi_{l,m}^{(n,\text{sqr})} e^{jl\alpha_n} e^{jm\alpha_{n-1}}. \quad (40)$$

By using (40), it can be verified (see e.g. [8]) that the Fourier coefficients of $f_{\theta_n|\theta_{n-1}}^{(p)} = \sqrt{f_{\theta_n|\theta_{n-1}}^{(p)}} \sqrt{f_{\theta_n|\theta_{n-1}}^{(p)}}$ from (26) are given by the following two-dimensional discrete convolution

$$\phi_{l,m}^{(n)} = \sum_{q=-\infty}^{\infty} \sum_{r=-\infty}^{\infty} \phi_{l-q,r}^{(n,\text{sqr})} \phi_{l,m-r}^{(n,\text{sqr})}, \quad (41)$$

$\forall l, m \in \mathbb{Z}$. In practice, the Fourier coefficients in the r.h.s. of (40) are usually computed numerically and the corresponding two-dimensional Fourier series is truncated. Therefore, the two-dimensional discrete convolution in (41) is finite, resulting in $\tilde{\phi}_{l,m}^{(n)}$ which approximates $\phi_{l,m}$, $\forall l = -D, \dots, D$, $m = -D, \dots, D$. For $|l| > D$ or $|m| > D$, $\phi_{l,m}$ is approximated by zero. The predicted pdf $f_{\theta_n|\mathbf{x}^{(n-1)}}$ is approximated by $f_{n|n-1,\text{FB}}^{(D)}: \Omega_\theta \rightarrow \mathbb{R}$ given by

$$f_{n|n-1,\text{FB}}^{(D)}(y) \triangleq \sum_{l=-D}^D \eta_{l,\text{FB}}^{(n|n-1,D)} e^{ily}, \quad (42)$$

where y is the argument of $f_{n|n-1,\text{FB}}^{(D)}$,

$$\eta_{l,\text{FB}}^{(n|n-1,D)} = 2\pi \sum_{m=-D}^D \eta_{-m,\text{FB}}^{(n-1,D)} \tilde{\phi}_{l,m}^{(n)}, \quad \forall l = -D, \dots, D. \quad (43)$$


Then, the Fourier series representation of $\sqrt{f_{n|n-1,\text{FB}}^{(D)}}$ is derived by applying the procedure, proposed in [44], for obtaining the Fourier coefficients of the square root of a pdf from the Fourier coefficients of the actual pdf.

REFERENCES

- [1] B. D. Anderson and J. B. Moore. Optimal filtering. Englewood Cliffs, NJ: Prentice-Hall, 1979.
- [2] M. S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp. A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking. IEEE Transactions on Signal Processing, 50, 2 (Feb. 2002), 174–188.
- [3] M. Azmani, S. Reboul, J.-B. Choquel, and M. Benjelloun. A recursive fusion filter for angular data. In Proceedings of the IEEE International Conference on Robotics and Biomimetics (ROBIO), Dec. 2009, 882–887.

- [4] M. Baccouche, F. Mamalet, C. Wolf, C. Garcia, and A. Baskurt
Sequential deep learning for human action recognition.
In *Human Behavior Understanding*, 2011, 29–39.
- [5] S. Basu and Y. Bresler
A global lower bound on parameter estimation error with periodic distortion functions
IEEE Transactions on Information Theory, 46, 3 (May 2000), 1145–1150.
- [6] I. Bilik and J. Tabrikian
Maneuvering target tracking in the presence of glint using the nonlinear Gaussian mixture Kalman filter.
IEEE Transactions on Aerospace and Electronic Systems, 46, 1 (Jan. 2010), 246–262.
- [7] D. Brunn, F. Sawo, and U. D. Hanebeck
Efficient nonlinear Bayesian estimation based on Fourier densities.
In *Proceedings of the IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems*, Sep. 2006, 317–322.
- [8] D. Brunn, F. Sawo, and U. D. Hanebeck
Nonlinear multidimensional Bayesian estimation with Fourier densities.
In *Proceedings of the 45th IEEE Conference on Decision and Control*, Dec. 2006, 1303–1308.
- [9] J. A. Carta, C. Bueno, and P. Ramírez
Statistical modelling of directional wind speeds using mixtures of von Mises distributions: case study.
Energy conversion and management, 49, 5 (2008), 897–907.
- [10] Z. Chen
Bayesian filtering: From Kalman filters to particle filters, and beyond.
Statistics, 182, 1 (2003), 1–69.
- [11] M. Costa, V. Koivunen, and H. V. Poor
Estimating directional statistics Using wavefield modeling and mixtures of von-Mises distributions.
IEEE Signal Processing Letters, 21, 12 (Dec. 2014), 1496–1500.
- [12] G. A. Einicke and L. B. White
Robust extended Kalman filtering.
IEEE Transactions on Signal Processing, 47, 9 (Sep. 1999), 2596–2599.
- [13] N. J. Gordon, D. J. Salmond, and A. F. M. Smith
Novel approach to nonlinear/non-Gaussian Bayesian state estimation.
IEE Proceedings F—Radar and Signal Processing, 140, 2 (Apr. 1993), 107–113.
- [14] S. R. Jammalamadaka and A. Sengupta
Topics in circular statistics.
World Scientific, 2001, vol. 5.
- [15] S. J. Julier and J. K. Uhlmann
Unscented filtering and nonlinear estimation.
Proceedings of the IEEE, 92, 3 (Mar. 2004), 401–422.
- [16] X. Kai, C. Wei, and L. Liu
Robust extended Kalman filtering for nonlinear systems with stochastic uncertainties.
IEEE Transactions on Systems, Man, and Cybernetics—Part A: Systems and Humans, 40, 2 (Mar. 2010), 399–405.
- [17] R. E. Kalman
A new approach to linear filtering and prediction problems.
Journal of Fluids Engineering, 82, 1 (1960), 35–45.
- [18] J. H. Kotecha and P. M. Djuric
Gaussian sum particle filtering.
IEEE Transactions on Signal Processing, 51, 10 (Oct. 2003), 2602–2612.
- [19] R. J. Kozick and B. M. Sadler
Maximum-likelihood array processing in non-Gaussian noise with Gaussian mixtures.
IEEE Transactions on Signal Processing, 48, 12 (Dec. 2000), 3520–3535.
- [20] G. Kurz, I. Gilitschenski, and U. D. Hanebeck
Recursive nonlinear filtering for angular data based on circular distributions.
In *Proceedings of the American Control Conference (ACC)*, June 2013, 5439–5445.
- [21] G. Kurz, I. Gilitschenski, and U. D. Hanebeck
Nonlinear measurement update for estimation of angular systems based on circular distributions.
In *Proceedings of the American Control Conference (ACC)*, June 2014, 5694–5699.
- [22] G. Kurz, I. Gilitschenski, and U. D. Hanebeck
Recursive Bayesian filtering in circular state spaces.
arXiv preprint arXiv:1501.05151 (2015).
- [23] G. Kurz, I. Gilitschenski, S. Julier, and U. D. Hanebeck
Recursive estimation of orientation based on the Bingham distribution.
In *Proceedings of the 16th International Conference on Information Fusion (FUSION)*, July 2013, 1487–1494.
- [24] G. Kurz, I. Gilitschenski, F. Pfaff, and L. Drude
libDirectional, 2015.
Homepage: <https://github.com/libDirectional>.
- [25] E. L. Lehmann and G. Casella
Theory of Point Estimation, 2nd edition.
Springer Texts in Statistics. New York: Springer, 1998.
- [26] J. Leitão and J. Moura
Acquisition in phase demodulation: application to ranging in radar/sonar systems.
IEEE Transactions on Aerospace and Electronic Systems, 31, 2 (Apr. 1995), 581–599.
- [27] J. Q. Li and A. R. Barron
Mixture density estimation.
In *Advances in Neural Information Processing Systems*, 12, (1999).
- [28] X. R. Li and V. P. Jilkov
Survey of maneuvering target tracking. Part V. Multiple-model methods.
IEEE Transactions on Aerospace and Electronic Systems, 41, 4 (Oct. 2005), 1255–1321.
- [29] F. Lindsten
Particle filters and Markov chains for learning of dynamical systems.
Ph.D. dissertation, Linköping University, Linköping, Sweden, 2013.
- [30] J. T.-H. Lo and A. S. Willsky
Estimation for rotational processes with one degree of freedom.
DTIC Document, Technical report, 1972.
- [31] B. C. Lovell and R. C. Williamson
The statistical performance of some instantaneous frequency estimators.
IEEE Transactions on Signal Processing, 40, 7 (July 1992), 1708–1723.
- [32] K. V. Mardia and P. E. Jupp
Directional Statistics.
Wiley Series in Probability and Statistics. Chichester: Wiley, 1999.
- [33] I. Markovic, F. Chaumette, and I. Petrovic
Moving object detection, tracking and following using an omnidirectional camera on a mobile robot.
In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, May 2014, 5630–5635.

- [34] I. Markovic and I. Petrovic
Bearing-only tracking with a mixture of von Mises distributions.
In Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Oct. 2012, 707–712.
- [35] R. G. McWilliam, B. G. Quinn, and I. V. L. Clarkson
Direction estimation by minimum squared arc length.
IEEE Transactions on Signal Processing, 60, 5 (May 2012), 2115–2124.
- [36] N. Nikolaidis and I. Pitas
Nonlinear processing and analysis of angular signals.
IEEE Transactions on Signal Processing, 46, 12 (Dec. 1998), 3181–3194.
- [37] E. Nitzan, T. Routtenberg, and J. Tabrikian
Cyclic Bayesian Cramér-Rao bound for filtering in circular state space.
In Proceedings of the 18th International Conference on Information Fusion (FUSION), July 2015, 734–741.
- [38] E. Nitzan, T. Routtenberg, and J. Tabrikian
A New Class of Bayesian Cyclic Bounds for Periodic Parameter Estimation.
IEEE Transactions on Signal Processing, 64, 1 (Jan. 2016), 229–243.
- [39] E. Nitzan, T. Routtenberg, and J. Tabrikian
Mean-cyclic-error lower bounds via integral transform of likelihood-ratio function.
Accepted to the 9th IEEE Sensor Array and Multichannel Signal Processing Workshop (SAM), July 2016.
- [40] E. Nitzan, J. Tabrikian, and T. Routtenberg
Bayesian cyclic bounds for periodic parameter estimation.
In Proceedings of the 5th IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP), Dec. 2013, 308–311.
- [41] M. Nyolt and T. Kirste
On resampling for Bayesian filters in discrete state spaces.
In Proceedings of the IEEE 27th International Conference on Tools with Artificial Intelligence (ICTAI), Nov. 2015, 526–533.
- [42] A. V. Oppenheim and R. W. Schaffer
Discrete Time Signal Processing, 3rd edition.
Prentice-Hall, 2010.
- [43] M. Orton and W. Fitzgerald
A Bayesian approach to tracking multiple targets using sensor arrays and particle filters.
IEEE Transactions on Signal Processing, 50, 2 (Feb. 2002), 216–223.
- [44] F. Pfaff, G. Kurz, and U. D. Hanebeck
Multimodal circular filtering using Fourier series.
In Proceedings of the 18th International Conference on Information Fusion (FUSION), July 2015, 711–718.
- [45] D. Rife and R. Boorstyn
Single tone parameter estimation from discrete-time observations.
IEEE Transactions on Information Theory, 20, 5 (Sep. 1974), 591–598.
- [46] T. Routtenberg and J. Tabrikian
MIMO-AR system identification and blind source separation for GMM-distributed sources.
IEEE Transactions on Signal Processing, 57, 5 (May 2009), 1717–1730.
- [47] T. Routtenberg and J. Tabrikian
Bayesian parameter estimation using periodic cost functions.
IEEE Transactions on Signal Processing, 60, 3 (Mar. 2012), 1229–1240.
- [48] T. Routtenberg and J. Tabrikian
Non-Bayesian periodic Cramér-Rao bound.
IEEE Transactions on Signal Processing, 61, 4 (Feb. 2013), 1019–1032.
- [49] T. Routtenberg and J. Tabrikian
Cyclic Barankin-type bounds for non-Bayesian periodic parameter estimation.
IEEE Transactions on Signal Processing, 62, 13 (July 2014), 3321–3336.
- [50] T. Routtenberg and J. Tabrikian
A fresh look at the cyclic Cramér-Rao bound for periodic parameter estimation.
Accepted to the 19th International Conference on Information Fusion (FUSION), July 2016.
- [51] S. Sarkka
On unscented Kalman filtering for state estimation of continuous-time nonlinear systems.
IEEE Transactions on Automatic Control, 52, 9 (Sep. 2007), 1631–1641.
- [52] H. W. Sorenson and D. L. Alspach
Recursive Bayesian estimation using Gaussian sums.
Automatica, 7, 4 (1971), 465–479.
- [53] G. Stienne, S. Reboul, M. Azmani, J.-B. Choquel, and M. Benjelloun
A multi-sensor circular particle filter applied to the fusion of the GPS-L2C channels.
In Proceedings of the 14th International Conference on Information Fusion (FUSION), July 2011, 1–8.
- [54] T. Tao
An introduction to measure theory.
American Mathematical Society, 2011.
- [55] P. Tichavsky, C. H. Muravchik, and A. Nehorai
Posterior Cramér-Rao bounds for discrete-time nonlinear filtering.
IEEE Transactions on Signal Processing, 46, 5 (May 1998), 1386–1396.
- [56] H. L. Van Trees and K. L. Bell
Bayesian Bounds for Parameter Estimation and Nonlinear Filtering/Tracking.
Wiley-IEEE Press, 2007.
- [57] A. S. Willsky
Fourier series and estimation on the circle with applications to synchronous communication—Part I: Analysis.
IEEE Transactions on Information Theory, 20, 5 (Sep. 1974), 577–583.
- [58] A. S. Willsky and J. T.-H. Lo
Estimation for rotational processes with one degree of freedom—Part II: Discrete-time processes.
IEEE Transactions on Automatic Control, 20, 1 (Feb. 1975), 22–30.
- [59] A. Xenaki, P. Gerstoft, and E. Fernandez-Grande
Sparse DOA estimation with polynomial rooting.
In Proceedings of the 3rd International Workshop on Compressed Sensing Theory and its Applications to Radar, Sonar and Remote Sensing (CoSeRa), June 2015, 104–108.
- [60] D. Zachariah, P. Wirfält, M. Jansson, and S. Chatterjee
Line spectrum estimation with probabilistic priors.
Signal Processing, 93, 11 (2013), 2969–2974.
- [61] A. Zygmund
Trigonometric series, 3rd edition.
Cambridge university press, 2002.



Eyal Nitzan received the B.Sc. (cum laude) and M.Sc. (cum laude) degrees in Electrical and Computer Engineering from the Ben-Gurion University of the Negev, Beer-Sheva, Israel, in 2012 and 2014, respectively. He is currently working toward the Ph.D. degree in the Dept. of Electrical and Computer engineering at Ben-Gurion University of the Negev. His research interests include estimation theory, statistical signal processing, and detection theory. He is the recipient of the Best Student Paper Award (1st prize) in IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP) 2013. He is the recipient of the Negev scholarship awarded to excellent Ph.D. candidates in 2014.



Tirza Routtenberg received the B.Sc. degree (magna cum laude) in bio-medical engineering from the Technion Israel Institute of Technology, Haifa, Israel in 2005 and the M.Sc. (magna cum laude) and Ph.D. degrees in electrical engineering from the Ben-Gurion University of the Negev, Beer-Sheva, Israel, in 2007 and 2012, respectively. She was a postdoctoral fellow with the School of Electrical and Computer Engineering, Cornell University, in 2012–2014. Since October 2014 she is a faculty member at the Department of Electrical and Computer Engineering at Ben-Gurion University of the Negev, Beer-Sheva, Israel. Her research interests include signal processing in smart grid, statistical signal processing, and estimation and detection theory. She was a recipient of the Best Student Paper Award in International Conference on Acoustics, Speech and Signal Processing (ICASSP) 2011 and in IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP) 2013 (coauthor). She was awarded the Negev scholarship in 2008, the Lev-Zion scholarship in 2010, and the Marc Rich foundation prize in 2011.



Joseph Tabrikian received the B.Sc., M.Sc., and Ph.D. degrees in Electrical Engineering from the Tel-Aviv University, Tel-Aviv, Israel, in 1986, 1992, and 1997, respectively. During 1996–1998 he was with the Department of Electrical and Computer Engineering, Duke University, Durham, NC as an Assistant Research Professor. He is now with the Department of Electrical and Computer Engineering at Ben-Gurion University of the Negev, Beer-Sheva, Israel. He served as Associate Editor of the IEEE Transactions on Signal Processing during 2001–2004 and 2011–2015 and as Associate Editor of the IEEE Signal Processing Letters during 2012–2015. Since 2015 he has been a Senior Area Editor of the IEEE Signal Processing Letters. He served as a member of the IEEE SAM technical committee during 2010–2015 and was the technical program co-chair of the IEEE SAM 2010 workshop. He is co-author of 5 award-winning papers in different IEEE conferences. His research interests include estimation and detection theory and radar signal processing.

Methods for Deterministic Approximation of Circular Densities

GERHARD KURZ
IGOR GILITSCHENSKI
ROLAND Y. SIEGWART
UWE D. HANEBECK

Circular estimation problems arise in many applications and can be addressed with methods based on circular distributions, e.g., the wrapped normal distribution and the von Mises distribution. To develop nonlinear circular filters, a deterministic sample-based approximation of these distributions with a so-called wrapped Dirac mixture distribution is beneficial. We present a closed-form solution to obtain a symmetric wrapped Dirac mixture with five components based on matching the first two trigonometric moments. Furthermore, we discuss the choice of a scaling parameter involved in this method and extend it by superimposing samples obtained from different scaling parameters. Finally, we propose an approximation based on a binary tree that approximates the shape of the true density rather than its trigonometric moments. All proposed approaches are thoroughly evaluated and compared in different scenarios.

Manuscript received December 4, 2015; revised June 17, 2016; released for publication July 26, 2016.

Refereeing of this contribution was handled by Chee-Yee Chong.

Authors' addresses: G. Kurz and U. Hanebeck, Intelligent Sensor-Actuator-Systems Laboratory (ISAS), Institute for Anthropomatics and Robotics, Karlsruhe Institute of Technology (KIT), Germany (E-mail: gerhard.kurz@kit.edu, uwe.hanebeck@ieee.org). I. Gilitschenski and R. Siegwart, Autonomous Systems Laboratory (ASL), Institute of Robotics and Intelligent Systems, Swiss Federal Institute of Technology Zurich, Switzerland (E-mail: igilitschenski@ethz.ch, rsiegwart@ethz.ch).

This is an extended version of the paper "Deterministic Approximation of Circular Densities with Symmetric Dirac Mixtures Based on Two Circular Moments" [32] published at the *17th International Conference on Information Fusion (Fusion 2014)*, which received the *Jean-Pierre Le Cadre Award for Best Paper*.

1557-6418/16/\$17.00 © 2016 JAIF

1. INTRODUCTION

Many estimation problems involve circular quantities, for example the orientation of a vehicle or the angle of a robotic joint. Since conventional estimation algorithms perform poorly in these applications, particularly if the angular uncertainty is high, circular estimation methods such as [31], [34], [35], [6], [54], [57], and [43] have been proposed. These methods use circular probability distributions stemming from the field of directional statistics [21], [42].

Circular estimation methods have been applied to a variety of problems in different fields. For example, many signal processing applications necessitate the consideration of circular quantities. Consider for instance phase estimation and tracking [39], [7], signal processing for global navigation satellite systems (GNSS) [54], [53], [26], and azimuthal speaker tracking [57]. In meteorology, estimation of the wind direction [11], [8] is of interest and in aerospace applications, the heading of an airplane may be estimated [35]. Through a suitable mapping, constrained object tracking problems on periodic one-dimensional manifolds can also be interpreted as circular estimation problems [29]. Finally, circular densities arise naturally in bearings-only tracking [12], [44].

To facilitate the development of nonlinear filters, sample-based approaches are commonly used. The reason is that samples, which we represent as Dirac delta distributions, can easily be propagated through nonlinear functions. We distinguish deterministic and nondeterministic approaches. In the noncircular case, typical examples for deterministic approaches include the unscented Kalman filter (UKF) [24] as well as extensions thereof [56], the cubature Kalman filter [4], [23], [22], and the smart sampling Kalman filter (S^2KF) [52]. Nondeterministic filters for the noncircular case are the particle filter [5], the Gaussian particle filter [25], and the randomized UKF [55].

We focus on deterministic approaches because they have several distinct advantages. First of all, as a result of their deterministic nature, all results are reproducible, i.e., for the same input (e.g., measurements and initial estimate), deterministic filters will always produce the same output.¹ Second, the samples are placed according to certain optimality criteria (i.e., moment matching [24], shape approximation [16], [50]), or a combination thereof [19], [17]. Consequently, a much smaller number of samples is sufficient to achieve a good approximation. Third, nondeterministic approaches usually have a certain probability of causing the filtering algorithm to fail due to a poor choice of samples. This is avoided in deterministic methods.

¹Randomized approaches can be made reproducible by choosing a fixed seed for the random number generator. However, this choice is completely arbitrary and affects the performance. Also, minor changes to the implementation, e.g., the order in which certain random numbers are drawn or the choice of the underlying random number generator, will affect the result.

In our previous publication [31], we presented a deterministic approximation for von Mises and wrapped normal distributions with three samples. This approximation is based on matching the first trigonometric moment. The first trigonometric moment is a complex number and a measure of both location and dispersion. This approximation has already been applied to constrained object tracking [29], sensor scheduling based on bearing-only measurements [12], as well as stochastic model predictive control [28].

The contributions of this paper can be summarized as follows. We present an extension of our previous approach [31] to match both the first and the second trigonometric moment, which was first discussed in [32]. This yields an approximation with five samples. Even though this approximation is slightly more complicated, it can still be computed in closed form and does not require any numerical computations or approximations. We have previously applied this method to the problem of heart phase estimation in [39].

The algorithm from [32] requires choosing a parameter $\lambda \in [0, 1]$. In this paper, we will show that the choice $\lambda = 0.8$ ensures good approximations even when the approximated distribution converges to a uniform distribution.

Furthermore, we present a novel superposition method that is able to combine sample sets with different choices of λ in order to obtain a larger number of samples while still maintaining the first and second trigonometric moment.

Finally, we also propose a new method based on the shape of the probability distribution function rather than its moments. This method creates a binary tree consisting of intervals in $[0, 2\pi)$ and distributes the samples in proportion to the probability mass contained in the interval. Unlike previous shape-based methods such as [17], the proposed method does not require numerical optimization. Thus, it is very fast, provided an efficient algorithm for calculating the cumulative distribution function of the respective density is available.

2. PREREQUISITES

In this section, we define the required probability distributions (see Fig. 1) and introduce the concept of trigonometric moments.

DEFINITION 1 (Wrapped Normal Distribution). A wrapped normal (WN) distribution [48] is given by the probability density function (pdf)

$$f(x; \mu, \sigma) = \frac{1}{\sqrt{2\pi}\sigma} \sum_{k=-\infty}^{\infty} \exp\left(-\frac{(x - \mu + 2k\pi)^2}{2\sigma^2}\right),$$

where $\mu \in [0, 2\pi)$ and $\sigma > 0$ are parameters for center and dispersion, respectively.

The WN distribution is obtained by wrapping a one-dimensional Gaussian density around the unit circle. It is of particular interest because it appears as a limit




Fig. 1. Probability density functions of WN, WC, and VM distributions with identical first trigonometric moment.

distribution on the circle, i.e., in a circular setting, it is reasonable to assume that noise is WN distributed. To see this, we consider i.i.d. random variables θ_i with $\mathbb{E}(\theta_i) = 0$ and finite variance. Then the sum

$$S_n = \frac{1}{\sqrt{n}} \sum_{k=1}^n \theta_k$$

converges to a normally distributed random variable if $n \rightarrow \infty$. Consequently, the wrapped sum ($S_n \bmod 2\pi$) converges to a WN-distributed random variable. Numerical computation of the pdf is discussed in [33].

DEFINITION 2 (Wrapped Cauchy Distribution). The wrapped Cauchy (WC) distribution [21], [42] has the pdf

$$f(x; \mu, \gamma) = \frac{1}{\pi} \sum_{k=-\infty}^{\infty} \frac{\gamma}{\gamma^2 + (x - \mu + 2k\pi)^2},$$

where $\mu \in [0, 2\pi)$ and $\gamma > 0$.

Similar to the WN distribution, a WC distribution is obtained by wrapping a Cauchy distribution around the circle. Unlike the WN distribution, it is possible to simplify the infinite sum in this case, yielding the closed-form expression

$$f(x; \mu, \gamma) = \frac{1}{2\pi} \frac{\sinh(\gamma)}{\cosh(\gamma) - \cos(x - \mu)}.$$

DEFINITION 3 (Von Mises Distribution). A von Mises (VM) distribution [58] is defined by the pdf

$$f(x; \mu, \kappa) = \frac{1}{2\pi I_0(\kappa)} \exp(\kappa \cos(x - \mu)),$$

where $\mu \in [0, 2\pi)$ and $\kappa > 0$ are parameters for location and concentration, respectively, and $I_0(\cdot)$ is the modified Bessel function of order 0.

According to [1, eq. 9.6.19], the modified Bessel function of integer order n is given by

$$I_n(z) = \frac{1}{\pi} \int_0^\pi \exp(z \cos \theta) \cos(n\theta) d\theta.$$

The von Mises distribution has a similar shape as a WN distribution and is frequently used in circular statistics.

DEFINITION 4 (Wrapped Dirac Distribution). A wrapped Dirac mixture (WD) distribution has the pdf

$$f(x; w_1, \dots, w_L, \beta_1, \dots, \beta_L) = \sum_{j=1}^L w_j \delta(x - \beta_j),$$

where L is the number of components, $\beta_1, \dots, \beta_L \in [0, 2\pi)$ are the Dirac positions, $w_1, \dots, w_L > 0$ are the weighting coefficients and δ is the Dirac delta distribution [31]. Additionally, we require $\sum_{j=1}^L w_j = 1$ to ensure that the WD distribution is normalized.

Unlike the continuous WN, WC and VM distributions, the WD distribution is a discrete distribution consisting of a certain number of Dirac delta components. These components can be seen as a set of samples and can be used to approximate a certain original density. WD distributions are useful for nonlinear estimation because they can easily be propagated through nonlinear functions [31], just as Dirac mixture densities in \mathbb{R}^n [52]. The WD distribution as defined above does not contain an infinite sum for wrapping, because wrapping a Dirac distribution results in a single component according to

$$\sum_{k=-\infty}^{\infty} \delta(x + 2\pi k - \beta) = \delta((x - \beta) \bmod 2\pi),$$

where $x \in [0, 2\pi)$. For consistency with the WN and WC distributions, we still refer to the WD distribution as *wrapped*, because it can be obtained by wrapping a Dirac mixture on \mathbb{R} , which results in taking all Dirac positions modulo 2π (see [27, Remark 3]). Thus, all properties of wrapped distributions apply to the WD distribution as well.

DEFINITION 5 (Trigonometric Moments). The n th trigonometric (or circular) moment of a random variable x with pdf $f(\cdot)$ is given by

$$m_n = \mathbb{E}(\exp(ix)^n) = \int_0^{2\pi} \exp(inx) f(x) dx,$$

where i is the imaginary unit [21], [42].

Trigonometric moments are the circular analogon to the conventional real-valued power moments $\mathbb{E}(x^n)$. Note, however, that $m_n \in \mathbb{C}$ is a complex number. For this reason, the first trigonometric moment already describes both location and dispersion of the distribution, similar to the first two conventional real-valued moments. The argument of the complex number is analogous to the mean whereas the absolute value describes the concentration.




Fig. 2. First trigonometric moment of wrapped normal, wrapped Cauchy, and von Mises distributions with zero mean plotted against their second trigonometric moment. The moments are real-valued in this case because $\mu = 0$.

LEMMA 1. *The trigonometric moments of WN, WC, VM, and WD distributions are given by*

$$m_n^{WN} = \exp(in\mu - n^2\sigma^2/2), \quad (1)$$

$$m_n^{WC} = \exp(in\mu - |n|\gamma), \quad (2)$$

$$m_n^{VM} = \frac{I_{|n|}(\kappa)}{I_0(\kappa)} \exp(in\mu), \quad (3)$$

$$m_n^{WD} = \sum_{j=1}^L w_j \exp(in\beta_j). \quad (4)$$

Derivations can be found in [27, Lemma 2]. The quotient of Bessel functions can be calculated numerically with the algorithm by [3]. Pseudocode for this algorithm can be found in [31, Fig. 4].

The parameters of WN, WC, and VM distributions are uniquely defined by the first trigonometric moment. However, WN, WC, and VM distributions with equal first moments significantly differ in their higher moments. This is illustrated in Fig. 1 and Fig. 2. This difference motivates the use of the second trigonometric moments in deterministic Dirac mixture approximations.

3. MOMENT-BASED DETERMINISTIC APPROXIMATION

In this section, we derive deterministic methods for computing Dirac approximation of WN, WC, and VM distributions. Without loss of generality, we only consider the case $\mu = 0$ in order to simplify the calculations. In the case of a location $\mu \neq 0$, the samples are computed for $\mu = 0$ and subsequently shifted by μ . The moment formulas (1)–(3) simplify to $m_n^{WN} = \exp(-n^2\sigma^2/2)$, $m_n^{WC} = \exp(-|n|\gamma)$, and $m_n^{VM} = I_{|n|}(\kappa)/I_0(\kappa)$. In particular, we find $\text{Im}(m_n^{WN}) = \text{Im}(m_n^{WC}) = \text{Im}(m_n^{VM}) = 0$, so there is no imaginary part and our calculations only involve real numbers. More

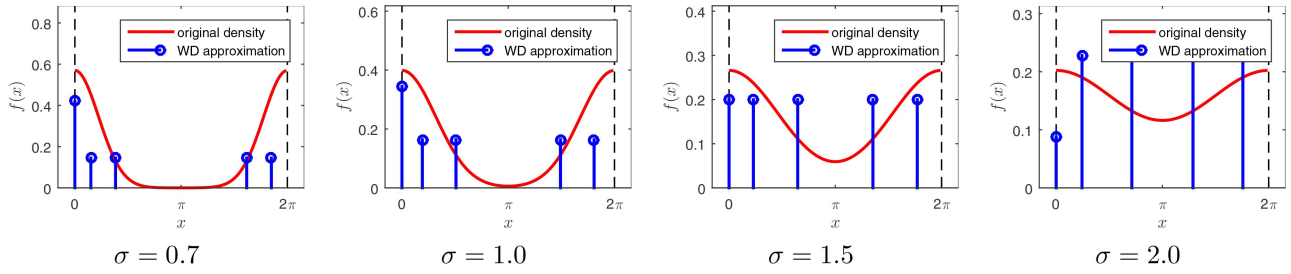


Fig. 3. WD approximations for WN distributions with different values for σ . In all cases, we use $\lambda = 0.5$.




Fig. 4. WD approximations for VM distributions with different values for κ . In all cases, we use $\lambda = 0.5$.

generally, for any circular distribution symmetric around $\mu = 0$, it holds that

$$\begin{aligned} \text{Imm}_n &= \int_0^{2\pi} \sin(nx) f(x) dx \\ &= \underbrace{\int_{-\pi}^{\pi} \sin(nx) f(x) dx}_{\text{odd function}} = 0. \end{aligned}$$

Note that this property only holds for symmetric circular distributions. In general, even for $\mu = 0$, only the first trigonometric moment is guaranteed to have no imaginary part, whereas higher moments are not necessarily real-valued.

3.1. First Circular Moment

First, we derive the approximation based on the first moment.

3.1.1. Two Components:

Obviously, one WD component is not sufficient to match a given first moment, because a single component only has a single degree of freedom, whereas the first moment has two degrees of freedom. For this reason, we propose a solution with $L = 2$ components, the minimum number possible. We use symmetric WD positions $\beta_1 = -\phi$, $\beta_2 = \phi$, and equal weights $w_1 = w_2 = \frac{1}{2}$. For the first moment, we have

$$m_1^{\text{WD}} = \sum_{j=1}^L w_j \exp(i\beta_j) = \cos(\phi).$$

Solving for ϕ results in $\phi = \arccos(m_1)$.

3.1.2. Three Components:

Now we extend the mixture with two components by adding a component at the circular mean. Con-

sider the WD distribution with $L = 3$ components, Dirac positions $\beta_1 = -\phi$, $\beta_2 = \phi$, $\beta_3 = 0$, and equal weights² $w_1 = w_2 = w_3 = \frac{1}{3}$. For the first moment, we have

$$m_1^{\text{WD}} = \sum_{j=1}^L w_j \exp(i\beta_j) = \frac{1}{3}(2\cos(\phi) + 1).$$

Now, we match with the first moment m_1 of a WN or VM distribution and obtain

$$\underbrace{\frac{1}{2}(3m_1 - 1)}_{=:c_1} = \cos(\phi).$$

Thus, we use $\phi = \arccos(c_1)$ to obtain a solution for the WD distribution.

This approximation method is closely related to the approach presented in [15], where a moment-based deterministic sampling scheme for the Bingham distribution on the unit hypersphere is proposed. In the circular case, the Bingham distribution can be seen as a rescaled von Mises distribution (see [38, Appendix B]), and the rescaled samples of the von Mises distribution proposed here can be shown to exactly match the samples of the Bingham distribution obtained by the method from [15]. Also, the samples are identical to those produced by the von Mises–Fisher sampling method [36] when it is applied to the circular case.

3.2. First Two Circular Moments

An approximation based on the first two trigonometric moments m_1 and m_2 is somewhat more involved. We consider a WD distribution with $L = 5$ components and

²A generalization of this method to nonequal weights is given in [27, Sec. 2.5.1-A].




Fig. 5. Feasible values for w_5 depending on a given concentration of the distribution. (a) WN distribution with parameter σ , (b) WC distribution with parameter γ , (c) VM distribution with parameter κ .

Dirac positions

$$\beta_1 = -\phi_1, \beta_2 = \phi_1, \beta_3 = -\phi_2, \beta_4 = \phi_2, \beta_5 = 0$$

that is symmetric around 0. As we will show later, moment matching does not allow a solution with an equally weighted Dirac mixture in general. Thus, we choose equal weights for the first four components

$$w_1 = w_2 = w_3 = w_4 = \frac{1 - w_5}{4}$$

and leave the weight w_5 of the component at $\beta_5 = 0$ to be determined. We will later derive constraints on the value of w_5 and see that $w_5 = \frac{1}{5}$, i.e., using equal weights for all components, does not always guarantee the existence of a solution.

For the first moment, we have

$$m_1^{WD} = 2 \frac{1 - w_5}{4} \cos(\phi_1) + 2 \frac{1 - w_5}{4} \cos(\phi_2) + w_5,$$

and obtain

$$\underbrace{\frac{2}{1 - w_5} (m_1 - w_5)}_{=: c_1} = \cos(\phi_1) + \cos(\phi_2). \quad (5)$$

Similarly, for the second moment, we have

$$m_2^{WD} = 2 \frac{1 - w_5}{4} \cos(2 \cdot \phi_1) + 2 \frac{1 - w_5}{4} \cos(2 \cdot \phi_2) + w_5.$$

At this point, we apply the trigonometric identity $\cos(2 \cdot x) = 2 \cos^2(x) - 1$. After a short calculation, we obtain

$$\underbrace{\frac{1}{1 - w_5} (m_2 - w_5) + 1}_{=: c_2} = \cos^2(\phi_1) + \cos^2(\phi_2). \quad (6)$$

By substituting $x_1 = \cos(\phi_1)$, $x_2 = \cos(\phi_2)$, we obtain a system of two equations

$$\begin{aligned} c_1 &= x_1 + x_2, \\ c_2 &= x_1^2 + x_2^2. \end{aligned}$$

We solve for x_1 and x_2 , which yields

$$x_1 = c_1 - x_2, \quad x_2 = \frac{2c_1 \pm \sqrt{4c_1^2 - 8(c_1^2 - c_2)}}{4}.$$

Obviously, there are two different solutions. Without loss of generality, we only consider the solution

$$x_1 = c_1 - x_2, \quad x_2 = \frac{2c_1 + \sqrt{4c_1^2 - 8(c_1^2 - c_2)}}{4} \quad (7)$$

because the other solution just swaps x_1 and x_2 , which is equivalent for our purposes. Finally, we obtain $\phi_1 = \arccos(x_1)$ and $\phi_2 = \arccos(x_2)$.

This leaves the question of choosing the weighting coefficient w_5 . The previous equations can only be evaluated if the conditions

$$-1 \leq x_i \leq 1, \quad i = 1, 2 \quad \text{and} \quad 4c_1^2 - 8(c_1^2 - c_2) \geq 0$$

hold. These conditions can be used to find a lower and an upper bound on w_5 . These bounds are

$$w_5^{\min} = \frac{4m_1^2 - 4m_1 - m_2 + 1}{4m_1 - m_2 - 3}, \quad (8)$$

$$w_5^{\max} = \frac{2m_1^2 - m_2 - 1}{4m_1 - m_2 - 3}. \quad (9)$$

In the following, we show that $w_5^{\min} \leq w_5^{\max}$ holds in all relevant cases.

LEMMA 2. (Existence of a Solution). *If $m_2 > -3 + 4m_1$, it holds that $w_5^{\min} \leq w_5^{\max}$, i.e., there exists a solution in $[w_5^{\min}, w_5^{\max}]$.*

PROOF Using $m_2 > -3 + 4m_1$, we know that the denominator in (8) and (9) is negative. Thus, we have

$$w_5^{\min} \leq w_5^{\max}$$

$$\Leftrightarrow 4m_1^2 - 4m_1 - m_2 + 1 \geq 2m_1^2 - m_2 - 1$$

$$\Leftrightarrow 2m_1^2 - 4m_1 + 2 \geq 0$$

$$\Leftrightarrow (m_1 - 1)^2 \geq 0,$$

which is always fulfilled.

It can be shown that the inequality $m_2 > -3 + 4m_1$ holds for WN, WC, and VM distributions [27, Lemma 18]. Consequently, for any $0 \leq \lambda \leq 1$,

$$w_5(\lambda) = w_5^{\min} + \lambda(w_5^{\max} - w_5^{\min}) \quad (10)$$




Fig. 6. WD approximations for WN distributions with different values for λ . In all cases, we use $\sigma = 1$. The periodic boundary is marked by a dashed line. Note that for $\lambda \approx 0$ and $\lambda \approx 1$ the mixture degenerates to three components.

is a feasible solution. Furthermore, weights $w_5(\lambda) < 0$ are invalid because negative weights violate Kolmogorov's axioms, i.e., the probability of any event has to be larger or equal to zero.³ The parameter λ has to be chosen accordingly. The range of admissible values is illustrated in Fig. 5. It is obvious from the figure that $w_5 = \frac{1}{5}$, i.e., equal weights for all WD components, is not necessarily contained in the region of feasible values. A good choice of the parameter λ is discussed below. The entire method is summarized in Algorithm 1 (see also [27, Sec. 2.5, Algorithm 3], [35, Algorithm 2]).

ALGORITHM 1 *Deterministic approximation with $L = 5$ components.*

Input: First circular moment m_1 , second circular moment m_2 ,

parameter $\lambda \in [0, 1]$ with default $\lambda = 0.5$

Output: $\mathcal{WD}(x; w_1, \dots, w_5, \beta_1, \dots, \beta_5)$

/* extract μ */

$\mu \leftarrow \text{atan2}(\text{Im}m_1, \text{Re}m_1)$;

$m_1 \leftarrow |m_1|$;

$m_2 \leftarrow |m_2|$;

/* compute weights */

$w_5^{\min} \leftarrow (4m_1^2 - 4m_1 - m_2 + 1)/(4m_1 - m_2 - 3)$;

$w_5^{\max} \leftarrow (2m_1^2 - m_2 - 1)/(4m_1 - m_2 - 3)$;

$w_5 \leftarrow w_5^{\min} + \lambda(w_5^{\max} - w_5^{\min})$;

$w_1, w_2, w_3, w_4 \leftarrow (1 - w_5)/4$;

/* obtain Dirac positions */

$c_1 \leftarrow \frac{2}{1 - w_5}(m_1 - w_5)$;

$c_2 \leftarrow \frac{1}{1 - w_5}(m_2 - w_5) + 1$;

$x_2 \leftarrow (2c_1 + \sqrt{4c_1^2 - 8(c_1^2 - c_2)})/4$;

$x_1 \leftarrow c_1 - x_2$;

$\phi_1 \leftarrow \arccos(x_1)$;

$\phi_2 \leftarrow \arccos(x_2)$;

/* shift Dirac positions by μ */

$(\beta_1, \dots, \beta_5) \leftarrow \mu + (-\phi_1, +\phi_1, -\phi_2, +\phi_2, 0) \bmod 2\pi$;

return $\mathcal{WD}(x; w_1, \dots, w_5, \beta_1, \dots, \beta_5)$;

³In practice, other filters such as the UKF [24] and the randomized UKF [55] are sometimes used with negative weights, which can give decent results, but does not have a proper probabilistic interpretation.

3.3. Properties of the Moment-based Approximation

There are several noteworthy properties of the presented approximation method. Obviously, it maintains the first and second trigonometric moment of the original density. Maintaining the first trigonometric moment guarantees that the conversion is reversible. If we take a WN, WC, or VM distribution and approximate it with a WD distribution, we can recover the original distribution by means of moment matching. In the case of a VM distribution, we can also obtain the original distribution by maximum likelihood estimation, which coincides with the result from moment matching [21, Remark 4.1].

Approximating not just the first, but also the second moment has the advantage of more accurately approximating the original distribution and producing a mixture with more components. As shown before, different types of distributions differ in their second moment, even if they are uniquely determined by their first moment (see Fig. 2). If we use a wrapped Dirac mixture to propagate a density through a nonlinear function, a larger number of mixture components captures the effect of the function more accurately.

One of the main advantages of the presented method is the fact that for WN and WC distributions, all required operations can be evaluated in closed form. The necessary formulas (1), (2), and (5)–(10) can be evaluated in constant time and are easily implemented even on embedded hardware with limited computational capabilities. In the case of a VM distribution, the calculation of the first and second trigonometric moment requires the evaluation of Bessel functions as given in (3), but all other steps (5)–(10) are still possible in closed form.

Examples for the approximation of both WN and VM densities with different concentrations are depicted in Fig. 3 and Fig. 4. These examples illustrate how the concentration of the density affects the placement and weighting of the samples. It can be seen that the weight of the middle samples increases as the density gets more concentrated. These figures also illustrate that the shape of the two densities is slightly different.

The influence of the parameter λ is illustrated in Fig. 6. If λ approaches 1, more and more weight is assigned to the Dirac component at zero whereas the other Dirac components have less influence. If, on




Fig. 7. Locations of the Dirac delta components β_1, \dots, β_5 depending on λ for a WN Distribution with parameters $\mu = \pi$ and different values of σ . Note that in the case of $\sigma = 2.5$, small values of λ are not valid because they yield a negative w_5 .

the other hand, λ approaches zero, two of the other components move towards the center Dirac component, effectively reducing the number of Dirac components to three. As both of these effects are undesirable, λ should not be chosen too close to either zero or one. It can be shown that for WN, WC as well as VM distribution, w_5 is positive for arbitrary concentrations if and only if $\lambda \geq 0.5$, which motivates a choice of $\lambda \in [0.5, 1]$.

LEMMA 3. (Condition for Positive Weights). For WN, VM, and WC distributions, w_5 is positive for arbitrary concentrations if and only if $\lambda \geq 0.5$.

PROOF We calculate

$$\begin{aligned}
 w_5(\lambda) &= w_5^{\min} + \lambda(w_5^{\max} - w_5^{\min}) \\
 &= \frac{4m_1^2 - 4m_1 - m_2 + 1}{4m_1 - m_2 - 3} \\
 &\quad + \lambda \left(\frac{2m_1^2 - m_2 - 1}{4m_1 - m_2 - 3} - \frac{4m_1^2 - 4m_1 - m_2 + 1}{4m_1 - m_2 - 3} \right) \\
 &= \frac{4m_1^2 - 4m_1 - m_2 + 1}{4m_1 - m_2 - 3} + \lambda \left(\frac{-2m_1^2 + 4m_1 - 2}{4m_1 - m_2 - 3} \right) \\
 &= \frac{4m_1^2 - 4m_1 - m_2 + 1 + \lambda(-2m_1^2 + 4m_1 - 2)}{4m_1 - m_2 - 3} \\
 &= \frac{(4 - 2\lambda)m_1^2 + (-4 + 4\lambda)m_1 - m_2 + 1 - 2\lambda}{4m_1 - m_2 - 3}.
 \end{aligned}$$

a) *WN*: From (1), we obtain the relation $m_2 = m_1^4$ and substitute accordingly.

$$\begin{aligned}
 w_5(\lambda) &= \frac{(4 - 2\lambda)m_1^2 + (-4 + 4\lambda)m_1 - m_1^4 + 1 - 2\lambda}{4m_1 - m_1^4 - 3} \\
 &= \frac{m_1^2 + 2\lambda + 2m_1 - 1}{m_1^2 + 2m_1 + 3}
 \end{aligned}$$

Because $m_1^2 + 2m_1 + 3 > 0$, we have

$$\begin{aligned}
 w_5(\lambda) &\geq 0 \\
 \Leftrightarrow m_1^2 + 2\lambda + 2m_1 - 1 &\geq 0 \\
 \Leftrightarrow \lambda &\geq \frac{1}{2} - 2m_1 - m_1^2 \xrightarrow{m_1 \rightarrow 0} \frac{1}{2}
 \end{aligned}$$

and $m_1 \in (0, 1)$ shows the claim.

b) *VM*: The property holds for VM distributions as well, but the proof is more tedious because of the involved Bessel functions. For this reason, we do not give a formal proof here.

c) *WC*: From (2), we obtain the relation $m_2 = m_1^2$

$$\begin{aligned}
 w_5(\lambda) &= \frac{(3 - 2\lambda)m_1^2 + (-4 + 4\lambda)m_1 + 1 - 2\lambda}{4m_1 - m_1^2 - 3} \\
 &= \frac{2\lambda m_1 - 2\lambda - 3m_1 + 1}{m_1 - 3}.
 \end{aligned}$$

Because $m_1 - 3 < 0$, we have

$$\begin{aligned}
 w_5(\lambda) &\geq 0 \\
 \Leftrightarrow 2\lambda m_1 - 2\lambda - 3m_1 + 1 &\leq 0 \\
 \Leftrightarrow \lambda(2m_1 - 2) &\leq -1 + 3m_1 \\
 \Leftrightarrow \lambda &\geq \frac{1}{2} \cdot \frac{1 - 3m_1}{1 - m_1} \xrightarrow{m_1 \rightarrow 0} \frac{1}{2}
 \end{aligned}$$

and $m_1 \in (0, 1)$ shows the claim.

Based on these results, we suggested to use $\lambda = 0.5$ in [32]. The value $\lambda = 0.5$ is a reasonable choice for both high and low concentrations and it guarantees $w_5 \geq 0$ in all cases.

However, the choice of $\lambda = 0.5$ has the disadvantage that the discrete approximation degenerates to four components as the distribution approaches a circular uniform distribution because the weight of the sample at the circular mean approaches zero. It would be preferable to obtain five equally weighted components in this case. For this purpose, we consider the limits

$$\lim_{m_1, m_2 \rightarrow 0} w_5^{\min} = -\frac{1}{3}$$

and

$$\lim_{m_1, m_2 \rightarrow 0} w_5^{\max} = \frac{1}{3}$$

of the minimum and maximum weights as the distribution approaches a circular uniform distribution (i.e., the moments approach zero). Then, we can solve for λ according to

$$\begin{aligned}
 w_5(\lambda) &= -\frac{1}{3} + \lambda \frac{2}{3} \stackrel{!}{=} \frac{1}{3} \\
 \Rightarrow \lambda &= 0.8.
 \end{aligned}$$

Thus, $\lambda = 0.8$ ensures convergence to a distribution with equally weighted components. According to Lemma 3, this choice also guarantees positive weights in all cases.

Even though we only presented an approximation for WN, WC, and VM distributions so far, the presented approach can easily be generalized to other symmetric circular probability distribution whose first and second trigonometric moments can be calculated and satisfy the condition given in Lemma 2.

4. SUPERPOSITION METHOD

In this section, we present an extension of the proposed method. For problems with strong nonlinearities, it may be desirable to obtain a deterministic approximation with a larger number of samples. However, generalizing the proposed method to more samples and/or higher trigonometric moments is not straightforward. For this reason, we propose the use of a superposition approach that combines several sample sets with different scaling parameters into one larger sample set, which also retains the first and second trigonometric moment. This approach is similar to the randomized UKF [55], where multiple UKF sample sets with different rotations and scalings are combined.

4.1. Proposed Method

To be specific, we consider the superposition of $q \in \mathbb{N}, q \geq 2$ approximations obtained using Algorithm 1 with different parameters $\lambda_1 < \dots < \lambda_q \in (0, 1)$. The resulting density is given by

$$\sum_{k=1}^q p_q \sum_{j=1}^L w_j^{(\lambda_k)} \delta(x - \beta_j^{(\lambda_k)}) \quad (11)$$

with weights $p_1, \dots, p_q > 0$ and $\sum_{k=1}^q p_q = 1$, where $w_1^{(\lambda_k)}, \dots, w_L^{(\lambda_k)}$ and $\beta_1^{(\lambda_k)}, \dots, \beta_L^{(\lambda_k)}$ are obtained using Algorithm 1 by choosing the parameter λ_k .

LEMMA 4. *For $n = 0, 1, 2$, the n th trigonometric moment of (11) is identical to the n th trigonometric moment of the original density.*

PROOF We use the linearity of integration and the fact that Algorithm 1 maintains the first and second moment as well as the normalization property, which allows us to obtain

$$\begin{aligned} & \int_0^{2\pi} \exp(inx) \sum_{k=1}^q p_q \sum_{j=1}^L w_j^{(\lambda_k)} \delta(x - \beta_j^{(\lambda_k)}) dx \\ &= \sum_{k=1}^q p_q \int_0^{2\pi} \exp(inx) \sum_{j=1}^L w_j^{(\lambda_k)} \delta(x - \beta_j^{(\lambda_k)}) dx \\ &= \sum_{k=1}^q p_q m_n \\ &= m_n. \end{aligned}$$

Because one sample is placed at the circular mean of the distribution regardless of λ , this method results in multiple samples at this location. In a practical implementation, these samples can be joined into a single sample with weight $\sum_{k=1}^q p_k w_5^{(\lambda_k)}$. If no other samples coincide, the approximation results in $4q + 1$ Dirac delta components. It should be noted that the weights $w_5^{(\lambda_k)}$ of individual samples at the center may be negative as long as the sum of the weights of all Dirac delta components at this location remains positive. This leads to the condition

$$\begin{aligned} & \sum_{k=1}^q p_k w_5^{(\lambda_k)} \geq 0 \\ \Leftrightarrow & \sum_{k=1}^q p_k (w_5^{\min} + \lambda_k (w_5^{\max} - w_5^{\min})) > 0 \\ \Leftrightarrow & w_5^{\min} + (w_5^{\max} - w_5^{\min}) \sum_{k=1}^q p_k \lambda_k \geq 0 \\ \Leftrightarrow & \sum_{k=1}^q p_k \lambda_k \geq \frac{w_5^{\min}}{w_5^{\min} - w_5^{\max}}. \quad (12) \end{aligned}$$

We will later make sure to that this condition always holds.

4.2. Choice of Parameters

While the approach discussed above is fairly straightforward, it is not obvious how to choose the parameters p_1, \dots, p_q and $\lambda_1, \dots, \lambda_q$. In principle, any choice fulfilling condition (12) would be valid, and the first and second trigonometric moments are maintained as shown in Lemma 4. However, for good performance in a circular filter, it is desirable that a large range is covered and that the Dirac delta components are as evenly spread as possible. It is possible to achieve this by defining a suitable cost function and using a numerical optimization procedure to obtain the parameters. Because this would incur a significant runtime cost, we instead suggest an intelligent heuristic that ensures good results in all cases and that can be calculated very efficiently.

For the heuristic, we choose uniform weights $p_k = 1/q$. When we consider the choice of the parameters $\lambda_1, \dots, \lambda_q$, we observe the influence of λ_k on the Dirac delta positions $\beta_1^{(\lambda_k)}, \dots, \beta_4^{(\lambda_k)}$ (see Fig. 7). It can be seen in the depicted examples that the outermost components at $\beta_1^{(\lambda_k)}$ and $\beta_2^{(\lambda_k)}$ move further away from the circular mean at first when λ_k increases, but start moving inward again as λ_k approaches 1. This behavior can formally be shown by the following Lemma.

LEMMA 5. *The angular distance between β_1 and the circular mean (as well as β_2 and the circular mean) is maximized for $\lambda^{\max} = 2\sqrt{2} - 2 \approx 0.828$.*

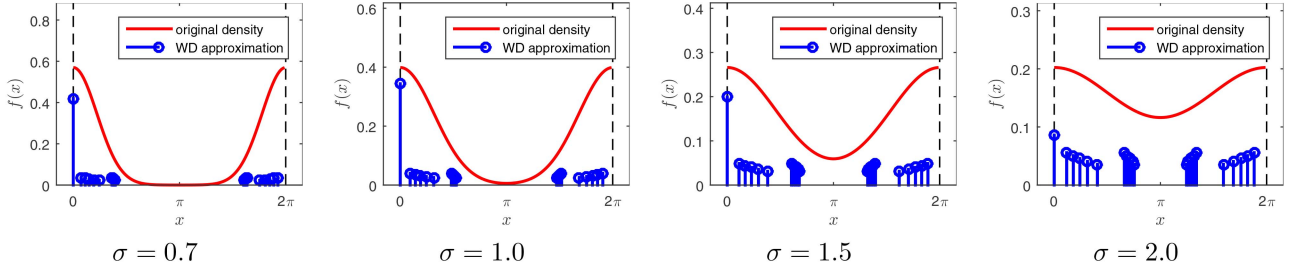


Fig. 8. WD approximations for WN distributions using superposition method with different values for σ . In all cases, we use $q = 5$.

PROOF We consider the partial derivative of x_1 with respect to λ , which yields

$$\begin{aligned}
\frac{\partial}{\partial \lambda} x_1 &= \frac{\partial}{\partial \lambda} (c_1 - x_2) \\
&= \frac{\partial}{\partial \lambda} \left(c_1 - \left(2c_1 + \sqrt{4c_1^2 - 8(c_1^2 - c_2)} \right) / 4 \right) \\
&= -\frac{\partial}{\partial \lambda} \left(c_1 + \sqrt{c_1^2 - 2(c_1^2 - c_2)} / 2 \right) \\
&= -\frac{\partial}{\partial \lambda} \left(c_1 + \sqrt{-c_1^2 + 2c_2} / 2 \right) \\
&= -\frac{\partial}{\partial \lambda} \left(\frac{2(m_1 - w_5)}{1 - w_5} \right. \\
&\quad \left. + \sqrt{-\left(\frac{2(m_1 - w_5)}{1 - w_5} \right)^2 + 2\left(\frac{m_2 - w_5}{1 - w_5} + 1 \right)} / 2 \right).
\end{aligned}$$

Furthermore, we obtain

$$\begin{aligned}
w_5 &= w_5^{\min} + \lambda(w_5^{\max} - w_5^{\min}) \\
&= \frac{4m_1^2 - 4m_1 - m_2 + 1 + \lambda(-2m_1^2 - 2 + 4m_1)}{4m_1 - m_2 - 3}.
\end{aligned}$$

Substituting this term for w_5 in the partial derivative x_1 , setting the result to zero, and solving for λ yields $2\sqrt{2} - 2$.

Note that this result does not depend on the moments m_1 and m_2 . To cover the maximum area, it is desirable to choose $\lambda_q = \lambda^{\max}$ and $\lambda_1, \dots, \lambda_{q-1} < \lambda^{\max}$. Because the Dirac delta components should be evenly distributed, we propose to define

$$\lambda_k = \lambda^{\min} + \frac{k}{q}(\lambda^{\max} - \lambda^{\min}), \quad k = 1, \dots, q,$$

i.e., the λ s are spread evenly between λ^{\min} and λ^{\max} , where λ^{\min} is yet to be determined.⁴ Obviously, it holds that $\lambda^{\min} \geq 0$, but for $\lambda^{\min} = 0$, condition (12) is not always guaranteed to hold. Thus, we determine λ^{\min} from the positivity condition (12).

⁴According to this definition, λ^{\max} is included in $\lambda_1, \dots, \lambda_k$, but λ^{\min} is not. This choice is motivated by the fact that the value of λ^{\min} that we derive in the following leads to a degenerate solution.

LEMMA 6. Condition (12) is satisfied if and only if

$$\lambda^{\min} \geq 2 \cdot \frac{q}{q-1} \cdot \frac{w_5^{\min}}{w_5^{\min} - w_5^{\max}} - \lambda^{\max} \cdot \frac{q+1}{q-1}.$$

PROOF First, we calculate a closed form solution for the sum of all λ s according to

$$\begin{aligned}
\sum_{k=1}^q \lambda_k &= \sum_{k=1}^q \lambda^{\min} + \frac{k}{q}(\lambda^{\max} - \lambda^{\min}) \\
&= q \cdot \lambda^{\min} + \frac{\lambda^{\max} - \lambda^{\min}}{q} \sum_{k=1}^q k \\
&= q \cdot \lambda^{\min} + \frac{\lambda^{\max} - \lambda^{\min}}{q} \cdot q \cdot (q+1)/2 \\
&= q \cdot \lambda^{\min} + (\lambda^{\max} - \lambda^{\min}) \cdot (q+1)/2 \\
&= \lambda^{\min} \cdot (q-1)/2 + \lambda^{\max} \cdot (q+1)/2.
\end{aligned}$$

Now, we use $p_1, \dots, p_q = 1/q$ and solve equation (12) for λ^{\min} , which leads to

$$\begin{aligned}
\lambda^{\min} \cdot (q-1)/2 + \lambda^{\max} \cdot (q+1)/2 &\geq \frac{q \cdot w_5^{\min}}{w_5^{\min} - w_5^{\max}} \\
\Leftrightarrow \lambda^{\min} \cdot (q-1)/2 &\geq \frac{q \cdot w_5^{\min}}{w_5^{\min} - w_5^{\max}} - \lambda^{\max} \cdot (q+1)/2 \\
\Leftrightarrow \lambda^{\min} &\geq 2 \cdot \frac{q}{q-1} \cdot \frac{w_5^{\min}}{w_5^{\min} - w_5^{\max}} - \lambda^{\max} \cdot \frac{q+1}{q-1}.
\end{aligned}$$

Based on this result, we choose

$$\lambda^{\min} = \max \left(0, 2 \cdot \frac{q}{q-1} \cdot \frac{w_5^{\min}}{w_5^{\min} - w_5^{\max}} - \lambda^{\max} \cdot \frac{q+1}{q-1} \right).$$

Examples of the proposed superposition approximation are depicted in Fig. 8. It can be seen that even though the λ s were chosen to cover the largest possible area and designed to be distributed as uniformly as possible, the Dirac delta components form four clusters within the possible ranges of β_1, \dots, β_4 . Also, there is a single central component located at β_5 , i.e., the circular mean, and that has a fairly large weight when the uncertainty is low. As a result, the superposition approximation is not very evenly spread and does not approximate the shape of the true density all that well. However, it is

very easy and fast to compute and is always guaranteed to maintain the first two trigonometric moments.

5. BINARY TREE APPROACH

In this section, we will introduce an approach for deterministic sampling of circular densities that is not based on trigonometric moment matching, but rather on approximation of the entire probability density function (as it is done in, e.g., [17]). The basic idea of the proposed approach is to construct a binary tree by recursively dividing the interval $[0, 2\pi)$ into smaller intervals while choosing the number of Dirac delta components in each interval proportionally to the contained probability mass.

5.0.1. Proposed Method:

Formally, the proposed algorithm (see Algorithm 2) works as follows. It considers the pdf $f(\cdot)$ on the interval $[l, r)$ that is to be approximated with a predefined number of $L \in \mathbb{N}^+$ samples. If $L = 1$, the algorithm simply returns the center of the interval.⁵ Otherwise, for $L > 1$, the interval is divided into two halves $[l, m)$ and $[m, r)$, where $m = (l + r)/2$ is the center of the interval. Thereby, two leaves are added to the binary tree, which can be processed recursively. The probability mass

$$p_{\text{left}} = \int_l^m f(x)dx, \quad p_{\text{right}} = \int_m^r f(x)dx$$

in each half is computed and the number of Dirac components in each half is chosen proportionally. Because the number of components in each half has to be an integer, we need to round the result. However, because of rounding errors, we may have corner cases where the number of Dirac components in both halves does not sum to L . For this reason, we round down first and then check if there is a remaining component. If that is the case, we assign it to the half where the rounding error is larger. In order to evaluate the probability mass within an interval, the cumulative distribution function (cdf) of the density $f(\cdot)$ is required. In the case of the WN, VM, and WC distributions, efficient methods for computing this function are available. For other distributions it may be necessary to fall back to numerical integration, which would typically be slower.

REMARK 1 (Cumulative Distribution Functions).

- 1) WN distribution: The cdf can be obtained by componentwise integration, which reduces the problem to the calculation of an infinite sum of Gaussian cdfs. The Gaussian cdfs are not available analytically but can be efficiently evaluated using the erf-function [1, Sec. 7.1]. Similar to the infinite sum in the pdf,

the infinite sum in the cdf can be truncated to just a few terms.

- 2) VM distribution: No analytic solution for the cdf of a VM distribution is known. However, Hill proposed an efficient approximation [18] with an accuracy of 12 decimal digits.
- 3) WC distribution: The cdf of a WC distribution can be computed analytically using the indefinite integral

$$\begin{aligned} & \int \frac{\sinh(\gamma)}{2\pi(\cosh(\gamma) - \cos(x - \mu))} dx \\ &= \frac{1}{\pi} \arctan(\coth(\gamma/2) \tan((x - \mu)/2)) \\ &+ \text{constant.} \end{aligned}$$

ALGORITHM 2 *approximateBT*

Input: left limit l , right limit r , number of samples L , pdf $f(\cdot)$
Output: sample positions $\underline{\beta}$
 $m = (l + r)/2;$
 /* one Dirac component remaining */
if $L = 1$ **then**
 | $\underline{\beta} = [m];$
 | **return** $\underline{\beta};$
end
 /* calculate integrals and distribute Diracs components proportionally */
 $p_{\text{left}} = \int_l^m f(x)dx;$
 $p_{\text{right}} = \int_m^r f(x)dx;$
 $L_{\text{left}} = \lfloor L \cdot p_{\text{left}} / (p_{\text{left}} + p_{\text{right}}) \rfloor;$
 $L_{\text{right}} = \lfloor L \cdot p_{\text{right}} / (p_{\text{left}} + p_{\text{right}}) \rfloor;$
 /* assign remaining Dirac delta component if necessary */
if $L_{\text{left}} + L_{\text{right}} = L - 1$ **then**
 | **if** $L \cdot p_{\text{left}} / (p_{\text{left}} + p_{\text{right}}) - L_{\text{left}} >$
 | $L \cdot p_{\text{right}} / (p_{\text{left}} + p_{\text{right}}) - L_{\text{right}}$ **then**
 | | $L_{\text{left}} = L_{\text{left}} + 1;$
 | **else**
 | | $L_{\text{right}} = L_{\text{right}} + 1;$
 | **end**
end
 /* Recursive calls for left and right half */
 $\underline{\beta} \leftarrow [];$
if $L_{\text{left}} \geq 1$ **then**
 | $\underline{\beta} \leftarrow [\underline{\beta}, \text{approximateBT}(l, m, L_{\text{left}})];$
end
if $L_{\text{right}} \geq 1$ **then**
 | $\underline{\beta} \leftarrow [\underline{\beta}, \text{approximateBT}(m, r, L_{\text{right}})];$
end
return $\underline{\beta};$

⁵A better alternative is to put the sample at the center of mass of the interval, i.e., at $\int_l^r x \cdot f(x)dx$. However, this typically necessitates numerical integration, and would thus significantly increase the computation time. If the integral can be computed analytically, this choice should be preferred, though.

5.1. Extensions

Although the binary tree approximation as presented above works quite well in most circumstances, some extensions are possible to improve it further.




Fig. 9. WD approximations for WN distributions using binary tree method with different values for σ (without moment correction). In all cases, we use $L = 25$.

5.1.1. Shifting Invariance:

In principle, the proposed algorithm can be applied to circular densities with arbitrary circular mean directly and it is not necessary to enforce $\mu = 0$. However, the binary tree approximation is not invariant under shifting, i.e., the results of the approximation depend on the location of the density. This is due to the choice of the initial interval $[l, r] = [0, 2\pi)$. In order to avoid introducing artifacts as a result of this issue, we recommend enforcing $\mu = 0$ before performing the binary tree approximation and shifting the samples by μ afterwards. This modification makes the proposed method invariant to shifting operations. Also, it is worth mentioning that for symmetric densities, this ensures symmetry around 0, which can be exploited to cut the computational effort in half if L is even. This is done by setting the initial interval as $[l, r] = [0, \pi]$, approximating with $L/2$ samples, and obtaining the samples in $[\pi, 2\pi]$ by mirroring.

5.1.2. Moment Correction:

Because of its shape-based approach, the binary tree approximation does not, in general, guarantee that any trigonometric moments are maintained. This has the disadvantage that the original continuous density cannot be exactly recovered from the samples using a moment-based estimator, not even for simple distributions such as the WN, VM, and WC distributions. Thus, propagation through an identity function or a simple shift operation is also not exact. This problem can be avoided by introducing a post-processing step to retroactively correct the first trigonometric moment. The downside of such a correction step is that the probability mass in each interval is not exactly retained anymore, i.e., we introduce a small error into the approximation of the shape in order to get an approximation that has the correct first trigonometric moment.

Correction of the first trigonometric moment is carried out in two steps. First, we correct the circular mean, i.e., we adjust the complex argument $\text{Arg}m_1$ to match the original distribution. This is achieved by calculating the circular mean of the original density and the circular mean of the binary tree approximation and then shifting all Dirac delta components by the difference. As long as the first trigonometric moment of the true density is

easy to calculate, this step does not require much computational effort, and thus, is always recommended.⁶

In order to correct the Euclidean norm $|\cdot|$ of the complex first trigonometric moment m_1 , i.e., the spread of the distribution, we propose the following approach. The basic idea is to perform a scaling operation that moves all Dirac delta component inwards or outwards with respect to the circular mean by scaling around the mean by a factor of $c > 0$, i.e., we obtain the wrapped Dirac mixture

$$\sum_{j=1}^L w_j \delta(x - (\text{Arg}m_1 + c \cdot \Phi(\beta_j - \text{Arg}m_1))),$$

where $\Phi(x) = (x + \pi \bmod 2\pi) - \pi$ is a function that switches to a parameterization on $[-\pi, \pi)$. The scaling parameter c is then obtained by solving the optimization problem

$$\arg \min_c \left| |m_1| - \sum_{j=1}^L \cos(c \cdot \Phi(\beta_j - \text{Arg}m_1)) \right|,$$

where we use an initial value of $c = 1$ (i.e., no scaling is performed). This one-dimensional optimization can be efficiently solved using a quasi-Newton algorithm [45, Sec. 6]. Typically, we obtain a result $c \approx 1$, i.e., only a slight scaling is necessary.

5.2. Properties of the Binary Tree Approximation

The runtime of the proposed method depends on the shape of the probability density that is to be approximated. In practical experiments, it is not quite as fast as the closed-form solution of the moment-based approximation discussed above, but it is still very fast compared to approaches that necessitate the solution of complicated multivariate nonlinear optimization problems (e.g., [17]). In particular, if the cdf can be computed efficiently as is the case for WN, VM, and WC distributions, the proposed method is very cheap to perform (see Sec. 6).

Unlike the moment-based approaches discussed above, we do not restrict ourselves to symmetric densities here. Thus, the binary tree method is also appli-

⁶For symmetric densities, the shifting invariance modification (Sec. 5.1.1) already guarantees that the circular mean is matched exactly. In this case, the correction step can be omitted.

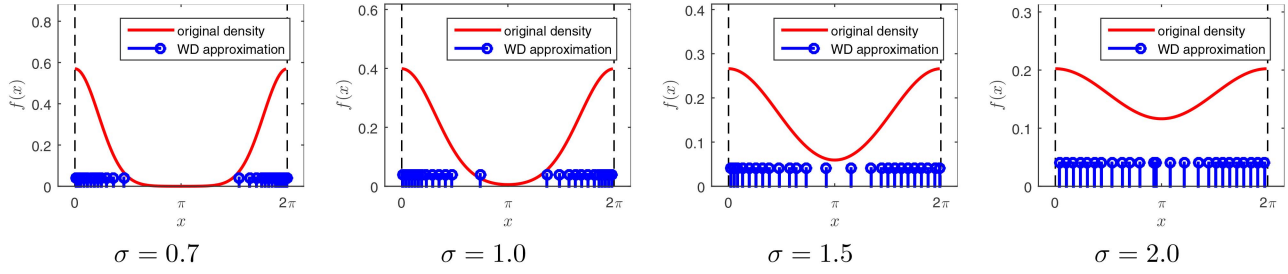


Fig. 10. WD approximations for WN distributions using binary tree method with different values for σ (with moment correction). In all cases, we use $L = 25$.

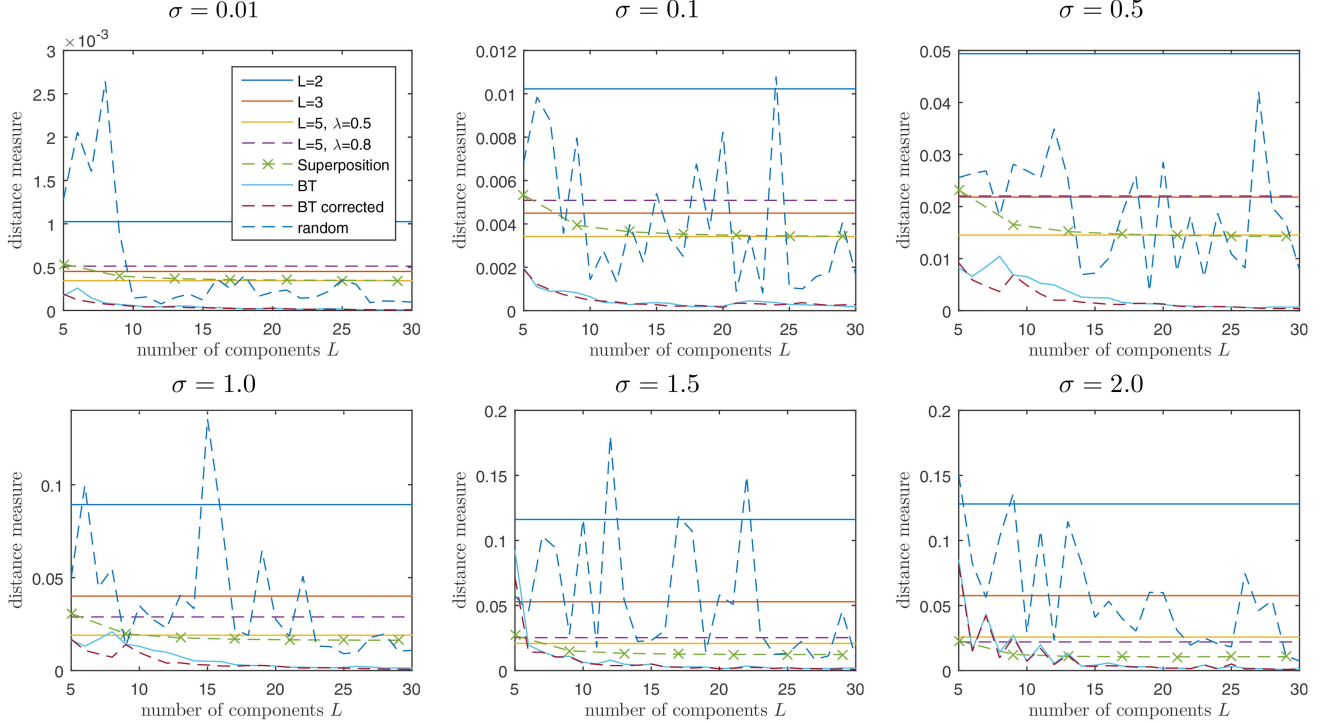


Fig. 11. Distance measure between true density $\mathcal{WN}(x; \pi, \sigma)$ and the different deterministic approximations.

cable to distributions with asymmetric densities such as the wrapped exponential distribution and the wrapped Laplace distribution [20]. More generally, it can also be applied to multi-modal circular densities, e.g., wrapped normal mixtures [2], von Mises mixtures [8], [44], [9], or Fourier densities [59], [10], [46], [47].

Furthermore, this method produces a uniformly weighted mixture. This is a desirable property because it has advantages when reweighting the samples as particle degeneration does not occur as quickly. Reweighting is a common technique used to derive nonlinear measurement updates, e.g., as part of a progressive filter [34], [35].

It should also be noted that the proposed approach can be generalized to higher dimensions, i.e., to probability distributions on the torus [30], [40], [51] and hypertorus [27], [41]. In that case, binary trees would be replaced by quadtrees, octrees, etc. Even though this generalization is straightforward, the higher-dimensional version is limited to a small number of dimensions, because it scales exponentially with respect

to the considered number of dimensions. Also, an efficient implementation of the multidimensional integral of the true density is required.

Examples of the binary tree approximation without moment correction and with moment correction are given in Fig. 9 and Fig. 10, respectively. It can be seen that the binary tree approximation represents the shape of the true distribution very well. The moment correction only changes the distribution very slightly, but exactly enforcing the moment constraint has significant advantages because it allows recovering the original density for distributions that are uniquely defined by their first trigonometric moment.

6. EVALUATION

In this section, we provide evaluations of the proposed methods according to different criteria. We consider a distance measure comparing the approximation to the true distribution and evaluate the accuracy when the approximations are applied for propagation of the density through a nonlinear function.

6.1. Approximation Accuracy Based on Distance Measure

To determine the approximation accuracy, we define a distance measure between the original continuous density and the discrete approximation. For two circular densities $f_1(\cdot)$ and $f_2(\cdot)$ with cumulative distributions⁷

$$F_1(x) = \int_0^x f_1(t)dt \quad \text{and} \quad F_2(x) = \int_0^x f_2(t)dt,$$

we define the distance measure according to

$$\int_0^{2\pi} (F_1(x) - F_2(x))^2 dx.$$

This measure is essentially a circular version of the distance used in [19], [49].

During the evaluation, we consider the density $\mathcal{WN}(x; \pi, \sigma)$ with several values of σ and approximated it using the methods discussed in this paper. The results are depicted in Fig. 11. It can be seen that the moment-based approximation with $L = 2$ components is always the worst method. The moment-based approximation with $L = 3$ components is usually not very accurate either. The moment-based approximation with $L = 5$ components is always quite acceptable for $\lambda = 0.5$. Setting $\lambda = 0.8$ decreases the accuracy of the approximation for small uncertainties, but improves the approximation quality for larger uncertainties. The superposition-based approximation seems similar to the $L = 5$, $\lambda = 0.5$ approximation for small noise, but outperforms it for larger noise as long as enough components are used. Both of the binary tree methods perform very well, especially if many components (say, $L \geq 15$) are used. The corrected version (see Sec. 5.1.2) is slightly better in some cases, but very similar to the uncorrected version in others. We also compare to random sampling, which can be seen to be very unreliable with such a small number of samples.

Moreover, we investigated the runtime performance of the proposed methods. All measurements were obtained on a laptop with an Intel Core i7-2640M @2.8GHz, 8GB RAM, and MATLAB 2015a. The results are given in Table I. As can be seen, all moment-based methods are very fast, including the superposition method. The binary tree method is more costly, particularly for the WN distribution whose cdf is the most expensive to calculate, at least in our current implementations. The retroactive moment correction procedure (see Sec. 5.1.2) somewhat increases the computation time, ir-

⁷It should be noted that the definition of the cumulative distribution on the circle is somewhat problematic because the starting point of the integration can be chosen anywhere in $[0, 2\pi)$, i.e., it is not invariant to shifting. The choice of the starting point also affects the distance measure used in this section. In particular if Dirac delta components are very close to the starting point, small changes to the starting point can cause large changes in the cdf, and consequently the distance. To minimize this effect, we use zero as a starting point and a density centered around π , such that the probability mass is concentrated far away from the starting point.




Fig. 12. Runtime of the binary tree approximation.

TABLE I
Average runtime for a single approximation.

method	time (s)		
	WN	VM	WC
$L = 2$	0.000166	0.000173	0.000150
$L = 3$	0.000166	0.000170	0.000152
$L = 5$	0.000246	0.000312	0.000223
Superposition ($q = 5$)	0.000334	0.000399	0.000315
BT ($L = 25$)	0.065933	0.004461	0.011233
BT corrected ($L = 25$)	0.081339	0.028380	0.033159
random ($L = 25$)	0.000269	0.000479	0.000254

respective of the underlying true distribution. Still, it can be seen that all proposed approximation algorithms are already suitable for many real-time applications even though our MATLAB implementation is not well optimized. Particularly, the moment-based approaches are very efficient and could even be used in situations where only little computational power is available.

In order to more closely investigate the runtime of the binary tree approximation, we provide a plot where we show the runtime as a function of the number of samples L in Fig. 12. It can be seen that the runtime scales approximately linearly in L , i.e., it can be applied efficiently even for many samples.⁸

6.2. Propagation Through a Nonlinear Function

Furthermore, we evaluate the proposed deterministic approximation methods by determining the error when propagating through a nonlinear function. For our evaluation, we consider the function $g_c : [0, 2\pi) \rightarrow [0, 2\pi)$ with

$$g_c(x) = \pi \cdot \left(\sin \left(\frac{\text{sign}(x - \pi) |x - \pi|^c}{2} \right) + 1 \right)$$

for some constant $c > 0$, which has previously been used for evaluation purposes in [14]. It can be shown that this

⁸It seems intuitive that the runtime is linear in most cases, but for pathological probability density functions it may not be possible to provide that guarantee.




Fig. 13. Propagation of a prior WN distribution with parameters $\mu = \pi$, $\sigma = 0.4$ through a nonlinear function $g_c(\cdot)$ by means of the proposed deterministic WD mixture approximation with five components. In this example, we use $c = 0.9$.




Fig. 14. The function $g_c(\cdot)$ for different values of c .

function is continuous for all $c > 0$. We have

$$g'_c(x) = \pi \cos\left(\frac{|x - \pi|^c}{2 \cdot \pi^{c-1}}\right) \cdot \frac{c \cdot |x - \pi|^{c-1}}{2 \cdot \pi^{c-1}}$$

which is positive for $x \in [0, 2\pi)$, i.e., the function $g_c(\cdot)$ is strictly increasing and furthermore bijective.⁹ Varying the value of c allows us to control how strong the nonlinearity is (see Fig. 14). The inverse function of $g_c(\cdot)$ can also be calculated analytically according to

$$g_c^{-1}(x) = \text{sign}\left(\arcsin\left(\frac{x}{\pi} - 1\right)\right) \cdot \left|2 \arcsin\left(\frac{x}{\pi} - 1\right) \pi^{c-1}\right|^{1/c} + \pi.$$

Now, we assume a random variable A is distributed according to a WN probability distribution $\mathcal{WN}(x; \mu_A, \sigma_A)$. We propagate A through the nonlinear function $g_c(\cdot)$ to get $B = g_c(A)$, and seek to obtain a WN distribution $f(x; \mu_B, \sigma_B)$ that approximates the distribution of $g_c(A)$. For $x \neq \pi$, the density of the exact distribution is given by

$$f^B(x) = \frac{f(g_c^{-1}(x); \mu, \sigma)}{g'_c(x)}.$$

⁹The proposed approach is not limited to injective or continuous functions. However, we use such a function because these properties simplify the calculation of the true posterior density.

This distribution is not a WN distribution, but can be approximated by one. To approximate the true distribution with a WN distribution, we proceed as follows. First, we deterministically approximate the prior distribution $\mathcal{WN}(x; \mu_A, \sigma_A)$ with a WD mixture $\sum_{j=1}^L w_j \cdot \delta(x - \beta_j)$ using one of the methods presented in this paper. Then, we propagate all samples through the nonlinear function $g_c(\cdot)$, which yields $\sum_{j=1}^L w_j \cdot \delta(x - g_c(\beta_j))$. Finally, we fit a WN distribution $\mathcal{WN}(x; \mu_B, \sigma_B)$ to the resulting WD mixture. This process is illustrated in Fig. 13.

We calculate the optimal WN approximation $\mathcal{WN}(x; \mu_{\text{opt}}, \sigma_{\text{opt}})$ of the posterior density f^B by matching the first trigonometric moment of f^B . This is achieved by using a grid of 20000 equidistant samples on the circle that are weighted using the prior pdf, propagated through the nonlinear function and then used to obtain the parameters of the optimal WN approximation. Then, we use the Kullback-Leibler divergence

$$D_{KL}(\mathcal{WN}(x; \mu_{\text{opt}}, \sigma_{\text{opt}}) \parallel \mathcal{WN}(x; \mu_B, \sigma_B)) = \int_0^{2\pi} \mathcal{WN}(x; \mu_{\text{opt}}, \sigma_{\text{opt}}) \log\left(\frac{\mathcal{WN}(x; \mu_{\text{opt}}, \sigma_{\text{opt}})}{\mathcal{WN}(x; \mu_B, \sigma_B)}\right) dx$$

between the $\mathcal{WN}(x; \mu_{\text{opt}}, \sigma_{\text{opt}})$ and the fitted distribution $\mathcal{WN}(x; \mu_B, \sigma_B)$ to quantify the information loss by this approximation. The results for $\mu = \pi$ and different values of the nonlinearity parameter c as well as several possible uncertainties σ are depicted in Fig. 15. In addition to the Kullback-Leibler divergence, we also consider the error in the first trigonometric moment

$$\left| \int_0^{2\pi} f^B(x) e^{ix} dx - \int_0^{2\pi} \sum_{j=1}^L w_j \cdot \delta(x - g_c(\beta_j)) e^{ix} dx \right|$$

of the approximation compared to the true density. Results obtained using this error measure are shown in Fig. 16.

In this evaluation, we consider the moment-based approaches with $L = 3$ and $L = 5$ samples ($\lambda = 0.5$), as well as the superposition approach with $q = 5$, the uncorrected and corrected versions of the binary tree method with $L = 25$ and also $L = 25$ random samples.




Fig. 15. Kullback-Leibler divergence between best WN that approximates the posterior and the WD-based WN approximation.

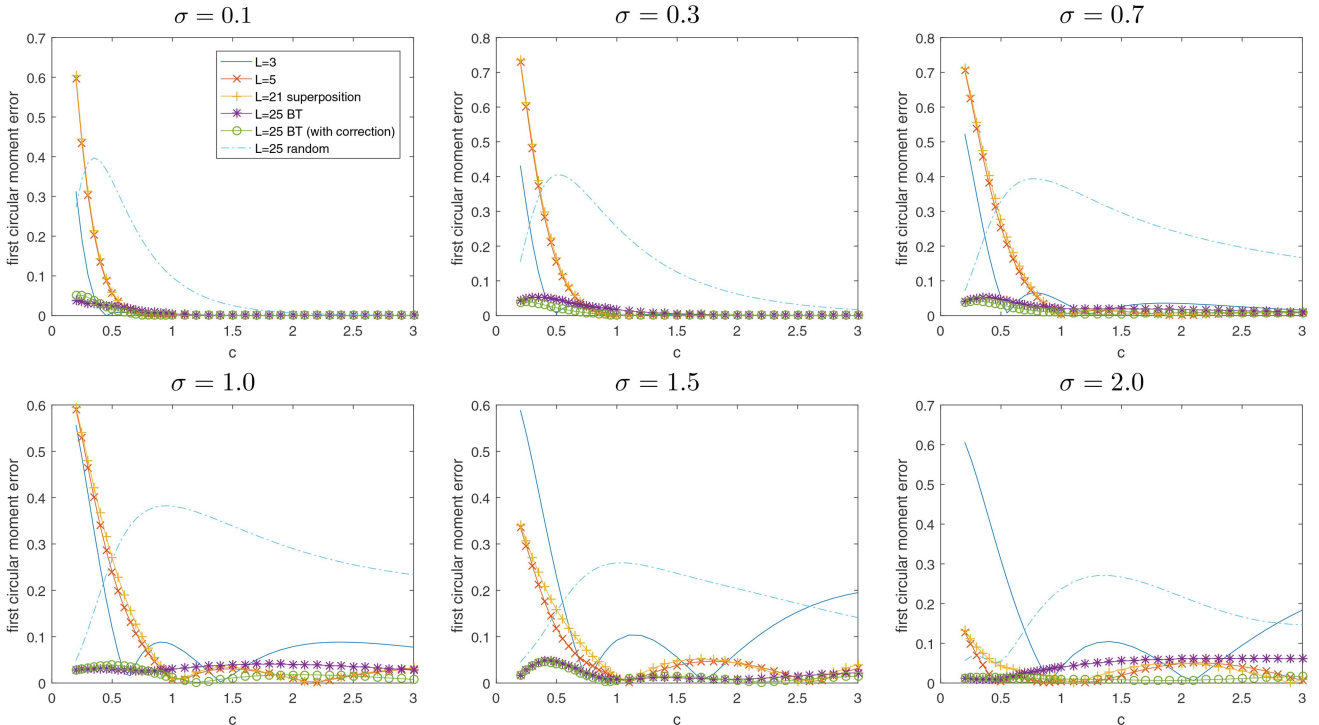


Fig. 16. Error in the first trigonometric moment between the true posterior and the WD-based approximation.

It can be seen that the random sampling generally performs very poorly compared to the deterministic sampling schemes, which is to be expected because random sampling generally needs a significantly larger number of samples to provide good results. As far as the moment-based approximations are concerned, the evaluation indicates that the method with $L = 5$ components

significantly outperforms the method with $L = 3$ components, particularly in scenarios with strong nonlinearity. The superposition method generally provides fairly similar results to the approximation with $L = 5$ components and does not seem to provide a significant advantage in the considered settings in spite of the increased number of Dirac delta components. Both versions of the

binary tree method generally provide very good results. In some cases, the uncorrected version suffers from the systematic error in the approximation, however, so the corrected version should be preferred if the slight increase in computation time is acceptable.

7. CONCLUSION

We presented several new methods to deterministically approximate circular distributions with wrapped Dirac mixtures. The proposed approaches are applicable to a variety of circular distributions, in particular the widely used wrapped normal and von Mises distributions.

First, we considered approximations with a fixed number of Dirac delta components based on one or two trigonometric moments. Because all expressions can be evaluated in closed form, these algorithms require little computational power and are suitable for implementation even on embedded devices. One might wonder if the presented algorithms can easily be generalized to higher moments, but such a generalization is nontrivial due to the fact that preserving n moments involves finding the roots of polynomials of degree n . Analytic solutions only exist for polynomials of degree ≤ 4 and are very complicated for $n = 3$ and $n = 4$.

Second, we presented a superposition method that is able to combine the results of multiple moment-based approximations to obtain a wrapped Dirac mixture with a larger number of components while still maintaining the first and second trigonometric moment. The resulting approximations can also be computed in closed form.

Third, we presented the binary tree method, a shape-based approximation with an arbitrary number of Dirac delta components. The resulting approximations represent the shape of the true continuous density much more accurately than those obtained with the superposition method, but they do not guarantee that the trigonometric moments are maintained. However, an additional correction step can be used to ensure that the first trigonometric moment is matched.

When considering a particular practical application, one of these approaches has to be chosen. The most appropriate choice mostly depends on the nonlinearity of the problem (within the range of the uncertainty) and on the available computation time. The moment-based approximation with five components is very fast, but may not be sufficient for strongly nonlinear problems. To handle nonlinear problems better, the superposition method can be used at a moderately higher computational cost. The best results are usually obtained by the binary tree method using a sufficiently large number of components, but it is computationally more costly and does not guarantee that the second moment is retained.

Future work includes the generalization of the proposed approaches to higher dimensions, e.g., the hypersphere, the hypertorus, partially wrapped spaces, and the

groups of rigid body motions in two and three dimensions. Some preliminary work in this area has already been published [15], [13], [36], but there are still many open questions, when it comes to deterministic sampling on periodic manifolds.

MATLAB implementations of the sampling methods presented in this paper are available as part of `libDirectional` [37], an open source library for directional statistics and directional estimation.

REFERENCES

- [1] Abramowitz, M., and Stegun, I. A. *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*, 10th ed. Dover, New York, 1972.
- [2] Agiomyrziannakis, Y., and Stylianou, Y. Wrapped Gaussian Mixture Models for Modeling and High-Rate Quantization of Phase Data of Speech. *IEEE Transactions on Audio, Speech, and Language Processing* 17, 4 (2009), 775–786.
- [3] Amos, D. E. Computation of Modified Bessel Functions and Their Ratios. *Mathematics of Computation* 28, 125 (1974), 239–251.
- [4] Arasaratnam, I., and Haykin, S. Cubature Kalman Filters. *IEEE Transactions on Automatic Control* 54, 6 (2009), 1254–1269.
- [5] Arulampalam, M., Maskell, S., Gordon, N., and Clapp, T. A Tutorial on Particle Filters for Online Nonlinear/Non-Gaussian Bayesian Tracking. *IEEE Transactions on Signal Processing* 50, 2 (2002), 174–188.
- [6] Azmani, M., Reboul, S., Choquel, J.-B., and Benjelloun, M. A Recursive Fusion Filter for Angular Data. In *IEEE International Conference on Robotics and Biomimetics (ROBIO 2009)* (2009), pp. 882–887.
- [7] Bucy, R. S., and Mallinckrodt, A. J. An Optimal Phase Demodulator. *Stochastics* 1, 1–4 (1975), 3–23.
- [8] Carta, J. A., Bueno, C., and Ramirez, P. Statistical Modelling of Directional Wind Speeds Using Mixtures of von Mises Distributions: Case Study. *Energy Conversion and Management* 49, 5 (May 2008), 897–907.
- [9] Costa, M., Koivunen, V., and Poor, H. Estimating Directional Statistics Using Wavefield Modeling and Mixtures of von-Mises Distributions. *IEEE Signal Processing Letters* 21, 12 (2014), 1496–1500.
- [10] Fernández-Durán, J. J. Circular Distributions Based on Nonnegative Trigonometric Sums. *Biometrics* 60, 2 (2004), 499–503.
- [11] Fisher, N. I. Problems with the Current Definitions of the Standard Deviation of Wind Direction. *Journal of Climate and Applied Meteorology* 26, 11 (1987), 1522–1529.
- [12] Gilitschenski, I., Kurz, G., and Hanebeck, U. D. Bearings-Only Sensor Scheduling Using Circular Statistics. In *Proceedings of the 16th International Conference on Information Fusion (Fusion 2013)* (Istanbul, Turkey, July 2013).
- [13] Gilitschenski, I., Kurz, G., and Hanebeck, U. D. A Stochastic Filter for Planar Rigid-Body Motions. In *Proceedings of the 2015 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI 2015)* (San Diego, California, USA, Sept. 2015).

- [14] Gilitschenski, I., Kurz, G., and Hanebeck, U. D. Non-Identity Measurement Models for Orientation Estimation Based on Directional Statistics. In *Proceedings of the 18th International Conference on Information Fusion (Fusion 2015)* (Washington D. C., USA, July 2015).
- [15] Gilitschenski, I., Kurz, G., Julier, S. J., and Hanebeck, U. D. Unscented Orientation Estimation Based on the Bingham Distribution. *IEEE Transactions on Automatic Control* 61, 1 (Jan. 2016), 172–177.
- [16] Hanebeck, U. D., and Klumpp, V. Localized Cumulative Distributions and a Multivariate Generalization of the Cramér-von Mises Distance. In *Proceedings of the 2008 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI 2008)* (Seoul, Republic of Korea, Aug. 2008), pp. 33–39.
- [17] Hanebeck, U. D., and Lindquist, A. Moment-based Dirac Mixture Approximation of Circular Densities. In *Proceedings of the 19th IFAC World Congress (IFAC 2014)* (Cape Town, South Africa, Aug. 2014).
- [18] Hill, G. W. Incomplete Bessel Function I_0 : The von Mises Distribution. *ACM Transactions On Mathematical Software* 3, 3 (Sept. 1977), 279–284.
- [19] Huber, M. F., and Hanebeck, U. D. Gaussian Filter based on Deterministic Sampling for High Quality Nonlinear Estimation. In *Proceedings of the 17th IFAC World Congress (IFAC 2008)* (Seoul, Republic of Korea, July 2008), vol. 17.
- [20] Jammalamadaka, S. R., and Kozubowski, T. J. New Families of Wrapped Distributions for Modeling Skew Circular Data. *Communications in Statistics—Theory and Methods* 33, 9 (2004), 2059–2074.
- [21] Jammalamadaka, S. R., and Sengupta, A. *Topics in Circular Statistics*. World Scientific, 2001.
- [22] Jia, B., and Xin, M. Adaptive Radial Rule based Cubature Kalman Filter. In *American Control Conference (ACC), 2015* (2015), pp. 3156–3161.
- [23] Jia, B., Xin, M., and Cheng, Y. High-degree cubature Kalman filter. *Automatica* 49, 2 (2013), 510–518.
- [24] Julier, S. J., and Uhlmann, J. K. Unscented Filtering and Nonlinear Estimation. *Proceedings of the IEEE* 92, 3 (Mar. 2004), 401–422.
- [25] Kotecha, J. H., and Djuric, P. Gaussian Particle Filtering. *IEEE Transactions on Signal Processing* 51, 10 (2003), 2592–2601.
- [26] Kucwaj, J.-C., Stienne, G., Reboul, S., Choquel, J.-B., and Benjelloun, M. Circular Multiple Change-points Estimation Applied to the GPS-L2C Phase Signal. In *Proceedings of the 17th International Conference on Information Fusion (Fusion 2014)* (Salamanca, Spain, July 2014).
- [27] Kurz, G. *Directional Estimation for Robotic Beating Heart Surgery*. PhD thesis, Karlsruhe Institute of Technology, Intelligent Sensor-Actuator-Systems Laboratory, Karlsruhe, Germany, 2015.
- [28] Kurz, G., Dolgov, M., and Hanebeck, U. D. Nonlinear Stochastic Model Predictive Control in the Circular Domain. In *Proceedings of the 2015 American Control Conference (ACC 2015)* (Chicago, Illinois, USA, July 2015).
- [29] Kurz, G., Faion, F., and Hanebeck, U. D. Constrained Object Tracking on Compact One-dimensional Manifolds Based on Directional Statistics. In *Proceedings of the Fourth IEEE GRSS International Conference on Indoor Positioning and Indoor Navigation (IPIN 2013)* (Montbeliard, France, Oct. 2013).
- [30] Kurz, G., Gilitschenski, I., Dolgov, M., and Hanebeck, U. D. Bivariate Angular Estimation Under Consideration of Dependencies Using Directional Statistics. In *Proceedings of the 53rd IEEE Conference on Decision and Control (CDC 2014)* (Los Angeles, California, USA, Dec. 2014).
- [31] Kurz, G., Gilitschenski, I., and Hanebeck, U. D. Recursive Nonlinear Filtering for Angular Data Based on Circular Distributions. In *Proceedings of the 2013 American Control Conference (ACC 2013)* (Washington D. C., USA, June 2013).
- [32] Kurz, G., Gilitschenski, I., and Hanebeck, U. D. Deterministic Approximation of Circular Densities with Symmetric Dirac Mixtures Based on Two Circular Moments. In *Proceedings of the 17th International Conference on Information Fusion (Fusion 2014)* (Salamanca, Spain, July 2014).
- [33] Kurz, G., Gilitschenski, I., and Hanebeck, U. D. Efficient Evaluation of the Probability Density Function of a Wrapped Normal Distribution. In *Proceedings of the IEEE ISIF Workshop on Sensor Data Fusion: Trends, Solutions, Applications (SDF 2014)* (Bonn, Germany, Oct. 2014).
- [34] Kurz, G., Gilitschenski, I., and Hanebeck, U. D. Nonlinear Measurement Update for Estimation of Angular Systems Based on Circular Distributions. In *Proceedings of the 2014 American Control Conference (ACC 2014)* (Portland, Oregon, USA, June 2014).
- [35] Kurz, G., Gilitschenski, I., and Hanebeck, U. D. Recursive Bayesian Filtering in Circular State Spaces. *IEEE Aerospace and Electronic Systems Magazine* 31, 3 (Mar. 2016), 70–87.
- [36] Kurz, G., Gilitschenski, I., and Hanebeck, U. D. Unscented von Mises-Fisher Filtering. *IEEE Signal Processing Letters* 23, 4 (Apr. 2016), 463–467.
- [37] Kurz, G., Gilitschenski, I., Pfaff, F., and Drude, L. libDirectional, 2015. <https://github.com/libDirectional>.
- [38] Kurz, G., and Hanebeck, U. D. Dynamic Surface Reconstruction by Recursive Fusion of Depth and Position Measurements. *Journal of Advances in Information Fusion* 9, 1 (June 2014), 13–26.
- [39] Kurz, G., and Hanebeck, U. D. Heart Phase Estimation Using Directional Statistics for Robotic Beating Heart Surgery. In *Proceedings of the 18th International Conference on Information Fusion (Fusion 2015)* (Washington D. C., USA, July 2015).
- [40] Kurz, G., and Hanebeck, U. D. Toroidal Information Fusion Based on the Bivariate von Mises Distribution. In *Proceedings of the 2015 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI 2015)* (San Diego, California, USA, Sept. 2015).
- [41] Mardia, K. V., Hughes, G., Taylor, C. C., and Singh, H. A Multivariate von Mises Distribution with Applications to Bioinformatics. *Canadian Journal of Statistics* 36, 1 (2008), 99–109.

- [42] Mardia, K. V., and Jupp, P. E. *Directional Statistics*, 1st ed. Wiley, Baffins Lane, Chichester, West Sussex, England, 1999.
- [43] Markovic, I., Cestic, J., and Petrovic, I. On Wrapping the Kalman Filter and Estimating with the SO(2) Group. In *Proceedings of the 19th International Conference on Information Fusion (Fusion 2016)* (Heidelberg, Germany, July 2016).
- [44] Markovic, I., and Petrovic, I. Bearing-Only Tracking with a Mixture of von Mises Distributions. In *Proceedings of the 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2012)* (Oct. 2012), pp. 707–712.
- [45] Nocedal, J., and Wright, S. *Numerical Optimization*, 2nd ed. Springer Series in Operations Research and Financial Engineering. Springer Science & Business Media, 2006.
- [46] Pfaff, F., Kurz, G., and Hanebeck, U. D. Multimodal Circular Filtering Using Fourier Series. In *Proceedings of the 18th International Conference on Information Fusion (Fusion 2015)* (Washington D. C., USA, July 2015).
- [47] Pfaff, F., Kurz, G., and Hanebeck, U. D. Nonlinear Prediction for Circular Filtering Using Fourier Series. In *Proceedings of the 19th International Conference on Information Fusion (Fusion 2016)* (Heidelberg, Germany, July 2016).
- [48] Schmidt, W. Statistische Methoden beim Gefügestudium kristalliner Schiefer. *Sitzungsberichte Akademie der Wissenschaften in Wien 126* (July 1917), 515–539.
- [49] Schrempf, O. C., Brunn, D., and Hanebeck, U. D. Density Approximation Based on Dirac Mixtures with Regard to Nonlinear Estimation and Filtering. In *Proceedings of the 2006 IEEE Conference on Decision and Control (CDC 2006)* (San Diego, California, USA, Dec. 2006).
- [50] Schrempf, O. C., Brunn, D., and Hanebeck, U. D. Dirac Mixture Density Approximation Based on Minimization of the Weighted Cramér-von Mises Distance. In *Proceedings of the 2006 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI 2006)* (Heidelberg, Germany, Sept. 2006), pp. 512–517.
- [51] Singh, H., Hnizdo, V., and Demchuk, E. Probabilistic Model for Two Dependent Circular Variables. *Biometrika* 89, 3 (2002), 719–723.
- [52] Steinbring, J., Pander, M., and Hanebeck, U. D. The Smart Sampling Kalman Filter with Symmetric Samples. *Journal of Advances in Information Fusion* 11, 1 (June 2016), 71–90.
- [53] Stienne, G. *Traitements des signaux circulaires appliqués à l’altimétrie par la phase des signaux GNSS*. PhD thesis, Université du Littoral Côte d’Opale, Dec. 2013.
- [54] Stienne, G., Reboul, S., Azmani, M., Choquel, J., and Benjeloun, M. A Multi-temporal Multi-sensor Circular Fusion Filter. *Information Fusion* 18 (July 2013), 86–100.
- [55] Straka, O., Dunik, J., and Simandl, M. Randomized Unscented Kalman Filter in Target Tracking. In *15th International Conference on Information Fusion (Fusion 2012)* (2012), pp. 503–510.
- [56] Straka, O., Dunik, J., and Simandl, M. Unscented Kalman Filter with Advanced Adaptation of Scaling Parameter. *Automatica* (2014).
- [57] Traa, J., and Smaragdis, P. A Wrapped Kalman Filter for Azimuthal Speaker Tracking. *IEEE Signal Processing Letters* 20, 12 (2013), 1257–1260.
- [58] von Mises, R. Über die “Ganzzahligkeit” der Atomgewichte und verwandte Fragen. *Physikalische Zeitschrift XIX* (1918), 490–500.
- [59] Willsky, A. S. Fourier Series and Estimation on the Circle with Applications to Synchronous Communication—Part I: Analysis. *IEEE Transactions on Information Theory* 20, 5 (1974), 584–590.



Gerhard Kurz received his diploma in computer science from the Karlsruhe Institute of Technology (KIT), Germany, in 2012. Afterwards, he obtained his Ph.D. in 2015 at the Intelligent Sensor-Actuator-Systems Laboratory, Karlsruhe Institute of Technology (KIT), Germany. His research interests include directional filtering, nonlinear estimation, and medical data fusion. He has authored multiple award-winning publications on these topics.



Igor Gilitschenski received a Ph.D. degree from the Intelligent Sensor-Actuator-Systems Laboratory, Karlsruhe Institute of Technology (KIT), Germany, in 2015. Before joining KIT, he obtained a diploma degree in mathematics from the University of Stuttgart, Germany. He is currently a Senior Researcher with the Autonomous Systems Laboratory, ETH Zürich, Zürich, Switzerland. His research interests include robotic perception and dynamic state estimation, with a focus on nonlinear systems and nonlinear domains.



Roland Siegwart received the M.Sc. and Ph.D. degrees in mechanical engineering from ETH Zürich, Zürich, Switzerland. He has been a Full Professor for Autonomous Systems with ETH Zürich, Zürich, Switzerland, since 2006, and the Vice President of Research and Corporate Relations from 2010 to 2014. From 1996 to 2006, he was an Associate Professor and later a Full Professor for Autonomous Microsystems and Robotics with the Ecole Polytechnique Federale de Lausanne, Lausanne, Switzerland. He leads a research group of around 30 people working on several aspects of robotics. Prof. Siegwart is a member of the Swiss Academy of Engineering Sciences. He has served as the Vice President for Technical Activities from 2004 to 2005, a Distinguished Lecturer from 2006 to 2007, and an AdCom member from 2007 to 2009 of the IEEE Robotics and Automation Society.

Uwe D. Hanebeck is a chaired professor of Computer Science at the Karlsruhe Institute of Technology (KIT) in Germany and director of the Intelligent Sensor-Actuator-Systems Laboratory (ISAS). Since 2005, he is the chairman of the Research Training Group RTG 1194 “Self-Organizing Sensor-Actuator-Networks” financed by the German Research Foundation.

Prof. Hanebeck obtained his Ph.D. degree in 1997 and his habilitation degree in 2003, both in Electrical Engineering from the Technical University in Munich, Germany. His research interests are in the areas of information fusion, nonlinear state estimation, stochastic modeling, system identification, and control with a strong emphasis on theory-driven approaches based on stochastic system theory and uncertainty models. Research results are applied to various application topics like localization, human-robot-interaction, assistive systems, sensor-actuator-networks, medical engineering, distributed measuring system, and extended range telepresence. Research is pursued in many academic projects and in a variety of cooperations with industrial partners.

Uwe D. Hanebeck was the General Chair of the “2006 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI 2006),” Program Co-Chair of the “11th International Conference on Information Fusion (Fusion 2008),” Program Co-Chair of the “2008 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI 2008),” Regional Program Co-Chair for Europe for the “2010 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2010),” and General Chair of the “19th International Conference on Information Fusion (Fusion 2016).” He is a Member of the Board of Directors of the International Society of Information Fusion (ISIF), Editor-in-chief of its Journal of Advances in Information Fusion (JAIF), and associate editor for the letter category of the IEEE Transactions on Aerospace and Electronic Systems (TAES). He is author and coauthor of more than 400 publications in various high-ranking journals and conferences.



Multitarget Tracking with the von Mises-Fisher Filter and Probabilistic Data Association

IVAN MARKOVIĆ
MARIO BUKAL
JOSIP ČESIĆ
IVAN PETROVIĆ

Directional data emerge in many scientific disciplines due to the nature of the observed phenomena or the working principles of a sensor. The problem of tracking with direction-only sensors is challenging since the motion of the target typically resides either in 3D or 2D Euclidean space, while the corresponding measurements reside either on the unit sphere or the unit circle, respectively. Furthermore, in multitarget tracking there is the need to deal with the problem of pairing sensors measurements with targets in the presence of clutter (the data association problem). In this paper we propose to approach multitarget tracking in clutter with direction-only data by setting it on the unit hypersphere, thus tracking the objects with a Bayesian estimator based on the von Mises-Fisher distribution and probabilistic data association. To achieve this goal we derive the probabilistic data association (PDA) filter and the joint probabilistic data association (JPDA) filter for the Bayesian von Mises-Fisher estimation on the unit hypersphere. The final PDA and JPDA filter equations are derived with respect to the Kullback-Leibler divergence by preserving the first moment of the hyperspherical distribution. Although the fundamental equations are given for the hyperspherical case, we focus on the filters on the unit 1-sphere (circle in \mathbb{R}^2) and the unit 2-sphere (surface of the unit ball in \mathbb{R}^3). The proposed approach is validated through synthetic data experiments on 100 Monte Carlo runs simulating multitarget tracking with noisy directional measurements and clutter.

Manuscript received February 2, 2016; released for publication March 10, 2016.

Refereeing of this contribution was handled by Igor Gilitschenski.

Authors' address: I. Marković, J. Česić, I. Petrović, Department of Computer and Control Engineering, University of Zagreb, Faculty of Electrical Engineering and Computing, Zagreb, Croatia. (E-mail: {ivan.markovic, josip.cesic, ivan.petrovic}@fer.hr). M. Bukal, Department of Applied Mathematics, University of Zagreb, Faculty of Electrical Engineering and Computing, Zagreb, Croatia. (E-mail: mario.bukal@fer.hr).

This is an extended version of the paper *Direction-only tracking of moving objects on the unit sphere via probabilistic data association* published at the *17th International Conference on Information Fusion (FUSION2014)*.

This work has been supported by the European Regional Development Fund under the project *Advanced Technologies in Power Plants and Rail Vehicles*.

1557-6418/16/\$17.00 © 2016 JAIF

I. INTRODUCTION

Directional data emerge in many scientific disciplines. Since the surface of the earth is approximately a sphere, such data arise readily in earth sciences, e.g. the location of the earthquake's epicenter, the paleomagnetic directions of the earth's magnetic pole etc. Furthermore, many astronomical observations are points on the celestial sphere and as such yield directional data. In multitarget tracking (MTT), it is not uncommon to work with sensors that can provide only directions to the objects in question. The measurement and estimation state space have a specific geometry of their own, which is different from the true trajectory space geometry.

The problem is challenging, because, although the motion of the target resides either in 3D or 2D Euclidean space, corresponding measurements reside either on the sphere or the circle, respectively. For example, if we are measuring and estimating only the direction to the target in 2D, i.e. the azimuth, the state and measurements will bear the non-Euclidean properties of angles. Applying standard filtering techniques employing the Gaussian distribution on \mathbb{R}^d would ignore the underlying geometry. In robotics, measurements from various sensors, due to their nature of operation, yield direction-only information of the objects of interest, e.g. microphone arrays measure the direction of the sound source, perspective and omnidirectional cameras measure directions of various features of interest in space.

Considering MTT, the goal of such a system is to estimate the multiple trajectories in scenarios with noisy measurements, clutter or false alarms (measurements that falsely appear to originate from moving objects). The duties of such a system are truly numerous, and in the past several seminal methods have been developed in order to tackle this problem wherein data association plays one of the crucial roles. To solve this problem the methods that can be used are the global nearest neighbor (GNN) which attempts to find the single most likely data association hypothesis at each scan [1], the probabilistic data association (PDA) filter for single target tracking and joint probabilistic data association (JPDA) filter for MTT where multiple hypotheses are formed after each scan and then these hypotheses are combined before proceeding further with the next scan [2], and the multiple hypothesis filter (MHT) [3] where multiple data association hypotheses are formed and propagated from scan to scan [1]. Furthermore, a group of methods based on random finite sets was developed with MTT in mind. An example is the probability hypothesis density (PHD) filter [4] which estimates the number of objects in the scene but does not solve the data association problem by itself, however, a solution has been presented in [5] for the Gaussian mixture PHD filter [6]. Further solutions within this framework were developed, such as the cardinalized PHD (CPHD) [7], [8], and multi-Bernoulli filters [9]–[12].

Estimation methods using the von Mises (vM) and the wrapped Gaussian distributions were recently discussed in [13]–[22]. Furthermore, inference on quaternions, namely the Bingham distribution (which actually models variables with 180° symmetry), was used in [23]–[25] and in [26] where, furthermore, a second-order filter was derived which included also the rotational velocity. These approaches, advocating the unit hypersphere as the appropriate filtering space, showed better performance of the Bingham filter compared to the extended Kalman filter. This approach was further extended to unscented orientation estimation based on the Bingham distribution and showed better performance than the unscented Kalman filter and better to equal performance than the particle filter (depending on the number of particles) [27]. However, the normalization constant of the Bingham distribution, hence its partial derivative, cannot be computed in closed form, but this can be surmounted by caching techniques and its relationship with the vM distribution [25]. On the contrary, the von Mises-Fisher (vMF) distribution does not require such techniques since the normalization constant and its partial derivative can be calculated in closed form, which will be of practical interest in the ensuing JPDA equations. Indeed, there are many choices for directional distributions, but for inferential purposes the vMF distribution is most widely used because of its exponential family structure [28].

Considering target tracking on the unit sphere (as in surface of the unit ball in \mathbb{R}^3), it was proposed in [29] to utilize the vMF distribution to model the system state and the sensors measurements after which a Bayesian estimator was developed for single object tracking. This method was used in our previous work [30] in order to track a single target detected by an omnidirectional camera on a mobile robot, thus not offering a consistent method for dealing with multiple moving objects nor false alarms/clutter. A global nearest neighbor (GNN), which in contrast to JPDA solves the data association by hard assignment, was applied in tracking multiple targets on the sphere in [31] and the Rényi α -divergence was used as a distance measure. In [32] we developed the vM mixture PHD filter and compared it to the Gaussian mixture PHD filter for MTT on the unit circle. The vM mixture PHD filter showed better performance since it was able to capture the circle geometry intrinsically, which proved important for the mixture PHD filter having components spread throughout the entire state space. In [33] the vMF filter was used to track multiple speakers. The authors also addressed modeling the object velocity with a Gaussian distribution by assuming that the posterior state distribution can factor to the product of the vMF and the Gaussian distribution. The authors compared the performance of the vMF filter to the constant velocity Kalman filter and a particle filter based on the vMF sampling. The vMF filter showed better performance than the Kalman filter, while the performance with respect to the vMF particle

filter depended on the number of particles (at the cost of increase in computational complexity). Multitarget tracking was solved by averaging observations prior to the update step, where the weights were computed based on the normalized vMF innovation term (also taking into account a uniform distribution to handle outliers). The authors conclude that the vMF filter strikes a good compromise between the efficiency of the Kalman filter and the statistical grounding of the vMF particle filter.

In [34] we proposed a probabilistic data association solution to the problem of tracking single and multiple targets in clutter with direction-only sensors. This paper serves as a foundation for the present paper and we are extending it in several ways. Instead of focusing on the unit sphere, in the present work the Bayesian filter is presented on the unit hypersphere, i.e. the $(d-1)$ -dimensional unit sphere, and it is shown how vM and vMF filters arise as special cases. We also provide in depth mathematical proofs of the moment-matching techniques used in the paper. Furthermore, validation gating is discussed in detail and experiments are extended through more thorough validation via MC runs of the vM and vMF PDA and JPDA filters.

Specifically, we pose the problem on the unit hypersphere and utilize a Bayesian tracking algorithm that is based on the vMF distribution on the $(d-1)$ -dimensional unit sphere. To solve the data association problem we derive the PDA and JPDA filter equations for the aforementioned Bayesian filter. This constituted (i) deriving the a posteriori probabilities of association events on the unit hypersphere which essentially weigh each hypothetical estimation and form a mixture of von Mises-Fisher densities, (ii) determining the final (single) component density as the result of the update in the PDA and JPDA filter by preserving the first moment of the hyperspherical distribution (which is optimal in the Kullback-Leibler sense), and (iii) modeling the false alarms as Poisson processes. The proposed algorithms were tested on 100 Monte Carlo runs of a synthetic data set comprising of single and multiple-object scenarios where direction-only measurements were corrupted with noise and clutter.

The paper is organized as follows. Section II presents the general mathematical background and formulae for tracking on the hypersphere with the von Mises-Fisher distribution. Section III describes the proposed PDA and JPDA filtering approaches based on the von Mises-Fisher distribution. Section IV presents the results and discussion of the synthetic data experiments, while Section V concludes the paper.

II. GENERAL BACKGROUND

When considering directions in d dimensions, i.e. unit vectors in d -dimensional Euclidean space \mathbb{R}^d , one can represent them as points on the $(d-1)$ -dimensional sphere \mathbb{S}^{d-1} of unit radius. Thus, in our notation, 1-sphere is the unit circle in \mathbb{R}^2 and the 2-sphere is

the surface of the unit ball in \mathbb{R}^3 . In the sequel we introduce all the necessary constituents and discuss the basic paradigm of a tracking system on the unit $(d-1)$ -sphere using von the Mises-Fisher distributions.

A. von Mises-Fisher distribution

Parametric probability distribution defined on the unit $(d-1)$ -dimensional sphere \mathbb{S}^{d-1} , whose probability density function (pdf) is given by

$$f(\mathbf{x}; \boldsymbol{\mu}, \kappa) = C_d(\kappa) \exp(\kappa \boldsymbol{\mu} \cdot \mathbf{x}), \quad \mathbf{x} \in \mathbb{S}^{d-1}, \quad (1)$$

is called *von Mises-Fisher distribution* with parameters $\kappa \geq 0$ and $\boldsymbol{\mu} \in \mathbb{S}^{d-1}$ denoting the concentration and the mean direction, respectively. Expression

$$C_d(\kappa) = \frac{\kappa^{d/2-1}}{(2\pi)^{d/2} I_{d/2-1}(\kappa)} \quad (2)$$

is the normalization constant with respect to the standard surface measure, while I_p denotes the modified Bessel function of the first kind and order p [35]. To enlighten many of its properties, it is worth of considering the vMF distributions as an exponential family. Recall, a parametric set of probability distributions parametrized by the natural parameter $\boldsymbol{\theta} \in \Theta \subseteq \mathbb{R}^d$ is called an *exponential family* if their probability densities admit the following canonical representation [36]

$$p(\mathbf{x}; \boldsymbol{\theta}) = \exp(T(\mathbf{x}) \cdot \boldsymbol{\theta} - F(\boldsymbol{\theta}) + C(\mathbf{x})), \quad (3)$$

where $T(\mathbf{x})$ is called sufficient statistics, $F(\boldsymbol{\theta})$ is the log-normalizing function and $C(\mathbf{x})$ denotes the carrier measure. Many familiar parametric distributions, like the Gaussian, Poisson, Gamma, Dirichlet etc., are exponential families [37].

One easily deduces from (1) that the vMF distribution constitutes an exponential family parametrized by the natural parameter $\boldsymbol{\theta} = \kappa \boldsymbol{\mu} \in \mathbb{R}^d$, the log-normalizing function given by

$$F_d(\boldsymbol{\theta}) = -\log C_d(\|\boldsymbol{\theta}\|), \quad (4)$$

and the trivial carrier measure. The minimal sufficient statistics is the identity map on \mathbb{S}^{d-1} , $T(\mathbf{x}) = \mathbf{x}$, hence, the vMF distribution is completely determined by the *directional (angular) mean*¹

$$\mathbb{E}[\mathbf{x}] = \int_{\mathbb{S}^{d-1}} \mathbf{x} f(\mathbf{x}; \boldsymbol{\mu}, \kappa) d\mathbf{x} = \nabla F_d(\boldsymbol{\theta}) = A_d(\kappa) \boldsymbol{\mu}, \quad (5)$$

where $A_d(\kappa)$ is the ratio of the following Bessel functions

$$A_d(\kappa) = \frac{I_{d/2}(\kappa)}{I_{d/2-1}(\kappa)}. \quad (6)$$

Please see Proposition A.1 in Appendix A for the proof. Let us mention two more distinctive properties of the vMF distribution: (i) density (1) is rotationally invariant around the mean direction and, (ii) analogously to

¹Note that the *directional mean* is defined by the integral (5), while the *mean direction* $\boldsymbol{\mu}$ is the parameter of the von Mises-Fisher distribution.

the multivariate Gaussian distribution, it is characterized by the *maximum entropy principle* in the following sense. Given any pdf on \mathbb{S}^{d-1} of prescribed directional mean $\boldsymbol{\eta}$, it is then the vMF distribution with the natural parameter $\boldsymbol{\theta} = (\nabla F)^{-1}(\boldsymbol{\eta})$, which maximizes the Boltzmann-Shannon entropy $-\int_{\mathbb{S}^{d-1}} p(\mathbf{x}) \log p(\mathbf{x}) d\mathbf{x}$ [29, Proposition 2.2].

In the present paper we present the vMF filter on \mathbb{S}^{d-1} , but due to the aimed application of MTT, we particularly show explicit relation in cases when $d=2$ and $d=3$, i.e. the vMF distributions on the unit 1-sphere (called the von Mises distribution) and the unit 2-sphere (also called the Fisher or Langevin distribution). For the former the above expressions amount to [35]

$$C_2(\kappa) = \frac{1}{2\pi I_0(\kappa)} \quad \text{and} \quad A_2(\kappa) = \frac{I_1(\kappa)}{I_0(\kappa)}, \quad (7)$$

while for the latter they simplify to

$$C_3(\kappa) = \frac{\kappa}{4\pi \sinh \kappa} \quad \text{and} \quad A_3(\kappa) = \frac{1}{\tanh \kappa} - \frac{1}{\kappa}. \quad (8)$$

An example of von Mises-Fisher distributions on the unit 1-sphere and the unit 2-sphere with different mean directions and concentration parameters are depicted in Fig. 1 and Fig. 2, respectively. Even though the application examples are shown for cases when $d=2$ and $d=3$, with the presented general approach the higher dimensional von Mises-Fisher filter could be applied in novel applications beyond MTT. For example, a higher dimensional von Mises-Fisher distribution was used in [38] to address large scale data mining applications, such as text categorization and gene expression analysis, which involve high-dimensional data that is also inherently directional in nature.

Please note that we will denote the 1-sphere Bayes filter as the von Mises filter, while with the practical slight abuse of terminology the 2-sphere Bayes filter will be denoted as the von Mises-Fisher filter and where necessary the general Bayes filter on the $(d-1)$ -sphere will be called the hyperspherical von Mises-Fisher filter. Furthermore, all the explicit formulae for the vMF distribution will be presented in angular variables with the following relation $\boldsymbol{\mu} = (\cos \alpha, \sin \alpha)$.

B. Motion model

In our model we assume that moving objects are relatively slow with respect to the sampling rate, i.e. changes in the objects's position between two consequent observations are relatively small. Mathematically, motion of such objects is then described by a wide-sense stationary stochastic processes, among which, the Wiener process (Brownian motion) is the standard choice [39]. These time-continuous processes are typically further approximated by a random walk of a fixed time step.

The Brownian motion distribution on \mathbb{S}^{d-1} with parameters $\boldsymbol{\mu}$ and κ is the distribution at time κ^{-1} of a random point which starts at $\boldsymbol{\mu}$ and moves on \mathbb{S}^{d-1} under




Fig. 1. Examples of the von Mises-Fisher distribution on the unit 1-sphere, i.e. the von Mises distribution, with equal mean directions and concentration parameters of 50 (red), 150 (green), and 500 (blue), which correspond approximately to standard deviations of 8.1°, 4.7° and 2.6°, respectively.




Fig. 2. Samples on the unit 2-sphere of the von Mises-Fisher distribution with different mean directions and concentration parameters of 50 (red), 150 (green), and 500 (blue), which correspond approximately to standard deviations of 10°, 5.7° and 3.1°, respectively.

an isotropic diffusion with infinitesimal variance \mathbf{I}_{d-1} (the identity matrix on the tangent space of \mathbb{S}^{d-1}) [28]. When $d = 2$ the Brownian motion distribution is the wrapped Normal distribution, which in turn is very close to and can be well approximated by the von Mises distribution with corresponding parameter transformations [16], [28]. This closeness of the densities also extends to the distributions on \mathbb{S}^{d-1} , i.e. to the Brownian motion distribution and the von Mises-Fisher distribution on \mathbb{S}^{d-1} , as we shall detail out in the sequel for $d = 3$. In practice, this means that in the prediction stage we can approximate Brownian motion well with the corresponding vMF distribution.

When $d = 3$, there are at least two viewpoints which motivate the vMF as a motion model. First, consider an isotropic Wiener process (Brownian motion) on \mathbb{S}^2 , which is defined as a temporally and spatially homogeneous Markov process [40]. It has been shown that such process exists and is unique. Moreover, the pdf of the distribution of a random point starting at $\hat{\mathbf{x}}^{k-1} \in \mathbb{S}^2$ and moving to $\hat{\mathbf{x}}^k$ at time $\tau > 0$ is given by

$$p(\hat{\mathbf{x}}^k | \hat{\mathbf{x}}^{k-1}) = \frac{1}{4\pi} \sum_{l=0}^{\infty} (2l+1) e^{-\tau l(l+1)/4} L_l(\hat{\mathbf{x}}^{k-1} \cdot \hat{\mathbf{x}}^k), \quad (9)$$

where L_l are Legendre polynomials of degree l . This formula is also obtained as the limit of the vanishing step size of a random walk on the sphere with all directions of movement being equally probable [41]. Furthermore, in the same paper the authors discuss the approximation of (9), which is the true \mathbb{S}^2 -analogue of the planar symmetric Gaussian distribution, by a von Mises-Fisher distribution, which unlike (9) features certain analyticity properties. They show that for small variance $\tau > 0$, one can approximate the Brownian motion on \mathbb{S}^2 by the von Mises-Fisher diffusion with large concentration parameter $\kappa_\tau = 2/\tau$, i.e.

$$p(\hat{\mathbf{x}}^k | \hat{\mathbf{x}}^{k-1}) \approx f(\hat{\mathbf{x}}^k; \hat{\mathbf{x}}^{k-1}, \kappa_\tau).$$

The second approach is from the applicational viewpoint probably even more relevant. Consider the isotropic Wiener process in \mathbb{R}^3 and corresponding time-discretization (random walk) of fixed time step $\tau > 0$. The transition probability density function of the process is given by the Gaussian density

$$p(\mathbf{x}^k | \mathbf{x}^{k-1}) = \frac{1}{(2\pi\sigma_\tau^2)^{3/2}} \exp(-\|\mathbf{x}^k - \mathbf{x}^{k-1}\|^2 / 2\sigma_\tau^2), \quad (10)$$

where $\sigma_\tau^2 := \sigma^2\tau$ and $\sigma > 0$ denotes the process strength. If we are confined to a measurement device which only measures direction $\hat{\mathbf{x}}^k \in \mathbb{S}^2$ of position vectors \mathbf{x}^k , then marginalizing (10) over the range, one obtains statistical model being the *angular Gaussian* density [42]

$$p(\hat{\mathbf{x}}^k | \mathbf{x}^{k-1}) = \frac{1}{C} \int_0^\infty s^2 \exp(-s^2/2\kappa_\tau + s\hat{\mathbf{x}}^{k-1} \cdot \hat{\mathbf{x}}^k) ds \quad (11)$$

with parameters $\hat{\mathbf{x}}^{k-1} = \mathbf{x}^{k-1}/\|\mathbf{x}^{k-1}\|$, $\kappa_\tau = \|\mathbf{x}^{k-1}\|^2/\sigma_\tau^2$, and normalization constant C . Following [29], for moderate or large values of κ_τ (practically most relevant cases), (11) can be well approximated by the vMF density $f(\hat{\mathbf{x}}^k; \hat{\mathbf{x}}^{k-1}, \kappa_\tau)$.

Usage of physically more realistic motion models like Ornstein-Uhlenbeck or Langevin processes [43], in place of simple Wiener process, requires more complex state representation manifolds and solving the Fokker-Planck equation to obtain the corresponding state transition densities. The latter typically needs to be further approximated by an appropriate parametric density which

will keep the model computationally tractable and statistically consistent with the remaining ingredients of the filtering algorithm: state distribution and measurement model.

C. Observation model

As already announced above, we assume that observation process consists of measuring a direction, where measurement disturbances are interpreted as random rotations, i.e. observed direction \mathbf{z} is a random rotation of the true direction \mathbf{x} . It is reasonable to statistically describe such model by a unimodal distribution which is rotationally invariant around the true direction \mathbf{x} .

Since our goal is to derive a Bayesian filter on \mathbb{S}^{d-1} , our choice for the measurement model will be the von Mises-Fisher distribution defined on the $(d-1)$ -sphere, represented by its density

$$p(\mathbf{z} | \mathbf{x}) = C_d(\kappa_o) \exp(\kappa_o \mathbf{x} \cdot \mathbf{z}), \quad \mathbf{z}, \mathbf{x} \in \mathbb{S}^{d-1}, \quad (12)$$

where the concentration parameter κ_o describes the measurement uncertainty. In the present paper specific examples will be given when $d = 2$ and $d = 3$, serving as a model for representing measurements of directions, i.e. angle-only measurements, in 2D and 3D Euclidean spaces.

D. Bayesian filter equations

With a Bayes filter we are striving to estimate the density $p(\mathbf{x}^k | \mathbf{z}^{1:k})$, i.e. the pdf of the state \mathbf{x}^k at time instant k given the history of all the measurements $\mathbf{z}^{1:k}$. This process can be decomposed in two steps—namely the prediction and the correction step. Let us assume that at time instant $k-1$ we have estimated the pdf of the targeted state, i.e. we have the available the posterior $p(\mathbf{x}^{k-1} | \mathbf{z}^{1:k-1})$. For the case at hand, filtering on \mathbb{S}^{d-1} , the prediction step amounts to solving the following integral

$$p(\mathbf{x}^k | \mathbf{z}^{1:k-1}) = \int_{\mathbb{S}^{d-1}} p(\mathbf{x}^k | \mathbf{x}^{k-1}) p(\mathbf{x}^{k-1} | \mathbf{z}^{1:k-1}) d\mathbf{x}^{k-1}, \quad (13)$$

where $p(\mathbf{x}^k | \mathbf{x}^{k-1})$ is the transition density or motion model of the tracked object. Then, after receiving the measurement at time instant k the correction step amounts to evaluating the Bayes rule

$$p(\mathbf{x}^k | \mathbf{z}^{1:k}) = \frac{p(\mathbf{z}^k | \mathbf{x}^k) p(\mathbf{x}^k | \mathbf{z}^{1:k-1})}{p(\mathbf{z}^k | \mathbf{z}^{1:k-1})}, \quad (14)$$

where $p(\mathbf{z}^k | \mathbf{x}^k)$ is the sensor model and the normalizer $p(\mathbf{z}^k | \mathbf{z}^{1:k-1})$ can be evaluated via

$$p(\mathbf{z}^k | \mathbf{z}^{1:k-1}) = \int_{\mathbb{S}^{d-1}} p(\mathbf{z}^k | \mathbf{x}^k) p(\mathbf{x}^k | \mathbf{z}^{1:k-1}) d\mathbf{x}^k. \quad (15)$$

The goal of the present paper is to employ the filtering equations (13), (14) and (15) when the underlying distributions are all of the vMF form. If these equations were to be solved for the Gaussian distribution, the result would be the Kalman filter [44].

If we assume that both the posterior at $k-1$ and the transition density are vMF distribution then the prediction step (13) would not yield another vMF distribution and the filtering equations would not stay in the domain of the same distribution. To solve this problem, a moment-matching technique is applied where the moments of the resulting distribution are matched to the moments of the corresponding vMF distribution. This procedure is also the optimal choice in the sense of the Kullback-Leibler divergence [45].

Let the state have the estimated position (direction) $\boldsymbol{\mu}^{k-1} \in \mathbb{S}^{d-1}$, conditioned upon all available measurements up to (and including) time $k-1$, which is statistically described by the density

$$p(\mathbf{x}^{k-1} | \mathbf{z}^{1:k-1}) = C_d(\kappa^{k-1}) \exp(\kappa^{k-1} \boldsymbol{\mu}^{k-1} \cdot \mathbf{x}^{k-1}).$$

Calculating the directional mean with respect to the prediction density

$$\begin{aligned} \mathbb{E}[\mathbf{x}^k | \mathbf{z}^{1:k-1}] &= \int_{\mathbb{S}^{d-1}} \mathbf{x}^k p(\mathbf{x}^k | \mathbf{z}^{1:k-1}) d\mathbf{x}^k \\ &= A_d(\kappa_o) A_d(\kappa^{k-1}) \boldsymbol{\mu}^{k-1}, \end{aligned}$$

and according to (5) we determine a unique vMF $f(\mathbf{x}^k; \bar{\boldsymbol{\mu}}^k, \bar{\kappa}^k)$ such that $\mathbb{E}[\mathbf{x}^k | \mathbf{z}^{1:k-1}] = A_d(\bar{\kappa}^k) \bar{\boldsymbol{\mu}}^k$. Thus, the prediction equations are

$$\bar{\boldsymbol{\mu}}^k = \boldsymbol{\mu}^{k-1}, \quad \bar{\kappa}^k = A_d^{-1}(A_d(\kappa_o) A_d(\kappa^{k-1})), \quad (16)$$

where A_d^{-1} denotes the inverse to the function A_d defined in (6). To compute the inverse, one must resort to numerical methods since the derived equations are transcendental. Please see Proposition A.2 in Appendix A for the proof. Similar equations were derived in [13], [15], [16], [28], [29] for the cases of $d = 2$ and $d = 3$, i.e. for the distributions on the unit 1-sphere and the unit 2-sphere.

Upon availability of the measurement \mathbf{z}^k at time k , and under the assumption that the sensor model follows the vMF density $f(\mathbf{z}^k; \mathbf{x}^k, \kappa_o)$, the posterior is found using the Bayes rule (14)

$$p(\mathbf{x}^k | \mathbf{z}^{1:k}) = C_d(\kappa^k) \exp(\kappa^k \boldsymbol{\mu}^k \cdot \mathbf{x}^k),$$

with corresponding update equations [29]

$$\begin{aligned} \kappa^k &= \|\kappa_o \mathbf{z}^k + \bar{\kappa}^k \bar{\boldsymbol{\mu}}^k\|, \\ \boldsymbol{\mu}^k &= \frac{\kappa_o \mathbf{z}^k + \bar{\kappa}^k \bar{\boldsymbol{\mu}}^k}{\kappa^k}. \end{aligned} \quad (17)$$

The Bayes normalizer (15) evaluates to

$$p(\mathbf{z}^k | \mathbf{z}^{1:k-1}) = \frac{C_d(\kappa_o) C_d(\bar{\kappa}^k)}{C_d(\kappa^k)}. \quad (18)$$

REMARK 2.1: The update equations of the 1-sphere Bayes filter, i.e. the von Mises filter [13], [15], [16], are

$$\kappa^k = \sqrt{\kappa_o^2 + (\bar{\kappa}^k)^2 + 2\kappa_o \bar{\kappa}^k \cos(\beta^k - \bar{\alpha}^k)}, \quad (19)$$

$$\alpha^k = \bar{\alpha}^k + \arctan \frac{\sin(\beta^k - \bar{\alpha}^k)}{\bar{\kappa}^k / \kappa_o + \cos(\beta^k - \bar{\alpha}^k)}, \quad (20)$$

where $\bar{\kappa}^k \bar{\boldsymbol{\mu}}^k = \bar{\kappa}^k (\cos \bar{\alpha}^k, \sin \bar{\alpha}^k)$ and $\kappa_o \mathbf{z}^k = \kappa_o (\cos \beta^k, \sin \beta^k)$. See Appendix A for the proof.

III. TRACKING IN CLUTTER WITH THE VON MISES-FISHER DISTRIBUTION

Target tracking in a cluttered environment requires, among other, to resolve the problem of measurement-to-target association. Moreover, in such approaches it is often practical to devise a validation gate so as to reject highly unlikely measurements. This way the computational complexity of the association procedure can be significantly lowered [1]. In this section we recall two basic probabilistic (nonbackscan) approaches, developed in seminal papers [46], [47] in the context of Poisson distributed clutter and models described by Gaussian distributions. Here we extend these approaches to directional (spherical) models described by vMF distributions.

A. Validation gating

In order to assess the validity of a measurement \mathbf{z}^k , we use the density (15) [48]. By inspecting (18) we can notice that the result of the integral does not have the form of a vMF distribution. If the goal was just to compute (18) for explicit values, we would not need to go further, but the validation gating assumes calculating confidence intervals which is inconvenienced by a non-standard density form of (18). To solve this issue, similar logic as in the prediction stage is used—in lieu of using the exact density, we will use the moment-matched vMF distribution. Given that, the directional mean with respect to (15) is

$$\begin{aligned} \mathbb{E}[\mathbf{z}^k | \mathbf{z}^{1:k-1}] &= \int_{\mathbb{S}^{d-1}} \mathbf{z}^k p(\mathbf{z}^k | \mathbf{z}^{1:k-1}) d\mathbf{z}^k \\ &= A(\bar{\kappa}^k) A(\kappa_o) \bar{\boldsymbol{\mu}}^k \end{aligned} \quad (21)$$

and according to (5) we determine a unique vMF $f(\mathbf{z}^k; \bar{\boldsymbol{\mu}}^k, \kappa_S^k)$ where

$$\kappa_S^k = A_d^{-1}(A_d(\kappa_o) A_d(\bar{\kappa}^k)). \quad (22)$$

Please confer Remark A.1 for the proof. Note that parameter κ_S^k has the role analogous to the Kalman innovation for linear models.

Having computed the required vMF density, we can determine the confidence intervals. When $d = 2$ the quantiles of the distribution with mean direction $\mu = 0$ are transferred to the linear interval $[-\pi, \pi]$ by cutting the unit circle at π . For an approximate validation region $100(1 - \alpha)\%$, e.g. 95% when $\alpha = 0.05$, the lower and upper tail area are respectively defined as [28]

$$\begin{aligned} \Pr(-180 < x < -180 + \delta) &= \alpha/2, \\ \Pr(180 - \delta < x < 180) &= \alpha/2. \end{aligned} \quad (23)$$

When $d = 3$, an approximate $100(1 - \alpha)\%$ validation region for \mathbf{z}^k is [28]

$$\{\mathbf{z}^k : \mathbf{z}^k \cdot \bar{\boldsymbol{\mu}}^k \geq \cos \delta\}, \quad (24)$$

which defines the intersection of the unit 2-sphere with the cone having vertex at the origin, axis the mean direction $\bar{\boldsymbol{\mu}}^k$ and semi-vertical angle δ which is defined by

$$\Pr(\mathbf{z} \cdot \bar{\boldsymbol{\mu}}^k \geq \delta) = \alpha, \quad (25)$$

with \mathbf{z} distributed according to vMF with parameters $\bar{\boldsymbol{\mu}}^k$ and κ_S . The tables linking desired validation regions defined by α , and for given concentration parameters κ the corresponding intervals defined by δ , can be found in [28], [49]. So the procedure is as follows, (i) specify the desired validation region, i.e. the α , (ii) given the computed κ_S from (22) and α by using tables we find the corresponding δ , (iii) we evaluate the obtained measurement \mathbf{z}^k with respect to the predicted mean $\bar{\boldsymbol{\mu}}^k$ to see if it satisfies the validation region constraint, and (iv) if so the measurement is assigned to the target in question, otherwise it is considered as clutter.

B. Probabilistic data association filter

First we assume a single target in track with multiple measurements where the number of false alarms is a Poisson distributed random variable. Let Z^k denote the set of all measurements that fall within the validation gate at time k

$$Z^k = \{\mathbf{z}_j^k : j = 1, \dots, m_k\},$$

and $Z^{1:k} = \{Z^1, \dots, Z^k\}$ the history of all measurements within the validation gate. We want to calculate the conditional probability density $p(\mathbf{x}^k | Z^{1:k})$ for all $k \geq 1$. Assume that at a given time $k - 1$, the object's direction is described by the vMF density

$$p(\mathbf{x}^{k-1} | Z^{1:k-1}) = C_d(\kappa^{k-1}) \exp(\kappa^{k-1} \boldsymbol{\mu}^{k-1} \cdot \mathbf{x}^{k-1}).$$

Obtaining measurements Z^k we build the following set of hypotheses:

$$\mathcal{H}_j = \{\mathbf{z}_j^k \text{ is the correct measurement}\}, \quad j = 1, \dots, m_k,$$

and

$$\mathcal{H}_0 = \{\text{none of the gated measurements are correct}\}.$$

Using the total probability formula, the posterior probability density at time k is given by

$$p(\mathbf{x}^k | Z^{1:k}) = \sum_{j=0}^{m_k} p(\mathbf{x}^k | \mathcal{H}_j, Z^{1:k}) P(\mathcal{H}_j | Z^{1:k}). \quad (26)$$

From the definition of \mathcal{H}_j and using the Bayes rule, for $j = 1, \dots, m_k$ we have

$$\begin{aligned} p(\mathbf{x}^k | \mathcal{H}_j, Z^{1:k}) &= p(\mathbf{x}^k | \mathcal{H}_j, Z^k, Z^{1:k-1}) \\ &= \frac{p(Z^k, \mathcal{H}_j | \mathbf{x}^k) p(\mathbf{x}^k | Z^{1:k-1})}{p(Z^k, \mathcal{H}_j | Z^{1:k-1})} \\ &= \frac{p(\mathbf{z}_j^k | \mathbf{x}^k) p(\mathbf{x}^k | Z^{1:k-1})}{p(\mathbf{z}_j^k | Z^{1:k-1})}. \end{aligned} \quad (27)$$

Assuming that the likelihood and the prior density are both vMF with respective parameters $\mathbf{z}_j^k, \kappa_o^k$, and $\bar{\boldsymbol{\mu}}^k, \bar{\kappa}^k$ given by (16), then the posterior density in (27) is also vMF with parameters analogous to those in (17)

$$\kappa_j^k = \|\kappa_o \mathbf{z}_j^k + \bar{\kappa}^k \bar{\boldsymbol{\mu}}^k\|, \quad (28)$$

$$\boldsymbol{\mu}_j^k = \frac{\kappa_o \mathbf{z}_j^k + \bar{\kappa}^k \bar{\boldsymbol{\mu}}^k}{\kappa_j^k}, \quad j = 1, \dots, m_k. \quad (29)$$

Clearly, for $j = 0$, $p(\mathbf{x}^k | \mathcal{H}_0, Z^{1:k}) = p(\mathbf{x}^k | Z^{1:k-1})$.

Let $w_j = P(\mathcal{H}_j | Z^{1:k})$ denote the a posteriori probabilities of each feature having originated from the object in track. According to calculations in [47, Appendix]

$$w_j = \frac{p(\mathbf{z}_j^k | \mathcal{H}_j, Z^{1:k-1})}{b + \sum_{i=1}^{m_k} p(\mathbf{z}_i^k | \mathcal{H}_i, Z^{1:k-1})}, \quad j = 1, \dots, m_k, \quad (30)$$

$$w_0 = \frac{b}{b + \sum_{i=1}^{m_k} p(\mathbf{z}_i^k | \mathcal{H}_i, Z^{1:k-1})}, \quad (31)$$

where $b = c(1 - p_G p_D)/p_D$, $c > 0$ is the clutter density, p_G is the probability that the correct feature will be inside the validation gate, and p_D is the probability that the correct feature will be detected. Density $p(\mathbf{z}_j^k | \mathcal{H}_j, Z^{1:k-1})$ denotes the probability density of a measurement conditioned upon past data and hypothesis that is correct, which is assumed to be known and in our case it is modeled by the vMF density

$$p(\mathbf{z}_j^k | \mathcal{H}_j, Z^{1:k-1}) = f(\mathbf{z}_j^k; \bar{\boldsymbol{\mu}}^k, \kappa_S^k), \quad (32)$$

where κ_S^k is given by (22).

Having defined and calculated all the ingredients, posterior density (26) becomes a mixture of vMF densities

$$p(\mathbf{x}^k | Z^{1:k}) = \sum_{j=0}^{m_k} w_j f(\mathbf{x}^k; \boldsymbol{\mu}_j^k, \kappa_j^k). \quad (33)$$

In order to estimate the object's direction $\boldsymbol{\mu}^k \in \mathbb{S}^{d-1}$, we calculate the directional mean

$$\mathbb{E}[\mathbf{x}^k | Z^{1:k}] = \int_{\mathbb{S}^{d-1}} \mathbf{x}^k p(\mathbf{x}^k | Z^{1:k}) d\mathbf{x}^k = \sum_{j=0}^{m_k} w_j A_d(\kappa_j^k) \boldsymbol{\mu}_j^k,$$

and, using (5), determine the unique vMF density $f(\mathbf{x}^k; \boldsymbol{\mu}^k, \kappa^k)$, which is the best approximation of (33) in

the sense of the Kullback-Leibler divergence by solving

$$\kappa^k = A_d^{-1} \left(\left\| \sum_{j=0}^{m_k} w_j A_d(\kappa_j^k) \boldsymbol{\mu}_j^k \right\| \right), \quad (34)$$

$$\boldsymbol{\mu}^k = \left(\sum_{j=0}^{m_k} w_j A_d(\kappa_j^k) \boldsymbol{\mu}_j^k \right) / A_d(\kappa^k). \quad (35)$$

The latter procedure is the analogon of computing the state estimate and covariance matrix from the mixture of Gaussians representing the posterior densities in [46], [47].

When $d = 2$, the expressions (34) and (35) need not be expressed in the vectorial form, but rather in the angular variables. Componentwise, they evaluate to the following formulae [50]

$$\begin{aligned} A_2^2(\kappa^k) &= \sum_{j=0}^{m_k} w_j^2 A_2^2(\kappa_j^k) \\ &+ 2 \sum_{\substack{j,i=1 \\ j < i}}^{m_k} w_j w_i A_2(\kappa_j^k) A_2(\kappa_i^k) \cos(\alpha_j^k - \alpha_i^k) \\ \tan \alpha^k &= \frac{\sum_{j=0}^{m_k} w_j A_2(\kappa_j^k) \sin \alpha_j^k}{\sum_{j=0}^{m_k} w_j A_2(\kappa_j^k) \cos \alpha_j^k}. \end{aligned} \quad (36)$$

C. Joint probabilistic data association filter

Next we consider the problem of tracking several interfering targets $\{\mathcal{O}_1, \dots, \mathcal{O}_N\}$, with the pertaining number being fixed to N . The main issue is how to appropriately assign features to targets in track. In principle, PDA filter approach could be applied for each object separately, but this would implicitly assume that all measurement features originated by another object in track are Poisson distributed clutter [47], and we would like to avoid such a rough assumption.

Let $\mathbf{X}^k = \{\mathbf{x}_1^k, \dots, \mathbf{x}_N^k\} \subset \mathbb{S}^{d-1}$ denotes the set of object's states (directions) at time k , and assume that at a given time $k - 1$ position of each object \mathcal{O}_i is described by the vMF density

$$p(\mathbf{x}_i^{k-1} | Z^{1:k-1}) = C_d(\kappa_i^{k-1}) \exp(\kappa_i^{k-1} \boldsymbol{\mu}_i^{k-1} \cdot \mathbf{x}_i^{k-1}).$$

Upon availability of a set of new measurements $Z^k = \{\mathbf{z}_j^k : j = 1, \dots, m_k\}$, the following set of hypotheses is built:

$$\mathcal{H}_{ij} = \{\mathbf{z}_j^k \text{ is caused by } \mathcal{O}_i\}, \quad j = 1, \dots, m_k,$$

and

$$\mathcal{H}_{i0} = \{\text{none of the measurements is caused by } \mathcal{O}_i\}.$$

Again, the total probability formula implies that the posterior density for object \mathcal{O}_i at time k is given by

$$p(\mathbf{x}_i^k | Z^{1:k}) = \sum_{j=0}^{m_k} p(\mathbf{x}_i^k | \mathcal{H}_{ij}, Z^{1:k}) P(\mathcal{H}_{ij} | Z^{1:k}), \quad (37)$$

where densities $p(\mathbf{x}_i^k | \mathcal{H}_{ij}, Z^{1:k})$ are computed following the same lines and assumptions as in the previous PDA filter approach. They are vMF densities $f(\mathbf{x}_i^k; \boldsymbol{\mu}_{ij}^k, \kappa_{ij}^k)$ with parameters

$$\kappa_{ij}^k = \|\kappa_o \mathbf{z}_j^k + \bar{\kappa}_i^k \bar{\boldsymbol{\mu}}_i^k\|, \quad (38)$$

$$\boldsymbol{\mu}_{ij}^k = \frac{\kappa_o \mathbf{z}_j^k + \bar{\kappa}_i^k \bar{\boldsymbol{\mu}}_i^k}{\kappa_{ij}^k}, \quad j = 1, \dots, m_k, \quad (39)$$

and $\kappa_{i0}^k = \bar{\kappa}_i^k$ and $\boldsymbol{\mu}_{i0}^k = \bar{\boldsymbol{\mu}}_i^k$.

The only difference between PDA filter and JPDA filter is in calculation of a posteriori association probabilities $w_{ij} = P(\mathcal{H}_{ij} | Z^{1:k})$, where JPDA filter takes into account measurement-to-object association events jointly across the set of objects. This means that hypothesis \mathcal{H}_{ij} consists of all *valid joint association events* \mathcal{E} which assign feature \mathbf{z}_j^k to object \mathcal{O}_i . By valid joint association events we consider those which assert that every feature lying within the validation gate region can originate from at most one object and every object can generate at most one feature. Thus, they partition the hypothesis \mathcal{H}_{ij} and

$$w_{ij} = \sum_{\mathcal{E} \in \mathcal{H}_{ij}} P(\mathcal{E} | Z^{1:k}), \quad j = 1, \dots, m_k, \quad (40)$$

$$w_{i0} = 1 - \sum_{j=1}^{m_k} w_{ij}. \quad (41)$$

In order to compute $P(\mathcal{E} | Z^{1:k})$, two auxiliary indicator functions are introduced: *measurement association indicator* $\varphi_j(\mathcal{E})$, which indicates whether in event \mathcal{E} measurement \mathbf{z}_j^k is associated with any object, and *target detection indicator* $\delta_i(\mathcal{E})$, which indicates whether in \mathcal{E} any measurement is associated with object \mathcal{O}_i . Following [47] and using vMF model instead of the Gaussian, we obtain

$$P(\mathcal{E} | Z^{1:k}) = B(\mathcal{E}) \prod_{\substack{l=1 \\ \varphi_l(\mathcal{E})=1}}^{m_k} f(\mathbf{z}_l^k; \bar{\boldsymbol{\mu}}_{i_l}^k, \kappa_{S,i_l}^k)$$

with $\kappa_{S,i_l}^k = A_d^{-1}(A_d(\kappa_o)A_d(\bar{\kappa}_{i_l}^k))$ analogous to (22), where i_l is the object index with which measurement \mathbf{z}_l^k is associated. Next,

$$B(\mathcal{E}) = \frac{c^{\phi(\mathcal{E})}}{p_G^{\alpha(\mathcal{E})} C} \prod_{\substack{i=1 \\ \delta_i(\mathcal{E})=1}}^N p_D^i \prod_{\substack{i=1 \\ \delta_i(\mathcal{E})=0}}^N (1 - p_D^i),$$

where $\phi(\mathcal{E})$ is the number of false features in joint event \mathcal{E} , which is assumed Poisson distributed, $\alpha(\mathcal{E}) = \sum_{j=1}^{m_k} \varphi_j(\mathcal{E})$ is the number of measurement-to-object associations in \mathcal{E} , p_G is the probability that the correct measurement will be inside the validation gate, p_D^i is the detection probability of object \mathcal{O}_i , and C is the normalization constant.

Posterior density (37) for object \mathcal{O}_i is again a mixture of vMF densities, and the estimated posterior direction $\boldsymbol{\mu}_i^k$ with uncertainty κ_i^k is calculated via

$$\kappa_i^k = A_d^{-1} \left(\left\| \sum_{j=0}^{m_k} w_{ij} A_d(\kappa_{ij}^k) \boldsymbol{\mu}_{ij}^k \right\| \right), \quad (42)$$

$$\boldsymbol{\mu}_i^k = \left(\sum_{j=0}^{m_k} w_{ij} A_d(\kappa_{ij}^k) \boldsymbol{\mu}_{ij}^k \right) / A_d(\kappa_i^k). \quad (43)$$

Here also for the case when $d = 2$, the resulting parameters can be computed as in (36).

IV. SYNTHETIC DATA EXPERIMENTS

In order to test the performance of the vMF PDA and JPDA filters, we have simulated the system for 250 time steps with maneuvering targets on the unit 1-sphere and the unit 2-sphere. The targets were uniformly spawned and their dynamics was described by a constant angular velocity model, where the disturbance acted as random noise in the angular acceleration. The number of spawned targets was a random integer from [3,5] for the JPDA case, while for the PDA case it was set to one. Since the JPDA filter assumes a constant and known number of targets in the scene, the originally spawned number of objects was kept constant during the simulation, i.e. there were no target births nor deaths. Please recall that the vM filter is the Bayes filter on the unit 1-sphere, while the vMF filter is the Bayes filter on the unit 2-sphere. The underlying motivation behind these simulations is MTT with directional sensors like microphone arrays and omnidirectional cameras. The former can determine directions to the sound sources based on microphone pair signal differences [15], [51], [52], while for the latter it has been shown that the image formation can be described by the unified spherical projection model yielding a representation of the omnidirectional image on the unit 2-sphere [30], [53], [54].

In order to make the simulations as realistic as possible (i) the trajectories were corrupted with the von Mises, i.e. the von Mises-Fisher noise, with concentration parameter $\kappa = 1500$,² (ii) the probability of detection was $p_D = 0.95$ and (iii) false alarms were simulated as a Poisson process on the unit spheres with the mean value $\lambda = \beta \mu(\mathbb{S}^{d-1})$, where μ denotes the area measure on \mathbb{S}^{d-1} , and the intensity $\beta = 0.25$ was defined as the number of false measurements per solid radian. For the 1-sphere case $\mu(\mathbb{S}^{d-1}) = 2\pi$, while for the 2-sphere case $\mu(\mathbb{S}^{d-1}) = 4\pi$. For example, on average we could expect $4\pi\beta$ false alarms per sensor frame sampled from a uniform distribution on the unit 2-sphere. For all the experiments the validation gate was computed for $\alpha = 0.01$, i.e. the validation region of 99% was used. The experiments involving the PDA and the JPDA filter were envisaged so as to simulate tracking of a single target and

²For the vM distribution this corresponds approximately to $\sigma = 1.5^\circ$ [16], while for the vMF distribution it is closer to $\sigma = 1.8^\circ$.




Fig. 3. An example of the experimental results of the tracking task for the von Mises PDA filter. The solid green line represents the estimated azimuth, while the solid black line is the ground truth. The gray circles represent false alarms. The mean error for this example was 0.97° .

multiple targets in clutter, both on the unit 1-sphere and the unit 2-sphere. In the end, we have performed 100 Monte Carlo runs of the previously described scenarios for both the vM and the vMF PDA and JPDA filters.

For the 1-sphere case the error Δ_{err} was computed as the absolute angular error by taking the periodicity into account

$$\Delta_{\text{err}} = |\text{mod}(\mathbf{x}_g - \mathbf{x}^k + \pi, 2\pi) - \pi|. \quad (44)$$

For the 2-sphere case the error Δ_{err} was computed as the great circle distance between the ground truth \mathbf{x}_g and the estimated state \mathbf{x}^k

$$\Delta_{\text{err}} = \arccos(\mathbf{x}_g \cdot \mathbf{x}^k). \quad (45)$$

The expression (45) would yield the same result as (44) if therein 2D unit vectors were used instead of angles.

For the PDA filter the error calculation is straightforward, since we have only one object in the scene. But for the JPDA filter the error calculation cannot be approached in the same manner since we are tracking multiple objects and a ground truth trajectory needs to be paired up with a vMF filter trajectory. In this paper we are assuming known and constant number of objects in the scene, and, furthermore, we are focusing on deriving the fundamentals for probabilistic data association techniques on the unit hyperspheres. Hence, more involved methods for track management are not discussed in the present paper and in error calculation we are not penalizing if filters switches tracks when the two tracks cross. Therefore, for the JPDA filter case, we first calculate errors between all the filters and the ground truth trajectories, after which we apply the Hungarian algorithm [55], [56] to optimally assign filters to the ground truth trajectories. However, in Section IV-A we discuss solutions






Fig. 4. An example of the experimental results of the tracking task for the vMF PDA filter. The solid green line represents the estimated direction, while the solid black line is the ground truth. The gray circles represent false alarms. The mean error for this example was 0.93° . (a) Trajectory on the unit 2-sphere. (b) Azimuth. (c) Elevation.

which could be applied for handling such multitarget tracking issues and metrics that can capture such errors.

Examples of the experimental results of the tracking task involving the PDA filter on the 1-sphere and the 2-sphere are shown in Figs. 3 and 4, respectively.






Fig. 5. An example of the experimental results of tracking five objects with the von Mises JPDA filter. The solid color lines represent the estimated azimuths, while the solid black lines are the ground truth. The gray circles represent false alarms. The mean absolute angular error of all the trajectories for this example was 0.99° .

The figures depict the ground truth, the estimated trajectory and false alarms. We can notice that the filters successfully manage to track the object in such a high clutter scenario, while yielding a mean error smaller than 1° . Figures 5 and 6 show an example of the JPDA filter tracking task. Examples with five objects in the scene are depicted, each color representing a single filter. We can see that in both cases filters manage to successfully track all the objects in the scene while maintaining overall mean error smaller than 1° . We can notice that the filters achieve mean error smaller than the measurement noise even when clutter is present. However, Fig. 5 deserves further comment. We can notice therein that during track crossing the filters switched tracks, e.g. at approximately 70 s the red and the cyan filter exchanged objects appearing as if they changed their course, while in truth they kept the same course during the whole simulation. As discussed previously, in the present paper we are not penalizing the track-switch, but we will address in Section IV-A how these issues can be alleviated.


Results of the 100 MC runs are depicted in Fig. 7. Therein statistics of the resulting errors is given both for the vM and vMF PDA and JPDA filters. The error for each run was calculated using (44) and (45) for the vM and the vMF filter, respectively. From the figure we can notice that the vM JPDA filter exhibited more outliers in comparison to others, but this was due to track coalescence. Namely, trajectories crossed much more often in the $[0, 2\pi]$ interval than they did on the unit 2-sphere. Overall, the median error was smaller than 1° .



(a)



(b)



(c)

Fig. 6. An example of the experimental results of the tracking task for the vMF JPDA filter. The solid color lines represent the estimated directions, while the solid black lines are the ground truth. The gray circles represent false alarms. The mean error of all the trajectories for this example was 0.89° . (a) Trajectories on the unit 2-sphere. (b) Azimuth. (c) Elevation.

A. Discussion

As previously mentioned, the JPDA filter assumes known and constant number of objects in the scene. However, in [57] the joint integrated probability data association (JIPDA) filter was proposed in order to al-




Fig. 7. Matlab's boxplot of the PDA and JPDA vM and vMF filter errors over 100 MC runs. The red lines are the median, the edges of the boxes are the 25th and 75th percentile, while the whiskers extend to the most extreme error not considered outliers and outliers are plotted as red pluses (errors 1.5 times larger than the difference between the percentiles).

leviate the JPDA assumption of the constant and known number of targets in the scene by including the probability of target existence within the framework. The results presented in this paper can be directly employed within the JIPDA framework. Furthermore, to handle issues like track switching and track coalescence, various other extensions of the PDA/JPDA filtering have been proposed [1], [58]. As long as the theoretical results derived in this paper are used, we believe that the aforementioned extensions can be applied to the vM and vMF JPDA filter as well.

In the present paper the likelihood for track switching is increased since the system state comprises of just the static component. If dynamical terms were included, the switch would be less likely to occur since terms like the estimated velocity would suggest the future motion of the object and would resolve situations like the one depicted in Fig. 5. However, principally including velocity components, which are in general Euclidean and not directional variables, into the same filtering framework is not a trivial task and is out of the scope of this paper. Examples of how the system state could be extended in this vein for the cases of the Bingham, vM and vMF distributions are given in [21], [26], [33], respectively.

Another solution for reducing the likelihood of track switch and coalescence is to include independent features for describing the tracked objects. For example, in multiple speaker tracking with a microphone array the fundamental frequency can be tracked [59] and used in the data association process. When tracking moving objects with an omnidirectional camera appearance-based features can be used to describe each object, like the HOG descriptor [60], which could also be used in the data association process. With these extensions at hand the

multitarget tracking performance with respect to track switch, track loss, track coalescence and similar issues could be evaluated using the CLEAR MOT metrics [61], the optimal subpattern assignment metric [62] and its extensions [63], [64].

V. CONCLUSION

In the present paper we have proposed methods for tracking single and multiple targets in clutter on the $(d-1)$ -sphere with the von Mises-Fisher distribution. The methods are based on Bayesian tracking and the data association logic of the PDA and JPDA filters. For single target tracking we have derived the PDA filter equations by assuming a moving object in a Poisson distributed clutter. This has resulted with a mixture of hypotheses represented as hyperspherical von Mises-Fisher densities which were weighted by the a posteriori probability that the selected measurement is correct. For multiple object tracking the JPDA filter was derived under similar assumptions which again resulted with a mixture of hypotheses represented as hyperspherical von Mises-Fisher densities, where each component was weighted by the a posteriori probability of the association event. The final single component estimate for each object in track, both in the PDA and JPDA filter case, was obtained by preserving the first moment of the distribution which is optimal in the Kullback-Leibler sense. For the cases of $d=2$ and $d=3$ the hyperspherical filter yields the vM PDA and JPDA and the vMF PDA and JPDA filter, respectively. In the end, the proposed methods were validated on synthetic data examples in 100 MC runs simulating scenarios of tracking a single and multiple targets in clutter on the unit 1-sphere and the unit 2-sphere.

APPENDIX A

PROPOSITION A.1 *The directional angular mean of the von Mises-Fisher density $f(\mathbf{x}; \boldsymbol{\mu}, \kappa)$ defined on \mathbb{S}^{d-1} equals*

$$\mathbb{E}[\mathbf{x}] = \nabla F_d(\boldsymbol{\theta}) = A_d(\kappa)\boldsymbol{\mu}, \quad (46)$$

where $A_d(\kappa) = I_{d/2}(\kappa)/I_{d/2-1}(\kappa)$.

PROOF The first equality in (46) follows from the identity

$$0 = \int_{\mathbb{S}^{d-1}} (\mathbf{x} - \nabla F_d(\boldsymbol{\theta})) e^{\mathbf{x} \cdot \boldsymbol{\theta} - F_d(\boldsymbol{\theta})} d\mathbf{x},$$

obtained by differentiating $1 = \int_{\mathbb{S}^{d-1}} e^{\mathbf{x} \cdot \boldsymbol{\theta} - F_d(\boldsymbol{\theta})} d\mathbf{x}$ with respect to $\boldsymbol{\theta}$. For the second equality, using $\boldsymbol{\theta} = \kappa\boldsymbol{\mu}$ and $\|\boldsymbol{\theta}\| = \kappa$, we straightforwardly calculate

$$\nabla F_d(\boldsymbol{\theta}) = -\frac{C'_d(\kappa)}{C_d(\kappa)}\boldsymbol{\mu}. \quad (47)$$

Using the basic recurrence relations for the modified Bessel functions one obtains

$$C'_d(\kappa) = -\frac{\kappa^{d/2-1}}{(2\pi)^{d/2}} \frac{I_{d/2}(\kappa)}{I_{d/2-1}^2(\kappa)},$$

which yields

$$-\frac{C'_d(\kappa)}{C_d(\kappa)} = \frac{I_{d/2}(\kappa)}{I_{d/2-1}(\kappa)} =: A_d(\kappa).$$

PROPOSITION A.2 *Let the prior and the transition density be defined as following von Mises-Fisher densities*

$$\begin{aligned} p(\mathbf{x}^k | \mathbf{x}^{k-1}) &= f(\mathbf{x}^k; \mathbf{x}^{k-1}, \kappa_\tau) \\ p(\mathbf{x}^{k-1} | \mathbf{z}^{1:k-1}) &= f(\mathbf{x}^{k-1}; \boldsymbol{\mu}^{k-1}, \kappa^{k-1}). \end{aligned} \quad (48)$$

Then the directional mean of the predicted density $p(\mathbf{x}^k | \mathbf{z}^{1:k-1})$ equals to

$$\mathbb{E}[\mathbf{x}^k | \mathbf{z}^{1:k-1}] = A_d(\kappa_\tau) A_d(\kappa^{k-1}) \boldsymbol{\mu}^{k-1}. \quad (49)$$

PROOF The result is obtained by applying the Fubini's theorem (rearranging the terms within the integrals) and using the previous proposition:

$$\begin{aligned} \mathbb{E}[\mathbf{x}^k | \mathbf{z}^{1:k-1}] &= \int_{\mathbb{S}^{d-1}} \mathbf{x}^k \left(\int_{\mathbb{S}^{d-1}} p(\mathbf{x}^k | \mathbf{x}^{k-1}) p(\mathbf{x}^{k-1} | \mathbf{z}^{1:k-1}) d\mathbf{x}^{k-1} \right) d\mathbf{x}^k \\ &= \int_{\mathbb{S}^{d-1}} p(\mathbf{x}^{k-1} | \mathbf{z}^{1:k-1}) \left(\int_{\mathbb{S}^{d-1}} \mathbf{x}^k p(\mathbf{x}^k | \mathbf{x}^{k-1}) d\mathbf{x}^k \right) d\mathbf{x}^{k-1} \\ &= A_d(\kappa_\tau) \int_{\mathbb{S}^{d-1}} \mathbf{x}^{k-1} p(\mathbf{x}^{k-1} | \mathbf{z}^{1:k-1}) d\mathbf{x}^{k-1} \\ &= A_d(\kappa_\tau) A_d(\kappa^{k-1}) \boldsymbol{\mu}^{k-1}. \end{aligned} \quad (50)$$

A von Mises-Fisher distribution $f(\mathbf{x}^k; \bar{\boldsymbol{\mu}}^k, \bar{\kappa}^k)$ with matched moments then has parameters

$$\bar{\boldsymbol{\mu}}^k = \boldsymbol{\mu}^{k-1}, \quad \bar{\kappa}^k = A_d^{-1}(A_d(\kappa_\tau) A_d(\kappa^{k-1})). \quad (51)$$

REMARK A.1: Observe that the same calculations as above yield the following. If the sensor likelihood and the predicted density are defined as the von Mises-Fisher densities

$$\begin{aligned} p(\mathbf{z}^k | \mathbf{x}^k) &= f(\mathbf{z}^k; \mathbf{x}^k, \kappa_o) \\ p(\mathbf{x}^k | \mathbf{z}^{1:k-1}) &= f(\mathbf{x}^k; \bar{\boldsymbol{\mu}}^k, \bar{\kappa}^k), \end{aligned} \quad (52)$$

then the directional mean of the Bayes normalizer $p(\mathbf{z}^k | \mathbf{z}^{1:k-1})$ evaluates to

$$\mathbb{E}[\mathbf{z}^k | \mathbf{z}^{1:k-1}] = A(\kappa_o) A(\bar{\kappa}^k) \bar{\boldsymbol{\mu}}^k. \quad (53)$$

PROOF OF REMARK 2.1. The equation (19) for κ^k follows directly from calculating the norm of the vector $\kappa_o \mathbf{z}^k + \bar{\kappa}^k \bar{\boldsymbol{\mu}}^k$. The proof of (20) is more involved, but also straightforward. First observe that

$$\alpha^k = \arctan \frac{\kappa_o \sin \beta^k + \bar{\kappa}^k \sin \bar{\alpha}^k}{\kappa_o \cos \beta^k + \bar{\kappa}^k \cos \bar{\alpha}^k}. \quad (54)$$

This form of the angle update is also frequently used [13], [16], but we find the expression (20) more expository from the prediction-correction standpoint [15]. Let us now transform the numerator in (54) as follows

$$\begin{aligned} \kappa_o \sin \beta^k + \bar{\kappa}^k \sin \bar{\alpha}^k &= \kappa_o \sin(\nu^k + \bar{\alpha}^k) + \bar{\kappa}^k \sin \bar{\alpha}^k \\ &= (\bar{\kappa}^k + \kappa_o \cos \nu^k) \sin \bar{\alpha}^k + \kappa_o \sin \nu^k \cos \bar{\alpha}^k, \end{aligned} \quad (55)$$

where $\nu^k = \beta^k - \bar{\alpha}^k$. Now using the trigonometric identity

$$A \sin x + B \cos x = \sqrt{A^2 + B^2} \sin \left(x + \arctan \frac{B}{A} \right),$$

we get the following formula

$$\begin{aligned} \kappa_o \sin \beta^k + \bar{\kappa}^k \sin \bar{\alpha}^k &= \sqrt{\kappa_o^2 + (\bar{\kappa}^k)^2 + 2\kappa_o \bar{\kappa}^k \cos \nu^k} \\ &\cdot \sin \left(\bar{\alpha}^k + \arctan \frac{\sin \nu^k}{\bar{\kappa}^k / \kappa_o + \cos \nu^k} \right). \end{aligned} \quad (56)$$

Analogous procedure can be performed for the denominator of (54), yielding an expression similar to (56), where the corresponding sines are replaced by cosines. Returning to (54) cancels the square root terms and a tangent is left due to the sine and cosine ratio—yielding finally the sought mean angle update formula (20).

REFERENCES

- [1] S. Blackman and R. Popoli *Design and Analysis of Modern Tracking Systems*. Artech House Publishers, 1999.
- [2] Y. Bar-Shalom and E. Tse *Sonar tracking of multiple targets using joint probabilistic data association filter*, *Automatica*, vol. 11, pp. 451–460, 1975.
- [3] D. Reid *An algorithm for tracking multiple targets*, *IEEE Transactions on Automatic Control*, vol. 24, no. 6, pp. 843–854, 1979.
- [4] R. P. S. Mahler *Multitarget Bayes filtering via first-order multitarget moments*, *IEEE Transactions on Aerospace and Electronic Systems*, vol. 39, no. 4, pp. 1152–1178, oct 2003.
- [5] K. Panta, D. E. Clark, and B.-N. Vo *Data association and track management for the Gaussian mixture probability hypothesis density filter*, *IEEE Transactions on Aerospace and Electronic Systems*, vol. 45, no. 3, pp. 1003–1016, 2009.
- [6] B.-N. Vo and W.-K. Ma *The Gaussian mixture probability hypothesis density filter*, *IEEE Transactions on Signal Processing*, vol. 54, no. 11, pp. 4091–4104, nov 2006.
- [7] R. Mahler *PHD filters of higher order in target number*, *IEEE Transactions on Aerospace and Electronic Systems*, vol. 43, no. 4, pp. 1523–1543, 2007.

- [8] B.-T. Vo, B.-N. Vo, and A. Cantoni
Analytic implementations of the cardinalized probability hypothesis density filter,
IEEE Transactions on Signal Processing, vol. 55, no. 7, pp. 3553–3567, jul 2007.
- [9] R. Mahler
Statistical Multisource-Multitarget Information Fusion.
Artech House, 2007.
- [10] B. T. Vo, B. N. Vo, and A. Cantoni
The cardinality balanced multi-target multi-Bernoulli filter and its implementations,
IEEE Transactions on Signal Processing, vol. 57, no. 2, pp. 409–423, 2009.
- [11] S. Reuter, B. T. Vo, B. N. Vo, and K. Dietmayer
The labeled multi-Bernoulli filter,
IEEE Transactions on Signal Processing, vol. 62, no. 12, pp. 3246–3260, 2014.
- [12] H. Deusch, S. Reuter, and K. Dietmayer
The labeled multi-Bernoulli SLAM filter,
IEEE Signal Processing Letters, vol. 22, no. 10, pp. 1561–1565, 2015.
- [13] M. Azmani, S. Reboul, J.-B. Choquel, and M. Benjelloun
A recursive fusion filter for angular data,
in *IEEE International Conference on Robotics and Biomimetics (ROBIO)*, dec 2009, pp. 882–887.
- [14] I. Marković and I. Petrović
Speaker localization and tracking with a microphone array on a mobile robot using von Mises distribution and particle filtering,
Robotics and Autonomous Systems, vol. 58, no. 11, pp. 1185–1196, 2010.
- [15] ———
Bearing-only tracking with a mixture of von Mises distributions,
in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2012, pp. 707–712.
- [16] G. Kurz, I. Gilitschenski, and U. D. Hanebeck
Recursive nonlinear filtering for angular data based on circular distributions,
in *American Control Conference (ACC)*, 2013, pp. 5439–5445.
- [17] G. Kurz, F. Faion, and U. D. Hanebeck
Constrained object tracking on compact one-dimensional manifolds based on directional statistics,
in *International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, oct 2013, pp. 1–9.
- [18] J. Traa and P. Smaragdis
A wrapped Kalman filter for azimuthal speaker tracking,
IEEE Signal Processing Letters, vol. 20, no. 12, pp. 1257–1260, 2013.
- [19] D. F. Crouse, R. W. Osborne, K. Pattipati, P. Willett, and Y. Bar-Shalom
Efficient 2D location estimation of sensor arrays using targets of opportunity,
Journal of Advances in Information Fusion, vol. 8, no. 1, pp. 73–89, 2013.
- [20] D. F. Crouse
Maximum-likelihood post-detection radar ambiguity resolution,
IEEE Transactions on Aerospace and Electronic Systems, vol. 50, no. 3, pp. 1876–1883, 2014.
- [21] G. Stienne, S. Reboul, M. Azmani, J. Choquel, and M. Benjelloun
A multi-temporal multi-sensor circular fusion filter,
Information Fusion, vol. 18, pp. 86–100, jul 2014.
- [22] F. Pfaff, G. Kurz, and U. D. Hanebeck
Multimodal circular filtering using Fourier series,
in *International Conference on Information Fusion (FUSION)*, 2015, pp. 711–718.
- [23] J. Glover, G. Bradski, and R. B. Rusu
Monte Carlo pose estimation with quaternion kernels and the Bingham distribution,
in *Proceedings of Robotics: Science and Systems*, Los Angeles, CA, USA, 2011.
- [24] G. Kurz, I. Gilitschenski, S. Julier, and U. D. Hanebeck
Recursive estimation of orientation based on the Bingham distribution,
in *International Conference on Information Fusion (FUSION)*, 2013, pp. 1–16.
- [25] ———
Recursive Bingham filter for directional estimation involving 180 degree symmetry,
Journal of Advances in Information Fusion, vol. 9, no. 2, pp. 90–105, 2014.
- [26] J. Glover and L. P. Kaelbling
Tracking the spin on a ping pong ball with the quaternion Bingham filter,
in *IEEE International Conference on Robotics and Automation (ICRA)*, 2014, pp. 4133–4140.
- [27] I. Gilitschenski, G. Kurz, S. J. Julier, and U. D. Hanebeck
Unscented orientation estimation based on the Bingham distribution,
IEEE Transactions on Automatic Control, vol. in press, p. 11, 2015.
- [28] K. V. Mardia and P. E. Jupp
Directional Statistics.
Wiley, 1999.
- [29] A. Chiuso and G. Picci
Visual tracking of points as estimation on the unit sphere,
in *The confluence of vision and control*, ser. Lecture Notes in Control and Information Sciences, D. Kriegman, G. Hager, and A. Morse, Eds. Springer London, 1998, vol. 237, pp. 90–105.
- [30] I. Marković, F. Chaumette, and I. Petrović
Moving object detection, tracking and following using an omnidirectional camera on a mobile robot,
in *International Conference on Robotics and Automation (ICRA)*, 2014.
- [31] J. Česić, I. Marković, and I. Petrović
Tracking of multiple moving objects on the unit sphere using a multiple-camera system on a mobile robot,
in *International Conference on Intelligent Autonomous Systems (IAS)*, 2014, p. 12.
- [32] I. Marković, J. Česić, and I. Petrović
Von Mises mixture PHD filter,
IEEE Signal Processing Letters, vol. 22, no. 12, pp. 2229–2233, 2015.
- [33] J. Traa and P. Smaragdis
Multiple speaker tracking with the factorial von Mises-Fisher filter,
in *IEEE International Workshop on Machine Learning for Signal Processing*, 2014.
- [34] I. Marković, M. Bukal, J. Česić, and I. Petrović
Direction-only tracking of moving objects on the unit sphere via probabilistic data association,
in *International Conference on Information Fusion (FUSION)*, 2014, pp. 1–7.
- [35] K. V. Mardia and P. E. Jupp
Directional statistics.
Wiley, 1999.
- [36] F. Nielsen and V. Garcia
“Statistical exponential families: A digest with flash cards,”
arXiv:0911.4863, 2009.
- [37] S. R. Jammalamadaka and A. Sengupta
Topics in Circular Statistics.
World Scientific, 2001.


- [38] A. Banerjee
Clustering on the unit hypersphere using von Mises-Fisher distributions,
Journal of Machine Learning Research, vol. 6, pp. 1–39, 2005.
- [39] A. H. Jazwinski
Stochastic processes and filtering theory.
New York: Academic press, 1970.
- [40] K. Yosida
Brownian motion on the surface of the 3-sphere,
Ann. Math. Statist., vol. 20, no. 2, pp. 292–296, 1949.
- [41] P. H. Roberts and H. D. Urse
Random Walk on a Sphere and on a Riemannian Manifold,
Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences, vol. 252, no. 1012, pp. 317–356, 1960.
- [42] G. S. Watson
Statistics on Spheres.
Wiley, 1983.
- [43] F. Beichlet
Stochastic Processes in Science, Engineering, and Finance. Chapman & Hall, 2006.
- [44] S. Thrun, W. Burgard, and D. Fox
Probabilistic Robotics.
The MIT Press, 2006.
- [45] O. Schwander and F. Nielsen
Learning mixtures by simplifying kernel density estimators,
in *Matrix Information Geometry*, F. Nielsen and R. Bhatia, Eds. Springer Berlin Heidelberg, 2013, pp. 403–426.
- [46] Y. Bar-Shalom and T. Edison
Tracking in a cluttered environment with probabilistic data association,
Automatica, vol. 11, no. 5, pp. 451–460, 1975.
- [47] T. Fortmann, Y. Bar-Shalom, and M. Scheffe
Sonar tracking of multiple targets using joint probabilistic data association,
IEEE Journal of Oceanic Engineering, vol. 8, no. 3, pp. 173–184, 1983.
- [48] T. Bailey, B. Upcroft, and H. Durrant-Whyte
Validation gating for non-linear non-Gaussian target tracking,
in *International Conference on Information Fusion*, 2006, pp. 1–6.
- [49] J. P. Marques de Sá
Applied Statistics using SPSS, Statistica, Matlab and R.
Springer-Verlag, 2007.
- [50] M. Bukal, I. Marković, and I. Petrović
Composite distance based approach to von Mises mixture reduction,
Information Fusion, vol. 20, pp. 136–145, 2014.
- [51] J.-M. Valin, F. Michaud, and J. Rouat
Robust localization and tracking of simultaneous moving sound sources using beamforming and particle filtering,
Robotics and Autonomous Systems, vol. 55, no. 3, pp. 216–228, 2007.
- [52] S. Argentieri, P. Danès, and P. Souères
“A survey on sound source localization in robotics: from binaural to array processing methods,”
Laboratory for Analysis and Architecture of Systems, The National Center for Scientific Research, Tech. Rep., 2014.
- [53] C. Geyer and K. Daniilidis
A unifying theory for central panoramic systems and practical implications,
in *European Conference on Computer Vision (ECCV)*, 2000, pp. 445–461.
- [54] P. J. Barreto and H. Araújo
Issues on the geometry of central catadioptric image formation,
in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2001, pp. 422–427.
- [55] H. W. Kuhn
The Hungarian method for the assignment problem,
Naval Research Logistics Quarterly, vol. 2, pp. 83–97, 1955.
- [56] J. Munkres
Algorithms for the assignment and transportation problems,
Journal of the Society for Industrial and Applied Mathematics, vol. 5, no. 1, pp. 32–38, 1957.
- [57] D. Mušicki and R. Evans
Joint Integrated Probabilistic Data Association—JIPDA,
IEEE Transactions on Aerospace and Electronic Systems, vol. 40, no. 3, pp. 1093–1099, 2004.
- [58] Y. Bar-Shalom and X. R. Li
Multitarget-Multisensor Tracking: Principles and Techniques.
Storrs, CT: YBS Publishing, 1995.
- [59] M. Murase, S. Yamamoto, J.-M. Valin, K. Nakadai, K. Yamada, K. Komatani, T. Ogata, and H. G. Okuno
Multiple moving speaker tracking by microphone array on a mobile robot,
in *INTERSPEECH*, 2005, p. 4.
- [60] A. A. Mekonnen, C. Briand, F. Lerasle, and A. Herbulot
Fast HOG based person detection devoted to a mobile robot with a spherical camera,
in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, nov 2013, pp. 631–637.
- [61] K. Bernardin and R. Stiefelhagen
Evaluating multiple object tracking performance: The CLEAR MOT metrics,
EURASIP Journal on Image and Video Processing, vol. 2008, p. 10, 2008.
- [62] D. Schuhmacher, B. T. Vo, and B. N. Vo
A consistent metric for performance evaluation of multi-object filters,
IEEE Transactions on Signal Processing, vol. 56, pp. 3447–3457, 2008.
- [63] B. Ristic, B.-N. Vo, D. Clark, and B.-T. Vo
A metric for performance evaluation of multi-target tracking algorithms,
IEEE Transactions on Signal Processing, vol. 59, no. 7, pp. 3452–3457, 2011.
- [64] T. Vu and R. Evans
A new performance metric for multiple target tracking based on optimal subpattern assignment,
in *International Conference on Information Fusion (FUSION)*, 2014, p. 8.



Ivan Marković received the B.Sc. and the Ph.D. degree in Electrical Engineering from the University of Zagreb, Faculty of Electrical Engineering and Computing (FER), Croatia in 2008 and 2014, respectively. He is currently employed at the FER, Zagreb, as a Postdoctoral fellow funded by the Ministry of Science, Education and Sport, Republic of Croatia. In 2014 he was awarded with the Silver Plaque “Josip Loncar” faculty award for outstanding doctoral dissertation and particularly successful scientific research. During his undergraduate studies, for outstanding academic achievements, he received the “Institute for Nuclear Technology Award” and the “Josip Loncar Award” in 2007 and 2006, respectively. In 2013 and 2014 he was a visiting researcher at Inria Rennes-Bretagne Atlantique in the Lagadic group directed by François Chaumette. He also serves as the coordinator of the technical editing team of the *Automatika* journal. His research interests are mobile robotics, estimation theory, especially detection and tracking of moving objects and speaker localization.



Mario Bukal received the M.S. degree in applied mathematics from the University of Zagreb, Faculty of Science in 2008 and the Ph.D. in applied mathematics from the Vienna University of Technology, Faculty of Mathematics and Geoinformation in 2012. From 2012 he is with the University of Zagreb, Faculty of Electrical Engineering and Computing, where he is a postdoctoral researcher working on several research projects. His research focus is in applied mathematics, in particular, mathematical and numerical analysis of diffusion equations and cross-diffusion systems, and homogenization and dimension reduction in nonlinear elasticity theory, but his research interests also include information fusion with applications to robotic systems.



Josip Česić received his B.Sc. and M.Sc. degree (Magna Cum Laude) in electrical engineering and information technology from University of Zagreb, Faculty of Electrical Engineering and Computing (FER) in 2011 and 2013, respectively. He finished the third semester of his master program at Chalmers University of Technology, Sweden, as an exchange student. He received several scholarships and awards for outstanding achievements during his undergraduate and graduate studies. He works as a research fellow at the Department of Control and Computer Engineering at FER since April 2013. His main research interests are in the areas of autonomous systems and mobile robotics, estimation theory and sensor processing.



Ivan Petrović received B.Sc. degree in 1983, M.Sc. degree in 1989 and Ph.D. degree in 1998, all in Electrical Engineering from the Faculty of Electrical Engineering and Computing (FER), University of Zagreb, Croatia. He had been employed as an R&D engineer at the Institute of Electrical Engineering of the Končar Corporation in Zagreb from 1985 to 1994. Since 1994 he has been with FER, where he is currently a full professor. He teaches a number of undergraduate and graduate courses in the field of control systems and mobile robotics. His research interests include various advanced control strategies and their applications to control of complex systems and mobile robots navigation. He has published more than 40 journal and 160 conference papers, and results of his research have been implemented in several industrial products. He is a member of IEEE, IFAC—TC on Robotics and FIRA—Executive committee. He is a member of the Croatian Academy of Engineering and Editor-in-Chief of the *Automatika* journal.

Density Estimation on the Rotation Group Using Diffusive Wavelets

NICOLAS LE BIHAN
JULIEN FLAMANT
JONATHAN H. MANTON

This paper considers the problem of estimating probability density functions on the rotation group $SO(3)$. Two distinct approaches are proposed, one based on characteristic functions and the other on wavelets using the heat kernel. Expressions are derived for their Mean Integrated Squared Errors. The performance of the estimators is studied numerically and compared with the performance of an existing technique using the De La Vallée Poussin kernel estimator. The heat kernel wavelet approach appears to offer the best compromise, with faster convergence to the optimal bound and guaranteed positivity of the estimated probability density function.

Manuscript received February 18, 2016; released for publication February 24, 2016.

Refereeing of this contribution was handled by Gerhard Kurz.

Authors' address: Nicolas Le Bihan: CNRS / Gipsa-Lab, 11 Rue des mathématiques, Grenoble Campus, BP 46, 38402 Saint Martin d'Heres Cedex, France (E-mail: nicolas.le-bihan@gipsa-lab.grenoble-inp.fr); Julien Flamant: Univ. Lille, Centre National de la Recherche Scientifique, Centrale Lille, UMR 9189, CRIStAL, Centre de Recherche en Informatique Signal et Automatique de Lille, 59000 Lille, France (E-mail: julien.flamant@phd.ec-lille.fr); Jonathan Manton: Department of Electrical and Electronic Engineering, The University of Melbourne, Victoria 3010, Australia (E-mail: jmanton@unimelb.edu.au).

Nicolas Le Bihan's research was supported by the REA, European Commission, through the International Outgoing Fellowship (IOF GeoSToSip 326176) program of the 7th PCRD.

1557-6418/16/\$17.00 © 2016 JAIF

I. INTRODUCTION

Statistics on Lie groups have become more and more popular in the last decade. Applications can be found in shape statistics [1], [2], medical imaging [3], multiple scattering processes [4], [5], crystallography [6], robotics and mechanics [7], [8] and many others. In directional statistics, the abundance of datasets taking values on spheres and homogeneous spaces has also motivated the study of random variables and processes on Lie groups [9]. Amongst all the matrix Lie groups, the most popular one is the rotation group in \mathbb{R}^3 , i.e., the Special Orthogonal group $SO(3)$. This is due to its predominant use in engineering problems [7], [10].

Even though the number of engineering challenges including random rotations has grown dramatically in the last years, the problem of probability density estimation for such variables was only considered in details recently in [6]. In parallel, the concept of wavelets on manifolds was transposed to the case of the rotation group, due to its relation to the 2-sphere which has attracted a lot of work since the 90s [11], [12], [13]. In [14], authors introduce diffusive wavelets on manifolds. The definition of a wavelet transform on a surface/manifold is conditioned by the possible definition of two operations on the manifold: scaling and translation. While translation is easily defined on Lie groups as it is based on the group action, the definition of scaling is less obvious. In [14], authors have chosen an intrinsic definition for scaling, whereas some extrinsic approaches had been proposed previously in [15]. The difference in these definitions resides in the way the mother wavelet is scaled: it is either firstly projected in the tangent plane before scaling and back projection (extrinsic); or scaled on the manifold directly (intrinsic). In this paper, we will make use of the intrinsic approach and study the ability of diffusive wavelets to define interesting estimators for densities on $SO(3)$.

The specificity of diffusive wavelets is that they are based on the heat kernel. As this kernel can be defined on manifolds, the diffusive wavelet approach overcomes the problem of "scaling" on manifold. Note that this definition problem was already pointed out by many authors [12], [15] when defining wavelets on the 2-sphere for example.

In this paper, we present an estimation technique for densities on $SO(3)$ based on the diffusive wavelet transform. Using wavelet estimators (linear or thresholded) is well known in non parametric estimation [16]. We propose the use of the linear¹ wavelet estimator to the case of $SO(3)$ -valued random variables and give some of its properties. In particular, the wavelet-based estimator is a characteristic kernel estimator. A comparison with other types of estimators (characteristic function and kernel) is provided.

¹Linear wavelet estimation refers to the standard estimation through wavelet coefficients estimation, as opposed to thresholded wavelet coefficients estimator sometimes referred as nonlinear [16].

The rest of the paper is organized as follows: Section II is dedicated to the presentation of harmonic analysis on $SO(3)$ and various concepts about random variables on this Lie group. In Section III, the diffuse wavelet formalism is introduced for the case of the rotation group. It is then used in Section IV to propose a probability density function estimator. Finally, Section V presents simulation results and comparison of the wavelet estimator with two other estimators.

II. CHARACTERISTIC FUNCTIONS FOR RANDOM VARIABLES ON $SO(3)$

We first review some basic concepts on the rotation group $SO(3)$ and random variables taking values on this well known Lie group. The concept of characteristic function for $SO(3)$ -valued random variables is of importance as it can be identified as the ‘‘Fourier transform’’ of the probability density of the random variable. It will also be of use in Section IV to provide a probability density estimator. The results presented here can be found in various textbooks, for example [17], [18].

A. The rotation group $SO(3)$

The set of rotations in the 3D space forms a compact Lie group denoted $SO(3)$. An element $x \in SO(3)$ can be parametrized in several ways [17, Chap. 3]. Using the so-called ZYZ convention of Euler angles parametrization, any element $x \in SO(3)$ can be associated to a matrix $R^x = R^x(\varphi, \theta, \psi)$ with $0 \leq \varphi, \psi < 2\pi$ and $0 \leq \theta \leq \pi$. With this convention, $R^x(\varphi, \theta, \psi)$ takes the form:

$$R^x(\varphi, \theta, \psi) = \begin{bmatrix} \cos \varphi & -\sin \varphi & 0 \\ \sin \varphi & \cos \varphi & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \cos \theta & 0 & \sin \theta \\ 0 & 1 & 0 \\ -\sin \theta & 0 & \cos \theta \end{bmatrix} \times \begin{bmatrix} \cos \psi & -\sin \psi & 0 \\ \sin \psi & \cos \psi & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

It is also possible to parametrize an element of $SO(3)$ in terms of its *rotation axis* and its *rotation angle* [17, Chap. 3]. The rotation axis η is a unit vector in \mathbb{R}^3 , i.e., $\eta \in \mathcal{S}^2$. The rotation angle, denoted $\omega(x)$, is given by:

$$\cos \omega(x) = \frac{\text{Tr}(R^x) - 1}{2} \quad (1)$$

Note that in this parametrization, the angle takes values in: $-\pi < \omega(x) \leq \pi$. It is also possible to express this angle $\omega(x)$ in terms of the Euler angles:

$$\omega(x) = 2 \arccos \left(\cos \frac{\theta}{2} \cos \frac{\varphi + \psi}{2} \right) \quad (2)$$

The *rotation angle* is of particular interest as it is a metric on $SO(3)$. In fact, one can define the distance between $x \in SO(3)$ and $y \in SO(3)$, denoted $d(x, y)$, as:

$$d(x, y) = |\omega(yx^{-1})| = \arccos \left(\frac{\text{Tr}(R^y(R^x)^{-1}) - 1}{2} \right) \quad (3)$$

In the sequel, we will make use of the notation $\mathbf{R}(\omega(x), \eta)$ for an element $x \in SO(3)$, keeping in mind that:

$$\mathbf{R}(\omega(x), \eta) = \exp(\omega(x)\mathbf{M}) \quad \text{with} \quad \mathbf{M} = \begin{pmatrix} 0 & -\eta_3 & \eta_2 \\ \eta_3 & 0 & -\eta_1 \\ -\eta_2 & \eta_1 & 0 \end{pmatrix} \quad (4)$$

where η_1, η_2, η_3 are the components of the vector η , i.e., $\eta = [\eta_1, \eta_2, \eta_3]^T$ and $\exp(\cdot)$ the matrix exponential [17].

B. Fourier series on $SO(3)$

We now consider the Fourier series expansion of functions taking values on the rotation group. Consider the set of square integrable functions $L^2(SO(3), \mathbb{R})$. By the Peter-Weyl theorem [18, Chap. III], a function $f \in L^2(SO(3), \mathbb{R})$ can be expressed as:

$$f(x) = \sum_{\ell \geq 0} \sum_{n=-\ell}^{+\ell} \sum_{m=-\ell}^{+\ell} (2\ell + 1) \hat{f}_{nm}^\ell D_{nm}^\ell(x) \quad (5)$$

for $x \in SO(3)$ and where $D_{nm}^\ell(x)$ are the Wigner-D functions [19] evaluated at position x . This infinite series expansion over ℓ has matrix coefficients \hat{f}^ℓ of dimension $(2\ell + 1) \times (2\ell + 1)$ with elements \hat{f}_{nm}^ℓ . The matrix entries \hat{f}_{nm}^ℓ are obtained by projection of f on the Wigner-D functions:

$$\hat{f}_{nm}^\ell = \langle f, D_{nm}^\ell \rangle_{SO(3)} = \int_{SO(3)} f(x) \overline{D_{nm}^\ell(x)} d\mu(x) \quad (6)$$

where $\langle f, h \rangle_{SO(3)}$ is the scalar product on $L^2(SO(3), \mathbb{R})$ and $d\mu(x)$ the bi-invariant Haar measure on $SO(3)$, i.e., $d\mu(x) = (8\pi^2)^{-1} \sin \theta d\varphi d\theta d\psi$ when using the ZYZ Euler angle parametrization (φ, θ, ψ) for elements $x \in SO(3)$. We also mention that the corresponding norm is:

$$\|f\|_2 = \sqrt{\langle f, f \rangle_{SO(3)}} \quad (7)$$

With the chosen parametrization of $SO(3)$, the previously introduced Wigner-D functions D_{nm}^ℓ take the form:

$$D_{nm}^\ell(\varphi, \theta, \psi) = e^{-in\varphi} P_{nm}^\ell(\cos \theta) e^{-im\psi} \quad (8)$$

where $P_{nm}^\ell(\cos \theta)$ are the generalized Legendre polynomials. We have used the fact that the Wigner-D functions D_{nm}^ℓ form a complete set of orthogonal functions:

$$\langle D_{nm}^\ell, D_{n'm'}^{\ell'} \rangle_{SO(3)} = \frac{1}{(2\ell + 1)} \delta_{nn'} \delta_{mm'} \delta_{\ell\ell'} \quad (9)$$

This shows clearly that the set of functions $\left\{ \sqrt{2\ell + 1} D_{nm}^\ell, \ell \geq 0, -\ell \leq n, m \leq \ell \right\}$ form an orthonormal basis for functions in $L^2(SO(3), \mathbb{R})$, allowing the decomposition of functions taking values on the rotation group using this basis. In the sequel, we may refer to \hat{f}_{nm}^ℓ (and abusively to \hat{f}^ℓ) as the *Fourier coefficients* of f .

C. Parseval identity

It is also well known that the Parseval identity holds for functions $f \in L^2(SO(3), \mathbb{R})$. Using our normalization convention, the following is true:

$$\|f\|_2^2 = \langle f, f \rangle_{SO(3)} = \sum_{\ell \geq 0} \sum_{n=-\ell}^{\ell} \sum_{m=-\ell}^{\ell} (2\ell+1) |\hat{f}_{nm}^\ell|^2 \quad (10)$$

The left hand side of equation (10) is commonly referred to as the *energy* and the Parseval identity simply states that the energy of f consists in the infinite sum of its modulus squared *Fourier coefficients*.

D. Zonal functions on $SO(3)$

Zonal functions, sometimes also called *conjugate invariant* functions, are radially symmetric functions with center $g_0 = e \in SO(3)$ where e is the identity element in $SO(3)$. See for example [20] for more details on radially symmetric functions on $SO(3)$. They will be at the heart of the diffusive wavelets construction in Section III. A function $f : SO(3) \rightarrow \mathbb{R}$ is called *zonal* iff it satisfies $\forall x, y \in SO(3)$:

$$f(yxy^{-1}) = f(x)$$

Equivalently, $f : SO(3) \rightarrow \mathbb{R}$ is *zonal* iff $f(x) = f(x')$ $\forall x, x' \in SO(3)$ such that $\omega(x) = \omega(x')$, i.e., x and x' have the same *angle*. In fact, it simply means that f , evaluated at $x \in SO(3)$, only depends on the rotation angle of x introduced in (1). It is known [6], [20] that the subspace of zonal functions is spanned by functions χ^ℓ , $\ell \in \mathbb{N}$, called the *characters* of $SO(3)$, and given as:

$$\chi^\ell(x) = \sum_{n=-\ell}^{\ell} D_{nn}^\ell(x) = \frac{\sin((\ell + \frac{1}{2})\omega(x))}{\sin \frac{\omega(x)}{2}} \quad (11)$$

$$= \mathcal{U}_{2\ell} \left(\cos \frac{\omega(x)}{2} \right) \quad (12)$$

where $\mathcal{U}_{2\ell}(\cdot)$ are the Chebychev polynomials of second kind and of (even) degree 2ℓ . For these (even degree) polynomials, the following orthogonality relation stands:

$$\int_{-1}^1 \mathcal{U}_{2\ell}(t) \mathcal{U}_{2\ell'}(t) \sqrt{1-t^2} dt = \frac{\pi}{2} \delta_{\ell\ell'} \quad (13)$$

In terms of the characters χ^ℓ , and using the notation $\omega(x) = \omega$ for simplicity, this integral becomes:

$$\int_0^{2\pi} \chi^\ell(\omega) \chi^{\ell'}(\omega) \sin^2 \left(\frac{\omega}{2} \right) d\omega = \pi \delta_{\ell\ell'} \quad (14)$$

where the link with the orthogonality relation for $\mathcal{U}_{2\ell}$ comes with $t = \cos(\omega/2)$. As a consequence, the characters χ^ℓ fulfill the following orthogonality relation:

$$\langle \chi^\ell, \chi^{\ell'} \rangle_{SO(3)} = \frac{1}{4\pi^2} \int_{S^2} d\eta \int_0^{2\pi} \chi^\ell(\omega) \chi^{\ell'}(\omega) \sin^2 \left(\frac{\omega}{2} \right) d\omega = \delta_{\ell\ell'} \quad (15)$$

where we used the expression of the Haar measure in terms of axis η and rotation angle ω for $SO(3)$. With this convention the Haar measure takes the form $d\xi = (1/4\pi^2) d\eta \sin^2(\omega/2) d\omega$. As a consequence, a zonal function $f \in L^2(SO(3), \mathbb{R})$ has a Fourier series expansion that can be written like:

$$f(x) = \sum_{\ell \geq 0} (2\ell+1) \hat{f}^\ell \chi^\ell(x) \quad (16)$$

where its Fourier coefficients \hat{f}^ℓ . These Fourier coefficients \hat{f}^ℓ are simply given by:

$$\hat{f}^\ell = \frac{1}{(2\ell+1)} \langle f, \chi^\ell \rangle_{SO(3)} = \frac{1}{\pi} \frac{1}{(2\ell+1)} \int_0^{2\pi} f(\omega) \chi^\ell(\omega) \sin^2 \left(\frac{\omega}{2} \right) d\omega \quad (17)$$

The Parseval identity for a zonal function f now reads:

$$\|f\|_2^2 = \sum_{\ell \geq 0} (2\ell+1)^2 |\hat{f}^\ell|^2 \quad (18)$$

Zonal functions will be used in Section III in the definition of wavelets on $SO(3)$.

E. Convolution

A very important feature of Fourier transformation is its behaviour with respect to the convolution product. In the case of functions taking values on $SO(3)$, the relation still holds. First, given $f, h \in L^1(SO(3), \mathbb{R})$, their convolution product is defined as:

$$(f * h)(x) = \int_{SO(3)} f(g) h(g^{-1}x) d\mu(g) \quad (19)$$

where the group operation stands naturally for translation. Now, if the Fourier coefficients of f and h are respectively \hat{f}^ℓ and \hat{h}^ℓ (in matrix format), then the following stands:

$$\widehat{f * h}^\ell = \hat{f}^\ell \hat{h}^\ell \quad (20)$$

Note that the right-hand side of the equation is a matrix product.

F. Characteristic function of $SO(3)$ -valued random variables

The characteristic function of a random variable is the Fourier transform of its probability transform. This well known result extends to random variables on the rotation group, thanks to the results on Fourier series expansion introduced in the previous Section. We now give some definitions and properties for characteristic functions of random variables on $SO(3)$ as it will be used to define a probability density function estimator in Section IV.

Consider the case of a random variable $X \in SO(3)$ with density f . Such random variables X taking values on the rotation group $SO(3)$ can simply be thought of as

random rotation matrices of dimension 3×3 , classically parametrized by Euler angles. Alternatively, one can think of these random variables as unit quaternions from the upper hemisphere of \mathcal{S}^3 .

1) *Definition:* Given a rotation random variable X with density f_X , the sequence $\Phi_X = \{\Phi_X(\ell)\}_{\ell \geq 0}$ of $(2\ell + 1) \times (2\ell + 1)$ matrices given by:

$$\Phi_X(\ell) = \mathbb{E}[\overline{D^\ell(X)}] \quad (21)$$

is the *characteristic function* of X . The elements of the matrix $\Phi_X(\ell)$ are denoted Φ_{nm}^ℓ and they read as:

$$\Phi_{nm}^\ell = \int_{SO(3)} f_X(x) \overline{D_{nm}^\ell(x)} d\mu(x) \quad (22)$$

Thus the density f has the following Fourier series expansion:

$$f_X(x) = \sum_{\ell \geq 0} \sum_{n,m=-\ell}^{\ell} (2\ell + 1) \Phi_{nm}^\ell D_{nm}^\ell(x) \quad (23)$$

One can see by comparison with Equation (5) that the characteristic function of f_X is the Fourier transform of its density f_X , i.e., $\Phi_X^\ell = \hat{f}_X^\ell$.

2) *Basic properties:* The following properties can easily be verified:

- Given two rotation random variables X and Y , then:

$$X = Y \quad \text{iff} \quad \Phi_X = \Phi_Y$$

- If X and Y are two independent rotation random variables and $Z = XY$, then:

$$\Phi_Z^\ell = \Phi_X^\ell \Phi_Y^\ell$$

- A rotation random variable $U \in SO(3)$ is uniformly distributed iff:

$$\Phi_U^\ell = 0 \quad \forall \ell > 0$$

- Consider n *i.i.d.* rotation random variables X_1, X_2, \dots, X_n , then the random variable consisting in the accumulated products of the X_n , denoted $Y = X_1 X_2 \dots X_n$ has the following characteristic function:

$$\Phi_Y^\ell = [\Phi_X^\ell]^n$$

where Φ_X^ℓ is the characteristic function shared by the X_n .

3) *Zonal invariance:* A rotation random variable X with density f_X is called *zonal invariant* if:

$$X \stackrel{d}{=} R X R^{-1} \quad \text{for all } R \in SO(3)$$

Now, if X is *zonal invariant*, then its characteristic function is:

$$\Phi_X^\ell = a_\ell I_\ell$$

where $a_\ell \in \mathbb{R}$ and I_ℓ is the $(2\ell + 1) \times (2\ell + 1)$ identity matrix. As a consequence, if X is *zonal invariant*, then its density f_X takes the form:

$$f_X(x) = \sum_{\ell \geq 0} (2\ell + 1) a_\ell \chi^\ell(x) \quad (24)$$

as detailed previously when considering zonal functions on $SO(3)$ in II-D.

This last expression shows that the characteristic function of *zonal invariant* random rotation variables are scalar coefficients a_ℓ , as opposed to matrix coefficients for random variables with no symmetries.

III. DIFFUSIVE WAVELETS ON $SO(3)$

We now introduce diffusive wavelets on the rotation group. They will be used in Section IV to propose a probability density estimator. Recently, Ebert and Wirth [21] introduced *diffusive wavelets* on groups. We detail in the sequel the special case of the rotation group in 3D, i.e., $SO(3)$.

A. Heat wavelet family on $SO(3)$

Here, we follow the construction given by Ebert [21]. First recall that on $SO(3)$ the heat kernel is given by:

$$\kappa_\rho(x) = \sum_{\ell \geq 0} (2\ell + 1) e^{-\ell(\ell+1)\rho} \chi^\ell(x) \quad (25)$$

for $x \in SO(3)$, and where χ^ℓ are the irreducible characters of $SO(3)$ introduced previously in Section II. Note that the Fourier series expansion of κ_ρ exhibits the fact that κ_ρ is a zonal function with Fourier coefficients $\hat{\kappa}_\rho^\ell = e^{-\ell(\ell+1)\rho}$.

Now, we introduce some of the properties of the heat kernel. First, $\kappa_\rho(x)$ is an *approximate identity* as it fullfils the following properties:

- $\|\hat{\kappa}_\rho^\ell\| \leq C \quad \forall \ell \in \mathbb{R}^+$
- $\lim_{\rho \rightarrow 0} \hat{\kappa}_\rho^\ell = Id \quad \forall \ell$
- $\lim_{\rho \rightarrow \infty} \hat{\kappa}_\rho^\ell = 0 \quad \forall \ell > 0$
- $-(\partial/\partial\rho)\hat{\kappa}_\rho^\ell$ is a symmetric positive definite matrix for all $\rho > 0$ and $\ell \geq 0$

where Id denotes the identity operator. As the heat kernel κ_ρ is an approximate identity, it follows that:

$$\kappa_\rho * h \xrightarrow{\rho \rightarrow 0} h \quad \forall h \in L^2(SO(3))$$

As detailed in [21], a family of wavelets corresponding to the heat kernel on $SO(3)$ is consequently of the form:

$$\Psi_\rho(x) = \frac{1}{\sqrt{\alpha(\rho)}} \sum_{\ell \geq 0} (2\ell + 1) \sqrt{\ell(\ell + 1)} e^{-\ell(\ell+1)/2\rho} \chi^\ell(x) \quad (26)$$

with $\alpha(\rho)$ a normalizing factor to be detailed below. Rephrasing this equation in terms of Fourier transform,

noticing that the heat wavelet family is zonal, one can write that:

$$\Psi_\rho(x) = \sum_{\ell \geq 0} (2\ell + 1) \hat{\Psi}_\rho \chi^\ell(x) \quad (27)$$

where the Fourier coefficients of Ψ_ρ are:

$$\hat{\Psi}_\rho = \frac{1}{\sqrt{\alpha(\rho)}} \sqrt{\ell(\ell+1)} e^{-\ell(\ell+1)/2\rho} \quad (28)$$

As can be seen in Equation (26), the choice for the normalization coefficients $\alpha(\rho)$ has to be made. In order to have a normalized/unitary wavelet family, we can use the Parseval identity to choose $\alpha(\rho)$. In fact, by (18), we have that:

$$\|\Psi_\rho\|_2^2 = \sum_{\ell \geq 0} (2\ell + 1)^2 |\hat{\Psi}_\rho|^2 \quad (29)$$

and imposing that $\|\Psi_\rho\|_2^2 = 1 \ \forall \rho$, one obtains that:

$$\alpha(\rho) = \sum_{\ell \geq 0} (2\ell + 1)^2 \ell(\ell+1) e^{-\ell(\ell+1)\rho} \quad (30)$$

This choice for $\alpha(\rho)$ will be made throughout the rest of the paper. Notice also that the heat wavelet family $\Psi_\rho \in L^2(SO(3))$ is a *diffusive wavelet family* as it satisfies the *admissibility condition*:

$$\begin{aligned} \kappa_\rho(x) &= \int_\rho^{+\infty} (\check{\Psi}_t * \Psi_t)(x) \alpha(t) dt \\ &= \int_\rho^{+\infty} \int_{SO(3)} \check{\Psi}_t(g) \Psi_t(g^{-1}x) d\mu(g) \alpha(t) dt \end{aligned} \quad (31)$$

in which $\kappa_\rho(x)$ is an diffusive approximate convolution identity and where we used the notation $\check{\Psi}_\rho(g) = \overline{\Psi_\rho(g^{-1})}$. The fact that $\kappa_\rho(x)$ is an approximate convolution identity ensures that the wavelet transform can be inverted. Note that the approximate convolution identity can be expressed using the Fourier coefficients. In the case of κ_ρ , it means that:

$$\lim_{\rho \rightarrow 0} \hat{\kappa}_\rho^\ell = 1 \quad \forall \ell$$

This means that an approximate convolution identity is characterized by constant Fourier coefficients in the limit.

With the diffusive wavelets introduced, we can now introduce the wavelet transform for functions on the rotation group.

B. Wavelet transform on $SO(3)$

Recall that we are interested in estimating probability density functions of $SO(3)$ -valued random variables using the diffusive wavelet transform. In the sequel, we consider probability density functions f belonging to $L^2(SO(3), \mathbb{R}) \cap L^1(SO(3), \mathbb{R})$ with the additional condition that $\int_{SO(3)} f d\mu(g) = 1$ and that they are non-negative. The diffusive wavelet transform for such densities is as follows. Given a diffusive wavelet family

$\Psi_\rho \in L^1(SO(3))$ as defined in III-A, then the Wavelet Transform (WT) of $f \in L^2(SO(3), \mathbb{R}) \cap L^1(SO(3), \mathbb{R})$ is:

$$WT : \begin{array}{ccc} f & \longrightarrow & WT_f \\ SO(3) & & \mathbb{R}^+ \times SO(3) \end{array}$$

with the following expression:

$$WT_f(\rho, g) = (f * \check{\Psi}_\rho)(g) = \int_{SO(3)} f(x) \check{\Psi}_\rho(x^{-1}g) d\mu(x) \quad (32)$$

and where we made use of the notation $\check{\Psi}(x) = \overline{\Psi(x^{-1})}$. Using scalar product on the rotation group, this expression can be written like:

$$WT_f(\rho, g) = \int_{SO(3)} f(x) \overline{\Psi_\rho(g^{-1}x)} d\mu(x) = \langle f, T_g^* \Psi_\rho \rangle_{SO(3)} \quad (33)$$

with T_g^* the following operator: $T_g^* : \Psi \rightarrow \Psi(g^{-1}\cdot)$. Equation (33) is a very general definition of wavelet transform on $SO(3)$. In the sequel, we will only make use of the heat wavelet family given in (26) to define the heat wavelet transform. One of the interesting properties of the wavelet transform is that it is invertible. The density f can thus be reconstructed in the following way:

$$\begin{aligned} f(x) &= \int_{\mathbb{R}^+} \int_{SO(3)} WT_f(t, g) \Psi_t(g^{-1}x) d\mu(g) \alpha(t) dt \\ &= \int_{\mathbb{R}^+} \int_{SO(3)} (f * \check{\Psi}_t)(g) \Psi_t(g^{-1}x) d\mu(g) \alpha(t) dt \\ &= \int_{\mathbb{R}^+} (f * \check{\Psi}_t * \Psi_t)(x) \alpha(t) dt \\ &= f * \int_{\mathbb{R}^+} (\check{\Psi}_t * \Psi_t)(x) \alpha(t) dt \\ &= f * \int_{\rho \rightarrow 0}^{+\infty} (\check{\Psi}_t * \Psi_t)(x) \alpha(t) dt \\ &= \lim_{\rho \rightarrow 0} (f * \kappa_\rho)(x) \\ &= f(x) \end{aligned} \quad (34)$$

where the last equality is obtained thanks to the fact that $\kappa_\rho(x)$ is a convolution identity. This can be verified in the Fourier domain, when denoting \hat{f}^ℓ the Fourier coefficients of f , by:

$$\begin{aligned} \lim_{\rho \rightarrow 0} (f * \kappa_\rho)(x) &= \lim_{\rho \rightarrow 0} \sum_{\ell \geq 0} (2\ell + 1) e^{-\ell(\ell+1)\rho} \hat{f}^\ell \chi^\ell(x) \\ &= \sum_{\ell \geq 0} (2\ell + 1) \hat{f}^\ell \chi^\ell(x) \\ &= f(x) \end{aligned} \quad (35)$$

An interesting property of the heat wavelet transform as defined above is that it is *unitary*. This property reads

$$\langle WT_{f_1}(\rho, g), WT_{f_2}(\rho, g) \rangle_{\mathbb{R}^+ \times SO(3)} = \langle f_1, f_2 \rangle_{SO(3)} \quad (36)$$

where we use the notation $\langle \dots \rangle_{\mathbb{R}^+ \times SO(3)}$ for the scalar product between wavelet transforms. This unitary condition makes possible to compare densities in the wavelet domain for example.

We have introduced the heat wavelet transform (also called diffusive wavelet transform) for probability densities of rotation random variables. In the next Section, we will make use of this transform to define an estimator for densities on $SO(3)$.

IV. ESTIMATION

In this section, we consider the problem of estimating the probability density function of a $SO(3)$ -valued random variable given an independent sample of size K : $\{X_1, X_2, \dots, X_K\}$. After presenting kernel estimators as defined in [6], we show that the characteristic function estimator and the heat wavelet estimator are actually kernel estimators. We provide the MISE of each of them and explain their differences.

A. Kernel estimators on $SO(3)$

First, we recall some of the results given in [6] for kernel estimators of probability density functions on $SO(3)$.

1) *Definition*: A kernel estimator, with kernel $\Xi \in L^2(SO(3))$ has the following expression:

$$\zeta_K(x) = \frac{1}{K} \sum_{k=1}^K \Xi(X_k^{-1}x) \quad (37)$$

This definition is the extension of the classical kernel estimator known for densities of random variables on the line (see [16] for details). Once again, we emphasize that translation is made through the group action. Before introducing different types of kernels $\Xi(\cdot)$, we introduce how to study their estimation performances.

2) *Bias and variance*: In order to characterize the kernel estimator, we provide here the expression of its Mean Integrated Square Error (MISE). This expression is general for kernel estimators and will be of use later to analyze the behaviour of the characteristic function and the wavelet estimators.

In the sequel, we will make use of the following notation in order to distinguish the different density estimator. Every estimator, based on a simple sample of size K , will be denoted $\zeta_K^\Delta(\cdot)$ with the superscript Δ taking the following values: $\Delta = \alpha$ for the ‘‘general’’ kernel estimator, $\Delta = \beta$ for the characteristic function estimator and $\Delta = \gamma$ for the heat wavelet estimator. Associate quantities will exhibit the α, β, γ values when needed.

As known in the classical case [16] and as given for the $SO(3)$ case in [6], the MISE of the kernel estimator $\zeta_K^\alpha(x)$ is made of a bias and a variance term:

$$\text{MISE}(\zeta_K^\alpha(x)) = \|f - \mathbb{E}[\zeta_K^\alpha]\|_2^2 + \mathbb{E}[\|\zeta_K^\alpha(x) - \mathbb{E}[\zeta_K^\alpha(x)]\|_2^2] \quad (38)$$

with the following fact:

$$\begin{aligned} \mathbb{E}[\zeta_K^\alpha(x)] &= \mathbb{E}\left[\frac{1}{K} \sum_{k=1}^K \Xi(X_k^{-1}x)\right] = \frac{1}{K} \sum_{k=1}^K \mathbb{E}[\Xi(X_k^{-1}x)] \\ &= \frac{1}{K} \sum_{k=1}^K \int_{SO(3)} \Xi(y^{-1}x) f(y) d\mu(y) = (\Xi * f)(x) \end{aligned} \quad (39)$$

This means that the mean value of the kernel consists in the convolution of the density with the kernel. One can also express the bias b_α and the variance σ_α^2 as follows:

$$b_\alpha^2 = \|f - \mathbb{E}[\zeta_K^\alpha]\|_2^2 = \|f - f * \Xi\|_2^2 \quad (40)$$

and

$$\sigma_\alpha^2 = \mathbb{E}[\|\zeta_K^\alpha(x) - \mathbb{E}[\zeta_K^\alpha(x)]\|_2^2] = \frac{1}{K} (\|\Xi\|_2^2 - \|\Xi * f\|_2^2) \quad (41)$$

Also, it is interesting to note that the following property holds:

$$\lim_{K \rightarrow \infty} \zeta_K^\alpha(x) = (\Xi * f)(x) \quad (42)$$

As a consequence, the MISE for the kernel estimator is:

$$\text{MISE}(\zeta_K^\alpha) = \|f - f * \Xi\|_2^2 + \frac{1}{K} (\|\Xi\|_2^2 - \|\Xi * f\|_2^2) \quad (43)$$

This general expression is useful for the study of kernel estimator. Using results from representation theory (Fourier series expansion), it is also possible to write the MISE in terms of the Fourier coefficients of Ξ and f . Remembering the Parseval identity and the convolution property of the Fourier series expansion on $SO(3)$, one gets, just like in [6] but with a slight $(2\ell + 1)$ factor due to our normalization choice, the following representation:

$$\begin{aligned} \text{MISE}(\zeta_K^\alpha) &= \sum_{\ell \geq 1} \left((2\ell + 1) (\hat{f}^\ell)^2 (1 - \hat{\Xi}^\ell)^2 + \frac{(2\ell + 1)}{K} (\hat{\Xi}^\ell)^2 (1 - \hat{f}^{\ell 2}) \right) \end{aligned} \quad (44)$$

where the following notation was used for the Fourier coefficients of the probability density function f :

$$(\hat{f}^\ell)^2 = (2\ell + 1)^{-1} \sum_{n,m=-\ell}^{\ell} |\hat{f}_{n,m}^\ell|^2 \quad (45)$$

This expression of the MISE of a kernel estimator on $SO(3)$ will be used with a specific kernel, namely the De La Vallée Poussin kernel, in Section V for comparison purposes with the characteristic function estimator and the wavelet estimator.

B. Characteristic function estimator on $SO(3)$

The characteristic function estimator presented here was used for example in [5] to provide a non-parametric

estimation of the density of a compound Poisson process on the rotation group $SO(3)$. Recall that our notation is $\zeta_K^\beta(\cdot)$ for this estimator. It is shown here that this estimator is in fact a kernel estimator and use can be made of the MISE expressions given in Section III-A.2. Note that we consider the truncated version of the estimator, with $\ell_{max} = L$, given by:

$$\zeta_K^\beta(x) = \sum_{\ell=0}^L \sum_{n=-\ell}^{\ell} \sum_{m=-\ell}^{\ell} (2\ell+1) \widetilde{f}_{nm}^\ell D_{nm}^\ell(x) \quad (46)$$

and where:

$$\widetilde{f}_{nm}^\ell = \frac{1}{K} \sum_{k=1}^K \overline{D_{nm}^\ell(X_k)} \quad (47)$$

The use of truncated version of the Fourier expansion is necessary for obvious computational reasons. Such a truncation has effects on the performance of the estimator as it only converges to a low-resolution version of the density. However, it is well-adapted to naturally band-limited functions. In the simulation in Section V, we will investigate the influence of the bandwidth on the estimation performances.

The definition of the characteristic kernel estimator leads by simple calculation to:

$$\zeta_K^\beta(x) = \frac{1}{K} \sum_{k=1}^K \sum_{\ell \geq 0} (2\ell+1) \chi^\ell(X_k^{-1}x) \quad (48)$$

This can be obtained thanks to the following property (see [22] for example):

$$\sum_{n=-\ell}^{\ell} \sum_{m=-\ell}^{\ell} \overline{D_{nm}^\ell(x)} D_{nm}^\ell(y) = \chi^\ell(x^{-1}y) \quad (49)$$

Note that this estimator indirectly estimates the density, as it is designed to estimate its Fourier coefficients. The estimator \widetilde{f}_{nm}^ℓ is the characteristic function estimator. From equation (48), one can see that the characteristic function estimator is a kernel estimator, with the specific kernel:

$$\Xi(\cdot) = \sum_{\ell \geq 0}^{\ell_{max}=L} (2\ell+1) \chi^\ell(\cdot)$$

From the expression of the MISE for a kernel estimator given in Equation (44), one can directly deduce the MISE for the characteristic function estimator. It suffices to replace the Fourier coefficients $\widehat{\Xi}^\ell$, remembering the linearity property of the Fourier expansion, by the sum of the unit coefficients up to $\ell_{max} = L$. This make the characteristic function estimator a very ‘‘simple’’ kernel with constant Fourier coefficients. Illustration of its behaviour compared to other estimators will be presented in Section V.

C. Linear diffusive wavelet estimator

Using the diffusive wavelet introduced in Section III-B, it is possible to build an estimator of the density f through its wavelet expansion. Recall that with our notation, the wavelet estimator is denoted $\zeta_K^\gamma(\cdot, \cdot)$. Note also that an extra scalar parameter is introduced when using the wavelet estimator. This is the scaling parameter of the wavelet transform. The estimator based on the wavelet transform consists in replacing the wavelet transform $WT_f(\rho, x)$ in the inversion formula given in Equation (34) by its empirical estimate, obtained from the data sample. The wavelet coefficients are thus estimated using:

$$\widetilde{WT}_f(\rho, x) = \frac{1}{K} \sum_{k=1}^K \check{\Psi}_\rho(X_k^{-1}x) \quad (50)$$

The estimated density $\zeta_K^\gamma(x)$ takes the following expression when plugging the estimated coefficients in the inversion formula:

$$\zeta_K^\gamma(x, t) = \frac{1}{K} \sum_{k=1}^K \int_t^{+\infty} \int_{SO(3)} \check{\Psi}_\rho(X_k^{-1}y) \Psi_\rho(y^{-1}x) d\mu(y) \alpha(\rho) d\rho \quad (51)$$

The ‘scaling’ coefficient ρ is a parameter for this estimator. Ideal range should be 0 to $+\infty$. Obviously it is not possible to reach the upper limit and one can only numerically tend to very high values. As can be seen in the Simulation section, it is not a limitation for the use of the estimator. In the sequel, we will only keep the $t \geq 0$ constraint and keep in mind the upper limit issue.

It can easily be shown that the wavelet estimator can be expressed using the kernel κ_t previously introduced in Equation (31). The following equality thus stands:

$$\begin{aligned} & \int_t^{+\infty} \int_{SO(3)} \check{\Psi}_\rho(X_k^{-1}g) \Psi_\rho(g^{-1}x) d\mu(g) \\ &= \int_t^{+\infty} \int_{SO(3)} \check{\Psi}_\rho(g') \Psi_\rho(g'^{-1}X_k^{-1}x) d\mu(g') \\ &= \int_t^{+\infty} (\check{\Psi}_\rho * \Psi_\rho)(X_k^{-1}x) \alpha(\rho) d\rho \\ &= \kappa_t(X_k^{-1}x) \end{aligned} \quad (52)$$

which is obtained by simple change of variable $g' = X_k^{-1}g$ and using the fact that the Haar measure is bi-invariant, inducing that $d\mu(X_k g') = d\mu(g')$. As a consequence, the linear wavelet estimator is:

$$\zeta_K^\gamma(x, t) = \frac{1}{K} \sum_{k=1}^K \kappa_t(X_k^{-1}x) \quad (53)$$

where one can obviously see that this estimator is a kernel estimator. In order to study this estimator, one can look at its MISE with the bias and variance term.

The bias term is:

$$\begin{aligned}
& \mathbb{E}[\zeta_K^\gamma(x, t)] \\
&= \mathbb{E} \left[\frac{1}{K} \sum_{k=1}^K \int_t^{+\infty} \int_{SO(3)} \check{\Psi}_\rho(X_k^{-1}y) \right. \\
&\quad \left. \times \Psi_\rho(y^{-1}x) d\mu(y) \alpha(\rho) d\rho \right] \\
&= \frac{1}{K} \sum_{k=1}^K \int_t^{+\infty} \int_{SO(3)} \int_{SO(3)} f(z) \check{\Psi}_\rho(z^{-1}y) \\
&\quad \times \Psi_\rho(y^{-1}x) d\mu(z) d\mu(y) \alpha(\rho) d\rho \\
&= f * \int_t^{+\infty} (\check{\Psi}_\rho \Psi_\rho)(x) \alpha(\rho) d\rho \\
&= (f * \kappa_t)(x) \tag{54}
\end{aligned}$$

Note that we have:

$$\lim_{t \rightarrow 0} \mathbb{E}[\zeta_K^\gamma(x, t)] = f(x) \tag{55}$$

Note that this is due to the fact that we used an approximate convolution identity. Thus, the bias of the linear wavelet estimator takes the following form:

$$b_{\gamma,t}^2 = \|f - \mathbb{E}[\zeta_K^\gamma(t)]\|_2^2 = \|f - f * \kappa_t\|_2^2 \tag{56}$$

It is noticeable that, asymptotically with t , the bias term vanishes:

$$\lim_{t \rightarrow 0} b_{\gamma,t}^2 = 0 \tag{57}$$

thanks to the definition of κ_t . This makes the wavelet estimator an unbiased estimator when t reaches 0. Now, for the variance term, notice first that:

$$\begin{aligned}
& \mathbb{E}[\zeta_K^\gamma(x, t) - \mathbb{E}[\zeta_K^\gamma(x, t)]]^2 = \text{var}[\zeta_K^\gamma(x, t)] \\
&= \frac{1}{K^2} \sum_{k=1}^K \text{var}[\kappa_t(X_k^{-1}x)] \tag{58}
\end{aligned}$$

It comes then naturally that the variance term takes the form:

$$\begin{aligned}
\sigma_{\gamma,t}^2 &= \int_{SO(3)} \text{var}[\zeta_K^\gamma(x, t)] d\mu(x) \\
&= \frac{1}{K^2} \sum_{k=1}^K \int_{SO(3)} (\mathbb{E}[\kappa_t^2(X_k^{-1}x)] - \mathbb{E}^2[\kappa_t(X_k^{-1}x)]) d\mu(x) \\
&= \frac{1}{K^2} \sum_{k=1}^K \int_{SO(3)} \int_{SO(3)} \kappa_t^2(z^{-1}x) f(z) d\mu(z) d\mu(x) \\
&\quad - \frac{1}{K^2} \sum_{k=1}^K \int_{SO(3)} (f * \kappa_t)^2(X_k) d\mu(X_k) \tag{59}
\end{aligned}$$

where we made use of the bi-invariance of the Haar measure and of the independence between the samples

X_k . Finally, it is possible to give the expression of the variance as:

$$\sigma_{\gamma,t}^2 = \mathbb{E}[\|\zeta_K^\gamma(x, t) - \mathbb{E}[\zeta_K^\gamma(x, t)]\|_2^2] = \frac{1}{K} (\|\kappa_t\|_2^2 - \|f * \kappa_t\|_2^2) \tag{60}$$

As was noticed earlier, the wavelet estimator is a kernel estimator. Using the general formula for the MISE of kernel estimators and the expressions of the bias and variance given above, one gets the following expression for the MISE:

$$\begin{aligned}
\text{MISE}(\zeta_K^\gamma, t) &= \sum_{\ell \geq 1} ((2\ell + 1) \hat{f}^\ell)^2 (1 - e^{-\ell(\ell+1)t})^2 \\
&\quad + \frac{(2\ell + 1)}{K} e^{-2\ell(\ell+1)t} (1 - \hat{f}^\ell)^2 \tag{61}
\end{aligned}$$

Using a first order approximation of the $\exp(\cdot)$ function for small t , one gets an expression of the MISE of the form:

$$\begin{aligned}
\text{MISE}(\zeta_K^\gamma, t) &= \sum_{\ell \geq 1} \left[\frac{(2\ell + 1)(1 - 2\ell(\ell + 1)t)}{K} \right. \\
&\quad \left. + \hat{f}^\ell)^2 \frac{(2\ell + 1)}{K} (2\ell(\ell + 1)t(1 + K) - 1) \right] \tag{62}
\end{aligned}$$

This expression will be used in Section V. It is of interest to note that the behaviour of the MISE with parameter t is as follows:

$$\lim_{t \rightarrow 0} \text{MISE}(\zeta_K^\gamma, t) = \sum_{\ell \geq 1} \frac{(2\ell + 1)}{K} (1 - \hat{f}^\ell)^2 \tag{63}$$

Also notice that as f is a probability density, then $f(x) \geq 0$ almost everywhere and we have that $\hat{f}^0 = 1$ because $\hat{f}^0 = \int_{SO(3)} f(x) d\mu(x)$.

The three studied estimators belong to the kernel estimator family, but they do have differences. The given expressions allows to study and compare their behaviors. The following section highlights the differences between them.

V. SIMULATIONS

In this section, we present some simulations that illustrate the differences between the three considered estimators studied in this paper. The comparison is performed in terms of the respective MISE computed in the context of estimation of a mixture of densities on the rotation group $SO(3)$.

A. Definition of a test function

First, we have to define a test function f on the rotation group, which we take similar to the one studied in [6]. This mixture takes the form:

$$\begin{aligned}
f(\mathbf{x}) &= 0.2 + 0.7 \psi_{\text{VP}}^{30}(\mathbf{R}(\mathbf{e}_1, \pi/6) \cdot \mathbf{x}) \\
&\quad + 0.1 \psi_{\text{VP}}^{45}(\mathbf{R}(\mathbf{e}_2, 4\pi/9) \cdot \mathbf{x}), \tag{64}
\end{aligned}$$




Fig. 1. Kernels used in the simulations. From left to right: Heat kernel, La Vallée Poussin kernel and characteristic function kernel. Kernels are displayed for different values of the “bandwidth” parameter.

where $\mathbf{R}(\eta, \omega)$ denotes the rotation of angle ω and axis η , and ψ_{VP}^κ is the de La Vallée Poussin kernel. This kernel has a closed form expression, which reads

$$\begin{aligned} \psi_{\text{VP}}^\kappa(x) &= \frac{(2\kappa + 1)2^{2\kappa}}{\binom{2\kappa + 1}{\kappa}} \cos^{2\kappa}\left(\frac{\omega(x)}{2}\right) \\ &= \binom{2\kappa + 1}{\kappa}^{-1} \sum_{\ell=0}^{\kappa} (2\ell + 1) \binom{2\kappa + 1}{\kappa - \ell} \chi^\ell(x). \end{aligned} \quad (65)$$

The parameter κ can be understood as the analog of the ρ parameter in the heat kernel wavelet approach, that is it plays the equivalent role of a bandwidth. Figure 2 illustrates the correspondence between different values of κ and ρ ($\log_2 \rho$ actually).

B. Computing the MISE

We recall that the MISE can be expressed in the Fourier domain for a kernel function Ξ by

$$\begin{aligned} \text{MISE}(\Xi) &= \sum_{\ell \geq 1} \left((2\ell + 1) (\hat{f}^\ell)^2 (1 - \hat{\Xi}^\ell)^2 + \frac{(2\ell + 1)}{K} (\hat{\Xi}^\ell)^2 (1 - (\hat{f}^\ell)^2) \right). \end{aligned} \quad (66)$$

The last expression of the MISE in terms of the Fourier coefficients of both the test function f and the chosen kernel Ξ allows us to compute the MISE in a simple way. The Fourier coefficients of the kernel are in general known, as it is the case here for the kernel used (de la Vallée Poussin, Heat kernel, and characteristic function kernel). However the coefficients of f have to be computed with a numerical implementation of the Fourier transform on the rotation group.

We used in our simulation an implementation of the FFT on $SO(3)$ as proposed by Kostelec and Rockmore [23]. This FFT is based on an equiangular sampling of the Euler angles. Since our test function f is a linear combination of de la Vallée Poussin kernels, it is bandlimited by the largest κ value chosen, that is in our case $\kappa = 45$. The FFT was thus performed up to degree $L = 49$, leading to 166650 complex valued Fourier coefficients \hat{f}_{nm}^ℓ . The energy per degree \hat{f}^{ℓ^2} can be computed like:

$$(\hat{f}^\ell)^2 = \frac{1}{(2\ell + 1)} \sum_{n,m=-\ell}^{\ell} |\hat{f}_{nm}^\ell|^2. \quad (67)$$

This expression, up to the maximum value of ℓ is then plugged into the MISE expression.

C. Results

In figure 1 we have plotted the kernels used in our simulations. We consider here only three kernel types, the de La Vallée Poussin kernel, the Heat kernel (sometimes called the Gauss-Weierstrass kernel) and the characteristic function kernel. For each of these kernels we looked at different bandwidth parameters, respectively κ , ρ , and L , the latter being the truncation order in the characteristic function estimator. The plots in Figure 1 emphasize the role of the respective bandwidth, as the concentration of the kernel functions increases with larger bandwidths. We also note that the de La Vallée Poussin and heat kernels are nonnegative kernels, whereas the characteristic function kernel exhibits negative values. Such behaviour has drawbacks, especially when the density to estimate exhibits narrow modes.

Before stepping into the MISE figures, we recall an important result from Hielscher [6, theorem 3], which gives the MISE optimal bound for the function f . It




Fig. 2. Bandwidth (number of non-null coefficients in the Fourier expansion) comparison for different values of κ (de la Vallée Poussin kernel) and ρ (Heat kernel).

reads explicitly as

$$\text{MISE}_{\text{opt}} = \sum_{\ell=1}^{\infty} (2\ell + 1) \frac{(\hat{f}^{\ell})^2 (1 - (\hat{f}^{\ell})^2)}{(K - 1)(\hat{f}^{\ell})^2 + 1} \quad (68)$$

where we have adapted the prefactor in $(2\ell + 1)$ due to our normalization choice. This optimal MISE is displayed on Figures 3, 4 and 5 (black solid line) for comparison with the studied estimators. Recall also that for comparison purpose, a comparison of parameters ρ and κ with respect to the bandwidth is displayed in Figure 2.

We investigate now the behavior of each kernel and its influence on the MISE for different values of bandwidth parameters. It is important to note that in each experiment, bandwidths were fixed, as opposed to what was done in [6] where the bandwidths were chosen as optimal with respect to the test function and the number of observations. Our approach is indeed motivated by a multiresolution approach, with the minimum amount of information about the function f being incorporated in the kernels.

In Figure 3 we have the MISE evaluated for de la Vallée Poussin kernel with bandwidth parameter ranging for $\kappa = 5, 60, 115, 170, 225, 280, 335, 390, 445$; along the theoretical lower bound. In Figure 4 we have the MISE evaluated for the heat kernel with bandwidth parameter ranging from $\rho = 2^{-1}$ to $\rho = 2^{-9}$, (which corresponds to a dyadic scaling) along the theoretical lower bound. Finally, in figure 5 we have the MISE evaluated for the characteristic function kernel with bandwidth parameter ranging from $L = 1$ to $L = 9$, along the theoretical lower bound.




Fig. 3. MISE computed for a de la Vallée Poussin kernel type, for bandwidth values $\kappa = 5, 60, 115, 170, 225, 280, 335, 390, 445$ (from light yellow to dark red). The optimal MISE bound given in Equation (68) is displayed in black for comparison.




Fig. 4. MISE computed for the heat kernel, for bandwidth values $\rho = 2^{-j}$ with $j = 1, 2, \dots, 9$ (from light yellow to dark red). The optimal MISE bound given in Equation (68) is displayed in black for comparison.

As expected with a multiresolution approach, for a given bandwidth the MISE reaches a lower threshold after a sufficient number of observations K . Moreover as the bandwidth increases the threshold is lowered. Comparing the three kernel types, the de La Vallée Poussin kernel seems to converge with various speeds for differ-




Fig. 5. MISE computed for the characteristic function kernel, for bandwidth values $L = 1, 2, \dots, 9$ (from light yellow to dark red). The optimal MISE bound given in Equation (68) is displayed in black for comparison.

ent bandwidths, while the heat or characteristic function kernel have a more regular rate of convergence with respect to the bandwidth. Therefore, the de La Vallée Poussin kernel seems slightly less appropriate for a multiresolution analysis. The heat kernel and characteristic function kernel are performing better, with a slight advantage to the characteristic function kernel as it comes closest to the MISE optimal bound. We note however that the characteristic function kernel does not have the nonnegative property of the heat kernel, which may be critical when the density to estimate possesses sharpened modes. Note that the nonnegativity can be handled during estimation procedures using for example “square root” estimators as in [24], and which thus require extra steps for the estimation. It is also noticeable that for small sizes of samples (from 10^1 to 10^2) the wavelet estimator gets closer to the optimal bound. This suggests that an adaptive choice of the parameter ρ could lead to very good estimates in most of the situations.

VI. CONCLUSION

We have demonstrated that the characteristic function estimator and the linear heat wavelet estimator on $SO(3)$ both belong to the larger family of kernel estimators for densities on $SO(3)$. The characteristic function estimator consists in a kernel estimator with constant Fourier coefficients up to a maximum bandwidth, while the wavelet estimator leads to a heat kernel with coefficients driven by the scaling parameter of the heat wavelet family. The MISEs of the introduced estimator have been presented and illustration of the differences between heat kernel, characteristic function and

De La Vallée Poussin kernels investigated. The diffusive wavelet based estimator combines the nice property of converging faster than the De La Vallée Poussin kernel and of being strictly positive (as opposed to the characteristic function kernel), allowing good performances in many configurations. The heat wavelet kernel thus combines naturally the advantages of both the De La Vallée Poussin and characteristic function kernels, by providing a nonnegative estimator with very good performances in a wide range of bandwidths. The simulation results presented demonstrate the advantage of using the heat wavelet transform for probability density estimation on the rotation group $SO(3)$. Future work should consist in studying the nonlinear (thresholded) version of the heat wavelet estimator which is known to perform better than linear wavelet estimator on the real line. Validation on real datasets should also be investigated to completely validate the proposed estimator.

REFERENCES

- [1] P. T. Fletcher, S. Joshi, C. Lu, and S. M. Pizer Gaussian distributions on lie groups and their application to statistical shape analysis, in *Information Processing in Medical Imaging*. Springer, 2003, pp. 450–462.
- [2] P. T. Fletcher, C. Lu, S. M. Pizer, and S. Joshi Principal geodesic analysis for the study of nonlinear statistics of shape, *Medical Imaging, IEEE Transactions on*, vol. 23, no. 8, pp. 995–1005, 2004.
- [3] P. T. Fletcher, C. Lu, and S. Joshi Statistics of shape via principal geodesic analysis on lie groups, in *Computer Vision and Pattern Recognition, 2003. Proceedings. 2003 IEEE Computer Society Conference on*, vol. 1. IEEE, 2003, pp. 1–95.
- [4] N. Le Bihan and L. Margerin Nonparametric estimation of the heterogeneity of a random medium using compound poisson process modeling of wave multiple scattering, *Physical Review E*, vol. 80, no. 1, p. 016601, 2009.
- [5] S. Said, C. Lageman, N. Le Bihan, and J. Manton Decomposing on compact lie groups, *Information Theory, IEEE Transactions on*, vol. 56, no. 6, pp. 2766–2777, June 2010.
- [6] R. Hielscher Kernel density estimation on the rotation group and its application to crystallographic texture analysis, *Journal of multivariate analysis*, vol. 119, pp. 119–143, 2013.
- [7] G. S. Chirikjian and A. B. Kyatkin *Engineering applications of noncommutative harmonic analysis: with emphasis on rotation and motion groups*. CRC press, 2000.
- [8] G. Chirikjian *Stochastic models, information theory and Lie groups*. Birkhäuser, 2012.
- [9] K. Mardia and P. Jupp *Directional statistics*. John Wiley & Sons Ltd, 2000.

- [10] J. T.-H. Lo and L. Eshleman
Exponential fourier densities on $SO(3)$ and optimal estimation and detection for rotational processes,
in *Decision and Control including the 15th Symposium on Adaptive Processes, 1976 IEEE Conference on*, 1976, pp. 134–140.
- [11] P. Schröder and W. Sweldens
Spherical wavelets: Efficiently representing functions on the sphere,
in *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*. ACM, 1995, pp. 161–172.
- [12] J.-P. Antoine and P. Vandergheynst
Wavelets on the 2-sphere: A group-theoretical approach,
Applied and Computational Harmonic Analysis, vol. 7, no. 3, pp. 262–291, 1999.
- [13] J.-L. Starck, Y. Moudden, P. Abrial, and M. Nguyen
Wavelets, ridgelets and curvelets on the sphere,
Astronomy & Astrophysics, vol. 446, no. 3, pp. 1191–1204, 2006.
- [14] S. Ebert and J. Wirth
Diffusive wavelets on groups and homogeneous spaces,
Proceedings of the royal society of Edinburgh, vol. 141A, pp. 497–520, 2011.
- [15] I. Bogdanova, P. Vandergheynst, J.-P. Antoine, L. Jacques, and M. Morvidone
Stereographic wavelet frames on the sphere,
Applied and Computational Harmonic Analysis, vol. 19, no. 2, pp. 223–252, 2005.
- [16] A. Tsybakov
Introduction to nonparametric estimation,
Springer, Ed., 2009.
- [17] S. L. Altmann
Rotations, Quaternions and double groups,
Dover, Ed. Dover, 2005.
- [18] H. Dym and H. McKean
Fourier series and integrals.
Academic Press, 1972.
- [19] A. Barut and R. Raczka
Theory of group representations and applications.
World Scientific, 1986, vol. 2.
- [20] R. Hielscher
“The inversion of the Radon transform on the rotational group and its application to texture analysis,”
Ph.D. dissertation, Freiberg University, 2007.
- [21] S. Ebert
“Wavelets on lie groups and homogeneous spaces,”
Ph.D. dissertation, Freiberg University, 2011.
- [22] R. Hielscher, J. Prestin, and A. Vollrath
Fast summation of functions on the rotation group,
Mathematical Geosciences, vol. 42, no. 7, pp. 773–794, 2010.
- [23] P. J. Kostelec and D. N. Rockmore
Ffts on the rotation group,
Journal of Fourier analysis and applications, vol. 14, no. 2, pp. 145–179, 2008.
- [24] F. Pfaff, G. Kurz, and U. Hanebeck
Multimodal circular filtering using Fourier series,
in *Proceedings of IEEE Conference on information fusion*, 2015, pp. 711–718.



Nicolas Le Bihan obtained his B.Sc. degree in physics from the Université de Bretagne Occidentale in Brest, France, in 1997. He received the M.Sc. and Ph.D. degrees in signal processing in 1998 and 2001, respectively, both from Grenoble INP. In 2011, he obtained the Habilitation degree from Grenoble INP. Since 2002, he has been a research associate at the Centre National de la Recherche Scientifique (CNRS) and is working at the Department of Images and Signals of the GIPSA-Lab (CNRS UMR 5083) in Grenoble, France. From 2013 to 2015, he was a visiting fellow at the University of Melbourne, holding a Marie Curie International Outgoing Fellowship from the European Union (IOF GeoSToSip 326176), ERA, 7th PCRD. His research interests include statistical signal processing on groups, noncommutative algebras and differentiable manifolds and its applications in polarized wave physics, waves in disordered media, and geophysics.



Julien Flamant obtained his M.Sc degree in Electrical Engineering and applied Physics from Ecole Normale Supérieure de Cachan, France in 2014, and completed his M.Phil degree in Electrical Engineering from The University of Melbourne in 2016. He is currently pursuing his PhD studies at CRISTAL, Lille, France. His research interests are centered around signal processing and its applications in physics.

Jonathan Manton holds a Distinguished Chair at the University of Melbourne with the title Future Generation Professor. He is also an Adjunct Professor in the Mathematical Sciences Institute at the Australian National University. Prof Manton is a Fellow of IEEE and a Fellow of the Australian Mathematical Society.

He received his Bachelor of Science (mathematics) and Bachelor of Engineering (electrical) degrees in 1995 and his Ph.D. degree in 1998, all from the University of Melbourne, Australia. From 1998 to 2004, he was with the Department of Electrical and Electronic Engineering at the University of Melbourne. During that time, he held a Postdoctoral Research Fellowship then subsequently a Queen Elizabeth II Fellowship, both from the Australian Research Council. In 2005 he became a full Professor in the Department of Information Engineering, Research School of Information Sciences and Engineering (RSISE) at the Australian National University. From July 2006 till May 2008, he was on secondment to the Australian Research Council as Executive Director, Mathematics, Information and Communication Sciences.

Prof Manton's traditional research interests range from pure mathematics (e.g. commutative algebra, algebraic geometry, differential geometry) to engineering (e.g. signal processing, wireless communications, systems theory). More recently, he has become interested in systems biology and systems neuroscience.



Uncertainty Propagation of Correlated Quaternion and Euclidean States Using the Gauss-Bingham Density

JACOB E. DARLING
KYLE J. DEMARS

A new probability density function, called the Gauss-Bingham density, is proposed and studied in the context of uncertainty propagation. The Gauss-Bingham density quantifies the correlation between a quaternion and Euclidean states on the cylindrical manifold on which these states naturally exist. The Gauss-Bingham density, including its canonical form, is developed. In order to approximate the temporal evolution of the uncertainty, an unscented transform for the Gauss-Bingham density is first developed. The sigma points are then transformed according to given (potentially) nonlinear system dynamics, and the maximum weighted log-likelihood parameters of the Gauss-Bingham density are recovered. Uncertainty propagation using the Gauss-Bingham density does not rely on a small angle assumption to project the uncertainty in the quaternion into a three parameter representation as does the predictor of the multiplicative extended Kalman filter, so its accuracy does not suffer when propagating large attitude uncertainties. Two simulations are presented to show the process and efficacy of uncertainty propagation using the Gauss-Bingham density and to compare it to the multiplicative extended Kalman filter.

Manuscript received March 23, 2016; released for publication March 29, 2016.

Refereeing of this contribution was handled by Igor Gilitschenski.

Authors' addresses: Department of Mechanical and Aerospace Engineering, Missouri University of Science and Technology, Rolla, MO, 65409 USA. (E-mail: {jed7w2,demarsk}@mst.edu).

1557-6418/16/\$17.00 © 2016 JAIF

I. INTRODUCTION

Consider the planar translation and rotation of a body, which are quantified by Cartesian position coordinates and heading angle, respectively. The position of the body is typically assumed Gaussian-distributed since it is not bounded to a given interval; however, the heading angle cannot be assumed Gaussian-distributed since it is required to be in the interval $[-\pi, \pi)$ and the support of the Gaussian density is infinite. A circular density, such as the wrapped normal density or von Mises density, which are defined on the interval $[-\pi, \pi)$, can be used to quantify the heading angle. The position and heading angle of the body are correlated in general, but they are properly quantified by two different densities; they must, therefore, be quantified under a common state density in order to properly represent their correlation.

Estimation approaches have been developed when the state consists of *only* a von Mises- or wrapped normal-distributed circular variable [1], [2]. The predictor step of these approaches, which propagates the state uncertainty in time, calculates or approximates the temporal evolution of the von Mises or wrapped normal density; however, they do not extend to a state with *both* a circular variable and other Euclidean (additive and unbounded) variables. Mardia and Sutton first proposed a state density to quantify a circular and Euclidean variable in which the state density is constructed as the product of a von Mises density and a Gaussian density conditioned on the von Mises-distributed variable [3]. The Gauss von Mises density is constructed in a similar manner to quantify a state with both a circular variable and other Euclidean variables as the product of a multivariate Gaussian density and von Mises density conditioned on the multivariate Gaussian-distributed variable [4].

A single circular variable can be used to quantify the heading angle, or one-dimensional attitude, of a body. Now consider the case when the three-dimensional attitude and other Euclidean states (such as position, velocity, angular velocity, etc.) of a body are quantified by an attitude quaternion and Cartesian coordinates, respectively. The attitude quaternion exists on the unit hypersphere and is antipodally symmetric; that is, opposing quaternions represent the same attitude. The attitude quaternion is a globally nonsingular attitude representation, and thus, it is a popular choice to represent the three-dimensional attitude of a body [5]–[7].

The Bingham density can be used to quantify the uncertainty of an attitude quaternion on the unit hypersphere [8], [9]. The Bingham density is a zero-mean Gaussian density conditioned on the unit hypersphere and is antipodally symmetric, so it is a proper probabilistic representation of the attitude quaternion because antipodal quaternions represent the same physical attitude and, therefore, must be equiprobable. Estimation approaches have been developed for a state that consists

only of an attitude quaternion [10]–[12]. In particular, [12] leverages an unscented transform to propagate the uncertainty of the Bingham-distributed attitude quaternion when the system dynamics are nonlinear. These approaches, however, do not quantify the correlation between the Bingham-distributed attitude quaternion and other Euclidean variables.

A state density that is similar to the Bingham density has been proposed to quantify the dual quaternion representing the pose (position and attitude) of a body [13]; however, this density does not extend to arbitrarily high dimensions to include the velocity, angular velocity, and other Euclidean states since it is constructed using a dual quaternion. The partially wrapped normal density has recently been proposed, which wraps m coordinates of an n -dimensional Gaussian density in order to quantify the correlation between m angular and $n - m$ Euclidean states [14]. This density applies to arbitrarily high m and n , so it can potentially be used to represent the uncertainty of a rotation sequence representing the three-dimensional attitude and other Euclidean states of a body. Because the temporal evolution of a rotation sequence is potentially singular [5]–[7], the temporal evolution of this uncertainty representation will be potentially singular as well.

In order to avoid this potential singularity, this work proposes a new state density, called the Gauss-Bingham density, which can be used to represent the uncertainty of an attitude quaternion and other Euclidean states of a body. The Gauss-Bingham density, which is born as the product of a Gaussian density and a Bingham density that is conditioned on the Gaussian-distributed random variable, quantifies the correlation between an antipodally symmetric s -dimensional unit vector and an r -dimensional vector of Euclidean states on their natural manifold, the unit hypercylinder. To approximate the temporal evolution of the Gauss-Bingham density given (potentially) nonlinear system dynamics, an unscented transformation [15]–[17] is developed for the Gauss-Bingham density. The sigma points generated by the unscented transform are transformed according to the nonlinear system dynamics, and the maximum weighted log-likelihood parameters of the Gauss-Bingham density are recovered from the sigma points. Uncertainty propagation using the Gauss-Bingham density is developed for arbitrary dimensions r and s , and s can then be specialized to $s = 1$ or $s = 3$ to represent the one- and three-dimensional attitude quaternion, respectively.

In order to present the uncertainty propagation of a Gauss-Bingham-distributed state vector, first a review of the attitude representations used is presented in Section II. A review of the Gaussian and Bingham densities is then presented in Section III. The Gauss-Bingham density, including the correlation structure between the Gaussian and conditional Bingham densities, is developed in Section IV. Uncertainty propagation of a Gauss-Bingham-distributed state vector is presented in Section V, which includes the unscented transform and

maximum weighted log-likelihood parameter estimation for the Gauss-Bingham density. Two simulations of uncertainty propagation using the Gauss-Bingham density are presented in Section VI. The first simulation propagates the uncertainty of the one-dimensional attitude quaternion and angular velocity of a body undergoing torque-free motion where the Gauss-Bingham density can easily be visualized in order to provide an intuitive example of this uncertainty propagation method. The second simulation propagates the uncertainty of the three-dimensional relative position, relative velocity, attitude quaternion, and angular velocity of a chase spacecraft with respect to a target spacecraft. A Monte Carlo approach and the predictor of the multiplicative extended Kalman filter are also simulated in order to evaluate the efficacy of uncertainty propagation using the Gauss-Bingham density and to compare it to more conventional methods.

II. ATTITUDE REPRESENTATIONS, KINEMATICS, AND DYNAMICS

Many different representations can be used to parameterize the attitude of a body. The following subsections provide a brief overview of the attitude representations, quaternion kinematics, and angular velocity dynamics used in the subsequent sections; [5]–[7] provide a comprehensive overview of these topics. The attitude matrix, which is a nine-parameter attitude representation, is used to fundamentally quantify the attitude of a body. In practice, the attitude matrix is difficult to quantify directly since it possesses six constraints, so a three or four parameter attitude representation is typically used to parameterize the attitude matrix in order to overcome this difficulty. Three parameter attitude representations provide a one-to-one representation; however, they are potentially singular. To avoid this singularity, a four parameter attitude representation, which is globally nonsingular, is typically used. Four parameter attitude representations are constrained in some way, since three parameters are sufficient to define attitude. Four parameter attitude representations also provide a two-to-one representation of attitude; that is, two sets of the same four parameter representation quantify the same physical attitude.

A. Attitude Matrix

The attitude of a body is fundamentally quantified by the attitude matrix, \mathbf{A} . The attitude matrix is defined as the orientation matrix that rotates the expression of a vector in the “ I ” coordinate-frame to its expression in the “ B ” coordinate-frame. This orientation matrix is denoted by T_I^B and is defined according to

$$\mathbf{x}^B = T_I^B \mathbf{x}^I \triangleq \mathbf{A} \mathbf{x}^I,$$

where \mathbf{x}^I and \mathbf{x}^B denote the physical vector $\mathbf{x} \in \mathbb{R}^3$ expressed in the I and B frames, respectively, and \mathbb{R}^n

represents n -dimensional Euclidean space. Typically, the I frame is taken to be an inertially-fixed frame and the B frame is taken to be a body-fixed frame of interest, but it is not required that the I frame be inertially fixed for these attitude representations to be valid.

Orientation matrices exist in the n -dimensional special orthogonal group, which is given by $SO(n) \triangleq \{T \in \mathbb{R}^{n \times n} : T^T T = I = T T^T, \det T = 1\}$, where $\det \cdot$ represents the determinant operator. Because orientation matrices are in $SO(n)$, they satisfy the following properties:

$$T_I^B = T_C^B T_I^C \quad \text{and} \quad T_B^I = [T_I^B]^{-1} = [T_I^B]^T. \quad (1)$$

These properties allow the rotation matrix relating a reference frame to the body frame, and thus quantifying the attitude error between these frames, to be expressed as

$$\delta A \triangleq T_R^B = T_I^B T_R^I = T_I^B [T_I^R]^T = A \hat{A}^T, \quad (2)$$

where the “ R ” coordinate frame represents the reference frame and \hat{A} represents the attitude matrix of the reference frame. As the attitude error approaches zero, $A \rightarrow \hat{A}$ and $\delta A \rightarrow I$, where I represents the identity matrix of appropriate dimension.

B. Axis-Angle

An intuitive four parameter representation of a rotation in three dimensions is the axis-angle representation. Euler’s theorem states that any rotation in three dimensions can be accomplished by a single rotation. The axis of this rotation is known as the Euler axis and is quantified by the unit vector \mathbf{e} . Define the corresponding angle of rotation about the Euler axis to be $\theta \in [-\pi, \pi)$. The axis-angle representation of this rotation is then given by the parameter set $\{\mathbf{e}, \theta\}$. If the sign of *both* the Euler axis *and* the rotation angle are changed, the rotation defined by the axis-angle representation is unaffected and the two-to-one representation of attitude using the axis-angle representation is apparent. The attitude matrix is given in terms of the Euler axis and rotation angle about the Euler axis by

$$A = I - \sin \theta [\mathbf{e} \times] + (1 - \cos \theta) [\mathbf{e} \times]^2, \quad (3)$$

where $[\mathbf{a} \times]$ represents the skew-symmetric cross product matrix of the arbitrary vector $\mathbf{a} \in \mathbb{R}^3$ and $[\mathbf{e} \times]^2 \triangleq [\mathbf{e} \times][\mathbf{e} \times]$.

C. Rotation Vector

A three parameter representation of a rotation in three dimensions is born as the product of the Euler axis and the rotation angle about the Euler axis as

$$\boldsymbol{\theta} = \theta \mathbf{e}. \quad (4)$$

This attitude parameterization is known as the rotation vector. The rotation vector is a one-to-one representation of attitude, which is apparent since $\{\mathbf{e}, \theta\}$ and $\{-\mathbf{e}, -\theta\}$, which quantify the same rotation, result in

the same rotation vector. The rotation vector is singular due to the potential discontinuity that can occur when propagating the rotation vector. The norm of the rotation vector is constrained to be no greater than π since $\theta \in [-\pi, \pi)$. During propagation, if the magnitude of the rotation vector is equal to π and has a positive temporal derivative, then the rotation vector must instantaneously change sign since $\{\mathbf{e}, \pm\pi\}$ and $\{-\mathbf{e}, \pm\pi\}$ represent equivalent rotations.

The axis-angle representation can be found from the rotation vector according to

$$\mathbf{e} = \frac{\boldsymbol{\theta}}{\|\boldsymbol{\theta}\|} \quad \text{and} \quad \theta = \|\boldsymbol{\theta}\|. \quad (5)$$

It is apparent from (5) that the Euler axis is undefined if the rotation angle is zero. This is not an issue, however, since this corresponds to a rotation angle of zero and $A = I$ in this case. The attitude matrix in terms of the rotation vector is given by substituting (3) into (5) to give the attitude matrix as

$$A = I - \sin \|\boldsymbol{\theta}\| \left[\frac{\boldsymbol{\theta}}{\|\boldsymbol{\theta}\|} \times \right] + (1 - \cos \|\boldsymbol{\theta}\|) \left[\frac{\boldsymbol{\theta}}{\|\boldsymbol{\theta}\|} \times \right]^2. \quad (6)$$

D. Attitude Quaternion

The attitude quaternion is a four parameter attitude representation and is defined by

$$\bar{\mathbf{q}} = \begin{bmatrix} \mathbf{q} \\ q \end{bmatrix} \in \mathbb{S}^3,$$

where $\mathbf{q} \triangleq [q_x \ q_y \ q_z]^T$ and q are the vector and scalar parts of the quaternion, respectively, and $\mathbb{S}^s \triangleq \{\mathbf{z} \in \mathbb{R}^{s+1} : \mathbf{z}^T \mathbf{z} = 1\}$ represents the s -dimensional unit hypersphere. Two important quaternion operations are multiplication and inversion, which are given for unit quaternions by

$$\bar{\mathbf{q}} \otimes \bar{\mathbf{r}} = \begin{bmatrix} r\mathbf{q} + q\mathbf{r} - \mathbf{q} \times \mathbf{r} \\ qr - \mathbf{q} \cdot \mathbf{r} \end{bmatrix} \quad \text{and} \quad \bar{\mathbf{q}}^{-1} = \begin{bmatrix} -\mathbf{q} \\ q \end{bmatrix},$$

where $\bar{\mathbf{q}} \otimes \bar{\mathbf{r}}$ represents the quaternion product of $\bar{\mathbf{q}}$ and $\bar{\mathbf{r}}$ and $\bar{\mathbf{q}}^{-1}$ represents the inverse of $\bar{\mathbf{q}}$. The quaternion multiplication is defined in this way such that successive rotations can be represented by multiplying quaternions in the same order as rotation matrices. Quaternion multiplication is not commutative, i.e. $\bar{\mathbf{q}} \otimes \bar{\mathbf{r}} \neq \bar{\mathbf{r}} \otimes \bar{\mathbf{q}}$ in general.

The attitude quaternion is the most widely used attitude representation because it is not singular and quaternion multiplication and inversion are used to represent sequential and opposite rotations. Equivalent properties as those given for orientation matrices in (1) can be expressed using quaternions according to

$$\bar{\mathbf{q}}_I^B = \bar{\mathbf{q}}_C^B \otimes \bar{\mathbf{q}}_I^C \quad \text{and} \quad \bar{\mathbf{q}}_B^I = [\bar{\mathbf{q}}_I^B]^{-1}.$$

The identity quaternion is defined as $\bar{\mathbf{p}} \triangleq [\mathbf{0}^T \ 1]^T$ and is the quaternion representing zero rotation, such that

$$\bar{\mathbf{q}} = \bar{\mathbf{q}} \otimes \bar{\mathbf{p}} = \bar{\mathbf{p}} \otimes \bar{\mathbf{q}} \quad \text{and} \quad \bar{\mathbf{p}} = \bar{\mathbf{q}} \otimes \bar{\mathbf{q}}^{-1} = \bar{\mathbf{q}}^{-1} \otimes \bar{\mathbf{q}}.$$

Using these relationships, the attitude error given in (2) can be expressed using quaternions according to

$$\delta\bar{\mathbf{q}} = \bar{\mathbf{q}} \otimes \hat{\mathbf{q}}^{-1},$$

where $\delta\bar{\mathbf{q}}$ is the quaternion representation of $\delta\mathbf{A}$ and $\hat{\mathbf{q}}$ is the reference quaternion. As the attitude error approaches zero, $\bar{\mathbf{q}} \rightarrow \hat{\mathbf{q}}$ and $\delta\bar{\mathbf{q}} \rightarrow \bar{\mathbf{p}}$.

The attitude quaternion is related to the axis-angle attitude representation and rotation vector according to

$$\bar{\mathbf{q}} = \begin{bmatrix} \mathbf{e} \sin \frac{\theta}{2} \\ \cos \frac{\theta}{2} \end{bmatrix} = \begin{bmatrix} \frac{\boldsymbol{\theta}}{\|\boldsymbol{\theta}\|} \sin \frac{\|\boldsymbol{\theta}\|}{2} \\ \cos \frac{\|\boldsymbol{\theta}\|}{2} \end{bmatrix}. \quad (7)$$

The unit-norm constraint imposed on the attitude quaternion is apparent in (7) due to trigonometric identities. It is also apparent in (7) that the attitude quaternion is a two-to-one and antipodal attitude parameterization since $\{\mathbf{e}, \theta\}$ and $\{-\mathbf{e}, -\theta\}$, and therefore $\bar{\mathbf{q}}$ and $-\bar{\mathbf{q}}$, represent equivalent rotations. Equation (7) is solved for the axis-angle attitude representation given the attitude quaternion according to

$$\theta = 2 \arccos q \quad \text{and} \quad \mathbf{e} = (1 - q^2)^{-1/2} \mathbf{q}.$$

After finding the axis-angle attitude representation, (4) can be used to find the equivalent rotation vector attitude representation according to

$$\boldsymbol{\theta} = \frac{2 \arccos q}{(1 - q^2)^{1/2}} \mathbf{q}. \quad (8)$$

E. Quaternion Kinematics and Attitude Dynamics

If $\bar{\mathbf{q}}$ is used to represent the equivalent rotation as the attitude matrix, \mathbf{A} , the temporal evolution of the attitude quaternion is given by

$$\dot{\bar{\mathbf{q}}} = \frac{1}{2} \begin{bmatrix} \boldsymbol{\omega} \\ 0 \end{bmatrix} \otimes \bar{\mathbf{q}}, \quad (9)$$

where $\boldsymbol{\omega}$ is angular velocity of the B frame with respect to the I frame expressed in the B frame. This is a kinematic relationship defining the temporal evolution of the attitude quaternion given the angular velocity of the body and is the rotational equivalent of the kinematic relationship between translational position and velocity.

If I is taken to be an inertial frame, the temporal evolution of the angular velocity for a rigid body is given by

$$\mathbf{J}^B \dot{\boldsymbol{\omega}} = \boldsymbol{\tau}^B - \boldsymbol{\omega} \times \mathbf{J}^B \boldsymbol{\omega}, \quad (10)$$

where \mathbf{J}^B is the inertia tensor of the body and $\boldsymbol{\tau}^B$ is the external torque acting on the body, both expressed in the B frame. This is the dynamic relationship relating the temporal evolution of the angular velocity to the inertia tensor of the body and the net external torque acting on the body. This relationship is the rotational analog to the dynamic relationship relating the temporal evolution of the translational velocity to the mass and net external force acting on the body.

III. PROBABILITY DENSITY FUNCTIONS

Before constructing the Gauss-Bingham density, the Gaussian and Bingham densities, which are used to construct the Gauss-Bingham density, are first presented in this section. Furthermore, some of their useful properties, including the applicability of probabilistically representing the attitude quaternion with the Bingham density, are presented.

A. Gaussian Density

The Gaussian density is given for a random vector $\mathbf{x} \in \mathbb{R}^r$ by

$$p_g(\mathbf{x}; \mathbf{m}, \mathbf{P}) = |\mathbf{P}|^{-1/2} \exp\{-\frac{1}{2}(\mathbf{x} - \mathbf{m})^T \mathbf{P}^{-1}(\mathbf{x} - \mathbf{m})\}, \quad (11)$$

where $\mathbf{m} \in \mathbb{R}^r$ is the mean and $\mathbf{P} = \mathbf{P}^T > \mathbf{0} \in \mathbb{R}^{r \times r}$ is the covariance of the Gaussian density.

B. Canonical Gaussian Density

The standard normal density, which is denoted as the canonical Gaussian density for consistent nomenclature with the canonical form of other densities, is introduced by substituting the transformation

$$\mathbf{x} = \mathbf{S}\mathbf{z} + \mathbf{m} \quad \text{where} \quad \mathbf{P} = \mathbf{S}\mathbf{S}^T \quad (12)$$

into (11), which yields the canonical Gaussian density as

$$\tilde{p}_g(\mathbf{z}) = p_g(\mathbf{z}; \mathbf{0}, \mathbf{I}) = (2\pi)^{-r/2} \exp\{-\frac{1}{2}\mathbf{z}^T \mathbf{z}\},$$

where the tilde notation is used to denote the canonical form of the density.

C. Bingham Density

The Bingham density is an antipodally symmetric density on the unit hypersphere that is a zero-mean Gaussian density conditioned on the unit hypersphere. Because antipodal attitude quaternions ($\bar{\mathbf{q}}$ and $-\bar{\mathbf{q}}$) represent the same attitude, the Bingham density rigorously quantifies the uncertainty in the quaternion representation of attitude without ambiguity between $\bar{\mathbf{q}}$ and $-\bar{\mathbf{q}}$. The Bingham density is defined for a random unit vector $\bar{\mathbf{q}} \in \mathbb{S}^s$ and is given by [8], [9]

$$p_b(\bar{\mathbf{q}}; \mathbf{M}, \mathbf{Z}) = F^{-1}(\mathbf{Z}) \exp\{\bar{\mathbf{q}}^T \mathbf{M} \mathbf{Z} \mathbf{M}^T \bar{\mathbf{q}}\}, \quad (13)$$

where $\mathbf{M} \in \text{SO}(s+1)$ is the orientation matrix describing the orientation of the density on the unit hypersphere, $\mathbf{Z} \in \mathbb{R}^{(s+1) \times (s+1)}$ is a diagonal matrix of concentration parameters with nondecreasing diagonal elements $Z_1 \leq \dots \leq Z_s \leq Z_{s+1} \stackrel{\Delta}{=} 0$, and $F(\mathbf{Z})$ is the normalization constant that ensures that $p_b(\bar{\mathbf{q}}; \mathbf{M}, \mathbf{Z})$ is a valid probability density function. The Bingham density possesses the property that $p_b(\bar{\mathbf{q}}; \mathbf{M}, \mathbf{Z}) = p_b(\bar{\mathbf{q}}; \mathbf{M}, \mathbf{Z} + c\mathbf{I})$ for all $c \in \mathbb{R}$; thus, Z_{s+1} can be defined to be zero with an appropriate choice of c for a given Bingham density without any change to the characteristics of the density.

An abuse of notation is used for $\bar{\mathbf{q}}$ because it is used to represent both a generic antipodally symmetric unit vector of arbitrary dimension s as well as the attitude quaternion; when $\bar{\mathbf{q}} \in \mathbb{S}^1$ or $\bar{\mathbf{q}} \in \mathbb{S}^3$, $\bar{\mathbf{q}}$ is a valid attitude quaternion representing the one- and three-dimensional orientation of a body, respectively. Whether $\bar{\mathbf{q}}$ represents a generic antipodally symmetric unit vector or the attitude quaternion is clear in the surrounding context.

The parameters of \mathbf{Z} control how tightly clustered the Bingham density is around its mean direction, while the orientation matrix, \mathbf{M} , specifies the mean direction itself. The normalization constant of the Bingham density is given by the hypergeometric function of a matrix argument according to

$$\begin{aligned} F(\mathbf{Z}) &= \int_{\mathbb{S}^s} \exp\{\bar{\mathbf{q}}^T \mathbf{Z} \bar{\mathbf{q}}\} d\mathbb{S}^s \\ &= |\mathbb{S}^s| {}_1F_1\left(\frac{1}{2}; \frac{s+1}{2}; \mathbf{Z}\right), \end{aligned} \quad (14)$$

where $|\mathbb{S}^s|$ represents the area of the unit hypersphere \mathbb{S}^s and ${}_1F_1(\cdot; \cdot; \cdot)$ represents the hypergeometric function of a matrix argument. The normalization constant is independent of the orientation matrix, which is intuitive since the orientation matrix simply changes the orientation of the density on the unit hypersphere. Many methods exist for calculating the normalization constant, including series expansions [18], saddle point approximations [19], [20], the holonomic gradient method [21], and interpolation of precomputed tabulated values [22]. In this work, the integral in (14) is numerically integrated directly in order to obtain the normalizing constant of the Bingham density.

Parallels between the parameters of the well-known and well-understood Gaussian density and the parameters of the Bingham density can be drawn in order to better understand the Bingham density. The Bingham density is a directional density; that is, it probabilistically quantifies the direction of a unit vector in \mathbb{S}^s . The orientation matrix, \mathbf{M} , is similar to the mean of the Gaussian density, \mathbf{m} , in that it specifies the mean *direction* of the Bingham density, while \mathbf{m} specifies the mean *location* of the Gaussian density. The matrix of concentration parameters of the Bingham density, \mathbf{Z} , is similar to the covariance matrix of the Gaussian density, \mathbf{P} , in that it specifies how tightly clustered the Bingham density is about its mean direction. Making the elements of \mathbf{Z} more negative leads to a more tightly clustered density about the mean direction for the Bingham density similarly to how decreasing \mathbf{P} leads to a more tightly clustered density about the mean for the Gaussian density. It is important to note that \mathbf{Z} is not the covariance of the Bingham density; however, they are directly related, as discussed in Section III-E.

Representing the uncertainty of an attitude quaternion using the Bingham density has three key advantages as compared to other methods of attitude uncertainty representation:

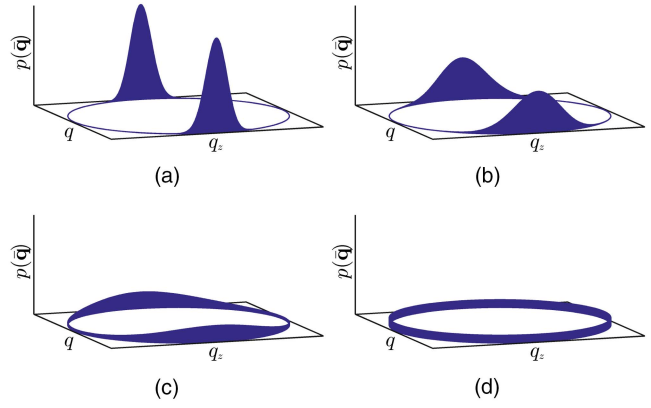


Fig. 1. Bingham densities on \mathbb{S}^1 for $\mathbf{M} = \mathbf{I}$ and varying values of Z_1 . (a) $Z_1 = -50$. (b) $Z_1 = -10$. (c) $Z_1 = -2$. (d) $Z_1 = 0$.

- The Bingham density is antipodally symmetric; thus, antipodal quaternions $\bar{\mathbf{q}}$ and $-\bar{\mathbf{q}}$ (which represent the same physical attitude) are equiprobable,
- the Bingham density quantifies the uncertainty of the attitude quaternion $\bar{\mathbf{q}}$ on its natural manifold \mathbb{S}^3 instead of projecting the attitude uncertainty into a local tangent space, which can potentially incur approximation errors, and
- the Bingham density possesses a simple representation of a uniformly distributed attitude quaternion on this manifold when the \mathbf{Z} matrix is null.

In order to visualize how the Bingham density represents the distribution of an attitude quaternion, the Bingham density is illustrated in \mathbb{S}^1 and \mathbb{S}^2 , where straightforward visualizations exist. The Bingham density is first shown for one-dimensional attitude uncertainty quantification, where the axis of rotation is defined to be the z -axis. In this case, the attitude quaternion simplifies to

$$\bar{\mathbf{q}} = \begin{bmatrix} 0 \\ 0 \\ q_z \\ q \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ \sin(\theta/2) \\ \cos(\theta/2) \end{bmatrix} \in \mathbb{S}^1, \quad (15)$$

where θ is the magnitude of the rotation about the z -axis. Fig. 1 shows the Bingham-distributed one-dimensional attitude quaternion for the identity orientation matrix and different values of Z_1 . Observation of Fig. 1 shows that the Bingham density is antipodally symmetric; that is, $\bar{\mathbf{q}}$ and $-\bar{\mathbf{q}}$ are equiprobable. Further observation highlights that as Z_1 becomes more negative, the uncertainty in the attitude quaternion decreases. Similarly, as Z_1 approaches zero, the uncertainty in the attitude quaternion increases until $Z_1 = 0$, in which case the attitude quaternion is uniformly-distributed.




Fig. 2. Bingham densities on \mathbb{S}^2 for varying values of Z_1 and Z_2 . (a) $Z_1 = -100, Z_2 = 0$. (b) $Z_1 = -100, Z_2 = -10$. (c) $Z_1 = -25, Z_2 = 0$. (d) $Z_1 = -25, Z_2 = -10$.

No valid attitude quaternion exists on \mathbb{S}^2 ; however, the Bingham density for the unit vector $\bar{\mathbf{q}} = [x_1 \ x_2 \ x_3]^T \in \mathbb{S}^2$ is illustrated in Fig. 2 for the orientation matrix

$$\mathbf{M} = \begin{bmatrix} \frac{\sqrt{2}}{2} & \frac{\sqrt{2}}{2} & 0 \\ 0 & 0 & 1 \\ \frac{\sqrt{2}}{2} & -\frac{\sqrt{2}}{2} & 0 \end{bmatrix}$$

and varying values of Z_1 and Z_2 to demonstrate how the \mathbf{Z} matrix affects the Bingham density in this dimension. When $Z_1 = Z_2 = 0$, all $\bar{\mathbf{q}}$ are equiprobable. When Z_2 is zero, the $\bar{\mathbf{q}}$ along a great circle defined by the orientation matrix \mathbf{M} are equiprobable, as observed in Figs. 2(a) and 2(c). In this case, the Z_1 parameter dictates how tightly clustered the probability density of $\bar{\mathbf{q}}$ is along the great circle. When Z_2 decreases from zero, the probability density of $\bar{\mathbf{q}}$ increases in the direction defined by the orientation matrix \mathbf{M} , as observed in Figs. 2(b) and 2(d).

No straightforward visualization of the Bingham density exists in \mathbb{S}^3 . Similar trends, however, are present as the entries of \mathbf{Z} change for the Bingham densities in \mathbb{S}^3 as they do for the Bingham density in \mathbb{S}^1 and \mathbb{S}^2 . A uniformly distributed attitude quaternion is given by the null \mathbf{Z} matrix, and as Z_1, Z_2 , and Z_3 decrease, the uncertainty in the attitude quaternion decreases. If any of the Z_1, Z_2 , and Z_3 are equal to zero, then the attitude quaternion becomes equiprobable around a

higher-dimensional circle or sphere, similar to Figs. 2(a) and 2(c) for the Bingham density in \mathbb{S}^2 .

D. Canonical Bingham Density

The canonical Bingham density is introduced by substituting the transformation

$$\bar{\mathbf{q}} = \mathbf{M}\bar{\mathbf{p}} \quad (16)$$

into (13), which yields the canonical Bingham density as

$$\tilde{p}_b(\bar{\mathbf{p}}; \mathbf{Z}) = p_b(\bar{\mathbf{p}}; \mathbf{I}, \mathbf{Z}) = F^{-1}(\mathbf{Z}) \exp\{\bar{\mathbf{p}}^T \mathbf{Z} \bar{\mathbf{p}}\}.$$

The canonical Bingham density still depends on the matrix of concentration parameters, \mathbf{Z} ; however, the elements of $\bar{\mathbf{p}}$ are uncorrelated.

While the canonical Bingham density is only defined on \mathbb{S}^s , it is still possible to express its mean and covariance in \mathbb{R}^{s+1} . First, define the expected value operator as

$$E_p\{f(\mathbf{x})\} = \int_{\Omega} f(\mathbf{x})p(\mathbf{x})d\mathbf{x},$$

where Ω represents the support of the probability density function $p(\mathbf{x})$ and $f(\mathbf{x})$ is an arbitrary (potentially) nonlinear function defined on Ω . Due to the antipodal symmetry of the canonical Bingham density, its mean in \mathbb{R}^{s+1} is

$$E_{\tilde{p}_b}\{\bar{\mathbf{p}}\} = \mathbf{0}.$$

Because the elements of $\bar{\mathbf{p}}$ are uncorrelated, the covariance of the canonical Bingham density in \mathbb{R}^{s+1} is given by (2.9) from [8] as

$$E_{\tilde{p}_b}\{\bar{\mathbf{p}}\bar{\mathbf{p}}^T\} = \text{diag}[f_1 \ f_2 \ \cdots \ f_{s+1}], \quad (17)$$

where the $\text{diag} \mathbf{v}$ operator constructs a matrix with diagonal entries defined by the arbitrary row vector \mathbf{v} and

$$f_i \triangleq F^{-1}(\mathbf{Z}) \frac{\partial F(\mathbf{Z})}{\partial Z_i}.$$

The f_i terms satisfy several important properties, given by

$$\sum_{i=1}^{s+1} f_i = 1 \quad (18a)$$

$$f_i > 0, \quad i = 1, 2, \dots, s+1 \quad (18b)$$

$$\lim_{Z_i \rightarrow -\infty} f_i = 0^+, \quad i = 1, 2, \dots, s \quad (18c)$$

$$\lim_{Z_1, Z_2, \dots, Z_s \rightarrow -\infty} f_{s+1} = 1^-. \quad (18d)$$

These properties will be exploited when determining the weights of sigma points of the unscented transform for the canonical Gauss-Bingham density in Section V-B. In order to obtain the $\partial F(\mathbf{Z})/\partial Z_i$ necessary to calculate the f_i , (14) is differentiated with respect to each Z_i , and the resulting integral expressions are evaluated numerically.

E. Mean and Covariance of the Bingham Density

Like the canonical Bingham density, the mean of the Bingham density is given in \mathbb{R}^{s+1} as

$$E_{p_b} \{\bar{\mathbf{q}}\} = \mathbf{0}$$

due to antipodal symmetry. The covariance of the Bingham density in \mathbb{R}^{s+1} is defined by $E_{p_b} \{\bar{\mathbf{q}}\bar{\mathbf{q}}^T\}$. In order to calculate this expected value, its argument is first pre- and post-multiplied by $\mathbf{M}\mathbf{M}^T = \mathbf{I}$, such that the covariance can be expressed as

$$E_{p_b} \{\bar{\mathbf{q}}\bar{\mathbf{q}}^T\} = E_{p_b} \{\mathbf{M}\mathbf{M}^T \bar{\mathbf{q}}\bar{\mathbf{q}}^T \mathbf{M}\mathbf{M}^T\}. \quad (19)$$

Introducing the transformation of variables defined in (16) and noting that \mathbf{M} is deterministic, (19) can be expressed as

$$E_{p_b} \{\bar{\mathbf{q}}\bar{\mathbf{q}}^T\} = \mathbf{M} E_{\tilde{p}_b} \{\bar{\mathbf{p}}\bar{\mathbf{p}}^T\} \mathbf{M}^T. \quad (20)$$

$E_{\tilde{p}_b} \{\bar{\mathbf{p}}\bar{\mathbf{p}}^T\}$ is simply the covariance of the canonical Bingham density, which is defined by (17). Substituting (17) into (20) yields the covariance of the Bingham density in \mathbb{R}^{s+1} as

$$E_{p_b} \{\bar{\mathbf{q}}\bar{\mathbf{q}}^T\} = \mathbf{M} \text{diag}[f_1 \quad f_2 \cdots f_{s+1}] \mathbf{M}^T,$$

which is seen to be a similarity transformation of the covariance of the canonical Bingham density according to the orientation matrix of the Bingham density.

IV. GAUSS-BINGHAM DENSITY

The Gauss-Bingham density quantifies a state vector composed of a Gaussian-distributed vector, $\mathbf{x} \in \mathbb{R}^r$, and a Bingham-distributed unit vector, $\bar{\mathbf{q}} \in \mathbb{S}^s$, on its natural manifold defined by $\mathbb{R}^r \times \mathbb{S}^s$. Before constructing the Gauss-Bingham density, first consider the motivating example of manipulating two jointly Gaussian-distributed random vectors given by \mathbf{x} and \mathbf{y} into the product of the density of \mathbf{x} and the density of \mathbf{y} conditioned on \mathbf{x} . The joint density of $[\mathbf{x}^T \ \mathbf{y}^T]^T$ is Gaussian and is given by

$$P_g \left(\begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix}; \begin{bmatrix} \mathbf{m}_x \\ \mathbf{m}_y \end{bmatrix}, \begin{bmatrix} \mathbf{P}_x & \mathbf{P}_{xy} \\ \mathbf{P}_{xy}^T & \mathbf{P}_y \end{bmatrix} \right),$$

where \mathbf{m} and \mathbf{P} represent the mean and covariance of their subscripted vector(s), respectively. The density of \mathbf{y} conditioned on \mathbf{x} is Gaussian-distributed and is given by [23]

$$\begin{aligned} & P_g(\mathbf{y} | \mathbf{x}; \mathbf{m}_{y|x}(\mathbf{x}), \mathbf{P}_{y|x}) \\ &= P_g(\mathbf{y} | \mathbf{x}; \mathbf{m}_y + \mathbf{P}_{xy}^T \mathbf{P}_x^{-1} (\mathbf{x} - \mathbf{m}_x), \mathbf{P}_y - \mathbf{P}_{xy}^T \mathbf{P}_x^{-1} \mathbf{P}_{xy}), \end{aligned}$$

where $\mathbf{m}_{y|x}(\mathbf{x})$ and $\mathbf{P}_{y|x}$ are the mean and covariance, respectively, of \mathbf{y} conditioned on \mathbf{x} . It is interesting to note the functional dependence of $\mathbf{m}_{y|x}$ on \mathbf{x} . From the definition of conditional probability, it follows that the

joint density of \mathbf{x} and \mathbf{y} can be expressed as

$$\begin{aligned} & P_g \left(\begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix}; \begin{bmatrix} \mathbf{m}_x \\ \mathbf{m}_y \end{bmatrix}, \begin{bmatrix} \mathbf{P}_x & \mathbf{P}_{xy} \\ \mathbf{P}_{xy}^T & \mathbf{P}_y \end{bmatrix} \right) = P_g(\mathbf{x}; \mathbf{m}_x, \mathbf{P}_x) \\ & \times P_g(\mathbf{y} | \mathbf{x}; \mathbf{m}_y + \mathbf{P}_{xy}^T \mathbf{P}_x^{-1} (\mathbf{x} - \mathbf{m}_x), \mathbf{P}_y - \mathbf{P}_{xy}^T \mathbf{P}_x^{-1} \mathbf{P}_{xy}). \quad (21) \end{aligned}$$

The conditional mean and covariance of $p(\mathbf{y} | \mathbf{x})$ are not restricted to be

$$\mathbf{m}_{y|x}(\mathbf{x}) = \mathbf{m}_y + \mathbf{P}_{xy}^T \mathbf{P}_x^{-1} (\mathbf{x} - \mathbf{m}_x) \quad (22a)$$

$$\mathbf{P}_{y|x} = \mathbf{P}_y - \mathbf{P}_{xy}^T \mathbf{P}_x^{-1} \mathbf{P}_{xy} \quad (22b)$$

for the definition of conditional probability to be valid; however, (22) must hold for the result to be Gaussian-distributed.

The left- and right-hand sides of (21) express the joint Gaussian density of $[\mathbf{x}^T \ \mathbf{y}^T]^T$ in two different forms. In the case where the vectors \mathbf{x} and \mathbf{y} are jointly Gaussian-distributed, as they are in this example, little if anything is gained by manipulating the left-hand side of (21) into the right-hand side of (21); however, in the case when one or both of the jointly distributed vectors are not Gaussian-distributed, correlation between the vectors can be introduced in a similar fashion to (21) by utilizing the definition of conditional probability. This allows the density of two jointly distributed random vectors, \mathbf{x}_1 and \mathbf{x}_2 , to be written as the product of the density of \mathbf{x}_1 and the density of \mathbf{x}_2 conditioned on \mathbf{x}_1 ; i.e.

$$p(\mathbf{x}_1, \mathbf{x}_2) = p(\mathbf{x}_1) p(\mathbf{x}_2 | \mathbf{x}_1).$$

Using the definition of conditional probability, the Gauss-Bingham density is constructed as the product of a Gaussian density and a Bingham density conditioned on the Gaussian-distributed random variable as

$$P_{gb}(\mathbf{x}; \mathbf{m}, \mathbf{P}, \mathbf{M}(\mathbf{z}), \mathbf{Z}) \triangleq P_g(\mathbf{x}; \mathbf{m}, \mathbf{P}) p_b(\bar{\mathbf{q}}; \mathbf{M}(\mathbf{z}), \mathbf{Z}), \quad (23)$$

where $\mathbf{x} = [\mathbf{x}^T \ \bar{\mathbf{q}}^T]^T$. The Bingham density is conditioned on the Gaussian-distributed random variable \mathbf{x} through the orientation matrix $\mathbf{M}(\mathbf{z})$ using the transformation of variables that defines the canonical Gaussian density, which is given by (12). The orientation matrix is expressed using \mathbf{z} (the random variable corresponding to the canonical Gaussian density) instead of \mathbf{x} for better numerical stability since \mathbf{z} is nondimensional. The functional dependence of the orientation matrix on \mathbf{z} is discussed in Section IV-A.

The Gauss-Bingham density possesses the following favorable properties for probabilistically quantifying the attitude quaternion (when $s = 1$ or $s = 3$) and other Euclidean states:

- The Gauss-Bingham density is antipodally symmetric in the attitude quaternion; thus, antipodal quaternions $\bar{\mathbf{q}}$ and $-\bar{\mathbf{q}}$ (which represent the same physical attitude) are equiprobable,




Fig. 3. Gauss-Bingham densities on $\mathbb{R}^1 \times \mathbb{S}^1$ for a linear and quadratic correlation structure. (a) Linear correlation. (b) Quadratic correlation. (c) Marginalized quaternion. (d) Marginalized quaternion.

- the Gauss-Bingham density quantifies the uncertainty of Euclidean states and the attitude quaternion on their natural manifold $\mathbb{R}^r \times \mathbb{S}^s$, and
- the Gauss-Bingham density possesses a simple representation of an equiprobable attitude quaternion for a given angular velocity when the \mathbf{Z} matrix is null.

In order to illustrate these properties, consider an application of the Gauss-Bingham density to quantify the uncertainty of the one-dimensional attitude quaternion and angular velocity of a body undergoing rotation about the z -axis. In this case, the state vector is defined as

$$\mathbf{x} = \begin{bmatrix} \omega \in \mathbb{R}^1 \\ \bar{\mathbf{q}} \in \mathbb{S}^1 \end{bmatrix} \in \mathbb{R}^1 \times \mathbb{S}^1, \quad (24)$$

where ω is the angular velocity about the z -axis and $\bar{\mathbf{q}}$ is the one-dimensional attitude quaternion of the body, which is defined by (15). No correlation structure for the orientation matrix, $\mathbf{M}(\mathbf{z})$, has yet been defined. Before formally defining this correlation structure, first consider two types of correlation, which are introduced into a set of parameters used to specify $\mathbf{M}(\mathbf{z})$: linear and quadratic. Figs. 3(a) and 3(b) show examples of the Gauss-Bingham density (with $Z_1 \neq 0$) for the linear and quadratic correlation structures, respectively. The marginalized attitude quaternion for the linear and quadratic correlation structures are shown in Figs. 3(c) and 3(d), respectively. It can be observed in these figures that the probability of the antipodal attitude quaternions is equal for any given angular velocity, which is a desirable property as these quaternions represent the same physical attitude.

When $\mathbf{Z} = \mathbf{0}$, the marginalized attitude quaternion is equiprobable regardless of the correlation structure




Fig. 4. Gauss-Bingham density on $\mathbb{R}^1 \times \mathbb{S}^1$ for $Z_1 = 0$. (a) Gauss-Bingham density. (b) Marginalized quaternion.

used. This is illustrated in Figs. 4(a) and 4(b), which show the Gauss-Bingham density in $\mathbb{R}^1 \times \mathbb{S}^1$ and the marginal density of the attitude quaternion when $Z_1 = 0$. This property of the Gauss-Bingham density is advantageous for representing the attitude quaternion as equiprobable when no prior attitude information is available.

A. Correlation Structure

In order to define the correlation structure for the orientation matrix $\mathbf{M}(\mathbf{z})$, it is important to note that $\mathbf{M}(\mathbf{z}) \in \text{SO}(s+1) \forall \mathbf{z} \in \mathbb{R}^r$. In order to ensure that this condition is met, the correlation structure is introduced into a minimum set of parameters necessary to specify the orientation matrix, denoted by $\phi(\mathbf{z})$, such that the orientation matrix is given by $\mathbf{M}(\phi(\mathbf{z}))$. A minimum parameter set, which is comprised of $s(s+1)/2 \triangleq n_\phi$ parameters, is necessary to define the orientation matrix; therefore $\phi(\mathbf{z}) \in \mathbb{R}^{s(s+1)/2}$ [24], [25]. The method for constructing the orientation matrix from the set of minimum parameters depends on s and the parameter set chosen. Methods for constructing the orientation matrix for dimensions $s = 1, 2, 3$ are presented in this section.

1) $s = 1$: First consider the Gauss-Bingham density specialized to $s = 1$. Only one parameter is necessary to specify the orientation matrix in this dimension since $n_\phi = 1$. This parameter is chosen to be the magnitude of the rotation about the known axis of rotation, which is given by $\theta(\mathbf{z})$, such that $\phi(\mathbf{z}) = \theta(\mathbf{z})$. The orientation matrix is then defined by

$$\mathbf{M}(\phi(\mathbf{z})) = \mathbf{M}(\theta(\mathbf{z})) = \begin{bmatrix} \cos \theta(\mathbf{z}) & \sin \theta(\mathbf{z}) \\ -\sin \theta(\mathbf{z}) & \cos \theta(\mathbf{z}) \end{bmatrix}.$$

The angle of rotation, $\theta(\mathbf{z})$, is defined on the interval $[-\pi, \pi)$ for all \mathbf{z} . Since $\mathbf{M}(\theta(\mathbf{z})) = \mathbf{M}(\theta(\mathbf{z}) + 2\pi k)$ for all $k \in \mathbb{Z}$, where \mathbb{Z} is the set of integers, $\theta(\mathbf{z})$ can be bounded to the interval $[-\pi, \pi)$ for all \mathbf{z} by adding the appropriate multiple of 2π .

2) $s = 2$: Now consider the Gauss-Bingham density specialized to $s = 2$. Three parameters are necessary to specify the orientation matrix in this dimension. These parameters are chosen to be the rotation vector, such

that $\phi(\mathbf{z}) = \theta(\mathbf{z})$. The orientation matrix is then given by (6) as

$$\begin{aligned} \mathbf{M}(\phi(\mathbf{z})) = \mathbf{M}(\theta(\mathbf{z})) = & \mathbf{I} - \sin \|\theta(\mathbf{z})\| \left[\frac{\theta(\mathbf{z})}{\|\theta(\mathbf{z})\|} \times \right] \\ & + (1 - \cos \|\theta(\mathbf{z})\|) \left[\frac{\theta(\mathbf{z})}{\|\theta(\mathbf{z})\|} \times \right]^2. \end{aligned}$$

If $\|\theta(\mathbf{z})\| = 0$, $\mathbf{M}(\mathbf{z}) = \mathbf{I}$, since the angle of rotation is zero. The norm of the rotation vector, $\|\theta(\mathbf{z})\|$, is defined on the interval $[-\pi, \pi)$ for all \mathbf{z} . Since $\mathbf{M}(\theta(\mathbf{z})) = \mathbf{M}(\theta(\mathbf{z}) + 2\pi k \theta(\mathbf{z}) / \|\theta(\mathbf{z})\|)$ for all $k \in \mathbb{Z}$, $\|\theta(\mathbf{z})\|$ can be bounded to the interval $[-\pi, \pi)$ for all \mathbf{z} by adding the appropriate multiple of $2\pi \theta(\mathbf{z}) / \|\theta(\mathbf{z})\|$.

3) $s = 3$: Finally, consider the Gauss-Bingham density specialized to $s = 3$. Six parameters are necessary to specify the orientation matrix in this dimension. These parameters are chosen to be two rotation vectors representing a left- and right-isoclonic rotation [26]. Let these rotation vectors be denoted by $\theta_L(\mathbf{z})$ and $\theta_R(\mathbf{z})$, respectively, such that $\phi(\mathbf{z}) = [\theta_L^T(\mathbf{z}) \ \theta_R^T(\mathbf{z})]^T$. The orientation matrix is then defined using left- and right-isoclonic rotations according to

$$\mathbf{M}(\phi(\mathbf{z})) = \mathbf{M} \left(\begin{bmatrix} \theta_L(\mathbf{z}) \\ \theta_R(\mathbf{z}) \end{bmatrix} \right) = \mathbf{L}(\bar{q}(\theta_L(\mathbf{z}))) \mathbf{R}(\bar{q}(\theta_R(\mathbf{z}))),$$

where

$$\begin{aligned} \mathbf{L}(\bar{q}) &= \begin{bmatrix} q & -q_x & -q_y & -q_z \\ q_x & q & -q_z & q_y \\ q_y & q_z & q & -q_x \\ q_z & -q_y & q_x & q \end{bmatrix} \\ \mathbf{R}(\bar{q}) &= \begin{bmatrix} q & -q_x & -q_y & -q_z \\ q_x & q & q_z & -q_y \\ q_y & -q_z & q & q_x \\ q_z & q_y & -q_x & q \end{bmatrix}, \end{aligned}$$

and $\bar{q}(\theta_L(\mathbf{z}))$ and $\bar{q}(\theta_R(\mathbf{z}))$ are defined by (7). The norm of each of these rotation vectors, $\|\theta_L(\mathbf{z})\|$ and $\|\theta_R(\mathbf{z})\|$, is defined on the interval $[-\pi, \pi)$ for all \mathbf{z} . Since

$$\mathbf{M} \left(\begin{bmatrix} \theta_L(\mathbf{z}) \\ \theta_R(\mathbf{z}) \end{bmatrix} \right) = \mathbf{M} \left(\begin{bmatrix} \theta_L(\mathbf{z}) + 2\pi k_L \theta_L(\mathbf{z}) / \|\theta_L(\mathbf{z})\| \\ \theta_R(\mathbf{z}) + 2\pi k_R \theta_R(\mathbf{z}) / \|\theta_R(\mathbf{z})\| \end{bmatrix} \right)$$

for all $k_L, k_R \in \mathbb{Z}$, $\|\theta_L(\mathbf{z})\|$ and $\|\theta_R(\mathbf{z})\|$ can be bounded to the interval $[-\pi, \pi)$ for all \mathbf{z} by adding the appropriate multiple of $2\pi \theta_L(\mathbf{z}) / \|\theta_L(\mathbf{z})\|$ and $2\pi \theta_R(\mathbf{z}) / \|\theta_R(\mathbf{z})\|$ to $\theta_L(\mathbf{z})$ and $\theta_R(\mathbf{z})$, respectively.

Now that the functional dependence of the orientation matrix, $\mathbf{M}(\phi(\mathbf{z}))$, on the minimal set of parameters, $\phi(\mathbf{z})$, has been defined for $s = 1, 2, 3$, the functional dependence of $\phi(\mathbf{z})$ on \mathbf{z} needs to be defined. Two choices for this functional dependence are considered: linear and quadratic. It is noted that the quadratic form of this functional dependence is not used after it is presented; however, it is a valid form of this functional

dependence and is presented to show the flexibility of the Gauss-Bingham density.

First, consider the quadratic dependence of $\phi(\mathbf{z})$ on \mathbf{z} , which is defined by

$$\begin{aligned} \phi(\mathbf{z}) = & \phi_0 + \beta \mathbf{z} \\ & + [\mathbf{z}^T \Gamma_1 \mathbf{z} \ \mathbf{z}^T \Gamma_2 \mathbf{z} \cdots \mathbf{z}^T \Gamma_{n_\phi} \mathbf{z}]^T, \end{aligned} \quad (25)$$

where $\phi_0 \in \mathbb{R}^{n_\phi}$, $\beta \in \mathbb{R}^{n_\phi \times r}$ and $\Gamma_i \in \{\mathbb{R}^{r \times r} : \Gamma_i = \Gamma_i^T\}$, $i = 1, \dots, n_\phi$ quantify the zeroth-, first-, and second-order correlation, respectively, of \mathbf{z} on $\phi(\mathbf{z})$. The choice of implementing \mathbf{z} instead of \mathbf{x} in the correlation structure results in nondimensional coefficients β and Γ_i , which is preferred for numerical stability. Noting that the orientation matrix $\mathbf{M}(\mathbf{z})$ is now explicitly defined by \mathbf{z} , ϕ_0 , β , $\Gamma_1, \dots, \Gamma_{n_\phi}$ and that \mathbf{z} is explicitly defined by \mathbf{x} , \mathbf{m} , and \mathbf{P} , the orientation matrix using the quadratic correlation structure is parameterized as

$$\mathbf{M}(\mathbf{z}) = \mathbf{M}(\mathbf{x}; \mathbf{m}, \mathbf{P}, \phi_0, \beta, \Gamma_1, \dots, \Gamma_{n_\phi}),$$

and the Gauss-Bingham density is given by the specialization of (23) as

$$\begin{aligned} p_{gb}(\mathbf{x}; \mathbf{m}, \mathbf{P}, \phi_0, \beta, \Gamma_1, \dots, \Gamma_{n_\phi}, \mathbf{Z}) = \\ p_g(\mathbf{x}; \mathbf{m}, \mathbf{P}) p_b(\bar{q}; \mathbf{M}(\mathbf{x}; \mathbf{m}, \mathbf{P}, \phi_0, \beta, \Gamma_1, \dots, \Gamma_{n_\phi}), \mathbf{Z}). \end{aligned}$$

The number of parameters necessary to quantify the quadratic correlation between \mathbf{z} and $\phi(\mathbf{z})$ is $\frac{1}{2} n_\phi (2 + 2r + r(r + 1))$, which increases quadratically with r . For one-dimensional attitude ($\mathbb{R}^1 \times \mathbb{S}^1$), three-dimensional attitude ($\mathbb{R}^3 \times \mathbb{S}^3$), and dynamic pose ($\mathbb{R}^9 \times \mathbb{S}^3$) quantification, 3, 60, and 330 unique parameters are needed to quantify the quadratic relationship between $\phi(\mathbf{z})$ and \mathbf{z} .

Now, consider the linear correlation structure for $\phi(\mathbf{z})$, which is given by a simplification of (25) as

$$\phi(\mathbf{z}) = \phi_0 + \beta \mathbf{z}.$$

Using the linear correlation structure, the orientation matrix is parameterized as

$$\mathbf{M}(\mathbf{z}) = \mathbf{M}(\mathbf{x}; \mathbf{m}, \mathbf{P}, \phi_0, \beta),$$

and the Gauss-Bingham density is given by the specialization of (23) as

$$\begin{aligned} p_{gb}(\mathbf{x}; \mathbf{m}, \mathbf{P}, \phi_0, \beta, \mathbf{Z}) \\ = p_g(\mathbf{x}; \mathbf{m}, \mathbf{P}) p_b(\bar{q}; \mathbf{M}(\mathbf{x}; \mathbf{m}, \mathbf{P}, \phi_0, \beta), \mathbf{Z}). \end{aligned}$$

The number of parameters necessary to quantify the linear correlation between \mathbf{z} and $\phi(\mathbf{z})$ is $n_\phi(1 + r)$, which increases linearly with r . For one-dimensional attitude, three-dimensional attitude, and dynamic pose quantification, 2, 24, and 60 unique parameters are needed to quantify the linear relationship between $\phi(\mathbf{z})$ and \mathbf{z} . Because the number of parameters necessary to quantify the linear correlation between \mathbf{z} and $\phi(\mathbf{z})$ increases linearly (as opposed to quadratically) with r , it is used in the remainder of this work.

B. Canonical Gauss-Bingham Density

The canonical Gauss-Bingham density is introduced by substituting the transformations

$$\mathbf{x} = \mathbf{S}\mathbf{z} + \mathbf{m} \quad \text{and} \quad \bar{\mathbf{q}} = \mathbf{M}(\mathbf{z})\bar{\mathbf{p}} \quad (26)$$

into (23), which yields the canonical Gauss-Bingham density as

$$\tilde{p}_{gb}(\mathbf{z}; \mathbf{Z}) = \tilde{p}_g(\mathbf{z})\tilde{p}_b(\bar{\mathbf{p}}; \mathbf{Z}),$$

where $\mathbf{z} = [\mathbf{z}^T \bar{\mathbf{p}}^T]^T$. The elements of \mathbf{z} are uncorrelated and zero mean, such that the covariance of \mathbf{z} is defined by \mathbf{I} and (17) and is given by

$$E_{\tilde{p}_{gb}}\{\mathbf{z}\mathbf{z}^T\} = \text{diag}[1 \cdots 1 \quad f_1 \cdots f_{s+1}]. \quad (27)$$

V. UNCERTAINTY PROPAGATION

In order to propagate the uncertainty of a given Gauss-Bingham-distributed state vector, an unscented transform is used and the weighted maximum log-likelihood parameters of the Gauss-Bingham density are found. The unscented transform generates a set of deterministically chosen sigma points that represent the given Gauss-Bingham density. Each sigma point is then transformed according to known (potentially) nonlinear system dynamics. The weighted maximum log-likelihood parameters of the propagated Gauss-Bingham density are then recovered from the transformed sigma points.

A. System Dynamics

Assume that discrete-time nonlinear system dynamics are given by

$$\mathbf{x}_k = \mathbf{f}(\mathbf{x}_{k-1}), \quad (28)$$

where $\mathbf{x} = [\mathbf{x}^T \bar{\mathbf{q}}^T]^T$, and define the state vector with the antipodal attitude quaternion as $\tilde{\mathbf{x}} = [\mathbf{x}^T \quad -\bar{\mathbf{q}}^T]^T$. Because $\bar{\mathbf{q}}$ quantifies an antipodally equivalent attitude representation in which $\bar{\mathbf{q}}$ and $-\bar{\mathbf{q}}$ represent the same physical attitude, the antipodal symmetry of $\bar{\mathbf{q}}$ must be preserved by the system dynamics; that is, if

$$\mathbf{x}_k = \mathbf{f}(\mathbf{x}_{k-1}) \quad \text{and} \quad \tilde{\mathbf{x}}_k = \mathbf{f}(\tilde{\mathbf{x}}_{k-1}), \quad (29)$$

then the quaternion elements of \mathbf{x}_k and $\tilde{\mathbf{x}}_k$ remain antipodal. Equation (29) defines an important property of the system dynamics, \mathbf{f} . This property states that the system dynamics preserve the antipodal symmetry of the quaternion, which is exploited to reduce the amount of computation necessary to propagate the sigma points representing the Gauss-Bingham density.

B. Unscented Transform

In order to select a set of weights and locations for the sigma points of the unscented transform for the canonical Gauss-Bingham density, the zeroth, first, and second moments between the canonical Gauss-Bingham density and the sigma points are matched in $\mathbb{R}^r \times \mathbb{R}^{s+1}$. The moments will be matched in $\mathbb{R}^r \times \mathbb{R}^{s+1}$; however, the sigma points will be parameterized such that they remain on the manifold $\mathbb{R}^r \times \mathbb{S}^s$. After finding sigma

points for the canonical Gauss-Bingham density, (26) is then used to convert the locations of the sigma points from the canonical Gauss-Bingham to the given Gauss-Bingham density.

In order to reduce the number of sigma points, only one of each pair of antipodal sigma points is considered and propagated since the system dynamics preserve the antipodal symmetry of the sigma points as shown by (29). To illustrate this concept, consider the following example antipodal sigma points in $\mathbb{R}^1 \times \mathbb{S}^1$ at t_{k-1} , \mathcal{X}_{k-1} and $\tilde{\mathcal{X}}_{k-1}$, that are antipodal in $\bar{\mathbf{q}}$ and given by

$$\mathcal{X}_{k-1} = \begin{bmatrix} 3 & \frac{1}{\sqrt{2}} & \frac{-1}{\sqrt{2}} \end{bmatrix}^T \quad \text{and} \\ \tilde{\mathcal{X}}_{k-1} = \begin{bmatrix} 3 & \frac{-1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{bmatrix}^T.$$

These sigma points are propagated by some (potentially) nonlinear function, \mathbf{f} , that satisfies the property given by (29). Assume that this propagation transforms the sigma points to

$$\mathcal{X}_k = [4 \quad 0 \quad 1]^T \quad \text{and} \quad \tilde{\mathcal{X}}_k = [4 \quad 0 \quad -1]^T,$$

which are still antipodal in $\bar{\mathbf{q}}$; thus, the computational expense can be lowered by considering only \mathcal{X}_{k-1} . \mathcal{X}_{k-1} is transformed according to \mathbf{f} to obtain \mathcal{X}_k , and antipodal symmetry can be used to obtain $\tilde{\mathcal{X}}_k$, if desired.

In order to generate the sigma points for the Gauss-Bingham density, motivation is drawn from the $2n + 1$ unscented transform. The $2n + 1$ unscented transform for the canonical Gaussian density places two sigma points at equal but opposite deviations from zero for each of the $n = r$ canonical Gaussian states. A central sigma point is then placed at the origin. When generating sigma points for the canonical Gauss-Bingham density, which considers only one of each pair of antipodal points in the attitude quaternion, a similar approach to that of the $2n + 1$ unscented transform for the canonical Gaussian density is used.

In order to generate the sigma points for the canonical Gauss-Bingham density, first a set of sigma points that quantify deviations from the origin in each state of \mathbf{z} are introduced as $\pm\delta$ while $\bar{\mathbf{p}}$ is held constant as the identity quaternion. The locations of these $2r$ sigma points are given by

$$\mathcal{Z}^{(1),(2)} = \begin{bmatrix} \overbrace{[\pm\delta \quad 0 \cdots 0]}^{\in \mathbb{R}^r} & \overbrace{[0 \cdots 0 \quad 1]}^{\in \mathbb{S}^s} \end{bmatrix}^T \\ \mathcal{Z}^{(3),(4)} = [0 \quad \pm\delta \cdots 0 \quad 0 \cdots 0 \quad 1]^T \\ \vdots \\ \mathcal{Z}^{(2r-1),(2r)} = [0 \cdots 0 \quad \pm\delta \quad 0 \cdots 0 \quad 1]^T,$$

with corresponding weights given by

$$w^{(1),(2)} = w^{(3),(4)} = \cdots = w^{(2r-1),(2r)} = \frac{w_g}{4r},$$

where $\mathcal{Z}^{(i),(j)}$ and $w^{(i),(j)}$ represent the locations and weights of the i th and j th sigma points, respectively, representing the canonical Gauss-Bingham density and w_g is a parameter used to specify the weights of these sigma points. The braces are used to denote the portions of \mathcal{Z} which are the Euclidean and quaternion states, respectively.

Next, angular deviations are introduced into the quaternion state as $\pm\alpha_\ell$ for $\ell = 1, 2, \dots, s$ while \mathbf{z} is held constant at zero in order to guarantee that the perturbed attitude quaternion remains on the unit hypersphere. These $2s$ sigma points are given by

$$\begin{aligned}\mathcal{Z}^{(2r+1),(2r+2)} &= \left[\overbrace{0 \dots 0}^{\in \mathbb{R}^r} \quad \overbrace{\pm S_{\alpha_1} \quad 0 \dots 0 \quad C_{\alpha_1}}^{\in \mathbb{S}^s} \right]^T \\ \mathcal{Z}^{(2r+3),(2r+4)} &= [0 \dots 0 \quad 0 \quad \pm S_{\alpha_2} \dots 0 \quad C_{\alpha_2}]^T \\ &\vdots \\ \mathcal{Z}^{(2r+2s-1),(2r+2s)} &= [0 \dots 0 \quad 0 \dots 0 \quad \pm S_{\alpha_s} \quad C_{\alpha_s}]^T,\end{aligned}$$

with corresponding weights given by

$$\begin{aligned}w^{(2r+1),(2r+2)} &= \frac{w_{b_1}}{4} \\ w^{(2r+3),(2r+4)} &= \frac{w_{b_2}}{4} \\ &\vdots \\ w^{(2r+2s-1),(2r+2s)} &= \frac{w_{b_s}}{4},\end{aligned}$$

where w_{b_ℓ} , for $\ell = 1, 2, \dots, s$, are parameters used to specify the weights of these sigma points and S_α and C_α represent the sine and cosine of α , respectively.

Finally, a central sigma point is placed at $\mathbf{z} = \mathbf{0}$ and in the “zero” direction of $\bar{\mathbf{p}}$, which is the identity quaternion. This single sigma point is given by

$$\mathcal{Z}^{(N)} = \left[\overbrace{0 \dots 0}^{\in \mathbb{R}^r} \quad \overbrace{0 \dots 0 \quad 1}^{\in \mathbb{S}^s} \right]^T,$$

with corresponding weight given by

$$w^{(N)} = \frac{w_c}{2},$$

where w_c is a parameter used to specify the weight of this sigma point and $N = 2r + 2s + 1$ is the total number of sigma points.

In order to find the parameters δ , α_ℓ , w_c , w_g , and w_{b_ℓ} , where $\ell = 1, 2, \dots, s$, which fully define the weights and locations of the sigma points for the canonical Gauss-Bingham density, the zeroth, first, and second moments between the sigma points and the canonical Gauss-Bingham density are matched. The zeroth and first moments of the canonical Gauss-Bingham are 1 and $\mathbf{0}$, respectively. The second moment of the canonical Gauss-Bingham density is given by (27). While only one of each antipodal pair of sigma points is stored and propagated, it is important to note that both of the

antipodal sigma points, which are equally weighted, are considered when calculating the moments of the sigma points. After accounting for the antipodal symmetry of each of the sigma points, the first moment of the sigma points is zero for any choice of the parameters. Matching the zeroth and second moments of the sigma points with the canonical Gauss-Bingham density yields

$$\sum_{\ell=1}^s w_{b_\ell} + w_c + w_g = 1 \quad (30a)$$

$$\frac{\delta^2 w_g}{r} = 1 \quad (30b)$$

$$w_{b_\ell} \sin^2 \alpha_\ell = f_\ell, \quad \ell = 1, 2, \dots, s \quad (30c)$$

$$\sum_{\ell=1}^s w_{b_\ell} \cos^2 \alpha_\ell + w_c + w_g = f_{s+1}, \quad (30d)$$

where (30a) stems from the zeroth moment and (30b)–(30d) stem from the second moment. Summing (30c) for $\ell = 1, 2, \dots, s$ and (30d) while noting the properties in (18) yields (30a); thus, (30d) is redundant and may be neglected. Solving (30b) and (30c) for δ and α_ℓ gives the locations of the sigma points as a function of their weights as

$$\delta = \sqrt{\frac{r}{w_g}} \quad \text{and} \quad \alpha_\ell = \text{asin} \sqrt{\frac{f_\ell}{w_{b_\ell}}}, \quad (31)$$

where $\ell = 1, 2, \dots, s$. Now, the weights must be selected for the sigma points. In order for (31) to have real solutions, w_{b_ℓ} must be greater than or equal to f_ℓ for all $\ell = 1, 2, \dots, s$. In order to ensure that this condition is met, a somewhat nonintuitive choice for the weights of the sigma points for the canonical Gauss-Bingham density is chosen that parallels the choice of weights of the sigma points for the Bingham density presented in [12]. Noting the properties given in (18), the weights of the sigma points representing the Gauss-Bingham density which satisfy (30a) are chosen as

$$w_{b_\ell} = f_\ell + (1 - \lambda - \kappa) \frac{f_{s+1}}{s}, \quad \ell = 1, 2, \dots, s \quad (32a)$$

$$w_c = \lambda f_{s+1} \quad (32b)$$

$$w_g = \kappa f_{s+1}, \quad (32c)$$

where λ and κ are positive tuning parameters such that $\lambda + \kappa < 1$. While choosing the weights according to (32) is nonintuitive, this choice of weights satisfies (30a) and provides real locations for the sigma points. λ and κ are chosen such that the weights of all the sigma points approach an equal weight of $1/N$ as the uncertainty in the states corresponding to $\bar{\mathbf{q}}$ approaches zero; that is, $Z_1, Z_2, \dots, Z_s \rightarrow -\infty$. This choice of weights ensures that the sigma points possess nearly equal weights, and thus have nearly equal importance, when the uncertainty in the attitude quaternion is small. Using the properties in (18), the λ and κ that yield equal weights for the sigma

points as the uncertainty in the quaternion goes to zero are given by

$$\lambda = \frac{1}{N} \quad \text{and} \quad \kappa = \frac{2r}{N}. \quad (33)$$

The sigma points on the canonical Gauss-Bingham density, which are defined in terms of the parameters in (31), (32), and (33), are transformed from the canonical Gauss-Bingham density to the Gauss-Bingham density of interest defined by $p_{gb}(\mathbf{x}; \mathbf{m}, \mathbf{P}, \phi_0, \beta, \mathbf{Z})$ according to (26). These transformed sigma points and their associated weights are denoted by $\mathcal{X}^{(i)}$ and $w^{(i)}$, respectively, where $i = 1, 2, \dots, N$.

The sigma points representing the Gauss-Bingham density at t_{k-1} , \mathcal{X}_{k-1} , are then transformed according to the nonlinear system dynamics given by (28) to obtain the sigma points representing the Gauss-Bingham density at t_k , \mathcal{X}_k . If the dynamical system is governed by continuous-time dynamics, the nonlinear discrete-time function in (28), f , is given by the integration of \mathbf{x}_{k-1} from t_{k-1} to t_k to obtain \mathbf{x}_k .

C. Maximum Weighted Log-Likelihood Gauss-Bingham Parameters

In order to obtain the parameters of the best-fit Gauss-Bingham density given the set of sigma points and weights at t_k , the parameters of the Gauss-Bingham density that maximize the weighted log-likelihood of the sigma points are found. To illustrate why the maximum weighted log-likelihood parameters are sought, consider the case when the unscented transform is used for a state that exists in \mathbb{R}^r . Given the sigma points and weights from the unscented transform, the mean and covariance are recovered from the weighted sample mean and covariance of the sigma points. It can be shown that the weighted sample mean and covariance of the sigma points is the mean and covariance of the Gaussian density that maximizes the weighted log-likelihood of the sigma points.

In this spirit, the parameters of the Gauss-Bingham density are recovered from the sigma points according to

$$\mathbf{m}_k, \mathbf{P}_k, \phi_{0,k}, \beta_k, \mathbf{Z}_k = \underset{\mathbf{m}, \mathbf{P}, \phi_0, \beta, \mathbf{Z}}{\operatorname{argmax}} \sum_{i=1}^N w^{(i)} \ln p_{gb}(\mathcal{X}_k^{(i)}; \mathbf{m}, \mathbf{P}, \phi_0, \beta, \mathbf{Z}). \quad (34)$$

This maximization can be performed analytically for the mean and covariance of the Gaussian density, \mathbf{m} and \mathbf{P} . First, note that the sigma points can be decomposed into their Euclidean and quaternion portions according to $\mathcal{X}_k = [\mathcal{X}_{x,k}^T \quad \mathcal{X}_{q,k}^T]^T$. The mean and covariance of the Gaussian density that maximizes the weighted

log-likelihood of the sigma points is given by the sample mean and covariance of the Euclidean portion of the sigma points according to

$$\mathbf{m}_k = 2 \sum_{i=1}^N w^{(i)} \mathcal{X}_{x,k}^{(i)} \quad (35a)$$

$$\mathbf{P}_k = 2 \sum_{i=1}^N w^{(i)} (\mathcal{X}_{x,k}^{(i)} - \mathbf{m}_k)(\mathcal{X}_{x,k}^{(i)} - \mathbf{m}_k)^T, \quad (35b)$$

where the factor of two is included since only one of each antipodal pair of sigma points in the quaternion state is quantified.

After using (35) to determine \mathbf{m}_k and \mathbf{P}_k , (34) becomes

$$\phi_{0,k}, \beta_k, \mathbf{Z}_k = \underset{\phi_0, \beta, \mathbf{Z}}{\operatorname{argmax}} J(\phi_0, \beta, \mathbf{Z}), \quad (36)$$

where

$$J(\phi_0, \beta, \mathbf{Z}) = \sum_{i=1}^N w^{(i)} \ln p_{gb}(\mathcal{X}_k^{(i)}; \mathbf{m}_k, \mathbf{P}_k, \phi_0, \beta, \mathbf{Z}).$$

This maximization is carried out numerically to find ϕ_0 , β , and \mathbf{Z} . In order to perform this numerical maximization, it is first transformed into a root-finding problem according to the first derivative conditions of a maximum, i.e.

$$\frac{\partial J(\phi_{0,k}, \beta_k, \mathbf{Z}_k)}{\partial \phi_{0,k}} = \mathbf{0} \quad (37a)$$

$$\frac{\partial J(\phi_{0,k}, \beta_k, \mathbf{Z}_k)}{\partial \beta_k} = \mathbf{0} \quad (37b)$$

$$\frac{\partial J(\phi_{0,k}, \beta_k, \mathbf{Z}_k)}{\partial \mathbf{Z}_k} = \mathbf{0}, \quad (37c)$$

where the explicit expressions for the derivatives are omitted for brevity. A root-finding algorithm is used to find the $\phi_{0,k}$, β_k , and \mathbf{Z}_k that satisfy (37). To initialize the root-finding algorithm, $\phi_{0,k-1}$, β_{k-1} , and \mathbf{Z}_{k-1} are used. By initializing the root-finding algorithm in this way, if the propagation time step is chosen sufficiently small, $\phi_{0,k-1}$, β_{k-1} , and \mathbf{Z}_{k-1} remain close to the local maximum and a gradient-based root-finding algorithm will converge to $\phi_{0,k}$, β_k , and \mathbf{Z}_k without excessive iteration required or risk of diverging to a different root.

A number of root-finding algorithms can be used to find the $\phi_{0,k}$, β_k , and \mathbf{Z}_k that satisfy (37). The Levenberg-Marquardt algorithm, a well-known optimizer, is chosen to find these $\phi_{0,k}$, β_k , and \mathbf{Z}_k [27], [28]. This algorithm is used to find the roots of an arbitrary system of equations defined by $\mathbf{g}(\mathbf{x}) = \mathbf{0}$ by minimizing the cost function $\mathbf{g}^T(\mathbf{x})\mathbf{g}(\mathbf{x})$ using \mathbf{x} as the minimization variable. The Levenberg-Marquardt algorithm was chosen to find $\phi_{0,k}$, β_k , and \mathbf{Z}_k since the cost function will remain near the minimum if the time step is chosen sufficiently small and $\phi_{0,k-1}$, β_{k-1} , and \mathbf{Z}_{k-1} are used

to initialize the algorithm. Applying the Levenberg-Marquardt algorithm in this manner to find the roots of (37) was found to be more computationally efficient than applying it to the optimization problem in (36) directly.

D. Uncertainty Propagation Algorithm

In summary, the algorithm used to propagate the Gauss-Bingham density is given by

— Given:

- A Gauss-Bingham-distributed state vector at t_0 defined by $p_{gb}(\mathbf{x}; \mathbf{m}_0, \mathbf{P}_0, \phi_{0,0}, \beta_0, \mathbf{Z}_0)$.
- System dynamics that preserve the antipodal symmetry of the quaternion as defined by the property given by (29).
- A sequence of times to which to propagate the Gauss-Bingham density, t_1, t_2, \dots, t_f .

- 1) Generate the sigma points and weights according to $p_{gb}(\mathbf{x}; \mathbf{m}_0, \mathbf{P}_0, \phi_{0,0}, \beta_0, \mathbf{Z}_0)$.
- 2) Set the time counter to $k = 1$
- 3) Propagate the sigma points from t_{k-1} to t_k according to the given system dynamics.
- 4) Recover \mathbf{m}_k and \mathbf{P}_k according to (35).
- 5) Recover $\phi_{0,k}$, β_k , and \mathbf{Z}_k according to the root-finding problem defined by (37) using $\phi_{0,k-1}$, β_{k-1} , and \mathbf{Z}_{k-1} to initialize the root-finding algorithm.
- 6) If $t_k = t_f$, stop; if $t_k < t_f$, set $k = k + 1$ and return to step 3.

The sequence of times to which to propagate the Gauss-Bingham density, t_1, t_2, \dots, t_f , should be chosen such that the time step is small enough that $\phi_{0,k-1}$, β_{k-1} , and \mathbf{Z}_{k-1} are close to $\phi_{0,k}$, β_k , and \mathbf{Z}_k in order to ensure that the root-finding algorithm converges to the proper solution for $\phi_{0,k}$, β_k , and \mathbf{Z}_k . The size of the time step is a compromise between computational cost and ensuring that the root-finding algorithm converges to the correct root. Because the sigma points are not resampled at each time step, no approximation error is introduced by choosing the time step too small. Since the time step chosen is problem dependent, general guidelines for choosing this time step can not be imposed.

VI. SIMULATIONS

Two simulations are performed to illustrate uncertainty propagation using the Gauss-Bingham density. The first simulation propagates the uncertainty of the planar attitude and angular velocity of a body in $\mathbb{R}^1 \times \mathbb{S}^1$, where the Gauss-Bingham density can be visualized on the unit cylinder. This simulation provides an intuitive example of uncertainty propagation using the Gauss-Bingham density. The second simulation propagates the uncertainty in the dynamic pose (position, velocity, attitude, and angular velocity) of a chase spacecraft with respect to a target spacecraft. This simulation compares uncertainty propagation using the Gauss-Bingham

density to the predictor of the multiplicative extended Kalman filter and a Monte Carlo approach in order to show the efficacy of uncertainty propagation using the Gauss-Bingham density.

A. Planar Attitude and Angular Velocity

Consider the attitude quaternion and angular velocity representing the one-dimensional attitude motion of a body undergoing rotation about the z -axis. In this case, the state vector of the body is defined by (24). The angular velocity comprises the Gaussian-distributed portion of the state vector, with initial mean and covariance given by

$$\mathbf{m}_0 = 0 \quad \text{and} \quad \mathbf{P}_0 = (0.01^\circ/\text{s})^2, \quad (38)$$

respectively. The attitude quaternion comprises the conditional Bingham-distributed portion of the state vector, and is initially uncorrelated with the angular velocity (that is, $\beta_0 = 0$). The parameters defining the orientation and concentration of the conditional Bingham-distributed portion of the state vector are given by

$$\phi_0 = 0 \quad \text{and} \quad Z_{1,0} = -100,$$

respectively. The Gauss-Bingham density representing the initial attitude quaternion and angular velocity of the body, as well as the sigma points generated by the unscented transform, are shown in Fig. 5(a). The marginalized density of the initial attitude quaternion is shown in Fig. 5(c).

The body undergoes torque-free motion, that is, $\tau^B = 0$. The temporal evolution of the attitude quaternion and angular velocity are given by (9) and (10), respectively. The uncertainty propagation algorithm summarized in Section V-D is used to propagate the uncertainty of the attitude quaternion and angular velocity forward in time. A time step of one minute is used to propagate the uncertainty, which is small enough to ensure that the root-finding algorithm converges to the proper $\phi_{0,k}$, β_k , and \mathbf{Z}_k at each time step. Fig. 5 shows the evolution of Gauss-Bingham density and sigma points representing the attitude quaternion and angular velocity of the body, as well as the marginalized density of the attitude quaternion over time. Table I provides the corresponding parameters of the Gauss-Bingham density over time. It is observed that the mean and covariance of the angular velocity, \mathbf{m} and \mathbf{P} , respectively, remain constant, which is expected since (10) shows that the angular-velocity is constant under torque-free motion.

The concentration parameter of the conditional Bingham density, \mathbf{Z} , remains constant (within numerical accuracy of the root finding algorithm). The mean direction of the Gauss-Bingham density, ϕ_0 , remains at zero since the mean of the angular velocity is zero for all time; however the linear correlation parameter, β evolves in time in order to quantify the effect of the uncertain angular velocity on the attitude quaternion.

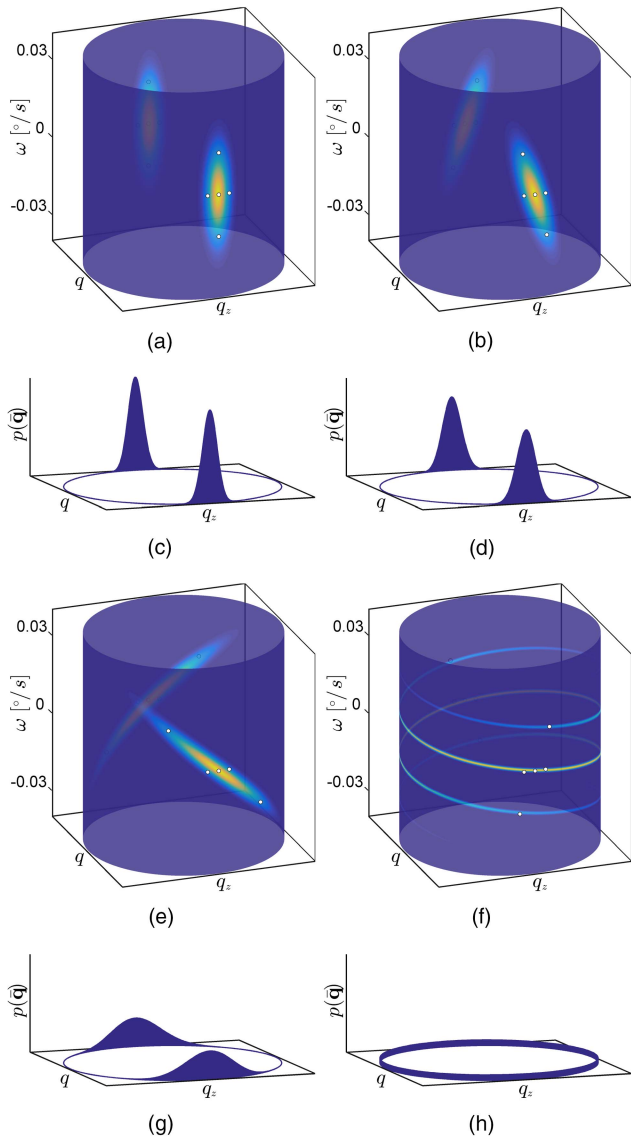


Fig. 5. Gauss-Bingham uncertainty propagation for one-dimensional attitude motion. (a) Initial state density. (b) State density at 15 minutes. (c) Initial quaternion density. (d) Quaternion density at 15 minutes. (e) State density at 1 hour. (f) State density at 6 hours. (g) Quaternion density at 1 hour. (h) Quaternion density at 6 hours.

TABLE I
Gauss-Bingham Parameters over time

Time [hours]	m [$^{\circ}/s$]	P [$(^{\circ}/s)^2$]	ϕ_0	β	Z_1
0	0	$(0.01)^2$	0	0.0000	-100
0.25	0	$(0.01)^2$	0	0.0785	-100
1	0	$(0.01)^2$	0	0.3142	-100
6	0	$(0.01)^2$	0	1.8850	-100

It is interesting to note that β evolves linearly in time for this problem. The Gauss-Bingham density eventually wraps around its cylindrical manifold as it is propagated, which causes the attitude quaternion to become equiprobable as time increases, and is apparent in Fig. 5(h). This is an expected result for a body undergoing

one-dimensional attitude motion with an uncertain angular velocity; as time increases, the uncertainty in the attitude quaternion of the body grows until the attitude quaternion becomes equiprobable.

Several important properties of the Gauss-Bingham density and its utility in uncertainty propagation can be observed in Fig. 5. The Gauss-Bingham density is antipodally symmetric in the quaternion state for all time, which is a necessary property to properly quantify the uncertainty in the attitude quaternion. Since this example quantifies the one-dimensional attitude motion in $\mathbb{R}^1 \times \mathbb{S}^1$, $N = 5$ sigma points are required to quantify the temporal evolution of the Gauss-Bingham density. The attitude quaternion becomes equiprobable as the uncertainty is propagated; however, the concentration parameter Z_1 does not approach zero. As the uncertainty is propagated, the attitude quaternion becomes equiprobable due to the wrapping of the Gauss-Bingham density around the cylinder, not because the concentration parameter approaches zero.

B. Spacecraft Relative Pose

Now consider an example in which a chase spacecraft is orbiting in close proximity to a target spacecraft. The state of the chase spacecraft is defined to be $[\omega^T \ \delta r^T \ \delta v^T \ \bar{q}^T]^T$, where \bar{q} and ω represent the attitude quaternion and angular velocity of the chase spacecraft, respectively, and δr and δv represent the relative position and velocity, respectively, of the chase spacecraft with respect to the target spacecraft. The chase spacecraft is taken to have an identity inertia tensor and undergoes torque-free motion, with the temporal evolution of the attitude quaternion and angular velocity given by (9) and (10), respectively. Since the body undergoes torque-free motion and has an identity inertia tensor, (10) shows that the angular velocity is constant in time.

In order to quantify the temporal evolution of the relative position and velocity, the Clohessy-Wiltshire equations are used [29]–[31]. The Clohessy-Wiltshire equations approximate the relative motion of the chase spacecraft with respect to the target spacecraft under the assumptions that the spacecraft are in close proximity and that the target spacecraft is in a circular orbit. If these assumptions are valid, the Clohessy-Wiltshire equations governing the temporal evolution of the relative position and velocity are

$$\begin{bmatrix} \delta \dot{r} \\ \delta \dot{v} \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 3n^2 & 0 & 0 & 0 & 2n & 0 \\ 0 & 0 & 0 & -2n & 0 & 0 \\ 0 & 0 & -n^2 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} \delta r \\ \delta v \end{bmatrix}, \quad (39)$$

where n is the mean motion of the target spacecraft and δr and δv are expressed in a rotating coordinate frame centered on the target spacecraft. The rotating

coordinate frame is defined by the position and velocity vectors of the target spacecraft. The target spacecraft is taken to be in a geostationary orbit with an orbital radius of 42164 km, which results in a mean motion of the target spacecraft of 7.2920×10^{-5} rad/s.

The Gauss-Bingham density is used to quantify the uncertainty of the state vector of the chase spacecraft. The Gaussian portion of the Gauss-Bingham density quantifies the uncertainty of the angular velocity, relative position, and relative velocity, with initial mean and covariance given by

$$\mathbf{m}_0 = \begin{bmatrix} 0.5^\circ/\text{s} \\ 0.8^\circ/\text{s} \\ 1.0^\circ/\text{s} \\ 0 \\ 10 \text{ km} \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} \quad \text{and} \quad \mathbf{P}_0 = \text{diag} \begin{bmatrix} 0.1^2 (\text{^\circ/s})^2 \\ 0.1^2 (\text{^\circ/s})^2 \\ 0.1^2 (\text{^\circ/s})^2 \\ 1 \text{ m}^2 \\ 1 \text{ m}^2 \\ 1 \text{ m}^2 \\ 0.01^2 (\text{m/s})^2 \\ 0.01^2 (\text{m/s})^2 \\ 0.01^2 (\text{m/s})^2 \end{bmatrix}^T, \quad (40)$$

respectively. The attitude quaternion of the chase spacecraft comprises the conditional Bingham-distributed portion of the state vector, and is initially uncorrelated with the angular velocity, relative position, and relative velocity (that is, $\beta_0 = \mathbf{0}$). The parameters defining the initial orientation and concentration of the conditional Bingham-distributed portion of the state vector are given by

$$\phi_0 = \mathbf{0} \quad \text{and} \quad Z_{1,0} = Z_{2,0} = Z_{3,0} = -5000,$$

respectively.

The uncertainty propagation algorithm summarized in Section V-D is used to propagate the uncertainty of the angular velocity, relative position, relative velocity, and attitude quaternion forward in time. A time step of fifteen seconds is used to propagate the uncertainty, which is small enough to ensure that the root-finding algorithm converges to the proper $\phi_{0,k}$, β_k , and \mathbf{Z}_k at each time step. Uncertainty propagation using the Gauss-Bingham density is compared to two other methods of uncertainty propagation to evaluate its efficacy: a Monte Carlo approach and the predictor step of the multiplicative extended Kalman filter (MEKF) [5], [32], [33]. 100,000 Monte Carlo samples are realized from the initial Gauss-Bingham density using an acceptance sampling method, and are propagated forward in time to quantitatively represent the true evolution of the initial Gauss-Bingham density.

The predictor step of the MEKF quantifies the “mean” using the attitude quaternion and relies on a small angle assumption to project the uncertainty in the

attitude quaternion into a three parameter attitude representation (the rotation vector is used in this analysis). Quotation marks are used around “mean” for the MEKF to indicate that it is not the mean as defined by the first moment of the state vector; rather, it is the “mean” quaternion as defined by one of the antipodal pair used to quantify the quaternion estimate. In order to find the equivalent “mean” and covariance for the MEKF given the initial Gauss-Bingham density, it is first noted that “mean” attitude quaternion is the identity quaternion since $\phi_0 = \mathbf{0}$ and $\beta_0 = \mathbf{0}$; thus, the mean for the MEKF is given by the concatenation of the mean given in (40) and the identity quaternion. The equivalent covariance of the MEKF state vector, which is expressed using the rotation vector instead of the attitude quaternion, is found by converting the quaternion portion of the initial Monte Carlo samples to their equivalent rotation vector according to (8), and calculating their sample covariance.

Since the angular velocity, relative position, and relative velocity are initially Gaussian-distributed, evolve according to linear dynamics, and their temporal evolution is not a function of the attitude quaternion, they remain Gaussian-distributed for all time. Because of this, both the Gauss-Bingham and MEKF uncertainty propagation methods perfectly capture the evolution of the uncertainty in these states, which is presented in Figs. 6–8 and shows the standard deviation of each component of these states quantified by both the MEKF and the Gauss-Bingham density over time. Furthermore, the mean of these quantities is constant for all time since their mean is a stationary solution to (10) and (39) under torque-free motion with an identity inertia tensor. Fig. 6 shows that the uncertainty of the angular velocity is constant, as expected since the angular velocity is constant. Fig. 7 shows that the uncertainty in the relative position grows as time increases. Fig. 8 shows that the uncertainty in the x - and y -components of the relative velocity increase, while the uncertainty in the z -component decreases. This decrease in uncertainty is expected due to the periodicity present in the Clohessy-Wiltshire equations. If the uncertainty is propagated for an entire orbit of the target spacecraft (approximately 24 hours), it would complete one cycle of its period.

Uncertainty propagation using the Gauss-Bingham density does not require that the system dynamics governing the Gaussian-distributed states be linear nor that their temporal evolution be functionally independent of the attitude quaternion. These conditions are used in this example to simplify the presentation and analysis of the results of the uncertainty propagation. If nonlinear system dynamics are used, or if the system dynamics are a function of the attitude quaternion, the best-fit Gaussian density that maximizes the weighted log-likelihood of the sigma points as defined by (35) is found.

Because the attitude uncertainty quantified by the Gauss-Bingham density and Monte Carlo samples are




Fig. 6. Gaussian-distributed angular velocity standard deviation quantified by the Gauss-Bingham density (black) and the MEKF (red).




Fig. 7. Gaussian-distributed relative position standard deviation quantified by the Gauss-Bingham density (black) and the MEKF (red).

expressed using the attitude quaternion and the uncertainty quantified by the MEKF predictor is expressed using the rotation vector, the uncertainty quantified by the Gauss-Bingham density and Monte Carlo samples are converted to rotation vector space in order to make a direct comparison. The rotation vector space is chosen for this comparison since it is a three parameter representation of attitude. In order to convert the uncertainty quantified by the Gauss-Bingham density and Monte Carlo samples from the attitude quaternion representation to the rotation vector representation, first, 100,000 samples of the Gauss-Bingham density are generated using an acceptance sampling method. The quaternion portion of the Gauss-Bingham samples, as well as the Monte Carlo samples, are then converted to their equivalent rotation vector according to (8). Expectation maximization [34] is then performed for each set of samples to fit a Gaussian mixture density to the x - y , y - z , and x - z projections of the rotation vector portion of the respec-




Fig. 8. Gaussian-distributed relative velocity standard deviation quantified by the Gauss-Bingham density (black) and the MEKF (red).

tive samples. This process is used *only* to visualize the uncertainty of the attitude quaternion quantified by the Gauss-Bingham density and the Monte Carlo samples in rotation vector space, and is not an element of the uncertainty propagation using the Gauss-Bingham density. Because the MEKF quantifies the mean and covariance of the rotation vector and not its density, the density is assumed to be Gaussian.

The attitude uncertainty quantified by the Gauss-Bingham density, Monte Carlo samples, and MEKF at a time of five minutes are presented in Fig. 9. Figs. 9(a) and 9(b) show the x - y projection of the rotation vector for 1,000 of the Monte Carlo and Gauss-Bingham samples, as well as the Gaussian mixture densities fit to these samples to show the agreement between the samples and the densities. These plots are repeated without the samples in Figs. 9(c) and 9(d) for clarity along with the uncertainty quantified by the MEKF in red in Fig. 9(d). Figs. 9(e) and 9(f) and Figs. 9(g) and 9(h) show the y - z and x - z projections, respectively, of the uncertainty quantified by the Gauss-Bingham density, true density (as approximated from the Monte Carlo samples), and MEKF. At the time of five minutes, the Gauss-Bingham density agrees very well with the true density. The MEKF quantifies the mean and covariance of the true density as well, which is attributed to the fact that the attitude uncertainty is still relatively small at this time.

Fig. 10 shows the uncertainty quantified by the Gauss-Bingham density, true density (as approximated from the Monte Carlo samples), and MEKF at a time of one hour in the same plots as Fig. 9. After propagating the uncertainty for one hour, the attitude uncertainty quantified by the MEKF does not agree with the true uncertainty, as it has outgrown the $\|\theta\| \in [-\pi, \pi]$ bound on the rotation vector. This uncertainty can potentially be wrapped such that it is expressed in the appropriate bounded region; however, this is not common practice when using the MEKF. The underlying small angle

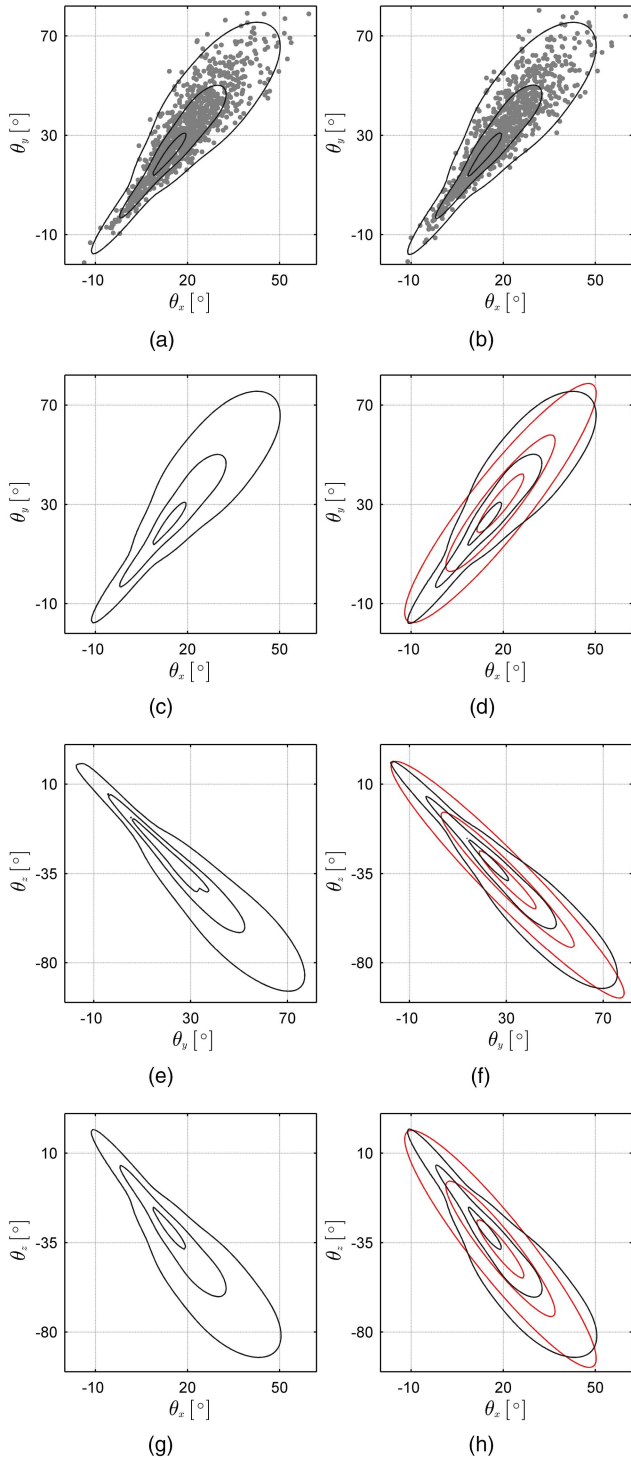


Fig. 9. True, Gauss-Bingham (GB), and MEKF attitude uncertainties expressed in rotation vector space at a time of five minutes. The MEKF density is shown in red. (a) True θ_x - θ_y density and samples. (b) GB θ_x - θ_y density and samples. (c) True θ_x - θ_y density. (d) GB and MEKF θ_x - θ_y densities. (e) True θ_y - θ_z density. (f) GB and MEKF θ_y - θ_z densities. (g) True θ_x - θ_z density. (h) GB and MEKF θ_x - θ_z densities.

assumption used to derive the predictor of the MEKF becomes invalid when the attitude uncertainties become large; thus, it is not well-suited to propagate large attitude uncertainties. This can be observed by noting that, even if the uncertainty were wrapped to the appropriate

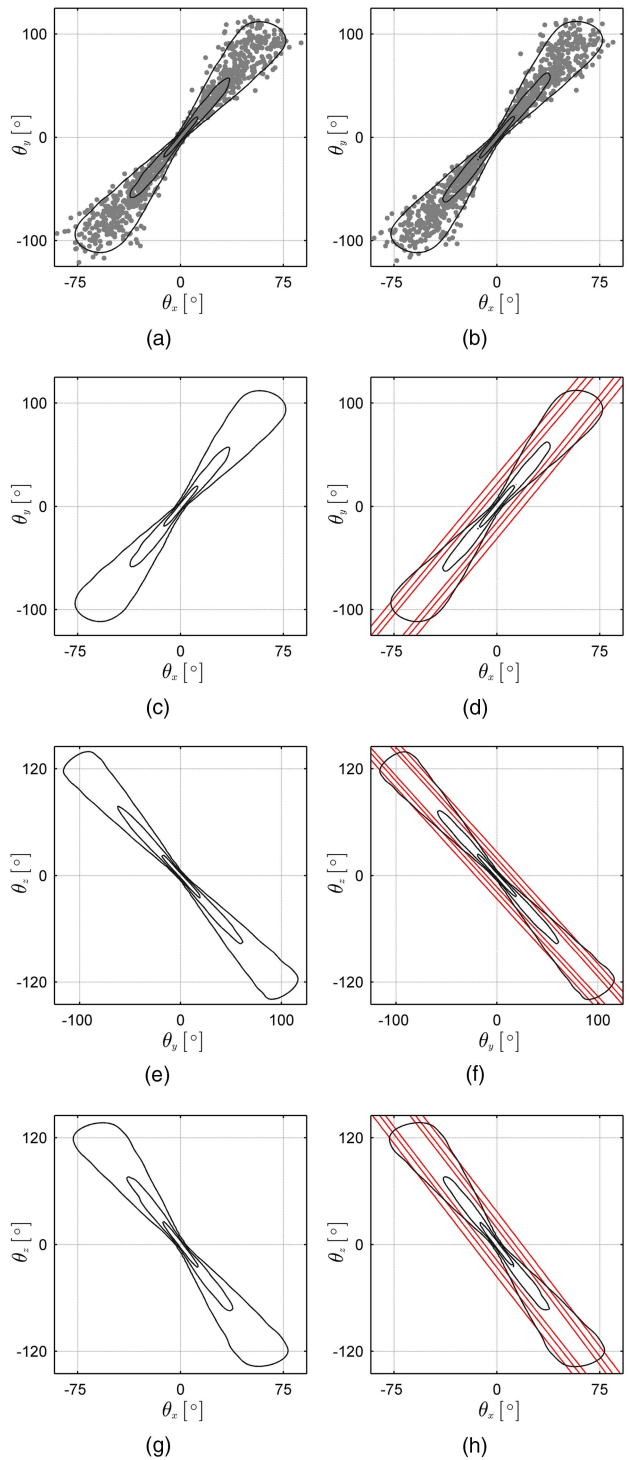


Fig. 10. True, Gauss-Bingham (GB), and MEKF attitude uncertainties expressed in rotation vector space at a time of one hour. The MEKF density is shown in red. (a) True θ_x - θ_y density and samples. (b) GB θ_x - θ_y density and samples. (c) True θ_x - θ_y density. (d) GB and MEKF θ_x - θ_y densities. (e) True θ_y - θ_z density. (f) GB and MEKF θ_y - θ_z densities. (g) True θ_x - θ_z density. (h) GB and MEKF θ_x - θ_z densities.

boundary, it still will not possess the appropriate hour-glass shape as the true density does.

After propagating the uncertainty for one hour, the Gauss-Bingham density is still in close agreement with the true density as is shown in Fig. 10. This is due

to the fact that the uncertainty propagation using the Gauss-Bingham density does not rely on a small angle assumption. The uncertainty is also quantified on the natural manifold of the attitude quaternion and the other Euclidean states, $\mathbb{R}^9 \times \mathbb{S}^3$, so the uncertainty cannot escape the bounded region on which it is defined. Because of these reasons, uncertainty propagation using the Gauss-Bingham density remains well-suited to quantify attitude uncertainty, even as the attitude uncertainty becomes large.

VII. CONCLUSIONS

A new probability density function, called the Gauss-Bingham density, was proposed that represents the uncertainty of an attitude quaternion and other Euclidean states of a body under a common probabilistic framework. The proposed Gauss-Bingham density is constructed as the product of a Gaussian density and a Bingham density that is conditioned on the Gaussian-distributed random variable, which quantifies the correlation between the quaternion and Euclidean states and exists on their natural manifold, $\mathbb{R}^r \times \mathbb{S}^s$. The Gauss-Bingham density quantifies a uniformly-distributed quaternion when its concentration matrix is null, which provides a convenient method to initialize the uncertainty when no attitude information is available. Uncertainty propagation using the Gauss-Bingham density was presented, which leverages an unscented transformation and recovers the maximum weighted log-likelihood parameters of the Gauss-Bingham density from the sigma points. The Gauss-Bingham density properly quantifies the uncertainty of the attitude quaternion and other Euclidean states without relying on a small-angle assumption to project the uncertainty in the attitude quaternion into a three parameter attitude representation, as does the predictor of the multiplicative extended Kalman filter (MEKF). Since the uncertainty is quantified on its natural manifold, it is not possible for the uncertainty propagated using the Gauss-Bingham density to outgrow its bounds like the predictor of the MEKF.

Two simulations were presented. The first simulation showed how the Gauss-Bingham density can be used to propagate the uncertainty of the one-dimensional attitude quaternion and angular velocity of a body where a convenient visualization of the uncertainty exists on $\mathbb{R}^1 \times \mathbb{S}^1$. As the uncertainty was propagated, it wraps around the cylinder defined by $\mathbb{R}^1 \times \mathbb{S}^1$, and the attitude quaternion becomes equiprobable, as expected. The second simulation showed how the Gauss-Bingham density can be used to propagate the uncertainty of the three-dimensional attitude quaternion, angular velocity, relative position, and relative velocity of a chase spacecraft about a target spacecraft. This simulation compared uncertainty propagation using the Gauss-Bingham to that of the predictor of the MEKF and a Monte Carlo approach. When the attitude uncertainty became large,

the uncertainty quantified by the Gauss-Bingham density remained in close agreement with the true density, where that of the MEKF did not.

REFERENCES

- [1] G. Stienne, S. Reboul, M. Azmani, J. B. Choquel, and M. Benjelloun
A multi-temporal multi-sensor circular fusion filter,
Information Fusion, vol. 18, pp. 86–100, 2014.
- [2] G. Kurz, I. Gilitschenski, and U. D. Hanebeck
Recursive Bayesian filtering in circular state spaces,
CoRR, vol. abs/1501.05151, 2015.
- [3] K. V. Mardia and T. W. Sutton
A model for cylindrical variables with applications,
Journal of the Royal Statistical Society. Series B (Methodological), vol. 40, no. 2, pp. 229–233, 1978.
- [4] J. T. Horwood and A. B. Poore
Gauss von Mises distribution for improved uncertainty realism in space situational awareness,
SIAM/ASA Journal of Uncertainty Quantification, vol. 2, pp. 276–304, 2014.
- [5] F. L. Markley and J. L. Crassidis
Fundamentals of Spacecraft Attitude Determination and Control.
Springer, 2014.
- [6] M. D. Shuster
A survey of attitude representations,
The Journal of the Astronautical Sciences, vol. 41, no. 4, pp. 439–517, October–December 1993.
- [7] P. C. Hughes
Spacecraft Attitude Dynamics.
John Wiley & Sons, 1986.
- [8] C. Bingham
An antipodally symmetric distribution on the sphere,
The Annals of Statistics, vol. 2, no. 6, pp. 1201–1225, 1974.
- [9] K. V. Mardia and P. E. Jupp
Directional Statistics.
John Wiley & Sons, 2000.
- [10] G. Kurz, I. Gilitschenski, S. Julier, and U. D. Hanebeck
Recursive Bingham filter for directional estimation involving 180 degree symmetry,
Journal of Advances in Information Fusion, vol. 9, no. 2, pp. 90–105, December 2014.
- [11] J. Glover and L. P. Kaelbling
Tracking the spin on a ping pong ball with the quaternion Bingham filter,
in *IEEE International Conference on Robotics and Automation*, 2014.
- [12] I. Gilitschenski, G. Kurz, S. J. Julier, and U. D. Hanebeck
Unscented orientation estimation based on the Bingham distribution,
IEEE Transactions on Automatic Control, vol. 61, no. 1, pp. 172–177, January 2016.
- [13] ———
A new probability distribution for simultaneous representation of uncertain position and orientation,
in *17th Annual International Conference on Information Fusion*, 2014.
- [14] G. Kurz, I. Gilitschenski, and U. D. Hanebeck
The partially wrapped normal distribution for SE(2) estimation,
in *IEEE International Conference on Multisensor Fusion and Information Integration*, 2014.
- [15] E. A. Wan and R. van der Merwe
The unscented Kalman filter for nonlinear estimation,
in *IEEE Adaptive Systems for Signal Processing, Communications, and Control Symposium*, 2000, pp. 153–158.

- [16] S. J. Julier and J. K. Uhlmann
Unscented filtering and nonlinear estimation,
Proceedings of the IEEE, vol. 92, no. 3, pp. 401–422, March 2004.
- [17] H. M. T. Menegaz, J. Y. Ishihara, G. A. Borges, and A. N. Vargas
A systematization of the unscented Kalman filter theory,
IEEE Transactions on Automatic Control, vol. 60, no. 10, pp. 2583–2598, 2015.
- [18] P. Koev and A. Edelman
Efficient evaluation of the hypergeometric function of a matrix argument,
Mathematics of Computation, vol. 75, no. 254, pp. 833–846, January 2006.
- [19] A. Kume and A. T. A. Wood
Saddlepoint approximations for the Bingham and Fisher-Bingham normalising constants,
Biometrika, vol. 92, no. 2, pp. 465–476, June 2005.
- [20] I. Gilitschenski, G. Kurz, S. J. Julier, and U. D. Hanebeck
Efficient Bingham filtering based on saddlepoint approximations,
in *IEEE International Conference on Multisensor Fusion and Information Integration*, 2014.
- [21] T. Sei and A. Kume
Calculating the normalising constant of the Bingham distribution on the sphere using the holonomic gradient method,
Statistics and Computing, vol. 25, no. 2, pp. 321–332, April 2013.
- [22] J. Glover
“Bingham statistics library,”
2009.
- [23] A. H. Jazwinski
Stochastic Processes and Filtering Theory.
Elsevier, 1970.
- [24] D. Mortari
On the rigid rotation concept in n-dimensional spaces,
The Journal of the Astronautical Sciences, vol. 49, no. 3, pp. 401–420, 2001.
- [25] A. J. Sinclair and J. E. Hurtado
Minimum parameter representations of N-dimensional principal rotations,
Journal of the Astronautical Sciences, vol. 53, no. 3, p. 317, June 2005.
- [26] F. Thomas and A. Pérez-Gracia
On Cayley’s factorization of 4D rotations and applications,
in *6th Conference on Applied Geometric Algebras in Computer Science and Engineering*, 2015.
- [27] K. Levenberg
A method for the solution of certain problems in least squares,
Quarterly of Applied Mathematics, vol. 2, pp. 164–168, 1944.
- [28] D. Marquardt
An algorithm for least-squares estimation of nonlinear parameters,
SIAM Journal Applied Mathematics, vol. 11, pp. 431–441, 1963.
- [29] W. H. Clohessy and R. S. Wiltshire
Terminal guidance system for satellite rendezvous,
Journal of Aerospace Science, vol. 27, no. 9, pp. 653–658, 1960.
- [30] K. T. Alfriend, S. R. Vadali, P. Gurfil, J. P. How, and L. S. Berger
Spacecraft Formation Flying.
Elsevier, 2010.
- [31] D. A. Vallado
Fundamentals of Astrodynamics and Applications.
Microcosm, 2001.
- [32] J. L. Crassidis, F. L. Markley, and Y. Cheng
Survey of nonlinear attitude estimation methods,
Journal of Guidance, Control, and Dynamics, vol. 30, no. 1, pp. 12–28, January–February 2007.
- [33] J. L. Crassidis and J. L. Junkins
Optimal Estimation of Dynamic Systems.
CRC press, 2011.
- [34] G. McLachlan and T. Krishnan
The EM Algorithm and Extensions.
John Wiley & Sons, 2008.



Jacob E. Darling is a Ph.D. candidate in Aerospace Engineering at Missouri University of Science and Technology in Rolla, MO, USA. He received his B.S. in Aerospace Engineering from Missouri University of Science and Technology in 2011. He is a Science, Mathematics, and Research for Transformation (SMART) fellow and will begin employment as a research Aerospace engineer for the United States Air Force Research Laboratory in early 2017. His research interests include dynamic pose estimation, attitude determination, spacecraft navigation, IMU-based navigation, and spacecraft relative motion.



Kyle J. DeMars is an assistant professor of Aerospace Engineering at Missouri University of Science and Technology in Rolla, MO, USA. Dr. DeMars received his Ph.D. degree in 2010, his M.S.E. degree in 2007, and his B.S. degree in 2004, all in Aerospace Engineering from The University of Texas at Austin. His research interests are in nonlinear uncertainty prediction and Bayesian inference, orbit determination, attitude dynamics and determination, data association, and autonomous sensor management.

Multivariate Angular Filtering Using Fourier Series

FLORIAN PFAFF
GERHARD KURZ
UWE D. HANEBECK

Filtering for multiple, possibly dependent angular variates on higher-dimensional manifolds such as the hypertorus is challenging as solutions from the circular case cannot easily be extended. In this paper, we present an approach to recursive multivariate angular estimation based on Fourier series. Since only truncated Fourier series can be used in practice, implications of the approximation errors need to be addressed. While approximating the density directly can lead to negative function values in the approximation, this problem can be solved by approximating the square root of the density. As this comes at the cost of additional complexity in the algorithm, we present both a filter based on approximating the density and a filter based on approximating its square root and closely regard the trade-offs. While the computational effort required for the filters grows exponentially with increasing number of dimensions, our approach is more accurate than a sampling importance resampling particle filter when comparing configurations of equal run time.

Manuscript received December 30, 2015; revised November 8, 2016; released for publication April 3, 2017.

Refereeing of this contribution was handled by William Blair.

Authors' address: Intelligent Sensor-Actuator-Systems Laboratory (ISAS), Institute for Anthropomatics and Robotics, Karlsruhe Institute of Technology (KIT), Germany (E-mail: florian.pfaff@kit.edu gerhard.kurz@kit.edu uwe.hanebeck@ieee.org).

1557-6418/16/\$17.00 © 2016 JAIF

I. INTRODUCTION

Periodic quantities are ubiquitous both in nature [1], [2], [3] and technology [4], [5], [6]. The most common periodic quantities are in the form of angles, such as the orientation in a two-dimensional space, but a variety of other periodic quantities exist, such as the phase of a signal [7], [8]. When dealing with orientations, we are usually only interested in the current orientation in our coordinate system and do not aim to count how often the object has revolved around the axis of rotation. While neglecting the latter simplifies the task, estimators need to be carefully crafted to properly account for the effects of periodicity.

For recursive Bayesian estimation, uncertainties in the system and measurement models have to be represented, e.g., via transition densities and likelihoods. Filters on linear domains that assume the support of the prior and posterior densities to be unbounded, such as the Kalman filter, have underlying assumptions that are incompatible with periodic manifolds. However, most periodic domains are locally similar to linear ones when regarding a very narrow region of the domain. Since the importance of a region to the estimation problem strongly depends on the probability mass in the region, densities that are concentrated on a very narrow region can still be handled with sufficient accuracy using approaches that rely on the linearity of the domain. Therefore, problems featuring very little uncertainty can still be handled properly using a modified Kalman filter or an unscented Kalman filter [9].

However, the wider the probability mass is spread on the domain, the less the estimation problem behaves like on a linear domain. For higher uncertainties, filters assuming linearity of the domain degrade and can become entirely misleading. In these cases, approaches based on directional statistics specifically crafted for periodic domains become a necessity. Directional statistics [10], [11] puts the focus on properly handling periodic manifolds by providing many analogues to concepts that are commonly used on linear manifolds. Fields in which directional statistics is applied include, e.g., geosciences [1], biology [2], [12], analysis of crystal structures [13], scattering theory [14], MIMO radar systems [15], robotics [4], and signal processing [5], [6], [16].

The main focus of directional statistics is on two classes of topologies. One of the two classes is the




Fig. 1. A mixture of two bivariate wrapped normal distributions shown as a heatmap. The two modes are shown in dark red.

hypersphere S^d , which is the surface of the unit ball in \mathbb{R}^{d+1} . The other important topology is the Cartesian product of multiple circles $S^1 \times \dots \times S^1$ called hypertorus, which will be the focus of this paper. The hypertorus is the proper topology to use when dealing with multiple, possibly stochastically dependent angles in the range of $[0, 2\pi)$. Examples featuring a hypertoroidal topology include angles at different points in time or connected rotatory robotic joints. If the variates of the random vector are stochastically independent, they can be treated as independent random variables and filters for circular topologies can be used. For the circle, several filters have been proposed [17], [18], [19], [20], [21], [22]. However, filters for multivariate angular problems are necessary if the random vector cannot be separated into multiple independent random variables. For example, if the random variables \mathbf{x} and \mathbf{y} are independent, the two variates of the vector $[\mathbf{x} \ \mathbf{x} + \mathbf{y}]^T$ are (usually, but not necessarily) dependent. They can, however, be easily transformed into a problem featuring independent variates that can be estimated independently. The obtained estimates can then be transformed back for use in the original problem featuring dependent variates. In practice, transforming the problem in a way that causes the components to be independent is usually nontrivial and frequently impossible.

On linear domains, an important filter is the Kalman filter that scales well with increasing number of variates and yields optimal estimation results if certain conditions are met. The Kalman filter can handle correlations and scales no more than cubically in the number of variates. For nonlinear system or measurement models, such an efficient and optimal solution does not exist in general. Therefore, even on linear domains, nonlinear estimation is a challenging task, especially for multivariate problems. On periodic domains, estimation problems are inherently nonlinear and thus challenging to deal with for higher numbers of variates.

For multivariate angular estimation, a recursive filter has been suggested in [23] for the two-dimensional torus. However, due to the approximation with a single bivariate wrapped normal distribution, multimodal bivariate posteriors, such as the one shown in Fig. 1, cannot be modeled adequately. As an alternative for toroidal and hypertoroidal estimation problems, the very general concept of a sampling importance resampling (SIR) particle filter [24] that is popular on linear domains can be adapted in a straightforward manner to periodic manifolds.

In this paper, we generalize a recursive Bayesian estimator based on Fourier series, which we proposed for univariate densities on the circle in [21]. Our new approach to multivariate angular estimation problems on the hypertorus is based on representing the density or its square root using a multidimensional Fourier series. Based on the representation used, we either call the filter the (angular) *Fourier identity filter* or the (angular)

Fourier square root filter and we abbreviate their names as IFF and SqFF. We show that the prediction and filter operations can be performed in computationally efficient ways, never exceeding an asymptotic complexity of $O(n \log n)$ for n Fourier coefficients and a fixed number of variates. However, for estimation problems with a higher number of variates, it is advisable to use more Fourier coefficients.

The rest of this paper is structured as follows. We give a brief explanation on how to perform recursive Bayesian estimation in general and lay out the required operations in the next section. We introduce the basics of Fourier series and directional statistics that are a prerequisite for our proposed approach in the third and fourth section. In the fifth section, we address related work and explicate the key idea of our proposed filters. In Sec. VI, we introduce an actual implementation of the Fourier filters. A comparison of our proposed filters and an evaluation comparing the Fourier filters with other approaches is given in Sec. VII. A conclusion and an outlook is provided in Sec. VIII. Finally, we present some useful properties of the Fourier series representations in the appendix.

II. RECURSIVE BAYESIAN ESTIMATION

While a lot of effort in estimation theory is geared towards estimating the mean of the posterior density, keeping track of the whole density is usually required for accurate results over multiple time steps. The focus on approximating the mean can be explained by the fact that on linear domains, the mean of the posterior density is the minimum mean squared error estimator [25, Sec. 10.3]. However, if the density is to be reused in future time steps, the mean of the posterior density itself does in general not suffice for an accurate calculation of the mean of the posterior density in future time steps. In the case of linear systems with Gaussian noise on linear domains, keeping track of the whole density is easy—the resulting Gaussian posterior density is precisely described by the mean vector and the covariance matrix.

For noise terms that are not exactly Gaussian distributed and especially for nonlinear models, describing the precise posterior density using a limited number of parameters is challenging or impossible even on linear domains. Many estimators focus on approximating the posterior density using a parameterized density of a prespecified family of densities. For example, Gaussian assumed density filters [26], [27], [28], [29] try to find a suitable Gaussian approximation for the true posterior density. To increase the accuracy of the approximation, Gaussian mixtures [26] can be used but mixtures entail further problems concerning component reduction [30]. Another popular approach is to use an SIR particle filter [24]. While simple particle filters can be used to asymptotically approximate moments such as the mean over multiple time steps, they are non-deterministic and do not directly provide a continuous approximation of the density.

In the subsections of this section, we introduce the general formulae that can be used for recursive Bayesian estimation with given likelihoods and transition densities and lay out the required operations for implementing a recursive Bayesian estimator. The formulae presented are suitable for manifolds with a topological group structure, such as \mathbb{R}^n and the hypertorus.

A. Prediction Step

The prediction step can be described by the Chapman–Kolmogorov equation

$$\begin{aligned} f_{t+1}^p(\mathbf{x}_{t+1} | \underline{z}_1, \dots, \underline{z}_t) \\ = \int_{\Omega_x} f_t^T(\mathbf{x}_{t+1} | \mathbf{x}_t) f_t^e(\mathbf{x}_t | \underline{z}_1, \dots, \underline{z}_t) d\mathbf{x}_t, \end{aligned}$$

in which Ω_x denotes the sample space—for hypertori, $[0, 2\pi)^d$ — f_t^T the transition density, f_t^e the posterior density, and f_{t+1}^p the prior density based on measurements up to the time step t . For identity models involving additive noise on linear domains, we can write

$$\mathbf{x}_{t+1} = \mathbf{x}_t + \mathbf{w}_t, \quad (1)$$

with \mathbf{w}_t distributed according to f_t^w . On periodic domains, this becomes

$$\mathbf{x}_{t+1} = (\mathbf{x}_t + \mathbf{w}_t) \bmod 2\pi,$$

which includes a nonlinear transformation. In the case of additive noise, we can simplify the formula for f_{t+1}^p to

$$\begin{aligned} f_{t+1}^p(\mathbf{x}_{t+1} | \underline{z}_1, \dots, \underline{z}_t) \\ = \int_{\Omega_x} f_t^w(\mathbf{x}_{t+1} - \mathbf{x}_t) f_t^e(\mathbf{x}_t | \underline{z}_1, \dots, \underline{z}_t) d\mathbf{x}_t \\ = (f_t^w * f_t^e)(\mathbf{x}_{t+1}) \end{aligned}$$

for both linear and circular domains and are able to perform the prediction step by calculating a continuous convolution.

B. Filter Step

If we obtain a measurement and know the corresponding measurement likelihood, we can use Bayes' formula for the filter step. This essential concept can be formulated as

$$\begin{aligned} f_t^e(\mathbf{x}_t | \underline{z}_1, \dots, \underline{z}_t) &= \frac{f_t^L(\underline{z}_t | \mathbf{x}_t) f_t^p(\mathbf{x}_t | \underline{z}_1, \dots, \underline{z}_{t-1})}{\int_{\Omega_x} f_t^L(\underline{z}_t | \mathbf{x}_t) f_t^p(\mathbf{x}_t | \underline{z}_1, \dots, \underline{z}_{t-1}) d\mathbf{x}_t} \\ &\propto f_t^L(\underline{z}_t | \mathbf{x}_t) f_t^p(\mathbf{x}_t | \underline{z}_1, \dots, \underline{z}_{t-1}), \end{aligned}$$

with the likelihood function f_t^L . It is important to note that the denominator is independent of \mathbf{x}_t and can thus be treated as a constant. Since we know that a proper pdf integrates to one, we can ignore the denominator if we have other means to normalize the density.

III. BASICS OF DIRECTIONAL STATISTICS

In this section, we introduce important concepts of directional statistics. In directional statistics, there are

counterparts to many important concepts used in the context of linear domains, some of which are addressed in this section. We always assume that our periodic region has a size of 2π along each dimension. While most formulae given in this paper do not explicitly depend on the precise region used, we say that our periodic quantities are always in $[0, 2\pi)$ in the scalar case and in $[0, 2\pi)^d$ in the d -variate case.

In this paper, we use the von Mises distribution as well as the wrapped normal distribution and generalize the latter to an arbitrary number of variates. We chose to use the multivariate wrapped normal distribution as it can be trivially generalized from its bivariate definition. Another important concept is that of trigonometric moments. While these moments are seldom used on linear domains, they are a useful concept to employ instead of power moments when dealing with periodic densities. We also introduce concepts for describing correlations between the variates of a random vector on hypertoroidal manifolds. For further reading, we recommend the two classic books about directional statistics [10], [11].

A. Von Mises Distribution

The von Mises distribution [10, Sec. 3.5], also called the circular normal distribution [11, Sec. 2.2.4], is a popular circular distribution. A useful property of this distribution is that the product of two von Mises distributions yields an (unnormalized) von Mises distribution again. The density of the von Mises distribution is given by

$$f_{\text{VM}}(x; \mu, \kappa) = \frac{e^{\kappa \cos(x-\mu)}}{2\pi I_0(\kappa)},$$

with $I_0(\cdot)$ being the modified Bessel function of the first kind, $\mu \in [0, 2\pi)$ being the location parameter, and $\kappa \geq 0$ describing its concentration.

B. Wrapped Normal Distribution

The wrapped normal distribution [10, Sec. 3.5] can be, visually speaking, obtained by wrapping a Gaussian distribution around the circle and summing up all probability mass at each point. The density is defined as

$$f_{\text{WN}}(x; \mu, \sigma^2) = \sum_{j \in \mathbb{Z}} \mathcal{N}(x + 2\pi j; \mu, \sigma^2),$$

in which we parameterize the density based on the mean $\mu \in [0, 2\pi)$ and the variance σ^2 of the underlying normal distribution. We use the variance instead of the standard deviation to provide a definition that is consistent with the multivariate case introduced in the next subsection.

C. Multivariate Wrapped Normal Distribution

The concept of a wrapped normal distribution can be generalized to higher dimensions, e.g., to the torus as a bivariate wrapped normal (also called wrapped bivariate normal) distribution [11, Sec. 2.3.2]. For the d -variate wrapped normal distribution, we wrap a

d -variate normal distribution onto the d -dimensional hypertorus. This leads to the formula

$$f_{\text{WN}}(\underline{x}; \underline{\mu}, \mathbf{C}_{\text{WN}}) = \sum_{\underline{j} \in \mathbb{Z}^d} \mathcal{N}(\underline{x} + 2\pi \underline{j}; \underline{\mu}, \mathbf{C}_{\text{WN}})$$

for the density of the d -variate distribution with the vector-valued mean $\underline{\mu} \in [0, 2\pi)^d$ and covariance matrix \mathbf{C}_{WN} .

D. Trigonometric Moments

One important concept on linear domains are power moments, usually simply referred to as moments. Moments describe useful properties of the distribution and are important for estimators. As previously mentioned, the first moment of the posterior density is the MMSE estimator on linear domains. Distributions of certain types can be parametrized by some of their moments, e.g., for the Gaussian distribution, the combination of the first moment and the second central moment (or its root) is the most commonly used parameterization.

On periodic manifolds, trigonometric (also called raw) moments feature some of the properties that power moments have on linear domains. The k th trigonometric moment ($k \in \mathbb{N}$) for scalar random variables is given by [11, Sec. 2.1]

$$m_k = \mathbb{E}(e^{ikx}) = \int_0^{2\pi} f(x) e^{ikx} dx.$$

It is also common practice to write trigonometric moments as vectors instead of as a complex number. For this representation, used for example in [10, Sec. 3.4.1], the parts represented by the real and imaginary part are calculated using separate integrals. However, using Euler's formula, it can be shown that the conversion from one representation to the other is straightforward.

Unlike power moments, trigonometric moments consist of two values—either the real and complex parts or the two components of the vector—and thus, a single moment can describe multiple properties. For example, the first trigonometric moment is not only a measure of the density's position but also of its dispersion. The parameters for some distributions, such as the wrapped normal and the von Mises distribution, can be derived from the first trigonometric moment [18]. As trigonometric moments express a lot about the density, several filters in the circular case approximate trigonometric moments and make use of them. Due to the close relationship of some of the Fourier coefficients to trigonometric moments, the Fourier filters are no exception.

For the k th moment of multivariate densities, we simply stack all k th moments of all variates. Thus, in

the d -variate case, the k th moment is given by

$$\underline{m}_k = \begin{bmatrix} m_{k,1} \\ m_{k,2} \\ \vdots \\ m_{k,d} \end{bmatrix} = \begin{bmatrix} \mathbb{E}(e^{ikx_1}) \\ \mathbb{E}(e^{ikx_2}) \\ \vdots \\ \mathbb{E}(e^{ikx_d}) \end{bmatrix} = \begin{bmatrix} \int_{[0,2\pi)^d} f(\underline{x}) e^{ikx_1} d\underline{x} \\ \int_{[0,2\pi)^d} f(\underline{x}) e^{ikx_2} d\underline{x} \\ \vdots \\ \int_{[0,2\pi)^d} f(\underline{x}) e^{ikx_d} d\underline{x} \end{bmatrix}. \quad (2)$$

E. Circular Mean Direction

The circular mean direction, which can be thought of as an analogue to the linear mean, only describes the density's location and can be calculated from the first trigonometric moment via [10, Sec. 2.2]

$$\mu = \text{atan2}(\mathcal{I}(m_1), \mathcal{R}(m_1)). \quad (3)$$

A useful property on linear domains is the linearity of the expected value, which does not hold for the circular mean direction [11, Sec. 2.2.1]. For the circular mean direction in the d -variate case, we simply calculate the circular mean direction for every component of the moment vector according to (3), yielding

$$\underline{\mu} = \begin{bmatrix} \text{atan2}(\mathcal{I}(m_{1,1}), \mathcal{R}(m_{1,1})) \\ \vdots \\ \text{atan2}(\mathcal{I}(m_{1,d}), \mathcal{R}(m_{1,d})) \end{bmatrix}.$$

F. Angular Correlation Coefficients

A variety of measures of correlation of two angular random variables have been introduced in the literature. Examples include the correlation coefficients by Jammalamadaka and Sarma [31], Johnson and Wehrley [32], and Jupp and Mardia [33]. One of the correlation coefficients is used by the only assumed density filter for bivariate toroidal problems [23] that we know of. As shown in Appendix D, a limited number of Fourier coefficients contain all information necessary to calculate all correlation coefficients mentioned.

Key to this is the close relationship of the correlation coefficients to a certain covariance matrix. In the bivariate case, all of these correlation coefficients can be calculated using entries of the covariance matrix

$$\Sigma = \mathbb{E} \left(\left(\begin{bmatrix} \cos(x_1) \\ \sin(x_1) \\ \cos(x_2) \\ \sin(x_2) \end{bmatrix} - \underline{\mu}_c \right) \left(\begin{bmatrix} \cos(x_1) \\ \sin(x_1) \\ \cos(x_2) \\ \sin(x_2) \end{bmatrix} - \underline{\mu}_c \right)^T \right)$$

with

$$\underline{\mu}_c = \mathbb{E}([\cos(x_1) \quad \sin(x_1) \quad \cos(x_2) \quad \sin(x_2)]^T),$$

which can be calculated via

$$\underline{\mu}_c = [\mathcal{R}(m_{1,1}) \quad \mathcal{I}(m_{1,1}) \quad \mathcal{R}(m_{1,2}) \quad \mathcal{I}(m_{1,2})]^T.$$

Thus, efficient calculation of this matrix allows efficient calculation of all correlation coefficients. It is easy to

extend this covariance matrix to arbitrary multivariate distributions by stacking the terms for all individual variates, yielding

$$\Sigma = \mathbb{E} \left(\left(\begin{pmatrix} \cos(x_1) \\ \sin(x_1) \\ \vdots \\ \cos(x_d) \\ \sin(x_d) \end{pmatrix} - \underline{\mu}_c \right) \left(\begin{pmatrix} \cos(x_1) \\ \sin(x_1) \\ \vdots \\ \cos(x_d) \\ \sin(x_d) \end{pmatrix} - \underline{\mu}_c \right)^T \right)$$

with

$$\underline{\mu}_c = \mathbb{E}([\cos(x_1) \quad \sin(x_1) \quad \cdots \quad \cos(x_d) \quad \sin(x_d)]^T).$$

IV. BASICS OF FOURIER SERIES

In this section, we give a brief introduction to multidimensional Fourier series, the second concept essential to this paper. For details regarding Fourier series, we refer the reader to the two-volume book series about trigonometric series by Zygmund [34] and books about harmonic analysis [35].

A. One-Dimensional Fourier Series

Using Fourier series, it is possible to approximate functions on $[0, 2\pi)$ using complex exponential functions. The set of functions

$$\{e^{ikx} \mid k \in \mathbb{Z}\}$$

is an orthogonal basis that can be used to represent any square-integrable (also called square-summable) complex function defined on $[0, 2\pi)$ using a square-summable sequence of Fourier coefficients [35, Sec. I-5]. Since densities encountered in practice are usually square-integrable, we can write their density $f(x)$ as a Fourier series

$$f(x) = \sum_{k=-\infty}^{\infty} c_k e^{ikx}, \quad (4)$$

where the Fourier coefficients $c_k \in \mathbb{C}$ fulfill

$$\sum_{k \in \mathbb{Z}} |c_k|^2 < \infty.$$

The Fourier coefficients can be calculated from the density according to

$$c_k = \frac{1}{2\pi} \int_0^{2\pi} f(x) e^{-ikx} dx.$$

For real functions, $c_{-k} = \bar{c}_k$ holds [34, Ch. I], causing imaginary parts to cancel out in (4). Furthermore, it is also possible to use real basis functions and coefficients to represent real functions. In this alternative representation, the series in (4) becomes a weighted sum of sine and cosine functions of different frequencies.

B. Higher-Dimensional Fourier Series

The straightforward generalization of one-dimensional Fourier series to the d -dimensional case is to use

functions of the orthogonal system

$$\{e^{i(k_1 x_1 + k_2 x_2 + \cdots + k_d x_d)} \mid \underline{k} \in \mathbb{Z}^d\}$$

to represent functions on the d -dimensional hypercube [34, Ch. XVII] (or in our case hypertorus) $[0, 2\pi)^d$. In the following, we write the basis functions described above using a dot product as $e^{i\mathbf{k} \cdot \mathbf{x}}$ and use a vector-valued index $\underline{k} \in \mathbb{Z}^d$ to specify individual entries $c_{\underline{k}}$ of the d -dimensional Fourier coefficient tensor. Using this notation, a multidimensional Fourier series can be written as

$$f(\mathbf{x}) = \sum_{\underline{k} \in \mathbb{Z}^d} c_{\underline{k}} e^{i\mathbf{k} \cdot \mathbf{x}}.$$

The individual Fourier coefficients can then be calculated via

$$c_{\underline{k}} = \frac{1}{(2\pi)^d} \int_{[0, 2\pi)^d} f(\mathbf{x}) e^{-i\mathbf{k} \cdot \mathbf{x}} d\mathbf{x}. \quad (5)$$

In this paper, our focus is on Fourier series for which only n specific coefficients are nonzero and we only consider sets of indices \mathcal{J} that are subsets of the integer lattice \mathbb{Z}^d with an equal subset of \mathbb{Z} in each dimension. If every subset of \mathbb{Z} ranges from $-k_{\max}$ to k_{\max} in each dimension, the total number of Fourier coefficients is $n = (2k_{\max} + 1)^d$. Similar to the one-dimensional case, $c_{-\underline{k}} = \bar{c}_{\underline{k}}$ holds for real functions.

REMARK 1. The formulae for the Fourier coefficients bear a close resemblance to the trigonometric moments. The k th trigonometric moment \underline{m}_k can be calculated from the Fourier coefficients via

$$\underline{m}_k = (2\pi)^d [c_{-k, 0, \dots, 0} \quad c_{0, -k, 0, \dots, 0} \quad \cdots \quad c_{0, \dots, 0, -k}]^T.$$

Thus, all trigonometric moments and especially the first trigonometric moment required for the calculation of the circular mean direction can be calculated efficiently from the Fourier coefficients. On the other hand, calculating arbitrary Fourier coefficients from the trigonometric moments is, in general, only possible in the one-dimensional case. For higher dimensions, many entries of the Fourier coefficient tensor do not have corresponding entries in the moment vectors as defined in (2). We visualize this for the two-dimensional case in Fig. 2. In Appendix D, we show that there is a relationship between the covariance matrix described in Sec. III-F and other entries of the Fourier coefficient tensor.

V. RELATED WORK AND KEY IDEA

For our approach, it is important to note that many important multivariate densities—such as the multivariate wrapped normal density—are square-integrable and thus lend themselves well to approximations using Fourier series. Mardia also states this observation for the univariate case in his book [10, Ch. 3–4], noting that if the Fourier coefficients are square-summable, the corresponding density is equal to the Fourier series almost everywhere.

Willsky discusses optimal filtering using non-truncated Fourier series with an infinite number of co-




Fig. 2. Visualization of the relationship between the Fourier coefficients and the trigonometric moments of a bivariate distribution. Fourier coefficients that are unrelated to all trigonometric moments are shown in blue, coefficients that are related to the first entry of a trigonometric moment vector are shown in green, and coefficients that are related to the second entry are shown in yellow. The zeroth coefficient in the middle determines both entries of the zeroth trigonometric moment vector and is identical for all (normalized) densities.

efficients in [6]. Since infinite series cannot be handled computationally, he discusses practical implementations in [16]. Willsky deems the performance of a filter working with truncated Fourier series to be insufficient for few coefficients and suggests making the assumption that the density is distributed according to a wrapped normal distribution, which we believe is too restrictive.

Fernández-Durán [36] observes that when approximating densities using Fourier series, the truncation of the coefficient vector can cause negative function values and suggests a computationally expensive way to ensure nonnegativity. In the context of nonlinear filtering for linear domains, Brunn et al. [37], [38] argue that approximating a transformed version of the pdf allows the reconstruction of a valid density with only nonnegative values in every time step.

Based on this idea, we have presented a filter for univariate periodic densities in [21] that ensures the validity of a density function

$$f : [0, 2\pi) \rightarrow \mathbb{R}_0^+$$

by approximating the square root of the density. Approximating the square root $g(x) = \sqrt{f(x)}$ of the density is reasonable as the square root of every density is square-integrable since

$$\int_{[0, 2\pi)} g(x)^2 dx = \int_{[0, 2\pi)} f(x) dx = 1$$

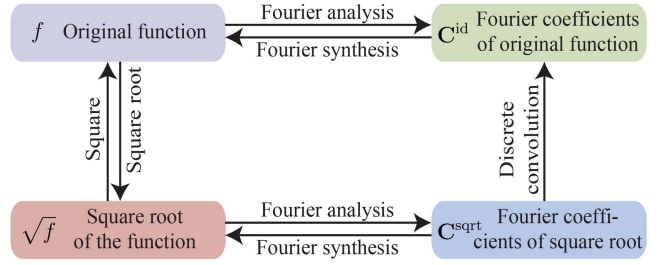


Fig. 3. Different representations that are used by the filters.

holds. Therefore, the Fourier coefficients are square-summable and thus converge to zero, facilitating the approximation via a Fourier series. Furthermore, if $g_{k_{\max}}$ denotes the truncated Fourier series respecting the coefficients from $-k_{\max}$ to k_{\max} , the convergence of

$$g_{k_{\max}} \xrightarrow{k_{\max} \rightarrow \infty} \sqrt{f} \text{ also implies } g_{k_{\max}}^2 \xrightarrow{k_{\max} \rightarrow \infty} f$$

almost everywhere. For this representation, we showed in [21] how prediction and filter steps can be calculated efficiently and accurately using the Fourier coefficients for univariate estimation problems. In a scenario featuring bimodality, the filter presented outperformed both a grid filter and a particle filter.

In this paper, we generalize the approach presented in [21] to higher dimensions and approximate multivariate periodic densities or their square root using multi-dimensional Fourier series. Based on these approximations, we describe how the prediction and filter steps of a recursive Bayesian estimator can be calculated efficiently in the multivariate case. Since run time is crucial for high numbers of variates, we put emphasis on comparing configurations of equal run time. Furthermore, we closely regard the IFF, which has disadvantages in theory but features a lower run time when using an identical number of coefficients.

VI. RECURSIVE BAYESIAN ESTIMATION BASED ON FOURIER SERIES

In this section, we present how a recursive Bayesian estimator can be implemented based on approximating the density or its square root using a Fourier series. For the SqFF, we are dealing with a total of four representations that can be used to obtain values of the probability density (see Fig. 3). These include the original function, the Fourier coefficients of the original function, the square root of the function, and the Fourier coefficients thereof. Only the first two representations are used in the IFF. Deriving a Fourier series representation for a function is commonly referred to as Fourier analysis whereas reconstructing the function is frequently called Fourier synthesis. As the different representations have implications for the algorithmic implementation of the operations, we will make a clear notational distinction between Fourier coefficient tensors representing the actual density \mathbf{C}^{id} and coefficient tensors representing the

square root of the density \mathbf{C}^{sqrt} and apply the same notation to the individual entries of the tensors denoted by c_k^{id} and c_k^{sqrt} with the vector-valued index \underline{k} .

For the upper two representations in Fig. 3, there are important dualities that we make use of. The convolution of two functions corresponds to a Hadamard (entrywise) product of their Fourier coefficient tensors. Furthermore, the multiplication of two functions represented as a d -dimensional Fourier series can be performed using a d -dimensional discrete convolution of the coefficient tensors.

In order to perform (or at least approximate) the prediction and filter steps described in Sec. II, we first need to be able to transform arbitrary densities and likelihoods into the two Fourier series representations. Second, we need to be able to perform multiplications and normalizations for the filter step in the respective representation. Third, convolutions are necessary to perform prediction steps for an identity model with additive noise (1). While exact results for both operations can be obtained in the IFF, an increase in the number of coefficients would be inevitable. While we may allow the number of coefficients to vary over time, parameter reduction becomes inevitable in the long run. Parameter reduction generally induces an approximation and is also necessary for the SqFF. For our implementation of the filters, we do not allow the number of coefficients to vary over time and truncate to an identical number of coefficients after each prediction and filter step.

In the following, we describe how the filter step, the prediction step, and the parameter reduction can be performed in $O(n \log n)$ for n Fourier coefficients. Further properties that are useful to applying the filter in practice but are not essential to the filter and prediction step are given in the appendix.

A. Transforming Multivariate Densities

An efficient approach to Fourier analysis was proposed by Cooley and Tukey [39]. This approach has become widespread for calculating the closely related discrete Fourier transform [40, Ch. 2] and is nowadays known as the fast Fourier transform (FFT) [41], which is also how we will refer to it and its higher-dimensional generalizations for the remainder of this paper.

The complexity of the FFT is $O(n \log n)$ for a total number of n coefficients. However, for a fixed k_{max} , the number of coefficients still grows exponentially with the dimensionality of the space. This is not a serious problem for low numbers of variates and we have verified good filter results with fast run times for densities with up to five variates. To obtain a Fourier series approximation of the square root of a density using the FFT, there is (aside from calculating the square root of each function value) no additional overhead involved. In the one-dimensional case, we were able to derive closed-form formulae for the coefficients for many important univariate densities [21]. In Appendix A of this

paper, we provide the formula for the Fourier coefficients of the multivariate wrapped normal distribution. While closed-form formulae can lead to lower run times, the cost of calculating n coefficients is always at best in $O(n)$ when at least n coefficients are required for an exact representation of the density.

For identity system and measurement models with independent, time-invariant additive noise terms, we can reduce the computational effort involved in obtaining the required Fourier coefficients. In these cases, it is not necessary to transform the density of the system noise and the likelihood in each time step. The density of the system noise simply stays identical while the likelihood is only influenced by the measurement via a shift.

For example, for an additive system noise that is distributed according to a multivariate wrapped normal distribution, the density of the system noise is

$$f_t^w(\underline{w}_t) = f_{\text{WN}}(\underline{w}_t; \underline{\mu}, \mathbf{C}_{\text{WN}})$$

in every time step. Similarly, we can avoid the need for transforming the likelihood multiple times for additive noise terms. If the likelihood is

$$f_t^L(\underline{z}_t | \underline{x}_t) = f_{\text{WN}}(\underline{z}_t; \underline{x}_t, \mathbf{C}_{\text{WN}}) = f_{\text{WN}}(\underline{x}_t; \underline{z}_t, \mathbf{C}_{\text{WN}})$$

we can initially transform the likelihood when setting $\underline{z} = \underline{0}$ (meaning, $f_{\text{WN}}(\underline{x}; \underline{0}, \mathbf{C}_{\text{WN}})$ in our case) and then calculate the Fourier coefficients of the actual likelihood respecting the current measurement \underline{z}_t from these coefficients. The individual Fourier coefficients c_k^{shifted} for a function shifted by \underline{z} can be calculated according to

$$c_k^{\text{shifted}} = c_k e^{-i\mathbf{k} \cdot \underline{z}}, \quad (6)$$

which is a straightforward generalization of the shifting operation for the scalar case (Theorem 1.1 (iv) in [34, Sec. II-1]).

B. Filter Step

To implement the filter step, two operations have to be performed for the two filters. The first operation is a multiplication of the prior density $f_t^p(\underline{x}_t | \underline{z}_1, \dots, \underline{z}_{t-1})$ and the likelihood $f_t^L(\underline{z}_t | \underline{x}_t)$. The coefficient tensors for the intermediate, unnormalized results will be called $\mathbf{C}_t^{e,\text{id}}$ and $\mathbf{C}_t^{e,\text{sqrt}}$. In the second step, the densities are normalized to yield the coefficient tensors $\mathbf{C}_t^{e,\text{id}}$ and $\mathbf{C}_t^{e,\text{sqrt}}$ to be used in the next prediction (or filter) step. The filter steps for the two Fourier filters are illustrated in Fig. 4 and the necessary operations are explained in detail in the following.

1) **Multiplication of Two Densities:** The first operation necessary to perform the filter step of our Bayesian filter is the multiplication operation. Let us now denote the Fourier coefficient tensor of $f_t^p(\underline{x}_t | \underline{z}_1, \dots, \underline{z}_{t-1})$ as $\mathbf{C}_t^{p,\text{id}}$ and refer to the Fourier coefficient tensor of $f_t^L(\underline{z}_t | \underline{x}_t)$ (for a fixed measurement \underline{z}_t , depending only on the state \underline{x}_t) as $\mathbf{C}_t^{L,\text{id}}$. For the IFF, we can then directly obtain




Fig. 4. Filter step of the IFF and the SqFF. (a) Illustration of the filter step of the IFF. (b) Illustration of the filter step of the SqFF.

the coefficient tensor $\check{C}_t^{e,id}$ for the new Fourier series representing the unnormalized multiplication of the two functions via

$$\check{C}_t^{e,id} = C_t^{p,id} * C_t^{L,id},$$

in which $*$ denotes the discrete convolution.

The discrete convolution can be performed in $O(n \log n)$ using FFT-based convolution approaches or by using alternative convolutions methods tailored to the multidimensional tensor convolution [42]. Since the discrete convolution of two tensors results in a larger tensor, parameter reduction as explained in Sec. VI-C becomes necessary.

The multiplication can be performed similarly for the SqFF. Owing to the fact that for all functions $f_t^p(\underline{x}_t | \underline{z}_1, \dots, \underline{z}_{t-1})$ and $f_t^L(\underline{z}_t | \underline{x}_t)$ and for all vectors \underline{x}_t and \underline{z}_t in the respective domains

$$\begin{aligned} & \sqrt{f_t^p(\underline{x}_t | \underline{z}_1, \dots, \underline{z}_{t-1})} \cdot \sqrt{f_t^L(\underline{z}_t | \underline{x}_t)} \\ &= \sqrt{f_t^p(\underline{x}_t | \underline{z}_1, \dots, \underline{z}_{t-1}) f_t^L(\underline{z}_t | \underline{x}_t)} \end{aligned}$$

holds, we can simply multiply the functions in the square root representation. Thus, the multiplication can be performed analogously to the IFF. For the Fourier coefficient tensors $C_t^{p,sqrt}$ of $\sqrt{f_t^p(\underline{x}_t | \underline{z}_1, \dots, \underline{z}_{t-1})}$ and $C_t^{L,sqrt}$ of $\sqrt{f_t^L(\underline{z}_t | \underline{x}_t)}$ (again, for a fixed measurement), we can calculate the unnormalized coefficient tensor $\check{C}_t^{e,sqrt}$ in the square root representation using

$$\check{C}_t^{e,sqrt} = C_t^{p,sqrt} * C_t^{L,sqrt}.$$

2) Normalization: The second operation necessary for the filter step is the normalization. We use $\check{c}_{t,\underline{k}}$ to refer to an entry of the coefficient tensor to be normalized, such as $\check{C}_t^{e,id}$ or $\check{C}_t^{e,sqrt}$ as obtained from the multiplication above. For the normalization, we need to integrate over the whole domain of the function, which is easy when a Fourier series is used to represent a real function. In the integral, all terms except the one stemming from the first coefficient integrate to zero. This is obvious since the exponential functions for all other

coefficients can be converted (pairwise) into sine and cosine functions of differing (but nonzero) frequencies using Euler's formula. Integrating these trigonometric functions over $[0, 2\pi)$ always yields zero. Therefore, the integral over the whole function can be calculated according to

$$\int_{[0, 2\pi)^d} \sum_{\underline{k} \in \mathbb{Z}^d} \check{c}_{t,\underline{k}}^{id} e^{i\mathbf{k} \cdot \mathbf{x}} d\mathbf{x} = \int_{[0, 2\pi)^d} \check{c}_{t,\underline{0}}^{id} d\mathbf{x} = (2\pi)^d \check{c}_{t,\underline{0}}^{id}$$

in the non-rooted representation. Thus, we can normalize the density by dividing all coefficients by $(2\pi)^d \check{c}_{t,\underline{0}}^{id}$, which ensures that the coefficient with index $\underline{0}$ of the new, normalized density is $c_{t,\underline{0}}^{id} = 1/(2\pi)^d$.

For the square root version, we can calculate the integral of the function by determining the Fourier coefficient with index $\underline{0}$ of the non-rooted representation. The coefficient $\check{c}_{t,\underline{0}}^{id}$ can be calculated from the coefficients $\check{c}_{t,\underline{k}}^{sqrt}$ of the square root representation via

$$\check{c}_{t,\underline{0}}^{id} = \sum_{\underline{k} \in \mathcal{J}} \check{c}_{t,\underline{k}}^{sqrt} \check{c}_{t,-\underline{k}}^{sqrt} = \sum_{\underline{k} \in \mathcal{J}} |\check{c}_{t,\underline{k}}^{sqrt}|^2.$$

Based on this insight, we can ensure that the density represented by the Fourier series integrates to one by dividing all $\check{c}_{t,\underline{k}}^{sqrt}$ by $\sqrt{(2\pi)^d \sum_{\underline{k} \in \mathcal{J}} |\check{c}_{t,\underline{k}}^{sqrt}|^2}$. Since calculating the sum of the square of the absolute values is in $O(n)$ and the division of all coefficients is also in $O(n)$, normalization is possible in $O(n)$ for both representations.

C. Parameter Reduction

The need for parameter reduction in the Fourier filters is reminiscent of nonlinear filters on linear domains. Except for very simple problems, representing the exact result of both convolution and multiplication operations usually results in an increase in the number of parameters. In the long run, parameter reduction becomes necessary and usually requires approximations. A popular example of nonlinear estimators are estimators based on Gaussian mixtures. For mixtures, component reduction is a non-trivial and expensive operation, in which the density has to be approximated using a lower number of components while preserving the shape of the density as well as possible [30].

Parameter reduction is also essential for our proposed approach using Fourier series. As described in Sec. VI-B.1, the multiplication operation requires calculating a discrete convolution of the coefficient tensors, resulting in an increase in the size of the tensor. As common densities and especially their square roots are square-integrable, the coefficients tend to zero in every dimension. Therefore, the influence of the coefficients on the shape of the function converges to zero for increasing $\|\mathbf{k}\|$ and it is thus reasonable to use truncation as an easy way for parameter reduction.

As can be seen in Sec. VI-B.2, renormalization is necessary after truncation for the SqFF as the integral




Fig. 5. Prediction step of the IFF.

depends on the sum of the squared absolute values of the coefficients. For the IFF, the integral only depends on $c_{t,0}^{\text{id}}$ and renormalization is thus not required.

D. Prediction Step

To perform the prediction step for an identity system model with additive noise (1), we only need to be able to convolve two densities. Convoluting two functions using their Fourier coefficient tensors is possible very efficiently as this only involves a Hadamard product. The prediction step for nonlinear system models usually requires more than a convolution operation. An approach for arbitrary transition densities that can easily be extended to the multivariate case is laid out in [43].

1) Prediction Step of the IFF: As illustrated in Fig. 5, we can directly use the rule for the convolution of two functions to obtain the Fourier coefficient tensor of the result. If $\mathbf{C}_t^{e,\text{id}}$ and $\mathbf{C}_t^{w,\text{id}}$ are the truncated Fourier coefficient tensors of $f_t^e(\underline{x}_t | \underline{z}_1, \dots, \underline{z}_t)$ and $f_t^w(\cdot)$, we can obtain the coefficient tensor $\mathbf{C}_{t+1}^{p,\text{id}}$ for

$$f_{t+1}^p(\underline{x}_{t+1} | \underline{z}_1, \dots, \underline{z}_t) = (f_t^e * f_t^w)(\underline{x}_{t+1})$$

via the use of the Hadamard product \odot according to

$$\mathbf{C}_{t+1}^{p,\text{id}} = \mathbf{C}_t^{e,\text{id}} \odot \mathbf{C}_t^{w,\text{id}}.$$

2) Prediction Step of the SqFF: Matters are more complicated when aiming to obtain the coefficient tensor $\mathbf{C}_t^{p,\text{sqrt}}$ representing the square root of the convolution of the densities. Since we intend to obtain the Fourier coefficients for

$$\sqrt{f_{t+1}^p(\underline{x}_{t+1} | \underline{z}_1, \dots, \underline{z}_t)} = \sqrt{(f_t^e * f_t^w)(\underline{x}_{t+1})},$$

we cannot simply use the Hadamard product of the coefficient tensors $\mathbf{C}_t^{p,\text{sqrt}}$ and $\mathbf{C}_t^{w,\text{sqrt}}$ as this would yield the coefficient tensor for the Fourier series representing

$$\sqrt{f_t^e(\underline{x}_t | \underline{z}_1, \dots, \underline{z}_t)} * \sqrt{f_t^w(\cdot)},$$

which is, in general, unequal to

$$\sqrt{(f_t^e * f_t^w)(\underline{x}_{t+1})}.$$

Instead, we first calculate the Fourier coefficient tensor for $(f_t^e * f_t^w)(\underline{x}_{t+1})$ and then determine the coefficients of the square root, which we lay out in the following and illustrate in Fig. 6. First, we derive the Fourier coefficients for the non-rooted densities $f_t^e(\underline{x}_t | \underline{z}_1, \dots, \underline{z}_t)$

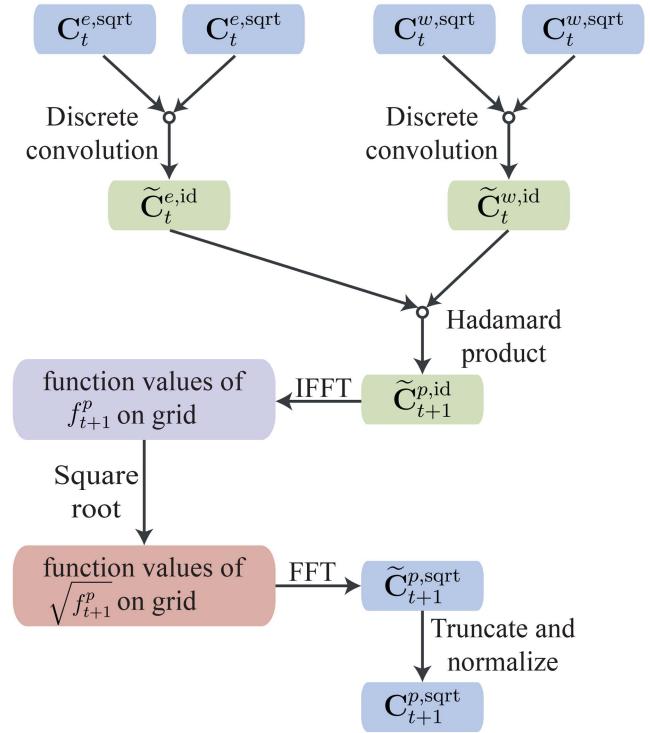


Fig. 6. Prediction step of the SqFF. The colors indicate the respective representation as laid out in Fig. 3.

and $f_t^w(\cdot)$ from the square root representations. These coefficient tensors are denoted by $\tilde{\mathbf{C}}_t^{e,\text{id}}$ and $\tilde{\mathbf{C}}_t^{w,\text{id}}$ to emphasize that they contain more coefficients than the original coefficient tensors $\mathbf{C}_t^{e,\text{sqrt}}$ and $\mathbf{C}_t^{w,\text{sqrt}}$. We can use the multiplication explained in Sec. VI-B.1 to derive the formulae

$$\tilde{\mathbf{C}}_t^{e,\text{id}} = \mathbf{C}_t^{e,\text{sqrt}} * \mathbf{C}_t^{e,\text{sqrt}} \quad \text{and} \quad \tilde{\mathbf{C}}_t^{w,\text{id}} = \mathbf{C}_t^{w,\text{sqrt}} * \mathbf{C}_t^{w,\text{sqrt}}.$$

Based on this, we can use the discrete convolution to obtain the Fourier coefficient tensor $\tilde{\mathbf{C}}_{t+1}^{p,\text{id}}$ for $f_{t+1}^p(\underline{x}_{t+1} | \underline{z}_1, \dots, \underline{z}_t)$ via

$$\tilde{\mathbf{C}}_{t+1}^{p,\text{id}} = \tilde{\mathbf{C}}_t^{e,\text{id}} \odot \tilde{\mathbf{C}}_t^{w,\text{id}}.$$

As the next step, we wish to obtain the Fourier coefficient tensor for the square root representation with a specified number of coefficients. For this, we first calculate n function values of $f_{t+1}^p(\underline{x}_{t+1} | \underline{z}_1, \dots, \underline{z}_t)$ on an equidistant grid, take the square root these function values, and then use the FFT to obtain the Fourier coefficient tensor. To obtain the function values on an equidistant grid, we could now naïvely evaluate the Fourier series with coefficient tensor $\tilde{\mathbf{C}}_{t+1}^{p,\text{id}}$ on all n grid points. However, to achieve a good run time behavior, we have to keep one more issue in mind. Unless there is a reason to do otherwise, it is reasonable that the number of coefficients to calculate for the result is linearly dependent on the number of coefficients used in the coefficient tensor $\tilde{\mathbf{C}}_{t+1}^{p,\text{id}}$. Since each function evaluation is in $O(n)$, n function evaluations would be in $O(n^2)$. Therefore, it is significantly cheaper to use the inverse

FFT with a complexity of $O(n \log n)$ to calculate the required n function values.

With the function values of $f_{t+1}^p(\underline{x}_{t+1} | \underline{z}_1, \dots, \underline{z}_t)$ at our disposal, we can calculate the square root of each value to obtain n function values of $\sqrt{f_{t+1}^p(\underline{x}_{t+1} | \underline{z}_1, \dots, \underline{z}_t)}$. These function values can then be used to approximate n Fourier coefficients via the FFT. Since the most expensive operations are the inverse FFT, the FFT, and the discrete convolution, all of which are in $O(n \log n)$, the total effort is in $O(n \log n)$.

In the step of calculating the square root, approximation errors are caused because taking the square root induces higher frequencies that are not accounted for. Furthermore, as the discrete convolutions used in the calculation of $\tilde{\mathbf{C}}_t^{e,\text{id}}$ and $\tilde{\mathbf{C}}_t^{w,\text{id}}$ result in an increase in the number of coefficients, we need to perform parameter reduction if the coefficients should not be allowed to increase with every prediction step. The parameter reduction can be performed as described in Sec. VI-C. Afterward, a normalization step is required due to the approximations performed.

VII. EXPERIMENTS AND EVALUATION

In this section, we compare our implementations of the two Fourier filters to each other and to other applicable filters. Both filters are available as part of libDirectional [44], a Matlab toolbox for directional statistics with a focus on recursive Bayesian estimation. In Sec. VII-A, we lay out theoretical benefits of the SqFF and describe an experiment to compare the filter steps of the Fourier filters in the univariate case. In the experiment, measurements are simulated regardless of how likely they are. This means we simulate both very likely measurements and unlikely measurements that we would rarely obtain if all underlying assumptions of the filter are correct. For all of these measurements, we evaluate the circular mean direction (3) and the approximation quality of the posterior density provided by the IFF and SqFF.

In the second subsection, we evaluate the error in the form of an angular distance for the two Fourier filters in two bivariate scenarios and one trivariate scenario with additive, wrapped normally distributed noise terms and compare the results with those of other applicable filters in these scenarios. While there is a bivariate wrapped normal filter for toroidal manifolds [23], the only approach to the knowledge of the authors that is applicable to arbitrary multivariate angular estimation problems is the particle filter. Since the number of coefficients used by the Fourier filters and the number of particles used by the particle filter has a major impact on the filter performance, we evaluated several possible configurations. In the scenarios in the second subsection, the likelihoods to be used in each time step can be efficiently calculated using an initial transformation of the likelihood that is shifted according to the measurement obtained.

In the third subsection, we simulate an application of the filters to estimating the angles of a robotic arm based on measurements of the position of the end effector. This application shows that we can use our filter as long as we have a likelihood function and do not require a periodic measurement space. However, this scenario requires more computational effort for the Fourier filters as a Fourier series approximation of the likelihood has to be performed in each time step.

Unlike our evaluation of the circular SqFF in [21], we do not regard the quality of the pdfs and cdfs for the multivariate cases. First, it is difficult to derive a continuous pdf from an SIR particle filter and comparing the cdf becomes increasingly difficult for higher dimensions, especially as a starting point of the integration has to be chosen on periodic domains. Second, we do not have any other filter to numerically approximate the ground truth with at our disposal. In our current evaluation, we also compare the run time of the filters and take the run times into account when assessing the individual filters.

A. Comparing the IFF and the SqFF for Varying Measurements

Allowing the function approximating the density to become negative has many inherent theoretic disadvantages as several useful concepts depend on valid pdfs.¹ Based on an approximation of the pdf with negative function values, no valid cdf can be derived. Furthermore, sampling the density using equally weighted samples may not be possible and sampling schemes such as Metropolis–Hastings sampling [46] do not work. Moreover, some measures of similarity between two densities, such as the Kullback–Leibler divergence [47, Sec. 8.5] and the Hellinger distance [48], cannot be calculated due to the logarithm or square root involved.

In most cases, we observed the approximation quality of the prior density or likelihood function to be already superior for the square root representation when compared with the non-rooted representation with an identical number of coefficients. An example of this is shown in Fig. 7. In the following, we regard the filter step of the Fourier filters and show how both filters perform when comparing the posterior densities and circular mean directions obtained.

Even if no truncation is performed in the filter operations described in Sec. II-B, the approximation errors in the prior density and the likelihood function cause errors in the filter result. The severity of this effect strongly depends on the actual prior density and likelihood function. One important factor influencing the quality of the approximation of the posterior density

¹Negative probabilities [45] can be a viable tool as long as they only appear as intermediate results or if implications are made for other properties that are unobservable simultaneously in the context of physics. In our case, they have no special semantic but stem from approximation errors and have a negative impact on our estimator.




Fig. 7. Von Mises distribution with $\mu = \pi$ and $\kappa = 10$ and corresponding Fourier series approximations using 7 coefficients.

is the overlap of regions of high and low function values of the prior density and the likelihood function.

To understand this effect, we have to take a closer look at the convergence of the Fourier series. A Fourier series converges in the L^2 distance, which is a measure of the squared absolute deviation. Due to the close relationship between the L^2 distance and the Hellinger distance, optimizing the L^2 distance is reasonable when regarding one density individually. However, when regarding the product of two functions, one also has to regard the relative deviation to make statements about the quality of the approximation of the (normalized) multiplication result. Since the Fourier series converges regarding the absolute deviation and not regarding the relative deviation, we expect to see a higher relative deviation in regions of low density than in regions of high density.

This has profound implications for the result of the multiplication of two functions. Let us assume both the prior density and the likelihood are unimodal and are close to zero except for values $\pm\pi/2$ around their modes, such as the von Mises distribution with $\kappa = 10$ that we show in Fig. 7. The prior distribution shall now have a mean of $\mu = \frac{3}{4}\pi$ and the likelihood of $\mu = \frac{3}{2}\pi$. Then, as can be seen in Fig. 8, the relative error of the approximation of the likelihood using 7 coefficients is very high in regions far from the mode, especially when approximating the function directly. When multiplying the prior with the likelihood, the small deviations visible in Fig. 7 are massively amplified in regions of high relative deviation of the likelihood (and also the other way around). This leads to a high deviation from the actual posterior in total. In this example, we can see that the approximation used in the SqFF is advantageous as the relative error is significantly less.

However, there is another effect occurring in the filter step that strongly differs for the two filters. While it is possible to efficiently ensure that a density represented by its Fourier coefficients integrates to one both for the non-rooted and the square root representation (see Sec. VI-B.2), the underlying pdfs that are integrated differ substantially. For the square root representation,

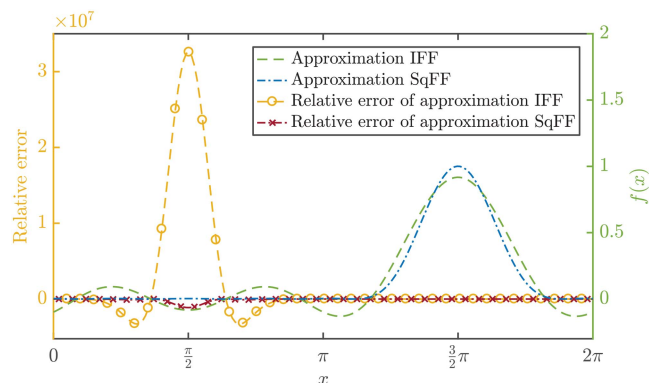


Fig. 8. Approximations of the likelihood and the relative error of the approximations when using Fourier series with 7 coefficients. The likelihood follows a von Mises distribution with $\mu = \frac{3}{2}\pi$ and $\kappa = 10$.

the integral is over nonnegative values, whereas negative function values are possible for the non-rooted version. Always normalizing to one induces a higher total deviation from the abscissa for functions with negative parts as positive and negative parts cancel out. Furthermore, if more than half of the supposed density is negative, normalization to one causes negative parts to become positive and vice versa.

In Fig. 9, we show the filter results for different distances between the modes when the prior density and the likelihood are von Mises distributions with $\kappa = 10$. Both the prior density and the likelihood function were approximated using five Fourier coefficients for both filters and the Fourier series for the posterior densities were also truncated to five coefficients. Fig. 9a and Fig. 9b show the posterior density for the two filters when the modes of the prior density and likelihood function are $\pi/2$ apart. While the normalized result of the Fourier identity filter shows highly negative parts and has the lowest probability density around the peak of the true posterior, the main peak of the SqFF matches that of the true posterior. Furthermore, the result of the SqFF captures the circular mean direction of the posterior density correctly, while the IFF is off by π . Fig. 9c and Fig. 9d show the results when the modes of the two functions are almost π apart. In this case, the true posterior is much flatter than the filter results that bear more of a resemblance to a mixture of the original densities. However, the mean is still correctly captured by the SqFF, while the IFF is off by π .

In Fig. 10, we provide an evaluation depending on the distance α between the two modes of the von Mises distributions. As the first criterion, we evaluate the quality of the posterior density. Since we are unable to use the Kullback–Leibler divergence and the Hellinger distance to compare the result of the IFF to the ground truth, we calculate the total variation [48] between the true densities and the filter results according to

$$d(h_1, h_2) = \int_{[0, 2\pi)} |h_1(x) - h_2(x)| dx,$$

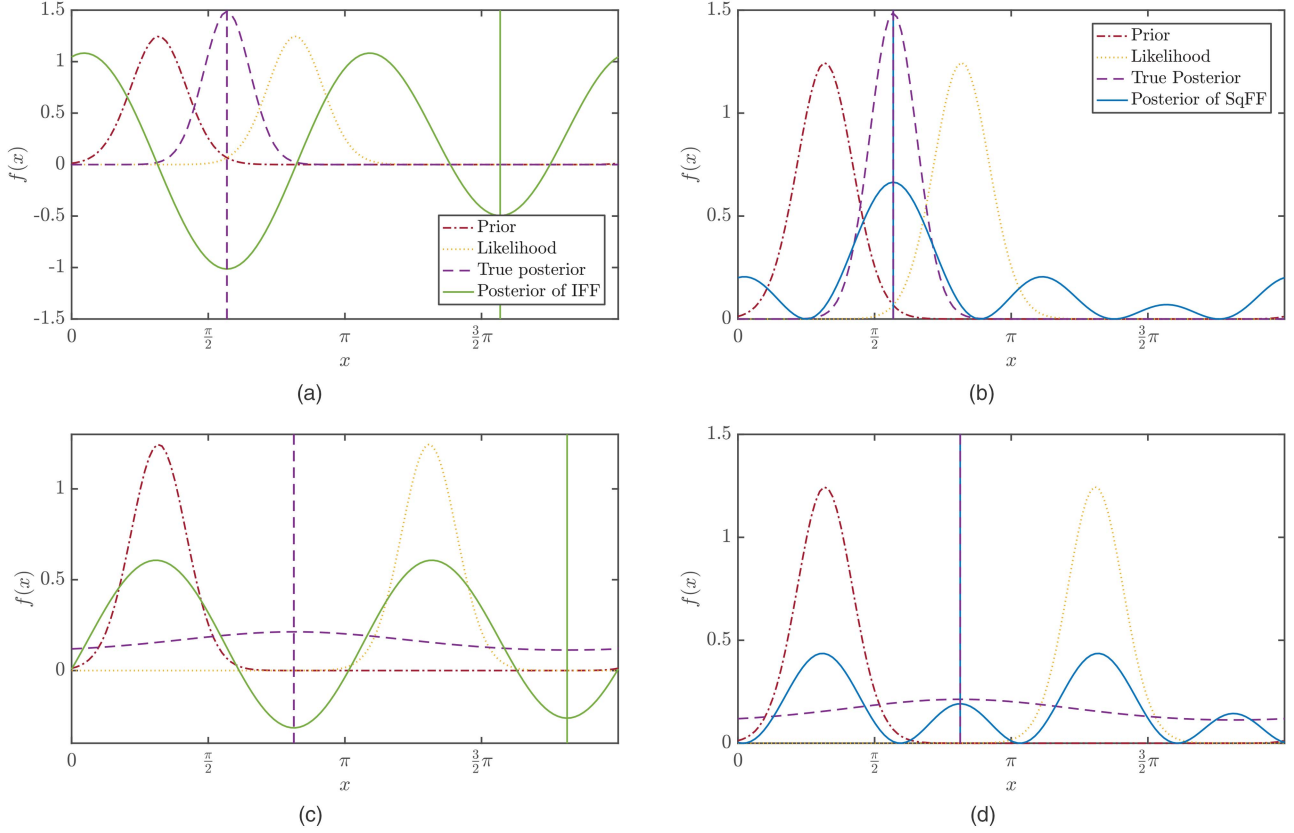


Fig. 9. Examples for posterior densities provided by the two Fourier filters after a single filter step when using five coefficients. The circular mean directions of the posterior densities are shown as a vertical line in the respective color. (a) Result obtained using the IFF when the densities are at a distance of $\pi/2$. (b) Result obtained using the SqFF when the densities are at a distance of $\pi/2$. (c) Result obtained using the IFF when the densities are at a distance of almost π . (d) Result obtained using the SqFF when the densities are at a distance of almost π .

in which $h_1 = f_1 f_2 / \|f_1 f_2\|_{L^1}$ denotes the true density after the filter step and h_2 is the result obtained by the respective variant of the filter. As the second criterion, we regard the quality of the point estimates. For this, we calculate the circular mean direction of the Fourier series approximation from the first trigonometric moment that can be easily obtained as described in Remark 1. Then, we calculate the shorter of the two arc lengths between the true circular mean direction β and the circular mean direction of the filter result γ via [11, Ch. 1.3.2]

$$d_{UV}(\beta, \gamma) = \min(\beta - \gamma, 2\pi - (\beta - \gamma)). \quad (7)$$

A more in-depth discussion of metrics on the circle is given in [49, Sec. 2.2.2].

In the experiment, we limit ourselves to the range $\alpha \in [0, \pi - 0.01]$ as the true posterior for $\alpha = \pi$ is the wrapped uniform distribution that does not have a circular mean direction. The plots clearly show the sensitivity of the filters to larger distances between the modes. The two filters tend to perform worse for $\alpha \in [\pi/2, \pi - 0.01]$ than for $\alpha \in [0, \pi/2]$ both regarding the circular mean direction and the total variation. Comparing the two filters with each other, the total variation between the filter result and the true posterior, as shown in Fig. 10a, can be seen to be higher in most cases when using the IFF. This was to be expected as negative function values are

possible for the result of the IFF. Both the negative parts and the highly positive parts required to cancel out the negative parts result in a higher total deviation. Therefore, deviations that are higher than the usual maximum of the total variation² can be observed, e.g., in the case shown in Fig. 9a. The circular mean direction shown in Fig. 10b is correct for most distances between the means for both filters. This is because the circular mean direction of the result is calculated only from the first trigonometric moment, which is approximated well. Errors in the trigonometric moment usually occur over multiple time steps as an incorrectly shaped density is used for the next filter step, resulting in a wrong circular mean direction. However, as apparent in Fig. 10b, the circular mean direction can also become totally off in a single filter step. Using this criterion, the IFF is also more susceptible to errors than the SqFF. All in all, our results suggest that the SqFF is more robust when unlikely measurements are obtained. Furthermore, we can see that results obtained by the IFF can become totally unlike the true posterior density and can strongly violate the properties of valid densities.

²The total variation usually cannot exceed the value of two for valid densities.




Fig. 10. Comparison of the performance of the IFF and the SqFF depending on the distance between the modes of the functions multiplied. (a) Total variation between the true posterior density and the densities obtained by the two filters. (b) Distance d_{UV} between the true circular mean direction and the circular mean directions provided by the filters.

REMARK 2. Approximating the square root of the density is not the only option. However, using the square root ensures nonnegativity of the reconstructed density while necessitating only few approximations to maintain this representation. An alternative would be to use the logarithm, but the logarithm is highly nonlinear and especially regions with density values close to zero (in which the logarithm approaches $-\infty$) need to be treated with caution. Nonetheless, we plan to consider other transformations than the square root in future work.

B. Evaluation of the Circular Mean Direction

Having compared the two proposed Fourier filters for one filter step in the previous subsection, we now evaluate the Fourier filters against other state-of-the-art approaches to multivariate angular filtering by comparing the estimation performance and run times when multiple filter and prediction steps are performed. In all scenarios, we initialized the filters using an approximation of the actual prior density used for the simulation. We have simulated the entire system behavior and the measurements for 50 time steps and performed alternating prediction and filter steps. A total of 1500 runs

were performed and the true states and the estimation results at each time step were saved for the calculation of our evaluation criterion. As the regarded scenarios feature multivariate estimation problems, we generalize the distance d_{UV} given in (7) to a vector of angles $\underline{\beta}$ and $\underline{\gamma}$ as


$$d_{MV}(\underline{\beta}, \underline{\gamma}) = \sqrt{\sum_{i=1}^d (\min(\beta_i - \gamma_i, 2\pi - (\beta_i - \gamma_i)))^2}.$$

This distance measure is similar to the Euclidean distance but takes periodicity into account. We calculate the distance between the ground truth vector and the estimate provided by the respective filter as a measure of the error of the estimator. To assess the estimation quality over all time steps and runs, we calculate the average of the errors over all time steps and all runs.


Comparing the two Fourier filters regarding the average error is also an important part of our evaluation. While the results regarding the approximation of the density in the previous subsection were much more promising for the SqFF, the error in the circular mean direction was identical when the modes of the prior density and the likelihood function were close, which is a very common case. Furthermore, as the SqFF requires more convolution and FFT operations in the prediction step, the IFF can be used with more coefficients at an identical run time. Therefore, it is not obvious a priori which Fourier filter performs better when evaluating the average error for configurations of equal run time.

All filters were implemented in Matlab without sophisticated optimizations and were compared on a laptop with an Intel Core i7-5500U processor, 12 GB of RAM, and Matlab 2016b running on Windows 10. The run times given are the average run times for each of the 50 time steps. For all filters, multiple different numbers of parameters were used. As k_{\max} was chosen to be equal in every dimension, the numbers of parameters for the Fourier filter were always odd integers taken to the second (bivariate scenarios) or third (trivariate scenario) power. In all evaluations in this subsection, the likelihood at each time step can be obtained by shifting an initially transformed likelihood and this was used to reduce the effort necessary for the Fourier filters.

In the following, we provide evaluation results for three scenarios. An identity model with additive noise (1) is used as the system and the measurement model in all scenarios. The first scenario features bivariate, unimodal transition densities and likelihood functions on the torus. The measurement and system noises are additive noise terms distributed according to multivariate wrapped normal distributions. This scenario is well suited for the use of the bivariate wrapped normal filter. The second scenario evaluated is also a bivariate scenario but features bimodal likelihoods instead of unimodal ones. The bimodality is introduced by using



(a)




(b)


Fig. 11. Average errors and run times for the different filters in the bivariate scenario with unimodal likelihoods. (a) Error of the different filters depending on the number of particles or coefficients used. (b) Run times of the different filters for one time step depending on the number of particles or coefficients used.

a mixture of two multivariate wrapped normal distributions at a distance of one radian in each dimension as additive measurement noise. Both mixture components are centered 0.5 radians away from the actual mean of the density along each axis. Approximating an arbitrary likelihood function using a (potentially unnormalized) wrapped normal distribution is no trivial matter and therefore we have not used the bivariate wrapped normal filter for this scenario and only compared our filters with the particle filter. The third scenario features trivariate densities on the three-dimensional hypertorus. The likelihoods are bimodal again and consist of two multivariate wrapped normal distributions at a distance of one radian in each dimension.

1) Bivariate, Unimodal Scenario: The results shown in Fig. 11 confirm our intuition that the bivariate wrapped normal filter is well suited to this scenario. However, as the results of the Fourier filters show, the performance of the wrapped normal filter can be slightly exceeded by the Fourier filters that approximate the whole posterior density more accurately, facilitating better performance in future time steps. While the Fourier



(a)



(b)

Fig. 12. Average errors and run times for the different filters in the bivariate scenario with bimodal likelihoods. (a) Error of the different filters depending on the number of particles or coefficients used. (b) Run times of the different filters for one time step depending on the number of particles or coefficients used.

filters already appear to reach their optimal accuracy using approximately 100 coefficients, the particle filter does not yet achieve as good of an accuracy for thousands of particles.

When taking the run times into account, the Fourier filters outperform the bivariate wrapped normal filter. The quality of the result of the bivariate wrapped normal filter can be surpassed with numbers of coefficients that still compare favorably in terms of the run time. While fast, the particle filter never comes close to the accuracy of the Fourier filters and the Fourier filters achieve better results using fast configurations with few coefficients. The IFF and the SqFF have almost identical estimation quality, but the IFF is faster. However, the filters achieve a quality that is close to its optimal estimation quality using only approximately 100 coefficients—a configuration at which the two filters differ only little in their run time.

2) Bivariate, Bimodal Scenario: In the results of this scenario, as shown in Fig. 12, the SqFF can be seen to provide better convergence when compared with

the IFF with identical numbers of coefficients. However, one can argue that the IFF is significantly faster and that, when comparing configurations with approximately equal run time, the IFF performs better than the SqFF.

The results of the particle filter do not reach the accuracy of the Fourier filters even for high numbers of coefficients. The particle filter using over 2000 particles is outperformed by the SqFF with less than 500 coefficients in both accuracy and run time. Similarly, this holds for the IFF, whose estimation quality quickly surpasses that of the particle filter and which is significantly faster, resulting in a better performance for configurations with comparable run time.

This scenario shows that the run time performance of the particle filter depends heavily on how fast the likelihood can be evaluated and even a mixture of two components instead of a single wrapped normal distribution can make a clear difference. For the Fourier filters, this does not hold for the identity model with additive noise as the likelihood is only transformed once and shifted afterward.

3) Trivariate, Bimodal Scenario: In the trivariate scenario, the possible configurations of the Fourier filters become more limited as we only use numbers of coefficients that can be written as an odd integer taken to the third power. However, as can be seen in Fig. 13, the superiority of the SqFF over the particle filter is evident. Its performance with corresponding numbers of parameters is better in terms of both the error and the run time, yielding a significant advantage when comparing configurations of comparable run time. Compared on a run time basis, the IFF also outperforms the particle filter since it can handle over 2000 coefficients with a run time that is lower than that of the particle filter with 200 particles.

The bad run time performance of the particle filter is caused by the high computational effort required for evaluating the trivariate wrapped normal density used as the likelihood. As the number of variates increases, the effort involved in evaluating the wrapped normal distribution grows exponentially. To the knowledge of the authors, there is, in general, no other way to calculate the density of a wrapped normal distribution with an arbitrary number of variates other than to sum up some of the addends of the infinite sum. The number of addends required to approximate the density with sufficient accuracy increases exponentially with the number of variates. Thus, the particle filter not only requires more particles for an increasing number of variates, the evaluation of the likelihood also becomes exponentially more expensive. If the filter is used for a large number of time steps, the problem is far less severe for the Fourier filters. For an arbitrary number of time steps, the likelihood and the noise density only have to be approximated once if the likelihood is shifted in the computationally efficient way given in (6) in each time

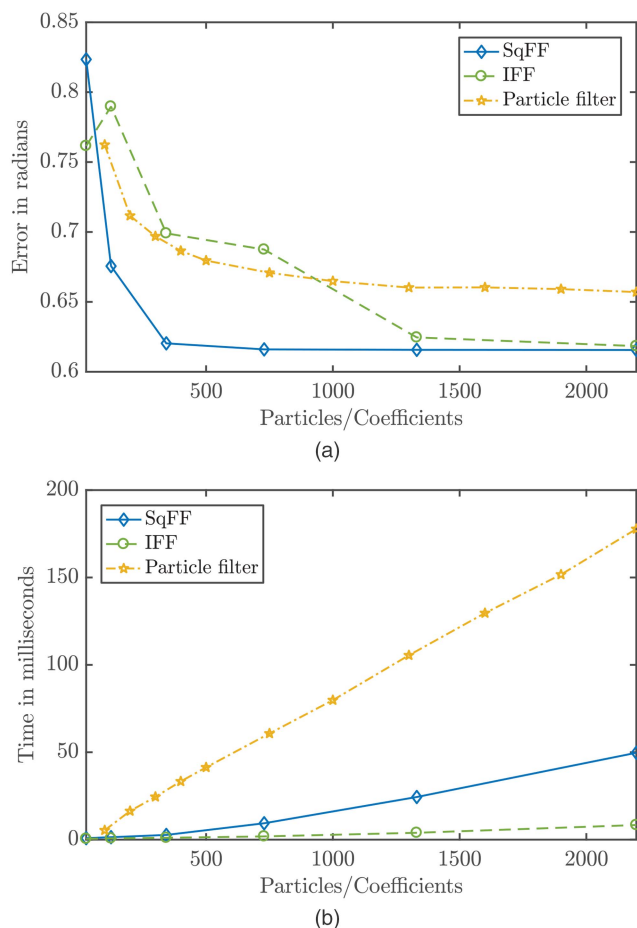


Fig. 13. Average errors and run times for the different filters in the trivariate scenario with bimodal likelihoods. (a) Error of the different filters depending on the number of particles or coefficients used. (b) Run times of the different filters for one time step depending on the number of particles or coefficients used.

step. Therefore, the expensive likelihood only has to be evaluated on a grid once at the beginning.

C. Estimating the Joint Angles of a Robotic Arm

In this subsection, we estimate the joint angles of a robotic arm based on measurements of the position of the end effector that are perturbed by multivariate Gaussian noise. Such a task could, e.g., arise when trying to validate the proper functionality of the robotic arm using external observations. The robotic arm, its joints, and the angles to be estimated are illustrated in Fig. 14. For simplicity, we assume that both joints can move freely and attain any angle. Only point measurements of the point in red on the end effector are obtained. As the end effector can only attain positions on a two-dimensional plane, the measurement is a vector comprising two components. To simplify the measurement equation, we set the zero coordinate of our coordinate system to the center of the first joint indicated in green in Fig. 14.

Given the kinematics of the system, we obtain the measurement equation

$$h(\underline{\alpha}) = \begin{bmatrix} \cos(\alpha_1) \\ \sin(\alpha_1) \end{bmatrix} l_1 + \begin{bmatrix} \cos(\alpha_1 + \alpha_2) \\ \sin(\alpha_1 + \alpha_2) \end{bmatrix} l_2$$




Fig. 14. Illustration of the robotic arm modeled in the simulation.

for $\underline{\alpha} = [\alpha_1 \ \alpha_2]^T$. In our evaluation, we set the joint lengths $l_1 = 2$ m and $l_2 = 1$ m. To simulate sensor noise, the measurements are generated according to

$$\underline{z}_t = h(\underline{\alpha}_t) + \underline{v}_t$$

with a multivariate Gaussian distributed noise term $\underline{v}_t \sim \mathcal{N}(\underline{v}; \underline{\mu}, \mathbf{C})$ with the parameters

$$\underline{\mu} = \begin{bmatrix} 0 \\ 0 \end{bmatrix} \text{ m} \quad \text{and} \quad \mathbf{C} = \begin{bmatrix} 0.2 & 0 \\ 0 & 0.2 \end{bmatrix} \text{ m}^2.$$

While the measurement noise is uncorrelated in the measurement space, the likelihood function

$$f^L(\underline{z}_t | \underline{\alpha}_t) = \mathcal{N}(\underline{z}_t; h(\underline{\alpha}_t), \mathbf{C})$$

is, as shown in Fig. 15 for $\underline{z}_t = [0 \ 2.3]^T$ m, asymmetric and the estimation problem thus cannot be trivially split up into univariate problems. Unlike in the previous subsection, the measurement equation is nonlinear and the likelihoods for differing measurements cannot be obtained by shifting an initial approximation of the likelihood. Therefore, a Fourier series approximation has to be performed in each time step, negatively affecting the run times of the Fourier filters.

As the system model, we use a periodic analogue to a random walk model. This means $\underline{\alpha}_t$ evolves according to

$$\underline{\alpha}_{t+1} = \underline{\alpha}_t + \underline{w}_t \quad \text{mod } 2\pi,$$

with a time-invariant, multivariate wrapped normally distributed additive system noise term \underline{w}_t and a modulo operator that ensures that the angles are always between 0 and 2π . The parameters $\underline{\mu}$ and \mathbf{C} of the system noise \underline{w}_t are identical to those of the measurement noise but the units are rad and rad^2 instead of m and m^2 .

The scenario was simulated for 50 time steps with alternating filter and prediction steps. As in Sec. VII-B, we determine the error d_{MV} in each time step and calculate the average over all 50 time steps and over 1500 runs. The results are depicted in Fig. 16 and are in line with the results obtained in the other bivariate scenarios. Both Fourier filters outperform the particle filter as configurations using only few coefficients provide better results than the particle filter using 2000 particles.






Fig. 15. Likelihood when the measurement $\underline{z}_t = [0 \ 2.3]^T$ m is obtained.



(a)



(b)

Fig. 16. Average errors and run times for the different filters in the scenario featuring a simulated robotic arm. (a) Error of the different filters depending on the number of particles or coefficients used. (b) Run times of the different filters for one time step depending on the number of particles or coefficients used.

However, as expected, the run time of the Fourier filters is slightly worse than in the previous scenarios as the approximation of the likelihood in every time step causes additional overhead. The IFF is faster than the SqFF with equal results, but the difference in the run time is not that pronounced when using the lowest number of coefficients necessary to obtain the highest accuracy achievable for the Fourier filters in this scenario.

VIII. CONCLUSION

In this paper, we have proposed filters based on Fourier series for multivariate angular estimation problems that fill a gap in recursive Bayesian estimation on hypertoroidal manifolds and allow for good estimation results even when likelihood functions or densities occurring are multimodal. In our evaluation of the error in the circular mean direction, the proposed IFF and SqFF achieve better results than the particle filter when comparing configurations of equal run time. The Fourier filters also outperform the bivariate wrapped normal filter in a scenario for which the latter is well suited.

Out of the proposed Fourier filters, the SqFF tends to perform better when comparing the error on a per coefficient basis. However, compared on a run time basis, the IFF is superior when only the estimation quality of the circular mean direction is evaluated. Since the SqFF prevents negative function values in the resulting approximation of the posterior density, it has significant theoretical advantages. As shown in Sec. VII-A, the SqFF is clearly the better choice if an accurate approximation of the posterior pdf is to be provided. Based on these experiments, we also recommend using the SqFF for additional robustness when the likelihood and the prior density share only little common regions of high function values (such as when an unlikely measurement is observed). All in all, we advise the users to employ the SqFF whenever feasible given the run time constraints to benefit from its higher robustness and the expressiveness of the pdf and to use the IFF when run time constraints are tight.

In future research, we intend to inspect the differences between the IFF and the SqFF more closely to be able to give recommendations for which filter to use given the likelihoods, transition densities, run time requirements, and measures of deviation to be minimized. Furthermore, automatically finding the lowest number of coefficients that results in close to optimal results may help users that want to utilize the filter with high estimation quality while saving computational effort. The number of parameters could be further reduced by using sparse representations of the Fourier coefficient tensors. Additional insights could also result from considering other transformations than the square root. Finally, using other basis functions or targeting other periodic manifolds will also be a subject of future research.

ACKNOWLEDGMENT

The IGF project 18798 N of the research association Forschungs-Gesellschaft Verfahrens-Technik e.V. (GVT) was supported via the AiF in a program to promote the Industrial Community Research and Development (IGF) by the Federal Ministry for Economic Affairs and Energy on the basis of a resolution of the German Bundestag. This work was also supported by the German Research Foundation (DFG) under grant HA 3789/13-1.

APPENDIX

In the following, we describe some useful formulae and properties that, despite not being essential to our filter, are important when working with filters and densities in general. We only present the formulae for the non-rooted representation. To use the formula for the Fourier coefficients of the multivariate wrapped normal distribution presented in Appendix A for the SqFF, the approach to derive the Fourier coefficients of the square root from the Fourier coefficients of the original density, as presented in the prediction step in Sec. VI-D.2, can be used. This, however, will not yield a higher accuracy than directly evaluating the square root of the density on a grid and then using the FFT. To use the formulae and properties in Appendix B–D for coefficient tensors representing the square root, we can simply derive the coefficients for the non-rooted representation in the computationally inexpensive way described in Sec. VI-B.1.

A. Fourier Coefficients for the Multivariate Wrapped Normal Distribution

Calculating Fourier coefficients for multivariate wrapped normal distributions in the non-rooted representation is possible using closed-form formulae. For this, we show that we can use the characteristic function of a regular multivariate normal distribution to derive the Fourier coefficients for the multivariate wrapped normal distribution. The characteristic function $\varphi_{\underline{x}}(\underline{k})$ of a random vector \underline{x} on \mathbb{R}^d with density $f_{\underline{x}}$ is defined as

$$\varphi_{\underline{x}}(\underline{k}) = \int_{\mathbb{R}^d} e^{i\underline{k} \cdot \underline{x}} f_{\underline{x}}(\underline{x}) d\underline{x}.$$

Now, let us assume that \underline{x} is normally distributed. Since $e^{i\underline{k} \cdot \underline{x}}$ is 2π -periodic in every dimension,

$$\begin{aligned} \varphi_{\underline{x}}(\underline{k}) &= \int_{\mathbb{R}^d} e^{i\underline{k} \cdot \underline{x}} \mathcal{N}(\underline{x}, \underline{\mu}, \mathbf{C}) d\underline{x} \\ &= \int_{[0, 2\pi)^d} e^{i\underline{k} \cdot \underline{x}} \sum_{\underline{j} \in \mathbb{Z}^d} \mathcal{N}(\underline{x} + 2\pi \underline{j}, \underline{\mu}, \mathbf{C}) d\underline{x} \end{aligned} \quad (8)$$

holds. Aside from a constant factor and a difference in a sign, the last line of the equation (8) is identical to the formula for the Fourier coefficients (5) of a multivariate wrapped normal distribution parametrized by $\underline{\mu}$ and \mathbf{C} . Now, we write the formula for the Fourier coefficients of the multivariate wrapped normal distribution depending on the characteristic function of the multivariate normal distribution [50, Table C] to obtain a closed-form solution. This leads to the formula

$$c_{\underline{k}} = \frac{1}{(2\pi)^d} \varphi_{\underline{x}}(-\underline{k}) = \frac{1}{(2\pi)^d} e^{-i\underline{k} \cdot \underline{\mu} - \underline{k}^T \mathbf{C} \underline{k} / 2}.$$

B. Integrating Fourier Series over Hyperrectangles

For a density given as a Fourier series, it is possible to efficiently calculate the integral over any axis-aligned

hyperrectangle directly from the Fourier coefficients. Let us first regard the one-dimensional case. In this case, we can rewrite the integral over a (truncated) Fourier series from l to r as

$$\int_l^r \sum_{k=-k_{\max}}^{k_{\max}} c_k e^{ikx} dx = \sum_{k=-k_{\max}}^{k_{\max}} c_k \underbrace{\int_l^r e^{ikx} dx}_{h_k}$$

and regard each addend h_k separately to obtain

$$h_0 = c_k [x]_l^r = c_k (r - l)$$

and

$$\begin{aligned} \forall k \neq 0, |k| \leq k_{\max} : h_k &= c_k \frac{1}{ik} [e^{ikx}]_l^r \\ &= -c_k \frac{i}{k} (e^{ikr} - e^{ikl}) \end{aligned}$$

via common integration rules. We now use $c_k = \bar{c}_{-k}$ and obtain

$$\begin{aligned} h_k + h_{-k} &= -c_k \frac{i}{k} (e^{ikr} - e^{ikl}) - \bar{c}_k \frac{i}{-k} (e^{-ikr} - e^{-ikl}) \\ &= \frac{i}{k} (-c_k (e^{ikr} - e^{ikl}) + \overline{c_k (e^{ikr} - e^{ikl})}) \\ &= \frac{2}{k} \mathcal{I}(c_k (e^{ikr} - e^{ikl})), \end{aligned}$$

with $\mathcal{I}(\cdot)$ denoting the imaginary part of the term. As expected, the sum of the pairs are real values. Based on this, we can calculate the integral via

$$\int_l^r \sum_{k=-k_{\max}}^{k_{\max}} c_k e^{ikx} dx = c_k (r - l) + \sum_{k=1}^{k_{\max}} \frac{2}{k} \mathcal{I}(c_k (e^{ikr} - e^{ikl})).$$

The integration formula provided can easily be extended to higher dimensions. If we use $\mathcal{J} \subset \mathbb{Z}^d$ to denote the index set comprising all indices of the nonzero Fourier coefficients,

$$\begin{aligned} \int_l^r \sum_{\underline{k} \in \mathcal{J}} c_{\underline{k}} e^{i\mathbf{k} \cdot \mathbf{x}} &= \sum_{\underline{k} \in \mathcal{J}} c_{\underline{k}} \int_l^r e^{i\mathbf{k}_1 x_1} \dots e^{i\mathbf{k}_d x_d} d\mathbf{x} \\ &= \sum_{\underline{k} \in \mathcal{J}} c_{\underline{k}} \left(\int_{l_1}^{r_1} e^{i\mathbf{k}_1 x_1} dx_1 \dots \int_{l_d}^{r_d} e^{i\mathbf{k}_d x_d} dx_d \right) \end{aligned}$$

holds and we can thus split the integration up and use the integration rule for the one-dimensional case to obtain the result in $O(nd)$.

C. Marginalizing Out Specific Dimensions

Calculating the Fourier coefficients for a marginalized density is computationally inexpensive. Without loss of generality, we marginalize the first dimension out. We denote the index set of all nonzero Fourier coefficients by \mathcal{J} and rewrite the integral

$$\begin{aligned} \int_0^{2\pi} \sum_{\underline{k} \in \mathcal{J}} c_{\underline{k}} e^{i\mathbf{k}_1 x_1} \dots e^{i\mathbf{k}_d x_d} dx_1 \\ = \sum_{\underline{k} \in \mathcal{J}} c_{\underline{k}} e^{i\mathbf{k}_2 x_2} \dots e^{i\mathbf{k}_d x_d} \int_0^{2\pi} e^{i\mathbf{k}_1 x_1} dx_1 \end{aligned}$$

and then use that the integral is always zero for $k_1 \neq 0$

$$\begin{aligned} &= \sum_{k_2=-k_{\max}}^{k_{\max}} \dots \sum_{k_d=-k_{\max}}^{k_{\max}} e^{i\mathbf{k}_2 x_2} \dots e^{i\mathbf{k}_d x_d} c_{0,k_2,\dots,k_d} \int_0^{2\pi} 1 dx_1 \\ &= \sum_{k_2=-k_{\max}}^{k_{\max}} \dots \sum_{k_d=-k_{\max}}^{k_{\max}} 2\pi e^{i\mathbf{k}_2 x_2} \dots e^{i\mathbf{k}_d x_d} c_{0,k_2,\dots,k_d}. \end{aligned}$$

Thus, we can calculate the new coefficient tensor by discarding all entries for which the index of the respective dimension is unequal to zero and then multiplying all remaining entries by 2π .

D. Calculating the Covariance Matrix

The covariance matrix mentioned in Sec. III-F is useful for calculating angular correlations and can be calculated efficiently for densities in a Fourier series representation. For higher dimensions, we first introduce a notational trick for Fourier coefficients representing a density after certain dimensions have been marginalized out. We write $c_{k_r, k_t} = (2\pi)^{d-2} c_{0,\dots,0,k_r,0,\dots,0,k_t,0,\dots,0}$ for the Fourier coefficients of the density with all dimensions unequal to $r \in \{1, 2, \dots, d\}$ and $t \in \{1, 2, \dots, d\}$ marginalized out and specify the exact coefficient via $c_{k_r=a, k_t=b}$ for $a \in \mathbb{Z}, b \in \mathbb{Z}$. We use a similar notation to index Fourier coefficients of densities with all dimensions except one marginalized out and write them as $c_{k_r} = (2\pi)^{d-1} c_{0,\dots,0,k_r,0,\dots,0}$ and denote specific entries via $c_{k_r=a}$. For additional brevity and clarity, we assign the real and the imaginary part of the first trigonometric moment the names $\underline{p} = \mathcal{R}(\underline{m}_1)$ and $\underline{q} = \mathcal{I}(\underline{m}_1)$.

The general term for entries of the covariance matrix Σ is

$$\mathbb{E}((\text{trig}_1(x_r) - u_r)(\text{trig}_2(x_t) - v_t)),$$

in which trig_1 and trig_2 can be sine or cosine functions. If trig_1 is cos, then $u_r = p_r$, if trig_1 is sin, then $u_r = q_r$ and the same applies to trig_2 and v_t . We present the formulae for the individual entries of the covariance matrix Σ without a derivation as the derivations are long and not essential to this paper.

As higher coefficients are lost in the integral involved in calculating the expectation value, only few coefficients are necessary for each entry. Furthermore, we can use the redundancy in the complex Fourier coefficients to arrive at even easier formulae depending on the real and complex parts of the Fourier coefficients. We further use

$\forall i \in \{1, \dots, d\} : p_i = 2\pi \mathcal{R}(c_{k_i=1})$ and $q_i = -2\pi \mathcal{I}(c_{k_i=1})$ to write the formulae in an even more compact manner.

For the formulae for the individual entries of the covariance matrix, keep in mind that Σ is a $2d \times 2d$ matrix. In the covariance matrix, an odd column index (i.e., $2r - 1$) indicates that trig_1 is the sine function, while an even column index (i.e., $2r$) indicates that trig_1 is the cosine function. The relationship is the same between trig_2 and the row index.

For all $r \neq t$, we obtain the formulae

$$\begin{aligned}\sigma_{2r-1,2t-1} &= \mathbb{E}((\cos(x_r) - p_r)(\cos(x_t) - p_t)) \\ &= 2\pi^2 \mathcal{R}(c_{k_r=1, k_t=1}) + 2\pi^2 \mathcal{R}(c_{k_r=1, k_t=-1}) \\ &\quad - p_r p_t\end{aligned}$$

and

$$\begin{aligned}\sigma_{2r-1,2t} &= \mathbb{E}((\cos(x_r) - p_r)(\sin(x_t) - q_t)) \\ &= -2\pi^2 \mathcal{I}(c_{k_r=1, k_t=1}) + 2\pi^2 \mathcal{I}(c_{k_r=1, k_t=-1}) - p_r q_t\end{aligned}\tag{9}$$

for the respective entries of the covariance matrix. Due to the symmetry of the covariance matrix Σ , we can use $\sigma_{2r,2t-1} = \sigma_{2t-1,2r}$ to calculate these entries using (9). For the terms with two sines, we obtain

$$\begin{aligned}\sigma_{2r,2t} &= \mathbb{E}((\sin(x_r) - q_r)(\sin(x_t) - q_t)) \\ &= -2\pi^2 \mathcal{R}(c_{k_r=1, k_t=1}) + 2\pi^2 \mathcal{R}(c_{k_r=1, k_t=-1}) - q_r q_t\end{aligned}$$

If $r = t$, other coefficients are involved in the calculation. In this case, we obtain

$$\begin{aligned}\sigma_{2r-1,2r-1} &= \mathbb{E}((\cos(x_r) - p_r)^2) \\ &= \pi \mathcal{R}(c_{k_r=2}) - p_r^2 + \frac{1}{2}, \\ \sigma_{2r-1,2r} &= \mathbb{E}((\cos(x_r) - p_r)(\sin(x_r) - q_r)) \\ &= \pi \mathcal{I}(c_{k_r=2}) - p_r q_r.\end{aligned}\tag{10}$$

Again, we can use the symmetry of Σ to obtain $\sigma_{2r,2r-1} = \sigma_{2r-1,2r}$, allowing us to use (10). Lastly, we get

$$\begin{aligned}\sigma_{2r,2r} &= \mathbb{E}((\sin(x_r) - q_r)^2) \\ &= -\pi \mathcal{R}(c_{k_r=2}) - q_r^2 + \frac{1}{2}.\end{aligned}$$

REFERENCES

- [1] K. V. Mardia
“Directional Statistics in Geosciences,”
Communications in Statistics—Theory and Methods, vol. 10, no. 15, pp. 1523–1543, 1981.
- [2] E. Batschelet
Circular Statistics in Biology, ser. Mathematics in Biology. London: Academic Press, 1981.
- [3] J. L. Crassidis, F. L. Markley, and Y. Cheng
“Survey of Nonlinear Attitude Estimation Methods,”
Journal of Guidance, Control, and Dynamics, vol. 30, no. 1, pp. 12–28, 2007.
- [4] J. Glover and L. P. Kaelbling
“Tracking the Spin on a Ping Pong Ball with the Quaternion Bingham Filter,”
in *Proceedings of the 2014 IEEE Conference on Robotics and Automation (ICRA 2014)*, Hong Kong, China, 2014.
- [5] J. Lo and A. Willsky
“Estimation for Rotational Processes with One Degree of Freedom—Part I: Introduction and Continuous-Time Processes,”
IEEE Transactions on Automatic Control, vol. 20, no. 1, pp. 10–21, Feb. 1975.
- [6] A. S. Willsky
“Fourier Series and Estimation on the Circle with Applications to Synchronous Communication—Part I: Analysis,”
IEEE Transactions on Information Theory, 1974.
- [7] G. Stienne, S. Reboul, M. Azmani, J. Choquel, and M. Benjelloun
“A Multi-Sensor Circular Particle Filter Applied to the Fusion of the GPS-L2C Channels,”
in *Proceedings of the 14th International Conference on Information Fusion (Fusion 2011)*, Chicago, Illinois, USA, 2011.
- [8] G. Kurz and U. D. Hanebeck
“Heart Phase Estimation Using Directional Statistics for Robotic Beating Heart Surgery,”
in *Proceedings of the 18th International Conference on Information Fusion (Fusion 2015)*, Washington D. C., USA, Jul. 2015.
- [9] S. J. Julier and J. K. Uhlmann
“Unscented Filtering and Nonlinear Estimation,”
Proceedings of the IEEE, vol. 92, no. 3, pp. 401–422, Mar. 2004.
- [10] K. V. Mardia and P. E. Jupp
Directional Statistics, 1st ed. Baffins Lane, Chichester, West Sussex, PO19 1UD England: Wiley, 1999.
- [11] S. R. Jammalamadaka and A. Sengupta
Topics in Circular Statistics. World Scientific, 2001.
- [12] W. Boomsma, K. V. Mardia, C. C. Taylor, J. Ferkinghoff-Borg, A. Krogh, and T. Hamelryck
“A Generative, Probabilistic Model of Local Protein Structure,”
Proceedings of the National Academy of Sciences, vol. 105, no. 26, pp. 8932–8937, 2008.
- [13] Y.-H. Chen, D. Wei, G. Newstadt, M. DeGraef, J. Simmons, and A. Hero
“Parameter Estimation in Spherical Symmetry Groups,”
IEEE Signal Processing Letters, vol. 22, no. 8, pp. 1152–1155, 2015.
- [14] F. Chatelain and N. Le Bihan
“Von Mises–Fisher Approximation of Multiple Scattering Process on the Hypersphere,”
in *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, 2013, pp. 6461–6465.
- [15] M. Akcakaya and A. Nehorai
“MIMO Radar Detection and Adaptive Design Under a Phase Synchronization Mismatch,”
IEEE Transactions on Signal Processing, vol. 58, no. 10, pp. 4994–5005, Oct 2010.
- [16] A. S. Willsky
“Fourier Series and Estimation on the Circle with Applications to Synchronous Communication—Part II: Implementation,”
IEEE Transactions on Information Theory, 1974.
- [17] M. Azmani, S. Reboul, J.-B. Choquel, and M. Benjelloun
“A Recursive Fusion Filter for Angular Data,”
in *2009 IEEE International Conference on Robotics and Biomimetics (ROBIO)*, Dec. 2009.
- [18] G. Kurz, I. Gilitschenski, and U. D. Hanebeck
“Recursive Nonlinear Filtering for Angular Data Based on Circular Distributions,”
in *Proceedings of the 2013 American Control Conference (ACC 2013)*, Washington D. C., USA, Jun. 2013.

- [19] ———
“Nonlinear Measurement Update for Estimation of Angular Systems Based on Circular Distributions,”
in *Proceedings of the 2014 American Control Conference (ACC 2014)*, Portland, Oregon, USA, Jun. 2014.
- [20] ———
“Recursive Bayesian Filtering in Circular State Spaces,”
IEEE Aerospace and Electronic Systems Magazine, vol. 31, no. 3, pp. 70–87, Mar. 2016.
- [21] F. Pfaff, G. Kurz, and U. D. Hanebeck
“Multimodal Circular Filtering Using Fourier Series,”
in *Proceedings of the 18th International Conference on Information Fusion (Fusion 2015)*, Washington D. C., USA, Jul. 2015.
- [22] G. Kurz, F. Pfaff, and U. D. Hanebeck
“Discrete Recursive Bayesian Filtering on Intervals and the Unit Circle,”
in *Proceedings of the 2016 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI 2016)*, Baden-Baden, Germany, Sep. 2016.
- [23] G. Kurz, I. Gilitschenski, M. Dolgov, and U. D. Hanebeck
“Bivariate Angular Estimation Under Consideration of Dependencies Using Directional Statistics,”
in *Proceedings of the 53rd IEEE Conference on Decision and Control (CDC 2014)*, Los Angeles, California, USA, Dec. 2014.
- [24] M. Arulampalam, S. Maskell, N. Gordon, and T. Clapp
“A Tutorial on Particle Filters for Online Nonlinear/Non-Gaussian Bayesian Tracking,”
IEEE Transactions on Signal Processing, vol. 50, no. 2, pp. 174–188, 2002.
- [25] S. Kay
Fundamentals of Statistical Signal Processing: Estimation Theory.
Prentice Hall, 1993.
- [26] M. F. Huber, F. Beutler, and U. D. Hanebeck
“Semi-Analytic Gaussian Assumed Density Filter,”
in *Proceedings of the 2011 American Control Conference (ACC 2011)*, San Francisco, California, USA, Jun. 2011.
- [27] U. D. Hanebeck
“PGF 42: Progressive Gaussian Filtering with a Twist,”
in *Proceedings of the 16th International Conference on Information Fusion (Fusion 2013)*, Istanbul, Turkey, Jul. 2013.
- [28] J. Steinbring and U. D. Hanebeck
“LRKF Revisited: The Smart Sampling Kalman Filter (S2KF),”
Journal of Advances in Information Fusion, vol. 9, no. 2, pp. 106–123, Dec. 2014.
- [29] J. Steinbring, M. Pander, and U. D. Hanebeck
“The Smart Sampling Kalman Filter with Symmetric Samples,”
Journal of Advances in Information Fusion, vol. 11, no. 1, pp. 71–90, Jun. 2016.
- [30] M. Huber, P. Krauthausen, and U. D. Hanebeck
“Superficial Gaussian Mixture Reduction,”
in *Proceedings of the IEEE ISIF Workshop on Sensor Data Fusion: Trends, Solutions, Applications (SDF 2011)*, Berlin, Germany, Oct. 2011.
- [31] S. R. Jammalamadaka and Y. Sarma
“A Correlation Coefficient for Angular Variables,”
Statistical Theory and Data Analysis II, pp. 349–364, 1988.
- [32] R. A. Johnson and T. Wehrly
“Measures and Models for Angular Correlation and Angular-Linear Correlation,”
Journal of the Royal Statistical Society. Series B (Methodological), vol. 39, no. 2, pp. 222–229, 1977.
- [33] P. E. Jupp and K. V. Mardia
“A General Correlation Coefficient for Directional Data and Related Regression Problems,”
Biometrika, vol. 67, no. 1, pp. 163–173, 1980.
- [34] A. Zygmund
Trigonometric Series,
3rd ed. Cambridge University Press, 2003, vol. 1 and 2.
- [35] Y. Katznelson
An Introduction to Harmonic Analysis,
3rd ed. Cambridge University Press, 2004.
- [36] J. J. Fernández-Durán
“Circular Distributions Based on Nonnegative Trigonometric Sums,”
Biometrics, vol. 60, no. 2, Jun. 2004.
- [37] D. Brunn, F. Sawo, and U. D. Hanebeck
“Efficient Nonlinear Bayesian Estimation Based on Fourier Densities,”
in *Proceedings of the 2006 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI 2006)*, Heidelberg, Germany, Sep. 2006, pp. 312–322.
- [38] ———
“Nonlinear Multidimensional Bayesian Estimation with Fourier Densities,”
in *Proceedings of the 2006 IEEE Conference on Decision and Control (CDC 2006)*, San Diego, California, USA, Dec. 2006, pp. 1303–1308.
- [39] J. W. Cooley and J. W. Tukey
“An Algorithm for the Machine Calculation of Complex Fourier Series,”
Mathematics of Computation, vol. 19, no. 90, pp. 297–301, 1965.
- [40] D. E. Dudgeon and R. M. Mersereau
Multidimensional Digital Signal Processing.
Prentice Hall, 1984.
- [41] P. Duhamel and M. Vetterli
“Fast Fourier Transforms: A Tutorial Review and a State of the Art,”
Signal Processing, vol. 19, no. 4, pp. 259–299, 1990.
- [42] S. P. Jan-P. Calliess, Michael Mai
“On the Computational Benefit of Tensor Separation for High-Dimensional Discrete Convolutions,”
Multidimensional Systems and Signal Processing, 2012.
- [43] F. Pfaff, G. Kurz, and U. D. Hanebeck
“Nonlinear Prediction for Circular Filtering Using Fourier Series,”
in *Proceedings of the 19th International Conference on Information Fusion (Fusion 2016)*, Heidelberg, Germany, Jul. 2016.
- [44] G. Kurz, I. Gilitschenski, F. Pfaff, and L. Drude
“libDirectional,”
2015. [Online]. Available: <https://github.com/libDirectional>.
- [45] R. P. Feynman
“Negative Probability,”
Quantum Implications: Essays in Honour of David Bohm, pp. 235–248, 1987.
- [46] W. K. Hastings
“Monte Carlo Sampling Methods Using Markov Chains and Their Applications,”
Biometrika, vol. 57, no. 1, pp. 97–109, 1970.
- [47] T. M. Cover and J. A. Thomas
Elements of Information Theory.
John Wiley & Sons, 2012.
- [48] F. Liese and I. Vajda
“On Divergences and Informations in Statistics and Information Theory,”
IEEE Transactions on Information Theory, 2006.
- [49] G. Kurz
“Directional Estimation for Robotic Beating Heart Surgery,”
Ph.D. dissertation, Karlsruhe Institute of Technology, Intelligent Sensor-Actuator-Systems Laboratory, Karlsruhe, Germany, 2015.
- [50] F. Oberhettinger
Fourier Transforms of Distributions and Their Inverses.
Academic Press, 1973.



Florian Pfaff started as a Ph.D. student at the Karlsruhe Institute of Technology (KIT) in 2014 after completing his diploma for which he received an award due to his outstanding performance. His research interests include directional estimation, distributed estimation, and multitarget tracking and its application to bulk material sorting.



Gerhard Kurz received his diploma in computer science from the Karlsruhe Institute of Technology (KIT), Germany, in 2012. Afterwards, he obtained his Ph.D. in 2015 at the Intelligent Sensor-Actuator-Systems Laboratory, Karlsruhe Institute of Technology (KIT), Germany. His research interests include directional filtering, nonlinear estimation, and medical data fusion. He has authored multiple award-winning publications on these topics.

Uwe D. Hanebeck is a chaired professor of Computer Science at the Karlsruhe Institute of Technology (KIT) in Germany and director of the Intelligent Sensor-Actuator-Systems Laboratory (ISAS). Since 2005, he is the chairman of the Research Training Group RTG 1194 “Self-Organizing Sensor-Actuator-Networks” financed by the German Research Foundation.

Prof. Hanebeck obtained his Ph.D. degree in 1997 and his habilitation degree in 2003, both in Electrical Engineering from the Technical University in Munich, Germany. His research interests are in the areas of information fusion, nonlinear state estimation, stochastic modeling, system identification, and control with a strong emphasis on theory-driven approaches based on stochastic system theory and uncertainty models. Research results are applied to various application topics like localization, human-robot-interaction, assistive systems, sensor-actuator-networks, medical engineering, distributed measuring system, and extended range telepresence. Research is pursued in many academic projects and in a variety of cooperations with industrial partners.

Uwe D. Hanebeck was the General Chair of the “2006 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI 2006),” Program Co-Chair of the “11th International Conference on Information Fusion (Fusion 2008),” Program Co-Chair of the “2008 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI 2008),” Regional Program Co-Chair for Europe for the “2010 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2010),” and will be General Chair of the “19th International Conference on Information Fusion (Fusion 2016).” He is a Member of the Board of Directors of the International Society of Information Fusion (ISIF), Editor-in-chief of its Journal of Advances in Information Fusion (JAIF), and associate editor for the letter category of the IEEE Transactions on Aerospace and Electronic Systems (TAES). He is author and coauthor of more than 300 publications in various high-ranking journals and conferences.



Particle Filtering with Observations in a Manifold: A Proof of Convergence and Two Worked Examples

SALEM SAID
JONATHAN H. MANTON

Particle filtering is currently a popular numerical method for solving stochastic filtering problems. This paper outlines its application to continuous time filtering problems with observations in a manifold. Such problems include a variety of important engineering situations and have received independent interest in the mathematical literature. The paper begins by stating a general stochastic filtering problem where the observation process, conditionally on the unknown signal, is an elliptic diffusion in a differentiable manifold. Using a geometric structure (a Riemannian metric and a connection) which is adapted to the observation model, it expresses the solution of this problem in the form of a Kallianpur-Striebel formula. The paper proposes a new particle filtering algorithm which implements this formula using sequential Monte Carlo strategy. This algorithm is based on an original use of the concept of connector map, which is here applied for the first time in the context of filtering problems. The paper proves the convergence of this algorithm, under the assumption that the underlying manifold is compact, and illustrates it with two worked examples. In the first example, the observations lie in the special orthogonal group $SO(3)$. The second example is concerned with the case of observations in the unit sphere S^2 .

Manuscript received February 28, 2016; released for publication March 11, 2016.

Refereeing of this contribution was handled by Gerhard Kurz.

Authors' addresses: S. Said, CNRS, Laboratoire IMS (UMR 5218), Université de Bordeaux; J. Manton, The University of Melbourne, Department of Electrical and Electronic Engineering.

1557-6418/16/\$17.00 © 2016 JAIF

1. INTRODUCTION

The language of differential geometry is increasingly being used by engineers. This is due to the realisation that, in many concrete situations, the most natural mathematical model involves a nontrivial differentiable manifold. Key differential geometric concepts (tangent vectors, geodesics, etc.) have allowed for tasks such as optimisation and stochastic modeling to be approached in a unified and intuitive way [1].

This trend has led to several variants of “particle filtering on manifolds” proposed in the literature, (for example, [2–9]). These involve discrete time filtering problems where either the unknown signal or the observation process lie in a differentiable manifold. The current paper outlines the application of particle filtering to continuous time filtering problems where the observation process takes its values in a differentiable manifold, and the unknown signal is a hidden Markov process. These are essentially different from problems where the unknown signal is a diffusion in a differentiable manifold and the observation process follows a classical additive white noise model; see Section 2.

The choice of continuous time over discrete time has technical and modeling advantages. It allows for the machinery of stochastic calculus and differential geometry to be applied. Also, the immense majority of physical models do not immediately discretise the time variable. Concretely, though, the difference between continuous and discrete time is a convenience. The question is to solve-then-discretise or discretise-then-solve. The final product in this paper, the algorithm of Section 4, is in discrete time and can be compared to any other algorithm developed directly in discrete time.

Predominantly, the differentiable manifolds appearing in engineering are classical matrix Lie groups and their symmetric spaces. The general stochastic filtering problem stated in Section 2 starts from an abstract differentiable manifold. All subsequent constructions are stated at this level of generality. It is hoped this will have the advantage of providing a deeper overall understanding. When (as a special instance of the general problem) the manifold is specified to be a matrix manifold, this allows for the complexity of differential geometric constructions to be reduced. This will be discussed again, using two concrete examples, in Section 5.

A subtheme of this paper is to realise a transfer of knowledge, from the mathematical field of stochastic differential geometry to the engineering community. Stochastic differential geometry was pioneered by Schwartz and Meyer. Among the fundamental texts, it is impossible to ignore the elegant and comprehensive account by Emery [10]. A more recent highly readable textbooks is the one by Hsu [11]. A recent account, written specifically for engineers, is [12].

Stochastic differential geometry begins with a manifold equipped with a connection. In differential geometry, connections are introduced to distinguish in

an invariant way those differentiable curves which are geodesics; in other words, zero acceleration lines. Stochastic differential geometry uses a connection to distinguish those pathwise continuous processes which are martingales; in other words zero drift processes.

For the purpose of filtering, it is very useful to think in terms of *antidevelopment*. The antidevelopment of a differentiable curve in a manifold is a differentiable curve in a Euclidean space, which could be identified with the tangent space at the base point of the curve. This can be visualised when the manifold is a two dimensional surface embedded in physical space, for instance a sphere. The antidevelopment of a differentiable curve drawn on the sphere is the trace that it leaves on a tangent plane while rolling without slipping (one could imagine the curve is drawn in ink so it leaves a trace on the plane). This visualisation is helpful for intuition, but it is important to keep in mind the notion of antidevelopment depends on a choice of connection. The picture of rolling without slipping corresponds to the connection inherited from the ambient space. Roughly speaking, the relation between a connection and the corresponding antidevelopment is that antidevelopment of a geodesic in a manifold is a straight line in Euclidean space. Similarly, antidevelopment of a martingale in the manifold is a martingale in Euclidean space. Required background from stochastic differential geometry is presented in Section 3.

In terms of stochastic filtering with observations in a manifold, the main problem considered in this paper, antidevelopment plays an essential role. While the observation process is a diffusion in a given differentiable manifold, its antidevelopment process takes its values in a Euclidean space. Regardless of the choice of connection, there is no loss of information in replacing the observation process by its antidevelopment. Moreover, see Proposition 1 of Section 3, an adequate choice of connection (roughly, one which is adapted to the observation model) leads to an antidevelopment process which depends on the unknown signal through a classical additive white noise model. Thus, antidevelopment can be thought of as a preprocessing, reducing the initial problem to a classical filtering problem defined by an additive white noise model. This idea of reduction was used in the engineering literature by Lo [13], in the case of observations in a matrix Lie group. In a recent paper [14], the authors were able to extend it to the general case of observations in a differentiable manifold.

The role of stochastic differential geometry in problems of stochastic filtering with observations in a manifold was exploited more extensively in the mathematical literature. Several authors have used it in deriving Zakai or filtering equations for these problems [15–17]. To the author’s knowledge, on the other hand, few papers have been devoted to their numerical solution. Note, however, the paper by Pontier [18], which will be quite important in the following. This proves the convergence of a

discrete time filter based on uniformly sampled observations to the solution of the continuous time problem.

The filtering problem is to compute the conditional distribution of the unknown signal given the observations. In section 4, the solution of this problem is expressed in the form of a Kallianpur-Striebel formula (see Proposition 4). This has a structure quite similar to that of the classical Kallianpur-Striebel formula. It is a Bayes formula where the prior distribution of the unknown signal is given by its Markov nature, i.e., initial distribution and transition kernel, and the likelihood function is an exponential functional of the observation process.

The proposed particle filtering algorithm implements the Kallianpur-Striebel formula using sequential Monte Carlo. There are at least two motivations for applying a sequential Monte Carlo approach. First, the fact that it is suitable for real time situations. Second, the computational stability which it provides in dealing with noisy observations over a longer time.

The algorithm follows sequential Monte Carlo strategy of sequential importance sampling with resampling. At a higher level, its main steps are exactly the same as for usual particle filtering. A fixed number of particles is used throughout. The interval of observation is subdivided into subintervals of equal length. Over each subinterval, the particles are propagated without interaction according to (an approximation of) the unknown signal’s transition kernel. They are subsequently given new weights corresponding to their likelihood and resampled to eliminate particles with lower weight. For a rigorous general discussion, see Del Moral’s monograph [19].

The specific role of stochastic differential geometry only appears in the computation of particle weights. This requires the use of *connector maps*. Under the assumption that the underlying manifold is compact, Proposition 3 of Section 3 shows that connector maps yield “increments” which are approximately normally distributed tangent vectors, conditionally on the unknown signal. The computation of weights then takes place as in the presence of additive normally distributed noise. Simply put, the proposed algorithm is a classical particle filter which employs the geometric concept of connector maps to locally linearise the observation process.

Proposition 5 of Section 4 states the convergence of this discrete algorithm to the solution of the continuous time problem. The convergence takes place as the number of particles tends to infinity and the length of each subdivision interval tends to zero. It depends on the quality of the approximation being used for the unknown signal transition kernel and also on the one described in Proposition 3. As several approximations are involved, the compromise between complexity and performance should be based on the most difficult one to realise. This leads to a kind of bottleneck effect. See discussion at the end of Section 4.

Section 5 presents two examples illustrating the implementation and performance of the proposed particle filtering algorithm. In the first example, see 5.1, the observation process is a left invariant diffusion with values in the special orthogonal group $SO(3)$; conditionally on the unknown signal. The fact that the observation model is compatible with the Lie group structure of $SO(3)$ leads to a certain simplification of the notions of antidevelopment and connector maps. Roughly speaking, these just amount to application of the group logarithm map (i.e., the matrix logarithm) and this can moreover be replaced by a linear approximation, which reduces computational complexity.

In the example of 5.2, the observation process has its values in the unit sphere S^2 . Here, a detailed discussion of the mathematical concepts introduced in the paper is provided and, as in the first example, it is shown how the particle filtering algorithm can be implemented in a way that reduces the complexity of geometric constructions.

These two examples are representative of problems where the observation process lies in a classical matrix Lie group or in a related symmetric space. With some adjustment, they could be extended in a straightforward way to deal with general matrix Lie groups and their symmetric spaces. This has not been possible in the current paper, essentially for a reason of space. However, for the case where the observation process takes its values in a Stiefel manifold, the reader is referred to [20].

The particle filtering algorithm proposed in this paper seems entirely new in the literature. It is encouraging that, building directly on existing results (from [14, 18] and [21]), the algorithm can be described and its convergence proved. However, the current treatment still suffers from certain drawbacks which should be addressed in future work. For instance, the restriction to compact manifolds is quite artificial and it should be possible to drop it with some additional care. Also, the convergence result of Proposition 5 does not explicitly provide a rate of convergence (see discussion after the proposition).

2. GENERAL FILTERING PROBLEM

To state a filtering problem, it is sufficient to define the signal and observation models [22]. In the following, the unknown signal will be a hidden Markov process x with values in some Polish (i.e., complete separable metric) space (S, \mathcal{S}) . Concretely, in most cases, this space S is either a finite set or a Euclidean space. The observation process Y will be a diffusion in a differentiable manifold \mathcal{M} of dimension d (see [23] for required background in differential geometry). Besides precisely giving the definition of x and Y , this section aims to put the resulting filtering problem into perspective. This is done by comparing it to other filtering problems, both classical and involving an observation process with values in a manifold, and by discussing how it can simulated numerically.

Both x and Y are defined in continuous time. One begins by introducing a complete probability space $(\Omega, \mathcal{A}, \mathbb{P})$ on which x and Y are defined.

The unknown signal x is a time invariant Markov process. Let $C_b(S)$ denote the space of bounded continuous functions $\varphi : S \rightarrow \mathbb{R}$. The generator of x is an operator A with domain $D(A)$, a dense subspace of $C_b(\varphi)$ which contains the constant function 1. The generator A is assumed to verify $A1 = 0$ and, for $\varphi \in D(A)$,

$$d\varphi(x_t) = A\varphi(x_t)dt + dM_t^\varphi, \quad (1)$$

where $t \geq 0$ is the time variable. The meaning of the above equation is that

$$\varphi(x_t) - \varphi(x_0) - \int_0^t A\varphi(x_s)ds = M_t^\varphi,$$

where M^φ is a martingale adapted to the augmented natural filtration of x , denoted \mathcal{X} [24]. The distribution of x is determined by its generator A and its initial distribution μ ; i.e., μ is the distribution of the initial value x_0 .

Two typical examples of the above definition are when x is a diffusion in some manifold \mathcal{N} of dimension l , possibly $\mathcal{N} = \mathbb{R}^l$, and when x is a finite state Markov process. The first example may appear when the problem is applied to tracking the pose of a rigid body. The second is generally considered within the framework of change detection. It is usual to refer to x as a hidden Markov process as it is unknown and only observed through Y .

In order to define the observation process Y , assume \mathcal{M} is a C^4 manifold. The signal x is encoded through the sensor function $H : S \times \mathcal{M} \rightarrow T\mathcal{M}$, where $T\mathcal{M}$ denotes the tangent bundle of \mathcal{M} . For $s \in S$, it is required the application $p \mapsto H(s, p)$ is a C^1 vector field on \mathcal{M} . That is, $H(s, p) \in T_p\mathcal{M}$ where $T_p\mathcal{M}$ is the tangent space to \mathcal{M} at p . Observation noise is introduced as follows. Let $(\Sigma_r; r = 1, \dots, m)$ be C^2 vector fields on \mathcal{M} and B a standard Brownian motion in \mathbb{R}^m , which is independent from x . The observation process Y is assumed nonexplosive (i.e., for $p \in \mathcal{M}$, if $Y_0 = p$ then Y_t is defined with values in \mathcal{M} for $t \geq 0$) and satisfying [25],

$$dY_t = H(x_t, Y_t)dt + \Sigma_r(Y_t) \circ dB_t^r. \quad (2)$$

Here and in all the following, summation convention is understood, (that is, a sum is understood over any repeated subscripts or superscripts). Since the Σ_r are vector fields on \mathcal{M} , the $\Sigma_r(Y_t)$ are tangent vectors to \mathcal{M} at Y_t . Moreover, the circle is the usual notation for the Stratonovich differential [26]. The filtering problem defined by (1) and (2) is more general than a classical filtering problem where Y depends on x through an additive white noise model. In (2) the observation noise B is “carried” by the vector fields Σ_r . In other words, it is introduced in a way which depends on the current observation. The model (2) reduces to a classical one when $\mathcal{M} = \mathbb{R}^d$ and i) $m = d$ with $\Sigma_r(p) = e_r$, where

$(e_r; r = 1, \dots, d)$ is a canonical basis of \mathbb{R}^d ; ii) $H(s, p) = H(s)$ is given by a function $H : S \rightarrow \mathbb{R}^d$.

Intuitively, the classical problem is a limit of problems where the observation process is sampled at times $k\delta$ where $k \in \mathbb{N}$ and δ the sampling step size. Then, since Y has values in \mathbb{R}^d , it is possible to consider discrete observations

$$(\text{for observations in } \mathbb{R}^d) \quad \Delta Y_k = Y_{(k+1)\delta} - Y_{k\delta}. \quad (3)$$

Each one of these increments is normally distributed conditionally on x . Therefore, the corresponding likelihood function is known analytically. This last aspect is what characterises a classical filtering problem. It is unaffected if x is a diffusion in a manifold \mathcal{M} . In practice, whatever the state space S , what matters is the ability to simulate the sample path distribution of x with sufficient precision.

The chief difference between the classical filtering problem, and the problem of filtering with observations in a manifold, given by the observation model (2), is that no formula similar to (3) is immediately available in the latter case. As a result, it is difficult to arrive at an analytic expression of the likelihood function. As the standard analysis, based on the assumption that Y depends on x through an additive white noise model, does not apply to the observation model (2), it becomes necessary to search for a well-defined generalisation of (3). This generalisation, will be introduced as of the beginning of Section 3, in the form of equation (6), based on an original use of the concept of *connector map*. The systematic use of connector maps constitutes the main ingredient of the new particle filtering algorithm proposed in the present paper.

In the mathematical literature, problems with observations in a manifold have been stated in two different forms. In [14, 15] the problem statement is the same as above, with the additional requirement that Y is an elliptic diffusion. In fact, this same requirement will be imposed in Section 3. In [16, 17], the diffusion Y is specified in terms of its horizontal lift. The difference between the two settings is that, in the latter one, there is a choice of connection already included in the problem statement. In the setting of (2), a metric and connection need to be constructed from the vector fields Σ_r . The current paper differs from [15] in the way this construction is defined. See Section 3.

In a real world application, the issue of how to simulate Y numerically is irrelevant. Indeed, Y is itself observed through some measurement device. However, in order to carry out a computer experiment, it is necessary to simulate Y given the model (2). In many cases, \mathcal{M} is embedded in a higher dimensional Euclidean space, $\mathcal{M} \subset \mathbb{R}^N$, and the vector fields H and Σ_r are restrictions of complete vector fields defined on all of \mathbb{R}^N . In particular, this is true for all the examples in Section 5. Then, numerical simulation of Y is a matter of solving a stochastic differential equation. It is possible to

use Euler, Milstein or a higher order stochastic Taylor scheme, according to desired precision [27]. All these numerical schemes will suffer from the same problem of producing an approximation which “falls off” the embedded manifold. A simple way of dealing with this is to project back onto the manifold once the approximation has been computed. When \mathcal{M} is presented as an abstract manifold, a generally applicable numerical scheme is the McShane approximation. This approximates Y by processes Y^δ where δ is a discretisation step size. On each interval $[k\delta, (k+1)\delta[$ for $k \in \mathbb{N}$, Y^δ is the solution of an ordinary differential equation,

$$\dot{Y}_t^\delta = H(x_t, Y_t^\delta) + \Sigma_r(Y_t) \Delta B_k^r, \quad (4)$$

where the dot denotes differentiation with respect to time and $\Delta B_k = \delta^{-1}(B_{(k+1)\delta} - B_{k\delta})$. When \mathcal{M} is a compact Riemannian manifold and Y_0^δ is taken to be the same as Y_0 , the rate of local uniform convergence of the paths of Y^δ to those of Y is given in [28],

$$\mathbb{P}(\sup_{t \leq T} d(Y_t^\delta, Y_t) > \varepsilon) = O(\delta), \quad (5)$$

where d is the Riemannian distance and $\varepsilon > 0$. Note that this depends on T so the same precision ε for larger T requires smaller δ .

The numerical approximation (4) gives an intuitive interpretation of the stochastic differential equation (2), whose rigorous definition may be found in [25]. Roughly, this approximation states that, over short time intervals of the form $[t, t + \delta]$, the observation process Y moves along the integral curves of a random vector field $H + \Sigma_r \Delta B^r$. This motion has a deterministic component H , which depends on the unknown signal, and a stochastic or “noise” component $\Sigma_r \Delta B^r$. These two components are not observed independently, but rather only through the observations Y_t and $Y_{t+\delta}$. From the point of view of filtering, it is natural to search for a transformation which takes these observations Y_t and $Y_{t+\delta}$ to a random vector of the form “ H +normally distributed noise.” This is the starting point of the following section.

3. STOCHASTIC DIFFERENTIAL GEOMETRY

The aim of this section is to generalise (3) so that it can be applied for the observation model (2). This is done by adopting the following point of view. In (3), ΔY_k is a tangent vector to $\mathcal{M} = \mathbb{R}^d$ at $Y_{k\delta}$, determined from two successive samples $Y_{k\delta}$ and $Y_{(k+1)\delta}$. Accordingly, in the general case of (2), a mapping $I : \mathcal{M} \times \mathcal{M} \rightarrow T\mathcal{M}$ will be constructed with the following properties. First, $I(p_1, p_2) \in T_{p_1}\mathcal{M}$ for all $p_1, p_2 \in \mathcal{M}$. Second, if ΔY_k is defined as

$$\Delta Y_k = I(Y_{k\delta}, Y_{(k+1)\delta}), \quad (6)$$

then, in the limit $\delta \downarrow 0$, the distribution of ΔY_k is approximately normal, conditionally on the unknown signal x . Such a mapping I is called a connector map, and it turns out that, it is always possible to choose a connector map

I which verifies the required properties. This is stated precisely in Proposition 3, which is the main result in this section.

Proposition 3 requires some basic concepts from stochastic differential geometry. These are recovered in 3.1 which, in particular, gives the definition of the antidevelopment process y of Y . The choice of connection leading to I is given in 3.2. Finally, Proposition 3 is stated and discussed in 3.3. Here, 3.1 and 3.2 are based on [14].

3.1. Stochastic antidevelopment

In the filtering problem defined by (1) and (2), the available information is a path of the observation process Y , taken over some time interval $t \leq T$. An observer is only able to compute functionals of the process Y . These are processes, in practice real or vector valued, adapted to the augmented natural filtration of Y ; which is denoted \mathcal{Y} . Two classes of such functionals are the building blocks for the following, the Stratonovich and Itô integrals along Y . The notion of antidevelopment is itself defined using these two kinds of integrals.

Assume a Riemannian metric $\langle \cdot, \cdot \rangle$ and a compatible connection ∇ are defined on \mathcal{M} . In the current paragraph, these are not specified and can be chosen arbitrarily.

To define the integrands involved in Stratonovich and Itô integrals along Y , let \mathcal{F} be the filtration where $\mathcal{F}_t = \mathcal{X}_\infty \vee \mathcal{B}_t$. Here, $\mathcal{X}_\infty = \bigvee_{t \geq 0} \mathcal{X}_t$ and \mathcal{B} is the augmented natural filtration of B . A stochastic integrand θ is an \mathcal{F} -adapted process with values in $T^*\mathcal{M}$ which is above Y ; this means $\theta_t \in T_{Y_t}^*\mathcal{M}$ for $t \geq 0$.

Let (θ_t, \cdot) denote the application of the linear form θ_t . The Stratonovich and Itô integrals of θ along Y are real valued \mathcal{F} -adapted processes with the following differentials,

$$(\theta_t, \circ dY) = (\theta_t, H)dt + (\theta_t, \Sigma_r) \circ dB_t^r, \quad (7)$$

$$(\theta_t, dY) = \{(\theta_t, H) + (1/2)(\theta_t, \nabla_{\Sigma_r} \Sigma_r)\}dt + (\theta_t, \Sigma_r)dB_t^r. \quad (8)$$

The Stratonovich differential does not involve the chosen connection. On the other hand, the connection appears explicitly in the Itô differential. Note that Y is called a ∇ -martingale if $H + (1/2)\nabla_{\Sigma_r} \Sigma_r \equiv 0$. In this case, (θ_t, dY) is a martingale differential, as is the case for a usual Itô differential.

Using the Riemannian metric $\langle \cdot, \cdot \rangle$, it is possible to define Stratonovich and Itô integrals of vector fields. A vector field along Y is a process E , \mathcal{F} -adapted with values in $T\mathcal{M}$ and which is above Y ; in the sense that $E_t \in T_{Y_t}\mathcal{M}$ for $t \geq 0$. A corresponding process θ in $T^*\mathcal{M}$ is then given by $(\theta_t, \cdot) = \langle E_t, \cdot \rangle$. The resulting differentials of (7) and (8) are written $\langle E_t, \circ dY \rangle$ and $\langle E_t, dY \rangle$.

In order to formulate the notion of antidevelopment, it is necessary to define what it means for E to be parallel (i.e., along Y). If the paths of Y were differentiable,

this would have the usual meaning from differential geometry. While this is not the case, due the presence of Brownian terms, it is still possible to introduce a stochastic covariant derivative of E and require this to vanish. Thus E is said to be parallel if

$$\nabla_{\circ dY} E_t = 0, \quad (9)$$

where the left hand side is the stochastic covariant derivative. In order to give this a precise meaning, consider first the case where G is a C^1 vector field on \mathcal{M} and $E = G(Y_t)$. Assuming the usual properties of a connection, (2) would give

$$\nabla_{\circ dY} G(Y_t) = \nabla_H G(Y_t)dt + \nabla_{\Sigma_r} G(Y_t) \circ dB_t^r. \quad (10)$$

This is extended to a general vector field E along Y by the following transformation. For a C^2 function f on \mathcal{M} , let $\nabla^2 f$ be the Hessian of f with respect to ∇ . By definition, $\nabla_H Gf = HGf - \nabla^2 f(H, G)$ and similarly for $\nabla_{\Sigma_r} G$. Now, $(\theta_t, \cdot) = \nabla^2 f(\cdot, E_t)$ is a process in $T^*\mathcal{M}$ as in (7). Thus, (10) can be generalised by writing

$$\nabla_{\circ dY} E_t f = dE_t f - \nabla^2 f(\circ dY, E_t), \quad (11)$$

which is taken as the definition of the stochastic covariant derivative. In integral notation, (11) defines a vector field $\int_0^t \nabla_{\circ dY} E_s$ along Y . For any C^2 function f on \mathcal{M} , this verifies

$$\int_0^t \nabla_{\circ dY} E_s f = E_t f - E_0 f - \int_0^t \nabla^2 f(\circ dY, E_s).$$

Recall the connection ∇ is compatible with the metric $\langle \cdot, \cdot \rangle$. Again, if the usual properties of a connection were assumed, one would expect the following. If E, K are vector fields along Y , then

$$d\langle E_t, K_t \rangle = \langle \nabla_{\circ dY} E_t, K_t \rangle + \langle E_t, \nabla_{\circ dY} K_t \rangle. \quad (12)$$

The proof that this identity indeed holds true is here omitted. The main point is that it is possible to perform calculations involving $\nabla_{\circ dY}$ just like in differential geometry, by treating $\circ dY$ as the tangent vector to the process Y .

Now, it is possible to define the antidevelopment of Y . This is a process y which has values in \mathbb{R}^d . A parallel orthonormal frame is a family $(E^i) \equiv (E^i; i = 1, \dots, d)$ of vector fields along Y which verify (9) and $\langle E_0^i, E_0^j \rangle = \delta_{ij}$. In this case, it follows from (12) that $\langle E_t^i, E_t^j \rangle = \delta_{ij}$ for $t \geq 0$. Given a parallel orthonormal frame (E^i) , the process y is defined by

$$y_t^i = \int_0^t \langle E_s^i, \circ dY \rangle. \quad (13)$$

It seems this definition is arbitrary, due to the issue of uniqueness of a parallel orthonormal frame. However, the classical uniqueness result for linear stochastic differential equations can be used to show y is essentially

unique. In fact, (E^i) is determined by (E_0^i) so that different choices of (E^i) only amount to y being multiplied by an orthogonal matrix.

Stratonovich and Itô integrals along Y can be written as classical Stratonovich and Itô integrals with respect to y . For the Stratonovich integral, this is a straightforward result of (13). For θ as in (7),

$$(\theta_t, \circ dY) = \theta_t(t) \circ dy_t^i, \quad (14)$$

where $\theta_t(t) = (\theta_t, E_t^i)$. This follows by the chain rule of Stratonovich calculus, since $(\theta_t, \cdot) = \theta_t(t) \langle E_t^i, \cdot \rangle$. Recall that (E^i) is a basis of the tangent space to \mathcal{M} at Y_t .

The case of the Itô integral is slightly more involved. Here, it is necessary to realise that (13) can also be written

$$y_t^i = \int_0^t \langle E_s^i, dY \rangle. \quad (15)$$

In other words, since the vector fields E^i are parallel, there is no difference between their Stratonovich and Itô integrals. Expressing (13) as in (7), (the arguments of H and Σ_r are dropped for space)

$$\langle E_t^i, \circ dY \rangle = \langle E_t^i, H \rangle + \langle E_t^i, \Sigma_r \rangle \circ dB_t^r.$$

In order to change the Stratonovich differential into an Itô differential, note from (9) and (12)

$$d\langle E_t^i, \Sigma_r \rangle = \langle E_t^i, \nabla_{\circ dY} \Sigma_r \rangle,$$

and from (10)

$$d\langle E_t^i, \Sigma_r \rangle = \langle E_t^i, \nabla_H \Sigma_r \rangle dt + \langle E_t^i, \nabla_{\Sigma_r} \Sigma_r \rangle \circ dB_t^r.$$

From this, it finally follows

$$\langle E_t^i, \Sigma_r \rangle \circ dB_t^r = \langle E_t^i, \Sigma_r \rangle dB_t^r + (1/2) \langle E_t^i, \nabla_{\Sigma_r} \Sigma_r \rangle dt, \quad (16)$$

so that (15) follows from (8). Given (15), it is possible to write

$$(\theta_t, dY) = \theta_t(t) dy_t^i, \quad (17)$$

using a reasoning similar to the one that lead to (14).

At the beginning of this paragraph, it was claimed that Stratonovich and Itô integrals along Y can be used to obtain functionals of this process. Even when the process θ is \mathcal{Y} -adapted, this is not clear from (7) and (8). For example, the right hand side in each of these two formulae contains a stochastic integral with respect to B . However, there is no reason to believe that \mathcal{B} is the same as \mathcal{Y} . Indeed, the definition (2) of Y involves both x and B .

It is possible to write (7) and (8) in an alternative form, which makes it evident that the resulting integrals are \mathcal{Y} -adapted as soon as θ is \mathcal{Y} -adapted. To do so, assume \mathcal{M} is embedded in a higher dimensional Euclidean space, $\mathcal{M} \subset \mathbb{R}^N$. There is little loss of generality in this assumption. Whitney's theorem asserts such an embedding exists if \mathcal{M} is C^∞ , connected and paracompact [10]. Such conditions are always verified in practice. Also, when \mathcal{M} is embedded in \mathbb{R}^N , it is possible to assume the vector fields H and Σ_r are restrictions

of vector fields defined on all of \mathbb{R}^N . Under these two assumptions, let η^1, \dots, η^N be canonical (i.e., rectangular) coordinates on \mathbb{R}^N and write $\theta(t) = \theta_\alpha(t) d\eta^\alpha$ where $\alpha = 1, \dots, N$. Replacing in (7) and using the chain rule of Stratonovich calculus,

$$(\theta_t, \circ dY) = \theta_\alpha(t) \circ dY_t^\alpha, \quad (18)$$

where Y^α are the coordinates of Y_t . In other words, the Stratonovich integral (7) is just the classical Stratonovich integral. Moving on to (8), a similar formula will be shown. Note that the coordinates η^α can be thought of as C^2 functions on \mathcal{M} . Using the notation $\nabla^2 \eta^\alpha(\Sigma_r, \Sigma_r) = H_{rr}^\alpha$,

$$(\theta_t, dY) = \theta_\alpha(t) \{dY_t^\alpha - (1/2) H_{rr}^\alpha(Y_t) dt\}. \quad (19)$$

Thus, the Itô integral (8) is the sum of the classical Itô integral, corresponding to $\theta_\alpha(t) dY_t^\alpha$, and of a correction term depending on the connection ∇ . It is possible that the latter correction term vanishes identically, so the Itô integral (8) is the same as the classical Itô integral. This is for instance the case when the connection ∇ is the one which \mathcal{M} inherits from \mathbb{R}^N . From the last two formulae, it is seen that the Itô and Stratonovich integrals are \mathcal{Y} -adapted as soon as the process θ is \mathcal{Y} -adapted. This follows from the same property for classical Itô and Stratonovich integrals.

The proof of (19) is as follows. Assume in (7) that $\theta_t = \tau(Y_t)$ for some differential form τ on \mathcal{M} . Passing from Stratonovich to Itô differentials,

$$(\theta_t, \Sigma_r) \circ dB_t^r = (\theta_t, \Sigma_r) dB_t^r + (1/2) \Sigma_r(\theta_t, \Sigma_r) dt.$$

Using the fact that ∇ is compatible with $\langle \cdot, \cdot \rangle$ and comparing to (8), the following general rule is found,

$$(\tau(Y_t), \circ dY) = (\tau(Y_t), dY) + (1/2) (\nabla_{\Sigma_r} \tau(Y_t), \Sigma_r) dt. \quad (20)$$

In order to obtain (19), it is enough to apply (20) with $\tau = d\eta^\alpha$ and recall the definition of $\nabla^2 \eta^\alpha$, (as cited before (11)).

In many cases, the embedding $\mathcal{M} \subset \mathbb{R}^N$ is known explicitly. For instance, \mathcal{M} is often directly defined as an embedded submanifold in Euclidean space. Then, formulae (18) and (19) can be used to compute the Itô and Stratonovich integrals. This is useful since the properties of classical stochastic integrals become available. This situation will apply throughout the examples of Section 5.

3.2. Le Jan-Watanabe connection

This paragraph describes the “right” choice for the metric $\langle \cdot, \cdot \rangle$ and the connection ∇ . When this choice is used to evaluate (13), the resulting antidevelopment y of Y depends on the unknown signal x through a classical additive white noise model. This reduces the initial filtering problem defined by (1) and (2) to a classical filtering problem. The precise statement is given in Proposition 1.

A usual simplifying assumption imposed on Y in the literature is that, conditionally on x , it is an elliptic diffusion in \mathcal{M} . This means that, for each $p \in \mathcal{M}$, the subspace of $T_p\mathcal{M}$ spanned by the $\Sigma_r(p)$ is equal to $T_p\mathcal{M}$. Under this assumption, elementary linear algebra implies there exists a unique Riemannian metric $\langle \cdot, \cdot \rangle$ such that

$$\langle E, K \rangle = \langle E, \Sigma_r(p) \rangle \langle K, \Sigma_r(p) \rangle, \quad (21)$$

for $E, K \in T_p\mathcal{M}$. This will be the choice of metric made in this following. In [15], the corresponding Levi-Civita connection is used in (13). Here, a different connection is used. Namely, the connection ∇ is taken to be the Le Jan-Watanabe connection. Based on [29], this is here defined as follows.

Let E be any C^1 vector field on \mathcal{M} . It results from (21) that this can be written $E = E^r \Sigma_r$, where $E^r = \langle E, \Sigma_r \rangle$. For $K \in T_p\mathcal{M}$, let

$$\nabla_K E = (KE^r) \Sigma_r(p). \quad (22)$$

This defines a connection ∇ compatible with $\langle \cdot, \cdot \rangle$. Precisely, for C^1 vector fields E, G on \mathcal{M} ,

$$\begin{aligned} K \langle E, G \rangle &= (KE^r) G^r + E^r (KG^r) \\ &= \langle \nabla_K E, G \rangle + \langle E, \nabla_K G \rangle, \end{aligned} \quad (23)$$

where (21) and (22) have been used.

With regard to the proof of Proposition 1, the only necessary property of ∇ is

$$\nabla_{\Sigma_r} \Sigma_r = 0. \quad (24)$$

To obtain (24), replace (21) in (22). Since Σ_r is a derivation,

$$\begin{aligned} \nabla_{\Sigma_r} \Sigma_r &= \\ \Sigma_r \langle \Sigma_r, \Sigma_v \rangle \Sigma_v &= \\ \Sigma_r \langle \Sigma_r, \Sigma_w \rangle \langle \Sigma_v, \Sigma_w \rangle \Sigma_v &+ \langle \Sigma_r, \Sigma_w \rangle \Sigma_r \langle \Sigma_v, \Sigma_w \rangle \Sigma_v. \end{aligned}$$

A simplification of the third line gives

$$\begin{aligned} \nabla_{\Sigma_r} \Sigma_r &= \\ \Sigma_r \langle \Sigma_r, \Sigma_w \rangle \Sigma_w &+ \Sigma_w \langle \Sigma_w, \Sigma_v \rangle \Sigma_v = \\ \nabla_{\Sigma_r} \Sigma_r &+ \nabla_{\Sigma_w} \Sigma_w, \end{aligned}$$

which immediately leads to (24).

PROPOSITION 1. *Let (E^i) be a parallel orthonormal frame and y given by (13), where the connection ∇ is defined by (22). Also, let \mathcal{Y} be the augmented natural filtration of y . Then, y has its values in \mathbb{R}^d . Moreover, for $t \geq 0$, $\mathcal{Y}_t = \mathcal{Y}_t$ and*

$$dy_t = h_t dt + d\beta_t, \quad h_t^i = \langle E_t^i, H \rangle, \quad (25)$$

where β is a Brownian motion in \mathbb{R}^d which is independent from x .

PROOF By definition, y has its values in \mathbb{R}^d . The proof of the second claim, i.e., $\mathcal{Y}_t = \mathcal{Y}_t$, is identical to that

of the analogous claim made in [15], (Theorem IV.3 on page 135). More generally, this claim holds for any connection ∇ compatible with $\langle \cdot, \cdot \rangle$.

In order to obtain (25) define first

$$\beta_t^i = \int_0^t \langle E_s^i, \Sigma_r \rangle dB_s^r. \quad (26)$$

Clearly, β is an \mathcal{F} -local martingale. Moreover, (21) implies

$$\int_0^t \langle E_s^i, \Sigma_r \rangle \langle E_s^j, \Sigma_r \rangle ds = \delta_{ij} \int_0^t ds.$$

By Lévy's characterisation, β is an \mathcal{F} -Brownian motion. Since $\mathcal{F}_0 = \mathcal{X}_\infty$, it follows β is independent from x .

Recall that y can be computed from (15). Replacing the definition (8) of this integral, it follows

$$dy_t^i = \{h_t^i + (1/2) \langle E_t^i, \nabla_{\Sigma_r} \Sigma_r \rangle\} dt + d\beta_t^i.$$

However, (24) states the second term is identically zero. This completes the proof of (25).

3.3. Connector maps

This paragraph is devoted to Proposition 3, which is the main result in the current section. This theorem states that the mapping I to be used in (6) is a connector map which verifies certain conditions, expressed in terms of the Le Jan-Watanabe connection introduced in (22) of the previous paragraph. Roughly, such a connector map becomes the discrete counterpart of antidevelopment. This is the content of Proposition 2 below.

As in 3.1, let $\langle \cdot, \cdot \rangle$ be a Riemannian metric and ∇ a compatible connection. Consider first the definition of a geodesic connector. Recall that, for each $p \in \mathcal{M}$ and $K \in T_p\mathcal{M}$ there exists a unique geodesic curve $\gamma :]-\epsilon, \epsilon[\rightarrow \mathcal{M}$ such that $\gamma(0) = p$ and $\dot{\gamma}(0) = K$, where $\epsilon > 0$. A geodesic curve is one satisfying the geodesic equation

$$\nabla_{\dot{\gamma}} \dot{\gamma}(t) = 0. \quad (27)$$

If the manifold \mathcal{M} is complete for the connection ∇ , then any geodesic curve γ can be extended to all $t \in \mathbb{R}$, (that is, ϵ can be taken arbitrarily large). Assuming this is the case, consider the exponential mapping $\exp : T\mathcal{M} \rightarrow \mathcal{M}$. For p and K as above, it is usual to write $\exp(K) = \exp_p(K)$ in order to distinguish the base point p . By definition, $\exp_p(K) = \gamma(1)$.

The exponential mapping is locally invertible. If two points $p, q \in \mathcal{M}$ are close enough to each other, then there exists a unique geodesic γ as above such that $\gamma(1) = q$. It is suitable to write $K = \log_p(q)$, since then "log is the inverse of exp." Note that both \exp_p and \log_p are C^2 mappings [23].

Let $d(\cdot, \cdot)$ denote the Riemannian distance associated to $\langle \cdot, \cdot \rangle$. It will be assumed that \mathcal{M} has a strictly positive radius of injectivity for the connection ∇ . That is, there exists $r > 0$ such that for all $p, q \in \mathcal{M}$ the inequality $d(p, q) < r$ implies $\log_p(q)$ is well defined. The two

assumptions of completeness of \mathcal{M} with respect to ∇ and of a strictly positive radius of injectivity are not very stringent. In particular, they are verified whenever \mathcal{M} is compact.

A geodesic connector is a mapping I defined following [15]. Let $u < r$ and $\phi : \mathbb{R}_+ \rightarrow [0, 1]$ a decreasing C^∞ function such that $\phi(d) = 1$ if $d \leq u$ and $\phi(d) = 0$ if $d \geq r$. Let I be given by

$$I(p, q) = \phi(d(p, q)) \log_p(q). \quad (28)$$

Now, I is essentially intended to be the mapping \log_p where p is the first argument. The function ϕ is only introduced as a cutoff, to avoid points q which lie beyond the injectivity radius.

In the following, let $\|\cdot\|$ denote Riemannian length, (e.g., $\|G\|^2 = \langle G, G \rangle$), and \mathbb{E} expectation with respect to \mathbb{P} .

PROPOSITION 2. *Assume \mathcal{M} is compact. Let G be a vector field along Y such that $t \mapsto \|G_t\|$ is bounded. Then, for $t \geq 0$, the Itô integral $R = \int_0^t \langle G_s, dY \rangle$ is square integrable (that is, $\mathbb{E}|R|^2 < \infty$). Moreover, if for $\delta > 0$*

$$R_\delta = \sum_{k\delta < t} \langle G_{k\delta}, I(Y_{k\delta}, Y_{(k+1)\delta}) \rangle, \quad (29)$$

then $\mathbb{E}|R_\delta - R|^2 \rightarrow 0$ as $\delta \downarrow 0$.

Proposition 2, (exactly the same as Theorem 3.4.2 on page 55 in [30]), generalises the definition of a classical Itô integral as a limit in the square mean of Riemann sums. In the classical definition, the increments of the integrating process are given by (3). Proposition 2 shows that for an Itô integral as in (8), at least when the manifold \mathcal{M} is compact, the limit continues to hold if the classical increments are replaced by those obtained from a geodesic connector. That is, from (6) where I is given by (28). Note that (3) is truly a special case of (28), since geodesics in $\mathcal{M} = \mathbb{R}^d$ are just straight lines.

The use of geodesic connectors seems natural from a theoretical point of view. However, the mapping I of (28) may be quite difficult to compute. It is important to note that Proposition 2 continues to hold, all other hypotheses being the same, if I is any mapping with the same property $I(p_1, p_2) \in T_{p_1}\mathcal{M}$, as long as (this is proved in Appendix A)

(I1) I is jointly C^2 (as a mapping $\mathcal{M} \times \mathcal{M} \rightarrow T\mathcal{M}$)

(I2) $dI(p, p)(V) = V$ for all $V \in T_p\mathcal{M}$

(I3) $\nabla^2 I(p, p)(V, V) = 0$ for all $V \in T_p\mathcal{M}$

Here, dI and $\nabla^2 I$ are the derivative and the Hessian of I with respect to its second argument. When the first argument p is fixed, I has its values in a fixed vector space $T_p\mathcal{M}$. Therefore, it is possible to speak of its derivative and Hessian. Conditions **(I1-I3)** are verified by the mapping of (28). Intuitively, any mapping I which verifies them is “geodesic connector up to second order.”

It is now possible to state Proposition 3.

PROPOSITION 3. *Assume \mathcal{M} is compact and the connection ∇ is given by (22). Moreover, let y be given by (13). For $\delta > 0$, let ΔY_k be given by (6) and $\Delta y_k = y_{(k+1)\delta} - y_{k\delta}$ where $k \in \mathbb{N}$. Then,*

$$\mathbb{E}|\Delta y_k^i - \Delta Y_k^i|^2 = o(\delta), \quad (30)$$

where $\Delta Y_k^i = \langle E_{k\delta}^i, \Delta Y_k \rangle$. In particular, as $\delta \downarrow 0$, the distribution of the vector ΔY_k^i conditionally on $\mathcal{F}_{k\delta}$ is approximately normal. Precisely, as $\delta \downarrow 0$,

$$\mathcal{L}\left\{\delta^{-1/2}(\Delta Y_k^i - \delta \times h_{k\delta}^i)\right\} \longrightarrow N_d. \quad (31)$$

Here, \mathcal{L} denotes the law of the random vector whose components are in brackets and N_d denotes a standard normal distribution in \mathbb{R}^d .

PROOF The order of magnitude given in equation (30) is essentially a result of Proposition 2.

For (31), note from Proposition 1

$$\Delta y_k^i = \int_{k\delta}^{(k+1)\delta} h_t^i dt + \{\beta_{(k+1)\delta}^i - \beta_{k\delta}^i\},$$

but then, as is clear by bounded convergence,

$$\mathcal{L}\{\delta^{-1/2}(\Delta y_k^i - \delta \times h_{k\delta}^i)\} \longrightarrow N_d.$$

The proposition follows from (30) by noting that $h_{k\delta}$ is measurable w.r.t. to $\mathcal{F}_{k\delta}$ and β is an \mathcal{F} -Brownian motion, so $\beta_{(k+1)\delta} - \beta_{k\delta}$ is independent of $\mathcal{F}_{k\delta}$ with the required covariance.

4. PARTICLE FILTERING ALGORITHM

This section is concerned with the numerical solution of the filtering problem defined by (1) and (2). In 4.1, Proposition 4 expresses the solution in closed form using a Kallianpur-Striebel formula. In 4.2, a particle filtering algorithm is proposed in order to numerically implement this formula. The convergence of this algorithm to the solution of the original filtering problem, (i.e., that of (1) and (2)), is the subject of Proposition 5 in 4.3.

4.1. Kallianpur-Striebel formula

A Kallianpur-Striebel formula may be thought of as an abstract Bayes formula which expresses in closed form the solution of a continuous time filtering problem. For classical filtering problems, where the observation process depends on the unknown signal through an additive white noise model, the Kallianpur-Striebel formula is quite well known [22]. The idea in this section is to profit from Proposition 1 in order to obtain a similar formula for the problem defined by (1) and (2).

The following Proposition 4 uses the notion of a copy of the unknown signal independent from the observation process. This means a process \tilde{x} with the same law as the signal x , given by (1), but which is independent from the observation process Y .

Let π_t be the conditional distribution of x_t , given past observations \mathcal{Y}_t . The aim of the filtering problem is precisely to find $\pi_t(\varphi)$ for any function $\varphi \in C_b(S)$. By definition, π is a *càdlàg* process with values in the space of probability measures on (S, \mathcal{S}) and such that

$$\pi_t(\varphi) = \mathbb{E}[\varphi(x_t) \mid \mathcal{Y}_t]. \quad (32)$$

From this, it is clear that the process π is \mathcal{Y} -adapted. Thus it should be possible to write it down in terms of some functional of the process Y (see the discussion at the beginning of 3.1).

PROPOSITION 4. *Let \tilde{x} be a copy of x independent from Y . Let \tilde{H} be the process where $\tilde{H}_t = H(\tilde{x}_t, Y_t)$. For $\varphi \in C_b(S)$,*

$$\pi_t(\varphi) = \frac{\rho_t(\varphi)}{\rho_t(1)}, \quad \rho_t(\varphi) = \mathbb{E}[\varphi(\tilde{x}_t)L_t \mid \mathcal{Y}_\infty], \quad (33)$$

where $\mathcal{Y}_\infty = \bigvee_{t \geq 0} \mathcal{Y}_t$ and

$$L_t = \exp \left(\int_0^t \langle \tilde{H}_s, dY \rangle - (1/2) \int_0^t \|\tilde{H}_s\|^2 ds \right). \quad (34)$$

Here the Itô differential $\langle \tilde{H}_s, dY \rangle$ is computed using the Le Jan-Watanabe connection defined by (22).

PROOF Recall that Proposition 1 states $\mathcal{Y} = \bar{\mathcal{Y}}$, where $\bar{\mathcal{Y}}$ is the augmented natural filtration of y . If y is computed using the Le Jan-Watanabe connection defined by (22), then y depends on x using a classical additive white noise model (25). The corresponding Kallianpur-Striebel formula is the same as above, but with L_t given by (this is the classical formula [22])

$$L_t = \exp \left(\int_0^t \tilde{h}_i(s) dy_s^i - \int_0^t \tilde{h}_i(s) \tilde{h}_i(s) ds \right).$$

Here, $\tilde{h}_i(s) = \langle E_s^i, \tilde{H} \rangle$, as in (25). Using the fact that $\mathcal{Y} = \bar{\mathcal{Y}}$, the Proposition follows by (17) and (32).

4.2. Algorithm description

Numerical implementation of the Kallianpur-Striebel formula, i.e., of (33) in Proposition 4, can proceed in at least two ways [22]. First, it is possible to derive from this formula the corresponding Zakai equation, whose solution can be attempted using a finite difference or spectral method. The Zakai equation is a stochastic differential equation satisfied by the unnormalised distribution ρ . It was indeed found in [14], but it seems its solution has not been specifically considered in the literature. Second, thinking of the Kallianpur-Striebel formula as an abstract Bayes formula, it is possible to implement it using sequential Monte Carlo strategy. Here, this latter option is pursued. The result takes the form of a particle filtering algorithm.

The proposed particle filtering algorithm and its convergence are easily understood in the framework of a discretised version of the original filtering problem of (1) and (2). Note that, in (33), \tilde{x} and Y are taken

to be independent. It is then natural to think of \tilde{x} as a computer simulated version of x . However, this is just an idealisation. A computer experiment can only yield an approximation of the model (1). Such an approximation is obtained for a given discretisation step size $\delta > 0$. It is a process denoted x^δ with its value at time $k\delta$ written $x_{k\delta}^\delta = x_k^\delta$. The particle filtering algorithm will be convergent when x^δ verifies the following conditions

- (x1) $\{x_k^\delta; k \in \mathbb{N}\}$ is a time invariant Markov chain
- (x2) x_0^δ has the same distribution as x_0 (noted μ)
- (x3) $\{x_k^\delta\}$ has strong order of convergence 0.5
- (x4) the transition kernel $q(s, ds')$ of $\{x_k^\delta\}$ is Feller

Conditions (x1-x2) are necessary for the statement of the algorithm. Conditions (x3-x4) do not appear in the algorithm, but are required in showing its convergence. The meaning of (x3) is

$$\mathbb{E}[d_S^2(x_t, x_k^\delta)] = O(\delta), \quad (35)$$

for k such that $t \in [k\delta, (k+1)\delta[$, where $d_S(\cdot, \cdot)$ is the metric in the state space S . The Feller condition (x4) is that for all $\varphi \in C_b(S)$, if $q\varphi$ is the function

$$q\varphi(s) = \int_S q(s, ds')\varphi(s'),$$

then $q\varphi \in C_b(S)$. Recall that $q(s, ds')$ is the transition probability for going from s to the “neighborhood” ds' .

If x is a diffusion in \mathbb{R}^l with bounded Lipschitz coefficients and x_0 square integrable, then the first order Euler scheme gives an approximation x^δ verifying (x1-x4). In fact, the assumptions on the coefficients of x can be weakened to those of [27], (Exercise 9.6.3 on page 326). Of course, in this case the metric $d_S(\cdot, \cdot)$ is just the Euclidean metric of \mathbb{R}^l .

If x is a diffusion in a compact Riemannian manifold \mathcal{N} , then conditions (x1-x4) are verified when x is approximated using successive geodesic steps.

In a purely heuristic way, consider now a discrete time filtering problem where the unknown signal is the Markov chain $\{x_k^\delta\}$ and the observations are the sequence of increments $\{\Delta Y_k\}$. Here, ΔY_k is obtained from (6) where the connection ∇ is that of (22).

Based on Proposition 3 and on (35), consider the likelihood for the observations ΔY_k to be

$$l(x_k^\delta, \Delta Y_k) = \exp \left(\langle H_k^\delta, \Delta Y_k \rangle - \frac{\delta}{2} \|H_k^\delta\|^2 \right), \quad (36)$$

where $H_k^\delta = H(x_k^\delta, Y_{k\delta})$. Applying the usual Bayes formula for a discrete time filtering problem obtained in this way, the resulting conditional distribution given the first $M+1$ observations is

$$\pi_M^\delta(\varphi) = \frac{\rho_M^\delta(\varphi)}{\rho_M^\delta(1)}, \quad \rho_M^\delta(\varphi) = \mathbb{E}[\varphi(\tilde{x}_M^\delta)L_M^\delta \mid \mathcal{Y}_\infty], \quad (37)$$

for $\varphi \in C_b(S)$, where $\{\tilde{x}_k^\delta\}$ is a copy of $\{x_k^\delta\}$ independent from Y and

$$L_M^\delta = \prod_{k=0}^M l(\tilde{x}_k^\delta, \Delta Y_k). \quad (38)$$

This is really just Bayes formula in discrete time. From (37),

$$\rho_M^\delta(\varphi) = \rho_{M-1}^\delta(\varphi_{M-}), \quad (39)$$

where

$$\varphi_{M-}(s) = \int_S l(s', \Delta Y_M) \varphi(s') q(s, ds'). \quad (40)$$

This is the usual ‘‘prediction-measurement update’’ formula, where prediction is according to the transition kernel $q(s, ds')$ and measurement update is based on the likelihood $l(s', \Delta Y_M)$.

It is important to note that π_M^δ is a true probability measure on S conditionally on the increments $\Delta Y_0, \dots, \Delta Y_M$. In particular, it makes sense to speak of sampling from π_M^δ once these increments are given. This is the aim of the proposed particle filtering algorithm.

The algorithm has as its input the sequence of increments $\{\Delta Y_k\}$ and is parameterised by the number of particles N . The output after $M+1$ observations is a family of N particles $\hat{x}_M^i, \dots, \hat{x}_0^i \in S$ which give a Monte Carlo approximation $\hat{\pi}_M^\delta$ of π_M^δ . This is

$$\hat{\pi}_M^\delta(\varphi) = (1/N) \sum_{i=1}^N \varphi(\hat{x}_M^i). \quad (41)$$

The implementation is the following,

- when ΔY_0 is available
 - 1 generate *i.i.d.* particles $\tilde{x}_0^1, \dots, \tilde{x}_0^N \sim \mu$
 - 2 compute normalised weights, $w_0^i \propto l(\tilde{x}_0^i, \Delta Y_0)$
 - 3 generate $(n_0^1, \dots, n_0^N) \sim \text{multinomial}(w_0^1, \dots, w_0^N)$ and replace \tilde{x}_0^i by n_0^i particles with same value
 - 4 relabel new particles $\hat{x}_0^1, \dots, \hat{x}_0^N$; set $w_0^i = 1/N$
- when ΔY_k is available ($k \geq 1$)
 - 1 generate particles $\tilde{x}_k^i \sim q(\hat{x}_{k-1}^i, ds)$
 - 2 compute normalised weights, $w_k^i \propto w_{k-1}^i l(\tilde{x}_k^i, \Delta Y_k)$
 - 3 generate $(n_k^1, \dots, n_k^N) \sim \text{multinomial}(w_k^1, \dots, w_k^N)$ and replace \tilde{x}_k^i by n_k^i particles with same value
 - 4 relabel new particles $\hat{x}_k^1, \dots, \hat{x}_k^N$; set $w_k^i = 1/N$

These steps are very much the same as in the classical bootstrap filter. The geometry of the observation process Y only appears through the use of a connector map I , which provides the increments ΔY_k .

The bootstrap filter is the simplest, but the least robust, particle filtering algorithm and there are many improvements upon it known in the literature. These can all be implemented in an equally direct way once the ΔY_k have been obtained.

To summarise, a connector map I leads to increments ΔY_k with approximate likelihood given by (36). Once this situation is accepted, it can be replaced into

any suitable algorithm. It is possible to say that the connector map serves to linearise the observation process Y locally, i.e., in the neighborhood of each sample $Y_{k\delta}$ as in (6).

4.3. Convergence

The convergence of the particle filtering algorithm proposed in the previous paragraph is here given in Proposition 5. Precisely, what is meant by this is the convergence of $\hat{\pi}_M^\delta$ to π_t as the step size δ goes to zero and the number of particles N goes to infinity, when M is taken of the order of t/δ . That is, the order of the number of increments ΔY_k which can be constructed from (6) up to time t .

An equally important question, not dealt with here, is the convergence as t goes to infinity of the conditional distribution π_t or of its Monte Carlo approximation obtained from the particle filtering algorithm. This is related to the eventual ergodicity or mixing properties of the unknown signal x .

Proposition 5 is based on two lemmas, which are first given. Lemma 1 states the convergence of π_M^δ to π_t . Lemma 2 states the convergence of $\hat{\pi}_M^\delta$ to π_M^δ for any given value of δ . This latter limit is not shown to be uniform in δ . Thus, in their form stated below, Lemmas 1 and 2 cannot be used to justify an approach where δ is taken proportional to $1/N$ (or some other function of N which converges to zero as N goes to infinity) and $\hat{\pi}_M^\delta$ is then computed for a large value of N .

It is possible to say that Proposition 5 only provides, in a satisfactory way, the consistency of the approximation $\hat{\pi}_M^\delta$. That is, the fact that it is possible to choose δ and N to make this approximation arbitrarily close to π_t .

For Lemma 1, the two following conditions are required.

- (H1) $\|H\|$ is bounded (as a function $S \times \mathcal{M} \rightarrow \mathbb{R}_+$)
- (H2) $\|H(s, p) - H(s', p)\| \leq C d_S(s, s')$ for all $p \in \mathcal{M}$, where the constant C does not depend on p

These are quite strong restrictions, however they allow for straightforward proofs. Replacing them by weaker conditions may lead to convergence in the square mean being replaced by convergence in probability, (see the statement of the lemma), but it would not introduce any more fundamental changes.

Assumption (H2) means that H is a Lipschitz continuous application from S to the space of continuous vector fields on \mathcal{M} , this latter space being equipped with its topology of uniform convergence with respect to the Riemannian metric.

LEMMA 1. Assume \mathcal{M} is compact and conditions (x1-x3), (H1-H2) hold. Let $\varphi \in C_b(S)$ be Lipschitz continuous, such that $\varphi \in D(A)$ and $A\varphi \in C_b(S)$, (recall the notation of (1)). Then, if M is the integer part of t/δ ,

$$\mathbb{E}|\pi_M^\delta(\varphi) - \pi_t(\varphi)|^2 = O(\delta).$$

PROOF The proof is identical to the one in [15], (Proposition 2.1 on page 292). Here, the main steps are indicated.

Note first that it is possible to consider $M\delta = t$. Indeed, the assumption that $\varphi \in D(A)$ and $A\varphi \in C_b(S)$ guarantees

$$\mathbb{E}|\pi_t(\varphi) - \pi_{M\delta}(\varphi)|^2 = O(\delta).$$

This is since π_t verifies the filtering equation given in [14], which has bounded coefficients under condition **(H1)**.

The proof follows from the identity

$$\begin{aligned} \pi_M^\delta(\varphi) - \pi_t(\varphi) &= \mathbb{E}[\varphi(x_M^\delta) - \varphi(x_t) \mid \mathcal{Y}_\infty] \\ &\quad + \mathbb{E}[(A-1)(\varphi(x_M^\delta) - \pi_M^\delta(\varphi)) \mid \mathcal{Y}_\infty], \end{aligned} \quad (42)$$

where $A = L_M^\delta/L_t$. This can be proved by using the Kallianpur-Striebel formula (33) to express the conditional expectations and then by developing the products.

Since φ is Lipschitz continuous, the square mean of the first term is bounded by

$$\mathbb{E}|\varphi(x_M^\delta) - \varphi(x_t)|^2 \leq C\mathbb{E}[d_S^2(x_t, x_k^\delta)] = O(\delta), \quad (43)$$

where the second inequality uses **(x3)**.

The bound for the second term is more delicate. The idea is to note

$$L_M^\delta = \exp\left(\sum_{k\delta < t} \langle H_k^\delta, \Delta Y_k \rangle - (\delta/2) \sum_{k\delta < t} \|H_k^\delta\|^2\right) \quad (44)$$

and compare this to L_t , given by (34), using Proposition 2. The detailed development requires condition **(H2)** and gives

$$\mathbb{E}|(A-1)(\varphi(x_M^\delta) - \pi_M^\delta(\varphi))|^2 \leq 4\|\varphi\|^2\mathbb{E}(A-1)^2 = O(\delta), \quad (45)$$

where $\|\varphi\|$ is the supremum of $|\varphi(s)|$ over $s \in S$.

The proof is completed by applying Minkowski's inequality, (43) and (45).

LEMMA 2. Assume conditions **(x1-x2)**, **(x4)** and **(H1)** hold. For all $\varphi \in C_b(S)$ and any values of δ and M

$$\mathbb{E}|\hat{\pi}_M^\delta(\varphi) - \pi_M^\delta(\varphi)|^2 \rightarrow 0 \quad \text{as } N \uparrow \infty.$$

PROOF It is clear from (41) that $\hat{\pi}_M^\delta(\varphi)$ is bounded. Note, moreover, that $\pi_M^\delta(\varphi)$ is square integrable. By dominated convergence, in order to show the Lemma, it is enough to show that $\hat{\pi}_M^\delta(\varphi)$ converges to $\pi_M^\delta(\varphi)$ almost surely.

Almost sure convergence is a direct application of a general theorem from [21], (Theorem 1, on page 742). This requires that the transition kernel be Feller, which is condition **(x4)**, and that the likelihood function is continuous, bounded and strictly positive.

Here, the likelihood function is the one corresponding to (36). That is

$$l(s, \Delta Y_{k\delta}) = \exp\left(\langle H(s, Y_{k\delta}), \Delta Y_k \rangle - \frac{\delta}{2} \|H(s, Y_{k\delta})\|^2\right).$$

That this is continuous, as a function of s , and strictly positive is immediate. Boundedness follows since the second term under the exponential is negative and

$$\langle H(s, Y_{k\delta}), \Delta Y_k \rangle \leq r \|H(s, Y_{k\delta})\|.$$

By Cauchy-Schwarz inequality, where r is as in the definition (28) of the mapping l .

Finally, it is possible to conclude by condition **(H1)**.

Now, it is possible to state Proposition 5 which combines Lemmas 1 and 2.

PROPOSITION 5. Assume that \mathcal{M} is compact and that conditions **(x1-x4)**, **(H1-H2)** hold. Let φ be as in Lemma 1. Then, if M is the integer part of t/δ

$$\lim_{\delta \downarrow 0} \lim_{N \uparrow \infty} \mathbb{E}|\hat{\pi}_M^\delta(\varphi) - \pi_t(\varphi)|^2 = 0.$$

PROOF The conditions of Lemmas 1 and 2 are united. By Lemma 1

$$\lim_{\delta \downarrow 0} \mathbb{E}|\pi_M^\delta(\varphi) - \pi_t(\varphi)|^2 = 0,$$

where the expression under the limit does not depend on N . By Lemma 2

$$\lim_{N \uparrow \infty} \mathbb{E}|\hat{\pi}_M^\delta(\varphi) - \pi_M^\delta(\varphi)|^2 = 0,$$

for any values of δ and M . Thus, the proposition follows by adding together these two limits and applying Minkowski's inequality.

Proposition 5 does not explicitly provide the rate at which $\hat{\pi}_M^\delta$ converges to π_t . Obtaining this rate of convergence requires a deeper analysis than provided here (in Lemmas 1 and 2). This can be carried out on the basis of the corresponding analysis for a classical filtering problem [22], (Chapter 9), but still has not been pursued in the literature. Clearly, this situation represents an important drawback for practical application.

The information lacking from Proposition 5, i.e., the rate of convergence, can eventually be recovered on a case by case basis. For the convergence of π_M^δ to π_t , Lemma 1 can be used to obtain the precise rate of convergence, (by expressing the various constants appearing in the proof). Then, for any required values of δ and M , the problem is to find a number of particles N sufficiently large for a given precision. It is well known that $\hat{\pi}_M^\delta$ converges to π_M^δ at a rate of the order of $1/N$ but where the involved constants depend on the observations $\{\Delta Y_k\}$. For an individual realisation of the observations, this can be made precise either through additional calculation or through computer experiments, based on the specific model being studied, (that is, on a particular instance of equations (1) and (2)).

Regarding the convergence of π_M^δ to π_t , stated in Lemma 1, it is useful to make the following remark. The order of convergence in condition (x3) can often be improved. For example [27], if x is a diffusion in \mathbb{R}^l , using the Milstein approximation instead of a first order Euler scheme gives strong order 1 instead of 0.5. However, this is not enough to improve the overall order of magnitude given in the lemma. As can be seen from the proof, this order of magnitude involves both (43) and (45). While improving the order of convergence in condition (x3) will accordingly improve the order of magnitude in (43), it has no similar effect on (45). The computer experiments presented in the following section show that the proposed particle filtering algorithm, when implemented with adequate values of δ and N , performs in a sensibly satisfactory way for the chosen examples.

5. EXAMPLES: OBSERVATIONS IN $SO(3)$ AND S^2

The stochastic filtering problem stated in Section 2 is of a quite general form. By specifying the state space S , the manifold \mathcal{M} and the various objects appearing in (1) and (2), it is possible to recover a wide range of problems. As already mentioned, these include the classical ones defined by an additive white noise model.

Several natural and important observation models also arise as special cases of the problem of Section 2. Of particular interest in engineering is the case where the manifold \mathcal{M} is a classical matrix manifold, (that is, a matrix Lie group or a related symmetric space), and the observation process Y , conditionally on the unknown signal x , is an invariant diffusion.

In signal and image processing, classical matrix manifolds such as Stiefel and Grassmann manifolds are of widespread use and importance. Therefore, it is natural to consider observation models compatible with their underlying structure. Roughly speaking, these are exactly the ones involving invariant diffusions.

To explain and optimise the implementation and performance of the particle filtering algorithm of Section 4 in the special case of classical matrix manifolds should be one of the main objectives for the present work. Mainly for a reason of space, the current section has a more modest scope dealing only with two individual examples.

In 5.1, the problem is considered where the observation process Y is an invariant diffusion in the special orthogonal group $SO(3)$. In 5.2 a similar problem is studied but where Y has its values in the unit sphere S^2 . These two examples serve as case studies. They show how, when faced with a problem of the kind given by (1) and (2), to carry out the various steps leading to a successful implementation of the particle filtering algorithm of Section 4. These include, at least, specification of the metric (21) and the connection (22), choice of the connector map I to be replaced in (6) and choice of the approximation $\{x_k^\delta\}$ of the signal x .

It should be noted none of these steps is known a priori, just by knowing the manifold \mathcal{M} . They are all carried out based on (1) and (2). In particular, the geometric structure given by the metric and connection of (21) and (22) is adapted to the observation model.

5.1. Observations in $SO(3)$

The first example considers the case where the observation process Y , conditionally on the unknown signal x , is a left invariant diffusion in the special orthogonal group $SO(3)$. This is the set of 3×3 real matrices g which are orthogonal and have unit determinant. That is,

$$g^{-1} = g^\dagger, \quad \det(g) = 1, \quad (46)$$

where \dagger denotes matrix transpose. As a subset of the vector space $\mathbb{R}^{3 \times 3}$ (space of 3×3 real matrices), $SO(3)$ is connected and compact. Furthermore, it is closed under matrix multiplication and inversion. Thus, $SO(3)$ is a compact connected Lie group [31]. The 3×3 identity matrix is denoted e ; clearly $e \in SO(3)$.

Let $\mathfrak{so}(3)$ denote the subspace of $\mathbb{R}^{3 \times 3}$, consisting of all antisymmetric matrices. That is, matrices $\sigma \in \mathbb{R}^{3 \times 3}$ such that $\sigma + \sigma^\dagger = 0$. It can be shown that, for $\sigma \in \mathbb{R}^{3 \times 3}$ and $\gamma(t) = \exp(t\sigma)$ where $t \in \mathbb{R}$, $\gamma(t) \in SO(3)$ for all $t \in \mathbb{R}$ if and only if $\sigma \in \mathfrak{so}(3)$. By definition, this means that $\mathfrak{so}(3)$ is the Lie algebra of the Lie group $SO(3)$. For this and other facts on compact Lie groups used in the following, see [31].

It is not surprising that, being defined by the differentiable constraints (46), $SO(3)$ is a differentiable manifold. Moreover, $\mathfrak{so}(3)$ can be identified with the tangent space $T_e SO(3)$.

The special orthogonal group $SO(3)$ is quite important in many applications. A matrix $g \in SO(3)$ defines an orientation preserving rotation in \mathbb{R}^3 . Furthermore, $SO(3)$ is often thought of as the typical example of a nontrivial compact connected Lie group. The presentation in the rest of this paragraph generalises to any compact connected Lie group with very minor changes.

In terms of the general filtering problem of Section 2, the example considered here makes no restriction on the signal model (that is, on (1)). The observation model (2) is specified in the following way.

The sensor function H and the vector fields Σ_r are defined in terms of left invariant vector fields on $SO(3)$. For each $\sigma \in \mathfrak{so}(3)$, there is a corresponding vector field Σ on $SO(3)$ where,

$$\Sigma(g) = g\sigma, \quad g \in SO(3). \quad (47)$$

This means that for each $g \in SO(3)$, there exists some differentiable curve $\gamma_g :]-\epsilon, \epsilon[\rightarrow SO(3)$ such that $\gamma_g(0) = g$ and $\dot{\gamma}_g(0) = \Sigma(g)$. In fact, it is quite straightforward to obtain such a curve. First, let γ_e be the curve $\gamma_e(t) = \exp(t\sigma)$. This is defined for all real t and has its values in $SO(3)$ as mentioned above. By elementary properties of the matrix exponential, $\gamma_e(0) = e$ and $\dot{\gamma}_e(0) = \sigma$. To obtain the curve γ_g for any $g \in SO(3)$, it is

enough to put $\gamma_g(t) = g\gamma_e(t)$. The meaning of the name “left invariant vector field” is precisely that γ_g can be obtained from γ_e by left multiplication.

Note that σ and Σ define each other uniquely. Any left invariant vector field Σ is of the form (47) where $\sigma = \Sigma(e)$.

It is required that for all $s \in S$, the application $g \mapsto H(s, g)$ is a left invariant vector field. It follows that there exists a mapping $h : S \rightarrow \mathfrak{so}(3)$ such that for $g \in SO(3)$

$$H(s, g) = gh(s). \quad (48)$$

The vector fields Σ_r are also taken to be left invariant vector fields. Let $\sigma_1, \sigma_2, \sigma_3$ be any basis of $\mathfrak{so}(3)$. This basis being fixed, let $\Sigma_1, \Sigma_2, \Sigma_3$ be the corresponding left invariant vector fields as in (47). In other words,

$$\Sigma_r(g) = g\sigma_r. \quad (49)$$

Note that the number of vector fields Σ_r is here equal to 3, the dimension of $SO(3)$ (which is also the dimension of $\mathfrak{so}(3)$). This was not assumed in the general problem of Section 2.

Now that H and $\Sigma_1, \Sigma_2, \Sigma_3$ have been defined, it is possible to write down the observation model (2). From (48) and (49), a simple rearrangement shows

$$dY_t = Y_t \circ \{h(x_t)dt + d\hat{B}_t\}. \quad (50)$$

Here \hat{B} is a process with values in $\mathfrak{so}(3)$ defined from a Brownian motion B in \mathbb{R}^d , which is independent from x , by

$$\hat{B}_t = B_t^1 \sigma_1 + B_t^2 \sigma_2 + B_t^3 \sigma_3. \quad (51)$$

The reader should be immediately aware equation (50) is just a linear matrix stochastic differential equation. It can be understood for each matrix element after writing down the usual formula for matrix product. A process Y satisfying an equation of this form, when x is assumed known, is called a left invariant diffusion in $SO(3)$. An alternative way of writing equation (50) involves the vector fields $\Sigma_1, \Sigma_2, \Sigma_3$. This is

$$dY_t = \Sigma_r(Y) \circ \{h^r(x_t)dt + dB_t^r\}, \quad (52)$$

where $h = h^r \sigma_r$ for some functions $h^r : S \rightarrow \mathbb{R}$. note the differentiability conditions of Section 2 are verified. Indeed, as functions of g , $H(s, g)$ and $\Sigma_r(g)$ are linear and therefore smooth (C^∞).

In [13], Lo considered the observation model (50) by itself (but for a general matrix Lie group, not just $SO(3)$). It was proposed that this can be reduced to a classical, additive white noise model by the following simple transformation

$$y_t = \int_0^t Y_s^{-1} \circ dY_s. \quad (53)$$

By the chain rule of Stratonovich calculus, it is clear that

$$dy_t = h(x_t)dt + d\hat{B}_t. \quad (54)$$

This result is strikingly similar to (25) in Proposition 1. It is now shown that, in effect, it is a special case of that proposition.

Following the approach of 3.2, the Le Jan-Watanabe connection is now introduced. Note first the condition of ellipticity is here verified. In fact, a sharper result holds since for each $g \in SO(3)$ the vectors $\Sigma_1(g), \Sigma_2(g), \Sigma_3(g)$ form a basis of $T_g SO(3)$.

Definition (21) amounts to introducing a Riemannian metric on $SO(3)$ such that this basis is orthonormal,

$$\langle \Sigma_r(g), \Sigma_v(g) \rangle = \delta_{rv}. \quad (55)$$

Using this metric, the Le Jan-Watanabe connection is defined by (22) which gives

$$\nabla_{\Sigma_r} \Sigma_v = 0. \quad (56)$$

This immediately implies (24). In fact, for any $g \in SO(3)$ and tangent vector $K \in T_g SO(3)$, it follows by linearity that $\nabla_K \Sigma_r = 0$. Here, one says the vector fields Σ_r form a global parallel frame.

For the following, it is important to note the metric (55) is left invariant. If $g \in SO(3)$ and $E, K \in T_g SO(3)$,

$$\langle E, K \rangle = \langle g^{-1}E, g^{-1}K \rangle. \quad (57)$$

That is, the left hand side is computed in the tangent space $T_g SO(3)$ and the right hand side in $T_e SO(3)$, which is $\mathfrak{so}(3)$. This can be shown by putting $\eta = g^{-1}E \in \mathfrak{so}(3)$ and $\kappa = g^{-1}K \in \mathfrak{so}(3)$ and considering the corresponding left invariant vector fields as in (47). It is then a straightforward result of (55).

Due to (57), the Riemannian metric (55) is completely determined by the basis $\sigma_1, \sigma_2, \sigma_3$. This can be chosen in a completely arbitrary way. It is clear that the following matrices form a basis of $\mathfrak{so}(3)$

$$\omega_1 = \begin{pmatrix} & & \\ & -1 & \\ 1 & & \end{pmatrix} \omega_2 = \begin{pmatrix} & 1 & \\ & & \\ -1 & & \end{pmatrix} \omega_3 = \begin{pmatrix} & -1 & \\ 1 & & \end{pmatrix}.$$

The general form of the basis $\sigma_1, \sigma_2, \sigma_3$ is therefore

$$\sigma_r = b_{vr} \omega_v, \quad (58)$$

where b is an invertible matrix.

It is possible that the Riemannian metric (55) will not be biinvariant (i.e., both left and right invariant). This is the case if and only if the matrix b is orthogonal. In practice, there is no reason why this should be the case. In rigid body mechanics [32], an orthogonal matrix b may be chosen when studying the motion of a spherically symmetric body. In general, b is given by the inertia matrix of the body, reflecting its shape and mass distribution. This can be far from spherical symmetry (consider an airplane). The metric (55) is introduced based on the observation model (52). In other words, it is adapted to the observation model.

Furthermore, the connection defined by (56) is known as the Cartan-Schouten (–)-connection. Since it

has nonzero torsion, it is not the Levi-Civita connection of any Riemannian metric (in particular, of the metric (55)). When the matrix b is orthogonal, the Levi-Civita connection of the metric (55) is known as the Cartan-Schouten (0)-connection (see [23, 33]).

In order to apply Proposition 1, it is necessary to compute a parallel orthonormal frame (E^i) . This turns out to be especially simple. Applying (10) and (56),

$$\nabla_{\circ dY} \Sigma_r(Y_t) = 0.$$

Therefore, by (9) and (12),

$$d\langle E^i, \Sigma_r \rangle = 0.$$

Since both (E^i) and (Σ_r) are orthonormal families, this implies the existence of an orthogonal matrix $a = (a^{ir})$ such that

$$E_t^i = a^{ir} \Sigma_r(Y_t).$$

Moreover, by a choice of initial condition, it is possible to take $a = e$ identity matrix.

Now, (13) gives

$$y_t^i = a^{ir} \int_0^t \langle \Sigma_r(Y_s), \circ dY_s \rangle. \quad (59)$$

To evaluate this Stratonovich integral, it is possible to apply (18). Note that this states the Stratonovich integral is the same as a classical Stratonovich integral. Applying this prescription is easier than changing the current notation to that of (18) and then changing back. In short, it follows from (49) and (57),

$$y_t^i = \langle a^{ir} \sigma_r, \int_0^t Y_s^{-1} \circ dY_s \rangle.$$

This shows that (54) is the same as (59) up to a change of basis. Thus (54) is indeed a special case of (25), Proposition 1. It is interesting to note the Brownian motion β appearing in Proposition 1 turns out, in the present case, to be the same as the original Brownian motion B .

The transformation (54) taking Y into y can be extended to any Lie group, not just compact and matrix. It is known as the Lie group stochastic logarithm and was generalised extensively by Estrade [34]. That it coincides with a stochastic antidevelopment was first pointed out in [35].

Going on with the programme of applying Section 4 in the current example, consider now the construction of connector maps. Recall this was described in (6) and (28). Moreover, it can be implemented using any mapping I verifying conditions **(I1-I3)**.

As in 3.3, the starting point is the notion of geodesic. Finding the geodesics of the connection (56) is straightforward. Indeed, the definition of this connection suggests geodesics are precisely the flow lines of left invariant vector fields. These are the curves of the form γ_g discussed after (47). This is easily checked to be the case by replacing (56) in the geodesic equation (27).

Thus, the mapping $\exp : TSO(3) \rightarrow SO(3)$ is related in a simple way to the matrix exponential. For $g \in SO(3)$ and $K \in T_g SO(3)$,

$$\exp_g(K) = g \exp(g^{-1}K). \quad (60)$$

Indeed, letting $\sigma = g^{-1}K$, it follows from (47) that $\Sigma(g) = K$. Then, it is clear the right hand side is $\gamma_g(1)$.

To specify the mapping I of (28) to the current context, an estimate of the radius of injectivity of the connection (56) would be needed. In (28), this is the role played by r which is needed in constructing the cut off function ϕ . With the considered geometry, the group $SO(3)$ is a manifold of constant (strictly) positive curvature. The radius of injectivity is then known from Riemannian geometry [23].

Here, a less elegant but simpler approach is taken. For each $g_1 \in SO(3)$ and $K \in T_{g_1} SO(3)$, let $g_2 = \exp_{g_1}(K)$. If $|g_1^{-1}g_2 - e| < 1$, where $|\cdot|$ stands for the Euclidean matrix norm, then $g_1^{-1}g_2$ has a unique matrix logarithm. Denote this $\log(g_1^{-1}g_2)$. Then by (60),

$$K = g_1 \log(g_1^{-1}g_2).$$

For any $g_1, g_2 \in SO(3)$ verifying $|g_1^{-1}g_2 - e| < 1$, define

$$\log_{g_1}(g_2) = g_1 \log(g_1^{-1}g_2). \quad (61)$$

According to (28), when it is possible, $I(g_1, g_2)$ should coincide with $\log_{g_1}(g_2)$. To take into account couples g_1, g_2 for which this expression is not well defined, let $\phi : SO(3) \times SO(3) \rightarrow [0, 1]$ be a C^∞ function such that $\phi(g_1, g_2) = 0$ if $|g_1^{-1}g_2 - e| \geq 1$ and $\phi(p, q) = 1$ if $|g_1^{-1}g_2 - e| \leq 1 - \lambda$. Here, $0 < \lambda < 1$ is fixed. Now, I can be defined as

$$I(g_1, g_2) = \phi(g_1, g_2) \log_{g_1}(g_2). \quad (62)$$

Computing a matrix logarithm, even for 3×3 matrices, is a relatively involved task. It is possible to propose an alternative mapping I , which does not involve a matrix logarithm. Recall the first order Taylor expansion of the matrix logarithm at e ,

$$\log(g_1^{-1}g_2) = g_1^{-1}g_2 - e + O(|g_1^{-1}g_2 - e|^2). \quad (63)$$

Let Π be a linear projection from $\mathbb{R}^{3 \times 3}$ to $\mathfrak{so}(3)$. For instance, $\Pi(\sigma) = (1/2)[\sigma + \sigma^\dagger]$ associates to the matrix σ its antisymmetric part. Instead of (62), it is possible to use

$$I(g_1, g_2) = g_1 \Pi(g_1^{-1}g_2 - e). \quad (64)$$

Of course, it is an abuse of notation to call both mappings (62) and (64) by the same name I . Still, this is done since they serve the same purpose. Roughly, the difference between (62) and (64) is that in the latter expression the matrix logarithm is replaced by the first term in its Taylor expansion. Moreover, in this same expression (64), there is no need for a cut off factor since all operations are well defined for $g_1, g_2 \in SO(3)$.

Recall that any chosen mapping I is here required to verify conditions **(I1-I3)**. It is straightforward to show this is the case for the mapping (64).

That condition **(I1)** is verified, follows from smoothness of matrix inversion and multiplication; in addition to the linear operation Π . To verify conditions **(I1-I2)**, note that in the same notation used to state these conditions,

$$dI(g_1, g_2)(\Sigma_r) = g_1 \Pi(g_1^{-1} \Sigma_r(g_2)),$$

and, using (56) and the definition of the Hessian,

$$\begin{aligned} \nabla^2 I(g_1, g_2)(\Sigma_r, \Sigma_v) &= \Sigma_r \Sigma_v I(g_1, g_2) - \nabla_{\Sigma_r} \Sigma_v I(g_1, g_2) \\ &= 0. \end{aligned}$$

These simply follow from the fact that (64) is linear in g_2 . The latter formula immediately gives condition **(I3)**, since $\Sigma_1(g_1), \Sigma_2(g_1), \Sigma_3(g_1)$ is a basis of $T_{g_1} SO(3)$. For condition **(I1)**, note by the same argument

$$dI(g_1, g_1)(K) = g_1 \Pi(g_1^{-1} K),$$

for all $K \in T_{g_1} SO(3)$. However, $\kappa = g_1^{-1} K \in \mathfrak{so}(3)$, so that $\Pi(\kappa) = \kappa$. Then,

$$dI(g_1, g_1)(K) = g_1 \kappa = K,$$

which is condition **(I1)**. The mapping I of (64) is especially easy to compute. Thanks to (46), the matrix inverse is the same as the transpose.

All that is needed in order to apply the particle filtering algorithm of Section 4 to the current example is the mapping I . At each step of the algorithm, as described in 4.2, instructions 1,3 and 4 do not involve the observation process Y . This only appears through the increments ΔY_k in instruction 2. For this instruction, each ΔY_k is computed using I as in (6) and replaced in (36) in order to find particle weights.

To illustrate the above discussion, a computer experiment is now presented. For the signal model (1), the experiment considers $S = \mathbb{R}^3$. The unknown signal x is taken to be an Ornstein-Uhlenbeck process,

$$dx_t = -\nu x_t dt + dv_t, \quad x_0 = (0, 0, 0), \quad (65)$$

where $\nu > 0$ and v is a Brownian motion in \mathbb{R}^3 with variance parameter $\sigma^2 = 0.5$. As discussed in 4.2, the first step is to choose an approximation $\{x_k^\delta\}$ of x . In the following, only one value of δ is considered, $\delta = 0.1$. The x_k^δ are constructed using a first order Euler scheme, which verifies conditions **(x1-x4)**.

The function $h: \mathbb{R}^3 \rightarrow \mathfrak{so}(3)$ is taken to be a linear isomorphism, mapping the canonical basis of \mathbb{R}^3 to the basis $\omega_1, \omega_2, \omega_3$ of $\mathfrak{so}(3)$. In other words,

$$h(s) = s^1 \omega_1 + s^2 \omega_2 + s^3 \omega_3, \quad (66)$$

for each $s = (s^1, s^2, s^3)$ in \mathbb{R}^3 . It should be noted that, with this choice for h , the sensor function H of (48) verifies condition **(H2)** but does not verify condition **(H1)**.

It was stated in 4.3 that conditions **(H1-H2)** are not essential for the overall behaviour of the particle filtering algorithm. The missing condition here is condition **(H1)**, which requires the sensor function H to be bounded. It can be shown by (48) and (57) that

$$\|H(s, g)\|^2 = \|h(s)\|^2 = |s|^2,$$

where, as before, $|\cdot|$ denotes the Euclidean norm. Thus, condition **(H1)** is equivalent to the condition that the function h be bounded. It is clear this does not hold for the function h in (66). However, it should be noted that the unknown signal x of (65) is a normal process. That is, the distribution of x_t is normal for each $t \geq 0$. Thus, the distribution of $h(x_t)$ has exponentially decreasing tails and finite moments of all orders. This compensates for condition **(H1)** being dropped, since $\|h(x_t)\|^2 = |x_t|^2$ and this has an exponentially small probability of being large.

In order to simulate a trajectory of the observation process Y , a slight modification of the McShane approximation (4) was used. Applied to (50), the McShane approximation gives an approximating process Y^δ which satisfies a linear ordinary differential equation

$$\dot{Y}_t^\delta = Y_t^\delta \{h(x_t) + \Delta \hat{B}_k\}, \quad (67)$$

on each interval $[k\delta, (k+1)\delta]$; for $k \in \mathbb{N}$ and where, as in (4), $\Delta \hat{B}_k = \delta^{-1}(\hat{B}_{(k+1)\delta} - \hat{B}_{k\delta})$.

Equation (67) is a linear ordinary differential equation with time dependent coefficients. This is due to the presence of $h(x_t)$ which contains the unknown signal. Without changing the convergence rate (5), it is possible to consider another approximating process \bar{Y}^δ which satisfies an equation with piecewise constant coefficients. This has the advantage of having a straightforward analytical solution. In effect, if \bar{Y}^δ satisfies the equation

$$\dot{\bar{Y}}_t^\delta = \bar{Y}_t^\delta \{h(x_{k\delta}) + \Delta \hat{B}_k\}, \quad (68)$$

on each interval $[k\delta, (k+1)\delta]$ for $k \in \mathbb{N}$, then

$$\bar{Y}_t^\delta = \bar{Y}_{k\delta}^\delta \exp[(t - k\delta)(h(x_{k\delta}) + \Delta \hat{B}_k)], \quad (69)$$

on the interval $[k\delta, (k+1)\delta]$. Equation (68) is the same as (67), but with $h(x_t)$ replaced by $h(x_{k\delta})$.

Formula (69) is the one used in simulating a trajectory of Y . The experiment was carried out with $b = e$ in (58). In this case (69) admits a simple interpretation since, for $s \in \mathbb{R}^3$, if $\gamma(t) = \exp[th(s)]$ then γ represents a uniform rotation with angular velocity s . Thus, \bar{Y}^δ consists in a sequence of successive uniform rotations where the angular velocity is $x_{k\delta} + \Delta B_k$ on the interval $[k\delta, (k+1)\delta]$.

Once a trajectory of Y has been simulated, the particle filtering algorithm of Section 4 can be applied immediately. The algorithm has as its input the sequence of increments $\{\Delta Y_k\}$. Here, these are obtained using the mapping I of (64). With the choice of projection Π

discussed above, $\Pi(\sigma)$ equal to the antisymmetric part of σ , this mapping becomes

$$I(g_1, g_2) = g_1[(1/2)(g_1^\dagger g_2 - g_2^\dagger g_1)], \quad (70)$$

where (46) was used in order to avoid matrix inversion. In the particle filtering algorithm, the mapping I is only used in instruction 2 which requires evaluating the likelihood (36) (computation of normalised weights). Here, it is important to note the factor g_1 before the bracket in the right hand side of (70) can be overlooked. This is seen by replacing the invariance property (57) and definition (70) in (36). As a result, an unnecessary matrix multiplication can be avoided.

Figure 1 below is now used to illustrate the performance of the particle filtering algorithm. With the values of δ and σ^2 mentioned above, this figure shows the estimation error $|\hat{x}_t - x_t|$ for $t \in [1, T]$ where $T = 10$. Here, in the notation of (41), \hat{x}_t is the estimate

$$\hat{x}_t = (1/N) \sum_{i=1}^N \hat{x}_M^i, \quad (71)$$

where M is the integer part of t/δ . This is the expectation of the Monte Carlo approximation $\hat{\pi}_M^\delta$ obtained from the particle filtering algorithm. Since $\hat{\pi}_M^\delta$ converges to π_t , (as described in Proposition 5 of 4.3), \hat{x}_t should also converge to the expectation of π_t , say x_t^* . That x_t^* exists follows from the fact that x_t is normal and thus square integrable, since π_t is defined in (32) as the conditional distribution of x_t given past observations.

The notation \hat{x}_t obscures dependence on δ and N . However, the values of δ and N essentially control the estimation error when (65) is determined; i.e., when ν is given. Figure 1 considers $\nu = 1$ and $\nu = 0.5$. The estimation error in the present case can be understood as combining a bias and a variance. By its very definition, x_t^* is an unbiased estimator of x_t , whereas \hat{x}_t is constructed as an approximation of x_t^* . The bias is then the difference (in the square mean) between \hat{x}_t and x_t^* . The convergence result given in this paper, Proposition 5, is only concerned with this difference and does not say anything about the variance part of the estimation error.

The variance part of the error turns out to be the variance of the conditional distribution π_t . This is due to the following important remark.

The filtering problem considered here (with x given by (65) and Y by (50)) has a finite dimensional solution. This is quite similar to a Kalman-Bucy filter. Proposition 1 states that the conditional distribution π_t given past observations of Y is the same as given past observations of y , (in the proposition, this is the statement that $\bar{\mathcal{Y}}_t = \mathcal{Y}_t$). Now, y is here given by (54). Replacing (66) for the function h shows it is just a linear additive white noise model. Moreover, it is clear from that




Fig. 1. Influence of ν and N on estimation errors

equation (65) satisfied by x is a linear stochastic differential equation. In particular, as already mentioned, x is normal.

Thus, the conditional distribution π_t is a normal distribution with mean x_t^* and covariance $P_t e$ where $P_t \geq 0$, (recall e is the 3×3 identity matrix). Moreover, x^* and P satisfy the Kalman-Bucy equations [36]

$$dx_t^* = -\nu x_t^* + P_t [dz_t - x_t^* dt], \quad (72)$$

and

$$\dot{P}_t = -P_t^2 - 2\nu P_t + \sigma^2, \quad (73)$$

where z is a process with values in \mathbb{R}^3 whose coordinates are $z_t^r = \langle y_t, \omega_r \rangle$.

The conditional distribution π_t is the exact solution of the current filtering problem. It can be computed by integrating the Kalman-Bucy equations (72) and (73). This can be done in a standard way after replacing dz_t in terms of $dy_t = Y_t^\dagger \circ dY_t$, which is (53). On the other hand, $\hat{\pi}_M^\delta$ is an approximation of π_t . In addition to being unbiased, x_t^* is an optimal estimator of x_t in the sense of mean square error. As \hat{x}_t converges to x_t^* , the bias part of the estimation error disappears and the optimal error P_t is achieved in the limit.

From a practical point of view, there is no need here to implement a particle filter. In fact, it is even computationally less expensive to integrate (72) and (73). However, as it is well know, the existence of finite dimensional solutions is the exception rather than the rule in real situations. The connection between the present example and rigid body mechanics has already been mentioned. If Y is used to represent the pose of a rigid body, then x is the angular velocity. The filtering problem appears as the problem of tracking angular velocity based on observations of the pose alone. In general, the angular velocity of a rigid body satisfies Euler's equation of rigid body mechanics, which is far more complicated than (65) and in particular nonlinear [32]. Thus, when a realistic model is used for x , using a particle algorithm or some other approximation becomes indispensable.

Before going on to the next example, note the behavior of the estimation error in Figure 1. While $x_0 = (0, 0, 0)$, the particles were initialised at $(4, 0, 0)$. Since $\nu > 0$, this initial error is quickly overcome. However,

the fact that $\nu > 0$ leads to nonzero asymptotic variance. As t goes to infinity P_t goes to $P_\infty = -\nu + (\sigma^2 + \nu^2)^{1/2}$. Thus, it is not possible to track x exactly. This problem appears since, when $\nu > 0$, the process x is ergodic. In particular, its asymptotic distribution is normal with mean 0 and variance $\sigma^2/2\nu$. This imposes a fundamental limit on estimation error.

5.2. Observations in S^2

In this second example, the observation process Y lies in the unit sphere S^2 . Of course, S^2 is the set of $p \in \mathbb{R}^3$ such that $|p| = 1$. In comparison with the previous example, the current one will raise several additional difficulties. Roughly, these are due to the fact that S^2 is not a Lie group but a symmetric space of the compact Lie group $SO(3)$.

The manifold structure of S^2 is most easily understood as inherited from \mathbb{R}^3 . Precisely, S^2 is a compact embedded submanifold of \mathbb{R}^3 . For $p \in S^2$, the tangent space $T_p S^2$ is the subspace of \mathbb{R}^3 consisting of those vectors K such that $(K, p) = 0$. Here, (\cdot, \cdot) denotes the standard Euclidean scalar product in \mathbb{R}^3 . In particular, $(p, p) = |p|^2$.

As in the previous example, no restriction is made on the signal model (1). The sensor function H and the vector fields Σ_r are defined in terms of the action of $SO(3)$ on S^2 . This is now briefly discussed.

For each $p \in S^2$, consider the two following linear mappings. First, the orthogonal projection $\Pi_p : \mathbb{R}^3 \rightarrow T_p S^2$. This is defined by

$$\Pi_p(v) = v - (p, v)p = p \times v \times p, \quad (74)$$

for $v \in \mathbb{R}^3$, where \times denotes vector product.

The second mapping is $\Sigma_p : \mathbb{R}^3 \rightarrow T_p S^2$ defined as follows. Let $\sigma_1, \sigma_2, \sigma_3$ be as in (58), with $b = e$. In other words, $\sigma_r = \omega_r$. Also, let $\sigma(v) = v^r \sigma_r$ where $v \in \mathbb{R}^3$ (this was called $h(v)$ in (66)). The mapping Σ_p is given by

$$\Sigma_p(v) = \sigma(v)p = v \times p. \quad (75)$$

This is related to the action of $SO(3)$ on S^2 in a simple way. Note that $\sigma(v) \in \mathfrak{so}(3)$ for $v \in \mathbb{R}^3$. If $\gamma(t) = \exp(t\sigma(v))$ for $t \in \mathbb{R}$ then $\gamma(t) \in SO(3)$ and

$$\Sigma_p(v) = \left. \frac{d}{dt} \right|_{t=0} \gamma(t)p.$$

In other words, $\Sigma_p(v)$ is the velocity of the point p when it is in uniform rotation with angular velocity v .

Unlike Π_p , the mapping Σ_p is not a projection. However, both mappings are surjective and have the same kernel,

$$\text{Ker}(\Pi_p) = \text{Ker}(\Sigma_p) = N_p S^2,$$

where $N_p S^2$ is the normal space to S^2 at p . This is a one dimensional subspace of \mathbb{R}^3 consisting of vectors λp where $\lambda \in \mathbb{R}$.

Using the mappings Π_p and Σ_p , the tangent bundle of the sphere can be described in a covariant way. For any $p, q \in S^2$ there exists $k \in SO(3)$ such that $kp = q$. In fact, there are an infinity of such k . The following relations hold

$$\Pi_{kp}(kv) = k\Pi_p(v), \quad \Sigma_{kp}(kv) = k\Sigma_p(v), \quad (76)$$

and can also be written

$$\Pi_q(v) = (k\Pi_p k^{-1})(v), \quad \Sigma_q(v) = (k\Sigma_p k^{-1})(v). \quad (77)$$

Returning to the filtering problem, a general observation model where the observations lie on S^2 can be defined using either the mappings Π_p or Σ_p for each $p \in S^2$. It is preferable to use Σ_p , since it is immediately related to the action of $SO(3)$ on S^2 .

The sensor function H will be assumed of the following form,

$$H(s, p) = \Sigma_p(h(s)), \quad (78)$$

where $h : S \rightarrow \mathbb{R}^3$. The vector fields Σ_r will be defined by

$$\Sigma_r(p) = \Sigma_p(e_r), \quad r = 1, 2, 3, \quad (79)$$

where e_1, e_2, e_3 is the canonical basis of \mathbb{R}^3 .

As in the previous example, all the differentiability conditions of Section 2 hold, since the operations used to define H and Σ_r are linear. Moreover, the condition of ellipticity, required to introduce the metric (21) and the Le Jan-Watanabe connection (22), is verified. This is because, by construction, the mapping Σ_p is surjective for each $p \in S^2$. Replacing the definition (75) of Σ_p in (2) gives the observation model

$$dY_t = -Y_t \times \{h(x_t)dt + \text{od}B_t\}, \quad (80)$$

where B is a standard Brownian motion in \mathbb{R}^3 .

In the current example, the number of vector fields Σ_r is equal to 3 whereas the dimension of S^2 is equal to 2. As a result, there is no simple formula similar to (53) that can be used to find the antidevelopment process y of Y . Rather, it is necessary to consider a parallel frame along Y . This is done after introducing the metric (21) and the Le Jan-Watanabe connection (22). It turns out these are the same as the Riemannian metric that S^2 inherits from \mathbb{R}^3 and its associated Levi-Civita connection [29].

It is straightforward to show the Euclidean scalar product (\cdot, \cdot) verifies (21). Let $p \in S^2$ and $K \in T_p S^2$. Note that

$$(K, \Sigma_r(p)) = (p \times K)^r,$$

where the right hand side is the r^{th} component of $p \times K$ in the canonical basis e_1, e_2, e_3 . Evaluating the right hand side of (21) (with $K = E$) gives

$$(K, \Sigma_r(p))(K, \Sigma_r(p)) = |p \times K|^2.$$

But this is

$$(p, p)(K, K) - (K, p)^2 = (K, K),$$

since $(p, p) = |p|^2 = 1$ and $(K, p) = 0$. Thus, the metric (21) is the same as (\cdot, \cdot) .

For the connection (22), note that by definition this is given by

$$\nabla_K E = K(E, \Sigma_r) \Sigma_r(p),$$

for $K \in T_p S^2$ and any C^1 vector field E on S^2 . This can also be written

$$\nabla_K E = \Sigma_p((p \times KE)),$$

where KE is the vector $KE = (KE^1, KE^2, KE^3)$. The last formula follows from (75) and (79) by linearity. If (75) is applied to it again, then it follows

$$\nabla_K E = p \times KE \times p = \Pi_p(KE),$$

which is the definition of the Levi-Civitation connection associated to (\cdot, \cdot) ; see [23].

In order to construct a parallel frame along Y , let $E_0^1, E_0^2 \in T_{Y_0} S^2$ be orthonormal. Also, let E^1, E^2 be vector fields along Y solving the equation of parallel transport (9). Here, this reads

$$\nabla_{\circ dY} E_t^i = \Pi_{Y_t}(dE_t^i) = 0,$$

for $i = 1, 2$ and with initial conditions E_0^1, E_0^2 . Another way of writing this equation, based on (74), is

$$dE_t^i = -Y_t(E_t^i, \circ dY_t). \quad (81)$$

To obtain this, it is enough to replace in (74) the fact that

$$d(E_t^i, Y_t) = (\circ dE_t^i, Y_t) + (E_t^i, \circ dY_t) = 0.$$

If $Y_0, E_0^1, -E_0^2$ is a positively oriented orthonormal basis in \mathbb{R}^3 , then $Y_t, E_t^1, -E_t^2$ will have this same property. It is now assumed this is the case.

From its definition (13), the antidevelopment process y of Y has its values in \mathbb{R}^2 and is given by

$$y_t^j = \int_0^t (E_s^j, \circ dY_s). \quad (82)$$

Let z be the process with values in \mathbb{R}^3 ,

$$z_t = h(x_t) dt + dB_t. \quad (83)$$

It follows from (80) that

$$dy_t^j = -(E_t^j, Y_t \times \circ dz_t) = -(E_t^j \times Y_t, \circ dz_t).$$

Given the chosen orientation for Y_t, E_t^1, E_t^2 , this yields

$$dy_t^j = (E_t^j, \circ dz_t). \quad (84)$$

From this, it is possible to recover (25) of Proposition 1. Namely,

$$dy_t^j = h_t^j dt + d\beta_t, \quad (85)$$

where $h_t^j = (E_t^j, h(x_t))$ and β is a standard Brownian motion in \mathbb{R}^2 .

From (84),

$$dy_t^j = h_t^j dt + (E_t^j, \circ dB_t).$$

Thus, it is enough to show

$$(E_t^j, \circ dB_t) = (E_t^j, dB_t) = d\beta_t^j.$$

That the Stratonovich differential can be replaced by an Itô differential follows from

$$(E_t^j, \circ dB_t) = (E_t^j, dB_t) + \frac{1}{2}(dE_t^j, dB_t),$$

where the last term denotes quadratic covariation. From (80) and (81), this is

$$(dE_t^j, dB_t) = -(Y_t, dB_t)(E_t^j, Y_t \times dB_t) = (E_t^j, Y_t \times Y_t) dt,$$

which is identically zero. That β is a Brownian motion follows from the fact that E^1, E^2 are orthonormal.

Equation (80) can be rewritten in terms of the parallel frame E^1, E^2 and the antidevelopment process y . Replacing in (80) the fact that Y, E^1, E^2 is an orthonormal basis, it follows from (82) that

$$dY_t = E_t^1 \circ dy_t^1 + E_t^2 \circ dy_t^2. \quad (86)$$

Similarly, (81) can be rewritten using (82),

$$dE_t^i = -Y_t \circ dy_t^i. \quad (87)$$

Now (86) and (87) form a system of linear stochastic differential equations which can be solved knowing the antidevelopment y . This shows that Y can be obtained if y is known. This is in spite of the fact that y has its values in \mathbb{R}^2 while Y has its values in \mathbb{R}^3 .

In order to apply the particle filtering algorithm of Section 4 to the current example, it is enough to specify a mapping I verifying conditions **(II-13)** of 3.3. Two such mappings are now considered. First, recall that I can be given by (28) as a geodesic connector. Here, the connection ∇ is the Levi-Civita connection corresponding to the Euclidean scalar product (\cdot, \cdot) . Thus, geodesics are to be understood in the usual meaning of large circles. Accordingly, for $p \in S^2$ and $K \in T_p S^2$ with $|K| \neq 0$,

$$\exp_p(K) = \cos |K| p + \sin |K| (K/|K|). \quad (88)$$

Moreover, when $p, q \in S^2$ and $(p, q) \neq \pm 1$,

$$\log_p(q) = \arcsin |\Pi_p(q)| (\Pi_p(q) / |\Pi_p(q)|). \quad (89)$$

Using this last formula, it is possible to implement (28). However, this involves several nonlinear operations. Another, simpler, mapping I can be guessed from (89). Consider the following

$$I(p, q) = \Pi_p(q). \quad (90)$$

This is the first order approximation of $\log_p(q)$ and is well defined for any $p, q \in S^2$. To see that it verifies conditions **(II-13)** note that

$$d\Pi_p(q)(V) = \Pi_p(V),$$

for all $V \in T_q S^2$. Thus, if $V \in T_p S^2$,

$$d\Pi_p(p)(V) = V,$$

which is condition **(I2)**. Condition **(I3)** reads

$$\nabla^2 \Pi_p(p)(V, V) = 0.$$

By definition of the connection ∇ , the left hand side is the projection on $T_p S^2$ of the acceleration at $t = 0$ of the geodesic

$$\gamma_p(t) = \cos(t)p + \sin(t)(V/|V|),$$

but this acceleration is equal to $-p$, so its projection on $T_p S^2$ is zero. Finally, condition **(I1)** is easily verified from (74).

The particle filtering algorithm of Section 4 can be applied as before. In particular, instruction 2 is carried out by replacing (78) and (90) into (36) in order to compute particle weights.

Precisely, the likelihood function based on successive samples $Y_{k\delta}$ and $Y_{(k+1)\delta}$ becomes

$$l(s) = \exp\left(\left(h(s) \times Y_{k\delta}, Y_{(k+1)\delta}\right) - \frac{\delta}{2} |h(s) \times Y_{k\delta}|^2\right). \quad (91)$$

By (78) and (90), the first term under the exponential is

$$(H(s, Y_{k\delta}), Y_{(k+1)\delta}) = (H(s, Y_{k\delta}), I(Y_{k\delta}, Y_{(k+1)\delta})),$$

just like in (36). The second term can also be found from (78).

The above discussion is now illustrated with a computer experiment. For the signal model, consider $S = \mathbb{R}^3$. The signal is simply a constant $x^* = (0, 0, 1)$. The function h of (78) is taken to be the identity function $h(s) = s$ for $s \in \mathbb{R}^3$. Just like in the previous example, it can be noted that condition **(H1)** is not verified. Again, since x is normal this condition can be overlooked.

With the function h chosen in this way, the observation model (80) becomes

$$dY_t = -Y_t \times \{x^* + \circ dB_t\}. \quad (92)$$

A trajectory of the observation process Y can be simulated using a formula similar to (68). Precisely, consider an approximating process \bar{Y}^δ where, (letting $\hat{x}^* = \hat{x}^r \sigma_r$),

$$\bar{Y}_t^\delta = \exp[(t - k\delta)(\hat{x}^* + \Delta \hat{B}_k)] \bar{Y}_{k\delta}^\delta, \quad (93)$$

on each interval $[k\delta, (k+1)\delta]$. The value of δ used here was $\delta = 0.1$, (the same as in the previous example).

In order to ensure the process \bar{Y}^δ has its values in $S^2 \subset \mathbb{R}^3$, it is enough to take $\bar{Y}_0^\delta \in S^2$. This was chosen to be $\bar{Y}_0^\delta = (0, 0, 1)$. The product appearing in (93) is between the matrix exponential on the left, which belongs to $SO(3)$, and the vector $\bar{Y}_{k\delta}^\delta$ on the right which belongs to S^2 .

Unlike the case of the previous example, (see (72) and (73)), there is here no known finite dimensional solution for optimal estimation of x^* . Application of particle filtering or of some other approximate solution is thus necessary.

Figure 2 below shows the distribution of $N = 1000$ particles in the (x_1, x_2) and (x_1, x_3) planes at times $T =$

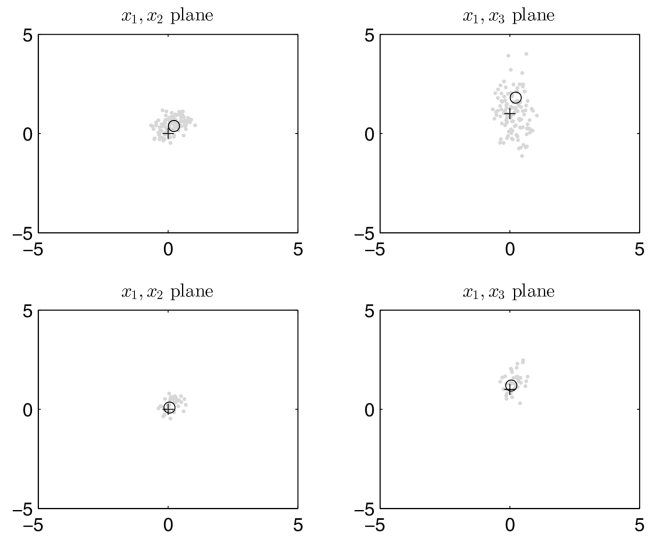


Fig. 2. Particles distribution (grey); estimate (o); true value (+)

1.5 (top row) and $T = 3$ (bottom row). Here, (x_1, x_2, x_3) are canonical coordinates in the basis e_1, e_2, e_3 . The position of x^* is designated by a + and the estimate \hat{x}_t (arithmetic average of the particles as in (71)) by o. The particles were initially generated from a normal distribution μ with mean $(0.5, 0.5, 1)$ and variance 1. A large value of N was chosen for visualisation. It is possible to use $N = 100$ with a similar performance.

Figure 2 shows the particle filtering algorithm of Section 4 is able to recover x^* within a relatively short time. It is interesting to note the larger variability of the particle distribution in the x_3 direction, apparent in the right column of the figure. This is because $Y_0 = (0, 0, 1)$ (the same as \bar{Y}_0^δ) so that, initially, the component of x^* along e_3 has no effect on the position of Y .

Due to the presence of noise B , the observation process Y rapidly explores a large area of S^2 (theoretically, Y is a recurrent process in S^2). This allows for the initial ambiguity in the x_3 direction to be overcome.

In the absence of noise, Y_t rotates uniformly around x^* . If x^* and Y_0 are parallel, $Y_t = Y_0$ for all $t \geq 0$. Then, Y contains no information regarding the magnitude of x^* . Otherwise, x^* can be recovered after an arbitrarily short time (knowing the model (92)). Such a situation cannot arise in the general case of noisy observations, since Y_t and x^* do not remain parallel.

The computer experiment presented here shows that the particle filtering algorithm of Section 4 is able to successfully handle a filtering problem which is not solvable by classical methods.

6. CONCLUSION

This paper considered continuous time filtering problems where the observation process, conditionally on the unknown signal, is an elliptic diffusion in a differentiable manifold. In order to numerically solve filtering problems of this kind, the paper proposed a particle filtering algorithm which it also proved to be

convergent under some additional technical conditions. Roughly, this algorithm combines the well known sequential Monte Carlo structure of a classical particle filter with the geometric construction of connector maps, used to locally linearise the observation process. To the author's knowledge, the proposed algorithm is entirely new in the literature.

The filtering problems considered in the paper are of a very general form. While this may have lead to some unnecessary abstraction, it also has clear advantages. When dealing with an applied problem, greater generality in mathematical formulation allows additional freedom in choosing a realistic observation model which includes sufficient a priori knowledge of the target application. Also, since most physical phenomena are naturally described in continuous time, the fact of starting from a continuous time formulation should accommodate the majority of physical models.

The particle filtering algorithm proposed in this paper leaves several choices open to the user wishing to implement it. These include the choice of an approximation of the hidden Markov structure and the choice of a connector map for local linearisation. This gives additional adaptability and allows for the trade-off between complexity and performance to be optimised according to applications. In any case, the paper gave precise conditions which the chosen implementation should satisfy in order to produce a consistent numerical solution.

This paper was only a first effort in the new direction of particle filtering with observations in a manifold. It was aimed at laying down a rigorous and adaptable general framework. Hopefully, additional papers strengthening convergence results and exploring in detail important engineering applications will be shortly submitted.

APPENDIX A

In 3.3, Proposition 2 was cited from [30]. However, soon after, a more general claim was made without proof. Namely, that Proposition 2 continues to hold if the mapping I of (28) is replaced by any other mapping which verifies conditions **(I1-I3)**.

The proof of this claim is a repetition of the one in [30], but does not seem to have been given explicitly in the literature. For completeness, it is here provided.

In preparation, consider the following generalised Itô formula. Let f be a C^2 function on M and replace $\tau = df$ in (20). This gives

$$df(Y_t) = (df, dY_t) + (1/2)\nabla^2 f(Y_t)(\Sigma_r, \Sigma_r)dt. \quad (94)$$

Let $I : \mathcal{M} \times \mathcal{M} \rightarrow T\mathcal{M}$ be any mapping which verifies conditions **(I1-I3)**. For $\delta > 0$ and any $k \in \mathbb{N}$, let

$$I_k(q) = I(Y_{k\delta}, q) \quad q \in \mathcal{M}.$$

Conditionally on $Y_{k\delta} = p$, this is a C^2 function on M with values in $T_p\mathcal{M}$. This is by condition **(I1)**. In [37],

it was shown that the Itô formula (94) can be applied so

$$I(Y_{k\delta}, Y_{(k+1)\delta}) = \int_{k\delta}^{(k+1)\delta} (dI_k, dY_t) + (1/2) \int_{k\delta}^{(k+1)\delta} \nabla^2 I_k(Y_t)(\Sigma_r, \Sigma_r)dt.$$

Here dI_k and $\nabla^2 I_k$ denote differentiation component by component of the vector valued function I_k after an arbitrary choice of basis.

It will be useful to rewrite this using (17),

$$I^i(Y_{k\delta}, Y_{(k+1)\delta}) = \int_{k\delta}^{(k+1)\delta} dI_k^{ij}(Y_t)dy_t^j + (1/2) \int_{k\delta}^{(k+1)\delta} \nabla^2 I_k^i(Y_t)dt,$$

where $I_k^i = \langle I_k, E^i \rangle$ and

$$dI_k^{ij} = (dI_k^i, E^j), \quad \nabla^2 I_k^i = \nabla^2 I_k^i(E^j, E^j).$$

On the other hand

$$y_{(k+1)\delta}^i - y_{k\delta}^i = \int_{k\delta}^{(k+1)\delta} dI_k^{ij}(Y_{k\delta})dy_t^j + (1/2) \int_{k\delta}^{(k+1)\delta} \nabla^2 I_k^i(Y_{k\delta})dt.$$

This is because, by conditions **(I2-I3)**,

$$dI_k^{ij}(Y_{k\delta}) = \delta_{ij}, \quad \nabla^2 I_k^i(Y_{k\delta}) = 0.$$

Note that, by (15),

$$dy_t^i = \langle E^i, H + (1/2)\nabla_{\Sigma_r} \Sigma_r \rangle dt + d\beta_t^i.$$

Here, the drift coefficient appearing before dt is uniformly bounded and β is a standard Brownian motion in \mathbb{R}^d .

Let, (this is the notation of (30) in Proposition 3),

$$\Delta y_k^i = y_{(k+1)\delta}^i - y_{k\delta}^i, \quad \Delta Y_k^i = I^i(Y_{k\delta}, Y_{(k+1)\delta}).$$

Then, by Itô isometry,

$$\begin{aligned} \mathbb{E}|\Delta y_k^i - \Delta Y_k^i|^2 &\leq C \max_j \int_{k\delta}^{(k+1)\delta} \mathbb{E}|dI_k^{ij}(Y_t) - dI_k^{ij}(Y_{k\delta})|^2 dt \\ &\quad + \int_{k\delta}^{(k+1)\delta} \mathbb{E}|\nabla^2 I_k^i(Y_t) - \nabla^2 I_k^i(Y_{k\delta})|^2 dt, \end{aligned}$$

where C is some positive constant (which does not depend on k).

By condition **(I1)** and the fact that the manifold \mathcal{M} is compact, dI_k^{ij} and $\nabla^2 I_k^i$ are bounded and continuous. Therefore, the expectations under the integral in each term tend to zero as $\delta \downarrow 0$.

This proves that

$$\mathbb{E}|\Delta y_k^i - \Delta Y_k^i|^2 = o(\delta), \quad (95)$$

by an extension of this result, it is straightforward to establish Proposition 2. This is now done.

As in the proposition, formula (29), let

$$R_\delta = \sum_{k\delta < t} \langle G_{k\delta}, \Delta Y_k \rangle = \sum_{k\delta < t} G_{k\delta}^i \Delta Y_k^i,$$

where $G_{k\delta}^i = \langle E^i, Gk\delta \rangle$.

By (17) and the definition of classical Itô integral, the integral $R = \int_0^t \langle G_s, dY \rangle$ is the limit in the square mean of

$$r_\delta = \sum_{k\delta < t} G_{k\delta}^i \Delta y_k^i.$$

Note as before

$$G_{k\delta}^i \Delta y_k^i = \int_{k\delta}^{(k+1)\delta} G_{k\delta}^i dI_k^{ij}(Y_t) dy_t^j + (1/2) \int_{k\delta}^{(k+1)\delta} G_{k\delta}^i \nabla^2 I_k^i(Y_t) dt,$$

and

$$G_{k\delta}^i \Delta y_k^i = \int_{k\delta}^{(k+1)\delta} G_{k\delta}^i dI_k^{ij}(Y_{k\delta}) dy_t^j + (1/2) \int_{k\delta}^{(k+1)\delta} G_{k\delta}^i \nabla^2 I_k^i(Y_{k\delta}) dt.$$

By summing over k and using Itô isometry and the fact that $\|G_t\|$ is bounded

$$\begin{aligned} \mathbb{E}|R_\delta - r_\delta|^2 &\leq C \sup_{s \leq t} \|G\|^2 \times \\ &\sum_{k\delta < t} \max_j \int_{k\delta}^{(k+1)\delta} \mathbb{E}|dI_k^{ij}(Y_t) - dI_k^{ij}(Y_{k\delta})|^2 dt \\ &\sum_{k\delta < t} \max_j \int_{k\delta}^{(k+1)\delta} \mathbb{E}|\nabla^2 I_k^i(Y_t) - \nabla^2 I_k^i(Y_{k\delta})|^2 dt, \end{aligned}$$

where C is some positive constant, possibly different from before. Using again condition **(I1)** and the fact that the manifold \mathcal{M} is compact it is seen that

$$\lim_{\delta \downarrow 0} \mathbb{E}|R_\delta - r_\delta|^2 = 0,$$

which, from the definition of r_δ , immediately gives Proposition 2. It is enough to write

$$\mathbb{E}|R_\delta - R|^2 \leq 2\mathbb{E}|R_\delta - r_\delta|^2 + 2\mathbb{E}|r_\delta - R|^2.$$

The first term has just been proved to converge to zero. The second term converges to zero by definition of r_δ .

As already mentioned, this proof is similar to the one in [30], but makes the additional remark that the only required properties for the mapping I are conditions **(I1-I3)**. Accordingly, there is no need to restrict I to being the geodesic connector mapping (28). In its above form, the proof explicitly uses compactness of \mathcal{M} in order to obtain convergence in the square mean. However, it is clear that this can be replaced by a milder assumption.

REFERENCES

- [1] J. H. Manton
On the role of differential geometry in signal processing, in *International Conference on Acoustics Speech and Signal Processing*, 2005.
- [2] Y. Liu, G. Li, and Z. Shi
Tracking deformable object via particle filtering on manifolds, in *Chinese Conference on Pattern Recognition*, 2008.

- [3] G. Wang, Y. Liu, and H. Shi
Covariance tracking via geometric particle filtering, in *Intelligent Computation Technology and Automation*, 2009, 2009.
- [4] Q. Rentmeesters, P. Absil, P. Van Dooren, K. Gallivan, and A. Srivastava
An efficient particle filtering technique on the Grassman manifold, in *International Conference on Acoustics Speech and Signal Processing*, 2010.
- [5] J. Kwon, M. Choi, C. Chun, and F. C. Park
Particle filtering on the Euclidean group, in *International Conference on Robotics and Automation*, 2007.
- [6] F. Tompkins and P. J. Wolfe
Bayesian filtering on the stiefel manifold, in *Advances in Multi-Sensor Adaptive Processing*, 2007.
- [7] J. Saboune and F. Charpillat
Using interval particle filtering for marker less 3d human motion capture, in *Tools with Artificial Intelligence*, 2005.
- [8] H. Snoussi and C. Richard
Monte carlo tracking on the riemannian manifold of multi-variate normal distributions, in *Digital Signal Processing Workshop and Signal Processing Education Workshop*, 2009.
- [9] A. Srivastava
A nonlinear filtering method for geometric subspace tracking, in *Sensor Array and Multichannel Signal Processing Workshop*, 2000.
- [10] M. Emery
Stochastic calculus in manifolds. Springer-Verlag (Universitext), 1989.
- [11] E. P. Hsu
Stochastic analysis on manifolds. Americal Mathematical Society (Graduate Studies in Mathematics), 2002.
- [12] J. Manton
A primer on stochastic differential geometry for signal processing, *Selected Topics in Signal Processing, IEEE Journal of*, vol. 7, no. 4, pp. 681–699, 2013.
- [13] J. T. Lo
Signal detection for bilinear systems, *Information Sciences*, vol. 9, pp. 249–378, 1975.
- [14] S. Said and J. H. Manton
On filtering with observation in a manifold: Reduction to a classical filtering problem, *SIAM Journal on Control and Optimisation*, vol. 51, no. 1, pp. 767–783, 2013.
- [15] M. Pontier and J. Szpirglas
Filtrage non linéaire avec observations sur une variété, *Stochastics*, vol. 15, no. 2, pp. 121–148, 1985.
- [16] S. K. Ng and P. E. Caines
Nonlinear filtering in Riemannian manifolds, *IMA Journal of Mathematical Control and Information*, vol. 2, pp. 25–36, 1985.
- [17] T. E. Duncan
Some filtering results in Riemannian manifolds, *Information and Control*, vol. 35, pp. 182–195, 1977.
- [18] M. Pontier
Approximation d'un filtre avec observation sur une variété, *Stochastics*, vol. 24, no. 4, pp. 285–304, 1988.
- [19] P. Del Moral
Feynman-Kac formulae: Geological and interacting particle systems with applications. Springer (Probability and Applications), 2004.

- [20] J. Boulanger, N. Le Bihan, S. Said, and J. Manton
Monte-carlo estimation from observation on stiefel manifold,
in *International Conference on Acoustics Speech and Signal Processing*, 2014.
- [21] D. Crisan and A. Doucet
A survey of convergence results on particle filtering methods for practioners,
IEEE Transactions on Signal Processing, vol. 50, no. 3, pp. 736–746, 2002.
- [22] A. Bain and D. Crisan
Fundamentals of stochastic filtering.
Springer, 2009.
- [23] I. Chavel
Riemannian geometry: A modern introduction.
Cambridge University Press, 2006.
- [24] L. C. G. Rogers and D. Williams
Diffusions, Markov processes and Martingales, Vol. 1.
Cambridge University Press, 2000.
- [25] N. Ikeda and S. Watanabe
Stochastic differential equations and diffusion processes.
North Holland, 1988.
- [26] L. C. G. Rogers and D. Williams
Diffusions, Markov processes and Martingales, Vol. 2.
Cambridge University Press, 2000.
- [27] P. E. Kloeden and E. Platen
Numerical solution of stochastic differential equations.
Springer (Stochastic modelling and applied probability), 1992.
- [28] M. Pontier and J. Szpirglas
Convergence des approximations de McShane d’une diffusion sur une variété compacte,
Séminaire de probabilités (Strasbourg), vol. 21, pp. 534–543, 1987.
- [29] K. D. ELworthy, Y. Le Jan, and X. M. Li
On the geometry of diffusion operators and stochastic flows.
Springer (Lecture notes in mathematics 1720), 1999.
- [30] R. W. R. Darling
“Martingales on manifolds and geometric Itô calculus,”
Ph.D. dissertation, The University of Warwick, 1982.
- [31] J. Faraut
Analysis on Lie groups: An introduction.
Cambridge University Press, 2008.
- [32] V. I. Arnold
Mathematical methods of classical mechanics.
Springer (Graduate Texts in Mathematics), 2010.
- [33] S. Helgasson
Differential geometry, Lie groups and symmetric spaces.
American Mathematical Society (Graduate Studies in Mathematics), 2001.
- [34] A. Estrade
Exponentielle stochastique et intégrale multiplicative discontinue,
Annales de l’Institut Henri Poincaré (B), probabilités et statistiques, vol. 28, no. 1, pp. 107–129, 1992.
- [35] M. Hakim-Dowek and D. Lépingle
L’exponentielle stochastique des groupes de lie,
Séminaire de probabilités (Strasbourg), vol. 20, pp. 352–374, 1986.
- [36] A. H. Jazwinski
Stochastic processes and filtering theory.
Academic Press, 1970.
- [37] H. Kunita
Some extensions of itô’s formula,
Séminaire de probabilités (Strasbourg), vol. 15, pp. 118–141, 1981.



Salem Said is a research scientist with CNRS since 2014, at IMS laboratory in Bordeaux, France. He currently holds an IdEx Bordeaux Chair of Installation. Previously, he held a research fellowship at the University of Melbourne school of engineering, under the supervision of Pr. Jonathan Manton. His current research focuses on information geometry and adaptive algorithms. His research also includes the areas of physical modeling and partial differential equations.

Jonathan Manton holds a Distinguished Chair at the University of Melbourne with the title Future Generation Professor. He is also an Adjunct Professor in the Mathematical Sciences Institute at the Australian National University. Prof Manton is a Fellow of IEEE and a Fellow of the Australian Mathematical Society.

He received his Bachelor of Science (mathematics) and Bachelor of Engineering (electrical) degrees in 1995 and his Ph.D. degree in 1998, all from the University of Melbourne, Australia. From 1998 to 2004, he was with the Department of Electrical and Electronic Engineering at the University of Melbourne. During that time, he held a Postdoctoral Research Fellowship then subsequently a Queen Elizabeth II Fellowship, both from the Australian Research Council. In 2005 he became a full Professor in the Department of Information Engineering, Research School of Information Sciences and Engineering (RSISE) at the Australian National University. From July 2006 till May 2008, he was on secondment to the Australian Research Council as Executive Director, Mathematics, Information and Communication Sciences.

Prof Manton's traditional research interests range from pure mathematics (e.g. commutative algebra, algebraic geometry, differential geometry) to engineering (e.g. signal processing, wireless communications, systems theory). More recently, he has become interested in systems biology and systems neuroscience.



Direct Position Determination for TDOA-based Single Sensor Localization

CHRISTIAN STEFFES
MARC OISPUU

In this paper, four different localization techniques based on TDOA-measurements for single sensor passive emitter localization are proposed. The use of signal structure information allows TDOA-based localization with a single moving sensor node. A direct position estimation scheme is derived for the single sensor TDOA localization problem. The feasibility of the proposed method is shown in simulations. The position estimation accuracy of the single sensor TDOA and the direct technique are compared using simulation results and the Cramér-Rao Lower Bound. Field experiments using an airborne sensor are conducted to prove the concept.

Manuscript received October 30, 2015; revised March 4, 2016; released for publication June 13, 2016.

Refereeing of this contribution was handled by Gerhard Kurz.

Authors' address: Dept. Sensor Data and Information Fusion, Fraunhofer Institute for Communication, Information Processing and Ergonomics FKIE, Fraunhoferstr. 20, 53343 Wachtberg, Germany. (E-mail: {christian.steffes, marc.oispuu}@fkie.fraunhofer.de).

1557-6418/16/\$17.00 © 2016 JAIF

I. INTRODUCTION

Passive emitter localization is a fundamental task encountered in various fields like wireless communication, radar, sonar, seismology, and radio astronomy. An airborne sensor platform is the preferable solution in many applications. The sensor is typically mounted e.g. on an aircraft, a helicopter, or an unmanned aerial vehicle (UAV). Airborne sensors provide in comparison to ground located sensors a far-ranging signal acquisition because of the extended radio horizon. Mostly for localization issues, sensors are installed under the fuselage or in the wings of the airborne sensor platform. In case of hard payload restrictions only compact sensors come into consideration.

Aspects of the two-dimensional and three-dimensional localization problem examined in the literature include numerous estimation algorithms, estimation accuracy, and target observability [16], [3]. Typical localization systems of interest obtain measurements like direction of arrival (DOA), frequency difference of arrival (FDOA), time difference of arrival (TDOA) or combinations of the aforementioned measurements [6].

Commonly, the desired source locations are determined in multiple steps: the signal processing step where the sensor data is computed from the raw signal data, and the sensor data fusion step where the localization and tracking task is performed. Alternatively, direct position determination (DPD) approaches have been proposed to compute the desired target parameters in a single step based on the raw signal data without explicitly computing intermediate measurements like DOA, FDOA, and TDOA [19], [18]. It has been shown that this kind of data processing offers a superior performance in scenarios with weak or closely-spaced sources but requires a higher computational burden in comparison to the standard multi-step processing. For example for TDOA-based localization, a direct approach based on the raw signal data has been proposed in [19], [1], and a standard approach based on TDOA/FDOA measurements has been proposed in [15], [16], respectively. In [10], a localization approach based on the complex ambiguity function (CAF) has been introduced which turned out to be a compromise between localization performance and computational burden. Analysis of emitter localization using a single moving observer based on frequency measurements with context knowledge has been introduced in [4]. The results show the advantage of using a priori knowledge concerning the emitter's altitude (either known or using a terrain model) on the performance of a localization system. In [2], a method for single platform geolocation using joint Doppler and AOA measurements is proposed. The combination of these heterogeneous measurements can allow more accurate position estimation. More recently, research on the single receiver TOA/TDOA-based localization using the periodicity of emitted signals has attracted attention. In [17], the single observer geolocation dealing

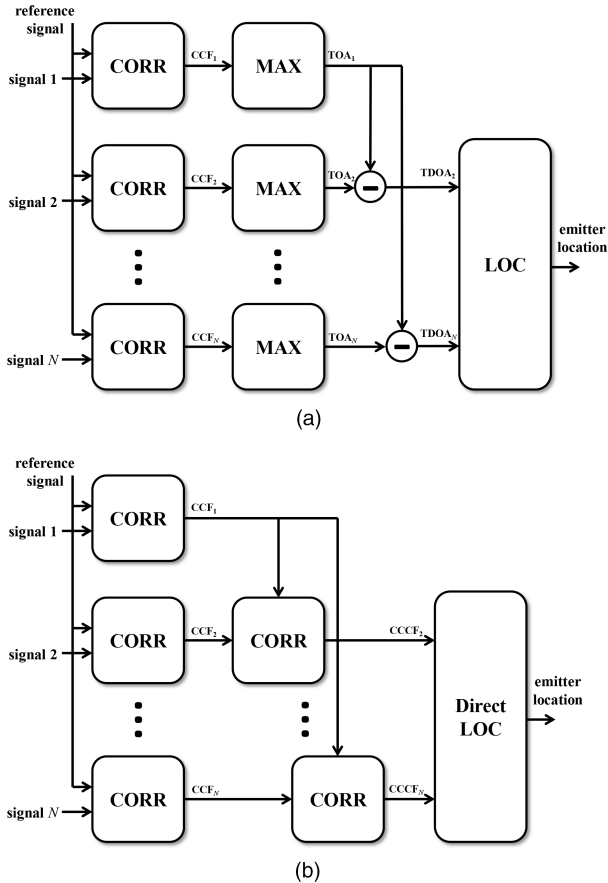


Fig. 1. Comparison of non-direct and direct S^4 TDOA approach. (a) S^4 TDOA approach. (b) DS^4 TDOA.

with oscillator instability is investigated. Experimental results using a moving observer and Kalman filters for the estimation of the local oscillator drift can be found in [7], [11].

In our previous work, we proposed a DPD approach for a moving antenna array sensor [9], [8]. Furthermore in [13], we introduced a *single element* TDOA localization approach using just a single omni-directional antenna. In the following, this approach is referred to as single sensor signal structure TDOA (S^4 TDOA) localization (Fig. 1(a)). Commonly, single-element approaches using a single directional antenna take the directions in which local maximum power is received to be the DOA estimates [12]. Since directional antennas cannot simultaneously scan in all directions, some transient signals can escape detection and fluctuations of the source signal strength and polarization during the sequential lobing process may have a significant impact on the DOA accuracy. However, these problems are circumvented by the technique proposed in [13] which is applicable when information about the signal structure is a priori known (e.g. communication and radar emitters). The method does not require knowledge of the contents of the emitted signal. The information that is needed, is that the emitter sends message bursts at a known repetition frequency. For the example of

GSM signals, the emitter sends data of a duration of $\approx 546.46 \mu\text{s}$ followed by a pause of $\approx 30.46 \mu\text{s}$. The knowledge of this repeated pattern of signal transmissions and pauses is used for the single sensor signal structure TDOA approach. In this paper, we assume the transmission on/off structure to be known but never assume the transmitted signal itself to be known to the estimator during the simulations and real data evaluation.

In [14], the single-element TDOA localization approach is extended by the key-idea of direct emitter localization. For an airborne scenario with a single stationary source, we introduce a novel direct localization approach based on the cross correlation function (CCF). Our simulation and experimental measurement results demonstrate that the proposed approach considerably outperforms the standard single-element localization approach. This approach is named as direct S^4 TDOA abbreviated with DS^4 TDOA (Fig. 1(b)).

The block diagram given in Fig. 1 depicts the different approaches in the measurement/localization steps. For the non-direct S^4 TDOA method, the signal received at each observation step is correlated with the reference signal. The maximum of this correlation yields the TOA of the signal. By differentiating TOAs of two observation steps, a TDOA measurement is obtained (first step: measurement step). These TDOA measurements are used in the localization step (2nd step). The direct method DS^4 TDOA omits the TDOA estimation step and the correlation functions are input to the localization algorithm (Direct Position Determination, single step localization).

This paper is based on the work presented in [14]. Two slightly different (D) S^4 TDOA are described. The newly introduced methods are called (D) S^4 TDOA* and do not rely on the explicit representation of the signal structure using a reference signal. All four (D) S^4 TDOA(*) approaches are compared using the Cramér-Rao Lower Bound CRLB and Monte-Carlo simulations.

This paper is organized as follows: In Section II, the considered localization problem is stated. The Cramér-Rao Lower Bounds on TDOA estimation and on TDOA-based emitter localization are described in Section III. In Section IV, we briefly review the S^4 TDOA localization approach based on the CAF [13] as well as the direct version DS^4 TDOA [14] and introduce the two novel (D) S^4 TDOA* approaches. Monte-Carlo simulations and the comparison to the Cramér-Rao Lower Bound are shown in Section V. Simulation results for a real data scenario comparing S^4 TDOA and DS^4 TDOA approaches are presented in Section VI. In Section VII, the experimental measurement results proof the concept. Finally, the conclusions are given in Section VIII.

The following notations are used throughout this paper: $f[k]$ is a discrete version of the function $f(t)$, $f^*[k]$ is the conjugate complex of the function $f[k]$,




Fig. 2. Three-dimensional localization scenario.

$f^{(\tau)}[k]$ denotes the sampled version of $f(t - \tau)$ and $(\cdot)^T$ denotes transpose.

II. PROBLEM FORMULATION

We consider an omni-directional antenna sensor mounted on an airborne platform moving along an arbitrary but known sensor path observing a single stationary, ground-located source at position \mathbf{x} . The target emits a coherent signal $s(t)$ which is built up by times, where information is transmitted and pause intervals between those transmissions. The duration of each transmissions and each pause intervals is assumed to be constant and known. For example in the case of a communication signals, the information is sent as bursts during the transmission time and the pause times are guard periods between consecutive bursts. The exact modulation method or the content of the transmission bursts doesn't need to be known as long as a certain level of signal-to-noise ratio results from the transmission.

During the movement, the sensor collects N signal data batches. The n th received signal at some measurement point reads

$$z_n(t) = a_n s(t - t_{e,n} - t_n) \exp(j\nu_n t) + w_n(t), \quad (1)$$

where a_n denotes a path attenuation factor, $t_{e,n}$ denotes the unknown signal emission time of the n th received signal, t_n denotes the time difference between signal emission and signal acquisition, ν_n is the signal Doppler shift induced by the movement of the own sensor platform, and w_n denotes some additional receiver noise, $n = 1, \dots, N$. The transmitted signal $s(t)$ and the received signals $z_n(t)$ are assumed to be complex base-band signals.

In practice, the sensor collects data samples from the received signal. In the considered scenario, the sampling rate is assumed to be high enough that the sensor location \mathbf{r}_n is approximately constant for each collected

data batch (Fig. 2). Then, t_n and ν_n are given by

$$t_n(\mathbf{x}) = \frac{\|\Delta \mathbf{r}_n(\mathbf{x})\|}{c}, \quad (2)$$

$$\nu_n(\mathbf{x}) = \frac{\mathbf{v}_n^T \Delta \mathbf{r}_n(\mathbf{x}) f_0}{\|\Delta \mathbf{r}_n(\mathbf{x})\| c}, \quad (3)$$

respectively, where $\Delta \mathbf{r}_n(\mathbf{x}) = \mathbf{x} - \mathbf{r}_n$ denotes the relative vector between sensor and source, \mathbf{v}_n is the sensor velocity vector, c is the signal propagation speed, and f_0 is the center frequency of the emitted signal.

Considering the time-discrete version of the received signal in (1), the k th data sample of the n th data batch is given by

$$z_n[k] = a_n s(k \Delta - \text{clk}_n - \tau_n) \exp(j\nu_n k \Delta) + w_n[k], \quad (4)$$

where Δ is the sample interval, clk_n is the known sensor clock of the n th measurement, τ_n is the signal time of arrival relative to the sensor clock. Please note that $\text{clk}_n + \tau_n = t_{e,n} + t_n$ holds. The additional receiver noise w_n is assumed to be temporally uncorrelated and zero-mean Gaussian.

For the single sensor TDOA estimation, a reference signal $\tilde{s}[k]$ is used which characterizes the repetition pattern of the transmitted signal. For the ideal case, the emitted signal would be known and thus $\tilde{s}[k] = s[k]$. Since usually, the emitted signal is unknown for almost all applications, the reference signal we employ throughout this paper only characterizes the transmission on/off pattern of the emitted signal, which basically results in a comparison of the amplitudes of the received signal with the reference signal. If the emitted signal were known, much better localization performance could be achieved. Throughout this paper, we assume the transmission and guard interval periods to be known. The method doesn't require knowledge of the contents of the emitted signal and we never assume the signal to be known during the simulations or real data evaluation (where in fact, we don't know the emitted signal).

Finally, the localization problem is stated as follows: Estimate the source location \mathbf{x} from the received signal data batches $\mathbf{z}_n = (z_n[1], \dots, z_n[K])^T$, $n = 1, \dots, N$.

III. CRAMÉR-RAO LOWER BOUND

The CRLB provides a lower bound on the estimation accuracy and its parameter dependencies reveal characteristic features of the estimation problem. The parameters to be estimated from the measurements $\mathbf{z} = (\mathbf{z}_1^T, \dots, \mathbf{z}_N^T)^T$ are given by the vector \mathbf{x} . In this case, the CRLB is related to the covariance matrix \mathbf{C} of the estimation error $\Delta \mathbf{x} = \mathbf{x} - \hat{\mathbf{x}}(\mathbf{z})$ of any unbiased estimator $\hat{\mathbf{x}}(\mathbf{z})$ as

$$\mathbf{C} = \text{E}\{\Delta \mathbf{x} \Delta \mathbf{x}^T\} \geq \mathbf{J}^{-1}(\mathbf{x}), \quad (5)$$

where the inequality means that the matrix difference is positive semidefinite and \mathbf{J} is the Fisher Information

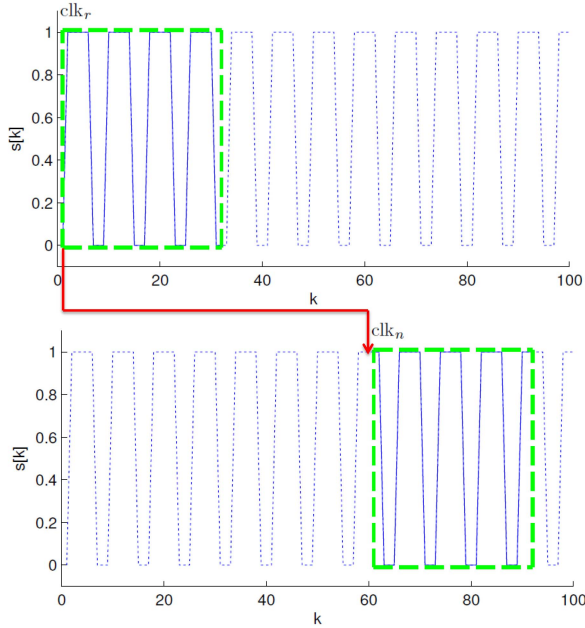


Fig. 3. Schematic representation of the received signals at measurement step r and n .

Matrix (FIM) given by

$$\mathbf{J}(\mathbf{x}) = \mathbb{E} \left\{ \left(\frac{\partial \mathcal{L}(\mathbf{z}; \mathbf{x})}{\partial \mathbf{x}} \right) \left(\frac{\partial \mathcal{L}(\mathbf{z}; \mathbf{x})}{\partial \mathbf{x}} \right)^T \right\}, \quad (6)$$

where \mathcal{L} denotes the log-likelihood function. If the estimator attains the CRLB then it is called *efficient*. The CRLB is given by the inverse Fisher Information.

In the following sections, CRLB for the TDOA estimation (Section III-A) as well as for the localization problem (Section III-B) are described.

A. CRLB on TDOA Estimation

The problem of estimating the TDOA τ of received signals can be stated as follows. From a set of signals $z_n(t)$, estimate the corresponding TDOAs. A lower bound on the achievable accuracy of this estimation process is essential for the calculation of the achievable localization accuracy described in Section III-B.

One of the most cited publication in this field is the work of Stein [15]. A CRLB on TDOA/FDOA estimation for the acoustic case is derived based on signal parameters like the bandwidth, the integration time and the signal-to-noise ratio of the received signals. The signals are assumed to be stationary Gaussian random processes. A more clear distinction between acoustic and electromagnetic signals was for example introduced in [5]. A more generalized bound for deterministic unknown signals was presented in [20]. The bound is determined using a given realization of the emitted signal itself. We use the Fisher Information \mathbf{J}_τ for the CRLB on emitter localization in the following section.

B. CRLB on Emitter Localization

The Fisher information for the TDOA-based localization problem is

$$\mathbf{J}(\mathbf{x}) = \sum_m \left(\left(\frac{\partial \tau_m}{\partial \mathbf{x}} \right) \mathbf{J}_\tau \left(\frac{\partial \tau_m}{\partial \mathbf{x}} \right)^T \right), \quad (7)$$

where m gives the TDOA measurement index and $\sigma_{\tau_m}^2$ is the variance of the corresponding TDOA estimation process modeled by the CRLB found in [20].

The CRLB for the position estimation accuracy is then given by

$$\sigma_{\mathbf{x}}^2 = \text{tr}(\mathbf{J}^{-1}(\mathbf{x})). \quad (8)$$

If TDOA measurements are temporally and spatially uncorrelated, the addition of the Fisher information of different measurement steps is possible.

IV. LOCALIZATION APPROACHES

In this section, approaches for the stated localization problem are presented i.e. the localization of a source with periodic coherent emission using a single moving sensor. Four measurement approaches for the single observer scenario are considered.

Firstly, an approach S^4 TDOA based on TOA measurements is presented (Section IV-A). Then, a derivation of this method called DS^4 TDOA based on the CCF and direct position determination is presented in Section IV-B. For both methods, approaches without the use of a representation of the reference signal $\tilde{s}[k]$ are introduced in Section IV-D. Those methods are called S^4 TDOA* and DS^4 TDOA* respectively. For the sake of simplicity, only TOA/TDOA measurements are considered in following and the Doppler is neglected. Nevertheless, the following techniques could be generalized to full CAF.

A. Two-step S^4 TDOA Approach [13]

Step 1: Commonly for a sensor network, TDOA measurements are extracted from the CCF

$$\text{CCF}(\Delta\tau) = \sum_{k=1}^K z_1^*[k] z_2^{(-\Delta\tau)}[k], \quad (9)$$

i.e. from the correlation of the two signals $z_1[k]$ and $z_2[k]$ in time domain. The TDOA estimates are calculated by detecting the peak in the CAF:

$$\hat{\Delta\tau} = \arg \max_{\Delta\tau} \text{CCF}(\Delta\tau). \quad (10)$$

However, since a single moving sensor is considered, the measurements are not taken simultaneously. Thus in the following, the signal processing for the single sensor case is presented (Fig. 1(a)). Due to the known signal structure, a quasi-TDOA measurement can be computed by considering the individual known sensor clock clk_n :

$$\hat{\tau}_n = \arg \max_{\tau} \text{CCF}_n(\tau) \quad (11)$$

with

$$\text{CCF}_n(\tau) = \sum_{k=1}^K z_n^*[k] \tilde{s}^{(-\text{clk}_n - \tau)}[k], \quad (12)$$

where $\tilde{s}[k]$ denotes a reference signal introduced in (4). Then similar to (10), a quasi-TDOA measurement can be calculated by taking the clock differences (Fig. 3)

$$\Delta \text{clk}_{n,r} = \left(\left\lfloor \frac{\text{clk}_n - \text{clk}_r}{T} \right\rfloor - \left\lfloor \frac{\text{clk}_n - \text{clk}_r}{T} \right\rfloor \right) T \quad (13)$$

into account, where the index r indicates some reference time. Then a quasi-TDOA measurement can be extracted from the TOA estimates by

$$\Delta \hat{\tau}_{n,r} = \hat{\tau}_n - (\hat{\tau}_r + \Delta \text{clk}_{n,r}), \quad (14)$$

i.e. by the difference of the individual estimated TOAs corrected by the clock difference. The correction of the clock difference is mandatory because the measurements are not taken simultaneously.

Step 2: The emitter localization problem can be solved by searching the emitter location that most likely explains the TDOA measurements calculated in (14). Therefore, the emitter location can be calculated by solving the following least-squares form:

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x}} \sum_{\substack{n=1 \\ n \neq r}}^N \frac{\|\Delta \hat{\tau}_{n,r} - \Delta \tau_{n,r}(\mathbf{x})\|^2}{\sigma_{\Delta \tau, n, r}^2}, \quad (15)$$

where $\Delta \tau_{n,r}(\mathbf{x})$ denotes the measurement function given analog to (14) by

$$\Delta \tau_{n,r}(\mathbf{x}) = t_n(\mathbf{x}) - (t_r(\mathbf{x}) + \Delta \text{clk}_{n,r}), \quad (16)$$

according to (2) and $\sigma_{\Delta \tau, n, r}^2$ denote the TDOA measurement variance, $n = 1, \dots, N$. The solution in (15) can be geometrically interpreted as the intersection of the hyperbolae represented by the individual TDOA measurements.

At this point it is worth to mention that in practice, the measurement variances are unknown and vary during the time. Consequently, the measurement variance have to be estimated because otherwise one could use an estimator with a reduced performance.

B. One-step DS⁴TDOA Approach [14]

The key-idea of direct localization approaches is to avoid the decision for one TOA/TDOA/AOA measurement in the first step of a localization algorithm. In the case of the S⁴TDOA method as described in the previous section, this decision is represented by the process of maximum determination of the CCF. The choice will always falls on the highest peak of the CCF, but when taking into account all measurement batches, this peak might be wrong. In this case, f.e. the second highest peak of the CCF would correspond to the sensor emitter geometry and fit all other measurement batches. Thus, leaving this decision open, allows the implicit evaluation

of multiple measurement hypotheses in one localization step.

The intention is to create a cost function that has to be optimized in the localization step, which takes into account all measurement batches at the same time without the explicit decision for TDOAs (Fig. 1(b)). By calculating the CCF of the CCF_r (CCF of \mathbf{z}_r and $s[k]$) and the CCF_n (CCF of \mathbf{z}_n and $s[k]$), the choice for an explicit TDOA can be postponed into the localization step. We call this approach direct single-sensor signal structure TDOA localization (DS⁴TDOA).

The choice of this approach is motivated by the scheme used for the multi-sensor TDOA localization, where the TDOA is not explicitly chosen in the first step but the TDOA measurement function is directly used as input for the localization step [1], [10]. Instead of a TDOA estimation from two received signals, DS⁴TDOA obtains the TDOA from two TOA measurements. The equivalent DS⁴TDOA then relies on the CCF of the TOA measurement functions, which are the CCFs of the received signals with the reference signal.

The proposed cost function (cross correlation of cross correlation functions) subject to the position \mathbf{x} is defined as

$$\text{CCCF}_{n,r}(\mathbf{x}) = \sum_{k=1}^K \text{CCF}_n^*[k] \text{CCF}_r^{(-\Delta \tau_n(\mathbf{x}))}[k], \quad (17)$$

with the CCF given in (12). The localization problem is then stated by

$$\hat{\mathbf{x}} = \arg \max_{\mathbf{x}} \sum_{\substack{n=1 \\ n \neq r}}^N \text{CCCF}_{n,r}(\mathbf{x}). \quad (18)$$

C. Discussion

The localization accuracy for both methods may degrade, if the distance between two observer positions is too big compared to the signal repetition duration T .

If $\Delta \tau_n(\mathbf{x}) \geq T/2$ the wrong peak may be chosen in the maximum determination of the CCFs in the case of S⁴TDOA. This choice has a direct effect on the localization accuracy using S⁴TDOA.

The influence on the localization for DS⁴TDOA is smaller if $\Delta \tau_n(\mathbf{x}) < T/2$ for almost all n . If $\Delta \tau_n(\mathbf{x}) \geq T/2$ for a significant number of measurements, the optimization of the localization function (18) may run into the maximum that corresponds to the wrong time slots. However this is unlikely, because the ambiguities that are due to $\Delta \tau_n(\mathbf{x}) \geq T/2$ are unlike to join in the same spatial position unless more than one emitter is present.

D. (D)S⁴TDOA without the use of $\tilde{s}[k]$

Both approaches described in the previous sections use an additional signal $\tilde{s}[k]$ representing the information on the signal structure. This allows data reduction for the localization step. If processing power and

data storage capacity and—in case of the use of the methods with multiple sensors—communication bandwidth is not an issue, the received signals can be stored and used for the localization process. In this case, instead of using TOA estimates calculated using $\tilde{s}[k]$ for S⁴TDOA and the CCCF for DS⁴TDOA, the cross correlation function of two received signals is used to estimate the TDOA or in the cost function of the direct method respectively. An additional shift factor according to the signal repetition interval and the observation time span has to be taken into account. We call these methods S⁴TDOA* and DS⁴TDOA*.

S⁴TDOA*:

The first step of the localization process of S⁴TDOA* is to calculate the maximum of the cross correlation function of two received signals at different time steps n, r :

$$\text{CCF}_{n,r}(\tau) = \sum_{k=1}^K z_n^*[k] z_r^{(-\text{clk}_{n,r}-\tau)}[k]. \quad (19)$$

The TDOA measurement is then given by

$$\hat{\tau}_{n,r} = \arg \max_{\tau} \text{CCF}_{n,r}(\tau). \quad (20)$$

In the second step, the emitter position is estimated by solving (15).

DS⁴TDOA*:

Similar to (17), the cost function for DS⁴TDOA* is given by

$$\text{CCF}_{n,r}(\mathbf{x}) = \sum_{k=1}^K z_n^*[k] z_r^{(-\Delta\tau_n(\mathbf{x}))}[k]. \quad (21)$$

The localization problem is then stated by

$$\hat{\mathbf{x}} = \arg \max_{\mathbf{x}} \sum_{\substack{n=1 \\ n \neq r}}^N \text{CCF}_{n,r}(\mathbf{x}). \quad (22)$$

For the evaluation of the real measurement data in this paper (sections VI and VII), S⁴TDOA* and DS⁴TDOA* are not applicable since processing power and storage capacity were limited. In the theoretical simulation and the CRLB evaluation (see section V), all four (D)S⁴TDOA^(*) methods are compared.

V. LOCALIZATION ACCURACY EVALUATION

A. Simulation Setup

To evaluate the four presented (D)S⁴TDOA^(*) localization approaches, Monte-Carlo simulations and CRLB analysis are conducted. A 2-dimensional scenario is investigated where one observer moves along a trajectory from west to east as depicted in Fig. 4. For each observation time step $n = 1 \dots 12$, a signal $s_n[k]$ that is emitted

from the target is simulated. We assume free space path loss

$$\text{FSPL}_{\text{dB}} = 10 \log_{10} \left(\frac{4\pi \|\Delta \mathbf{r}_n(\mathbf{x})\|}{\lambda} \right)^2 \quad (23)$$

and, by taking the receiver sensitivity S_{dB} into account, calculate the corresponding SNR

$$\text{SNR}_{\text{dB(n)}} = (P_E + G_E + G_R - \text{FSPL}_{\text{dB}}) - S_{\text{dB}}, \quad (24)$$

where λ is the wavelength of the signal, P_E is the transmitter power, G_E and G_R are antenna gain of the emitter and receiver antennas. The received signal is delayed by the time $t_n(\mathbf{x})$ the signal took to travel from the emitter to the observer according to (2).

The signals for each observation step are simulated as complex valued base-band signals at a sample rate of $f_s = 400$ kHz using the following parameters. The duration of each observed signal is $T = 1$ ms composed of repeated data transmission $T_{\text{data}} = 50 \mu\text{s}$ and guard periods with duration $T_{\text{guard}} = 10 \mu\text{s}$. During the time of data transmission, the emitted signal consists of a chirp signal with bandwidth $B = 200$ kHz. During the guard periods, no data is transmitted. White Gaussian noise is then added to the signal according to the SNR calculated using (24). The noise power is determined over the whole observation bandwidth of 400 KHz. The parameters for the path loss calculation are $G_E = 3$ dB, $G_R = 0$ dB, $S_{\text{dB}} = -90$ dBm at a center frequency of 1800 MHz. The transmission power P_E is varied for different evaluations. The received signal is then given by $z_n[k]$.

The reference signal $\tilde{s}[k]$ which is used by (D)S⁴TDOA has the same duration as the simulated received signal. The data transmission period starts with the first sample of $\tilde{s}[k]$ and the same reference signal is used at each time step. Since the TDOA and position estimation using (D)S⁴TDOA in this special realization rely on the amplitude comparison by correlating the received signals z_n with the reference signal \tilde{s} , the reference signal can be modeled as a real valued signal with $\tilde{s}[k] = 1$ during data transmission periods and $\tilde{s}[k] = 0$ during guard periods.

TDOA measurements are taken between consecutive observation steps resulting in a total of $N/2$ TDOA measurements. The measurement set for each measurement index $m = 1, \dots, N/2$ is given by $\{\hat{\tau}_{1,2}, \dots, \hat{\tau}_{2m-1,2m}\}$.

The cost functions of (D)S⁴TDOA^(*) are maximized using Nelder Mead simplex optimization under the assumption of constant TDOA measurement variance. The initial position estimate for the optimization process is calculated by evaluating the cost functions on a grid of possible emitter positions. The grid points are spaced by 500×500 m. For the first TDOA measurement, no position estimate is given, since the emitter location is not observable with only one TDOA measurement.

Simulations with 500 Monte-Carlo runs are conducted. For each run, the emitter position is chosen uniformly at random from an area of interest (AOI)_x =




Fig. 4. Scenario used for localization accuracy analysis including results of (D)S⁴TDOA. Zoom of target area shows only (D)S⁴TDOA* results. Transmission power $P_E = 30$ dBm. CRLB is depicted as 3σ ellipse.




Fig. 5. Comparison of localization approaches to CRLB (transmission power $P_E = 18$ dBm).

$-5000, \dots, 5000$ m, $AOI_y = 1000, \dots, 8000$ m). The position is estimated using all four (D)S⁴TDOA^(*) methods and the corresponding localization CRLB is calculated according to (8).

B. Results

The simulations are carried out for different transmission powers. Fig. 5 shows the results for $P_E = 18$ dBm. For many emitter positions throughout the area of interest, this results in low SNR values. Both direct localization approaches are more robust against low SNR, since ambiguities in the cross correlation functions have less effect on the localization. The two step localization methods need to chose one TDOA measurement in the first step independently of all other observation steps whereas the direct technique postpones




Fig. 6. Comparison of localization approaches to CRLB (transmission power $P_E = 30$ dBm).

this decision into the localization step, where all measurements are incorporated (see also Section IV-C). The similar accuracy of DS⁴TDOA and DS⁴TDOA* is due to the high repetition rate of transmission and guard periods. The correlation of the reference signal, having very high SNR, and the received signal with low SNR, still shows good cross correlation characteristics. The position estimation accuracy of S⁴TDOA is out of the scale of Fig. 5. The performance of S⁴TDOA* improves until measurement index 4 and then degrades again. This is due to the fact that the mean SNR for the given trajectory and randomized emitter positions from the area of interest is often lower at the last observation points and thus the probability of choosing a wrong peak of the CCF increases.

By increasing the signal transmission power to $P_E = 30$ dBm and thus having higher SNR, the performance of S⁴TDOA* is very similar to DS⁴TDOA* for all measurement steps. The results are depicted in Fig. 6. Again, DS⁴TDOA outperforms S⁴TDOA, which shows the lowest localization accuracy.

To show the distribution of the position estimates of all four (D)S⁴TDOA^(*) methods, a fixed emitter position is chosen. For this scenario, again 500 Monte-Carlo runs are conducted. The results of TDOA measurement index 6 are depicted in Fig. 4. The zoomed area shows only the estimates of S⁴TDOA* and DS⁴TDOA*. The CRLB is given by a 3σ error ellipse.

The advantages of (D)S⁴TDOA compared to (D)S⁴TDOA* are given by less need for storage space and a reduction of processing power (and communication requirements). A trade-off between localization accuracy and sensor requirements is possible using (D)S⁴TDOA.

VI. SIMULATION USING REAL DATA SCENARIO

The proposed localization approaches performances are evaluated in Monte-Carlo simulations for a given




Fig. 7. Simulation Scenario: Sensor trajectory and area of interest (green box).

scenario. GSM base stations are chosen as emitter with recurring signal structure. In this section, only S^4 TDOA and DS^4 TDOA algorithms are evaluated.

A. Simulation Setup

The desired signal is sent on the broadcast channel of a GSM base station and is divided into time slots. Each time slot has a duration of $576.92 \mu s$. A time slot is divided into data transmission time and guard period during which no transmission takes place. This time slot signal structure is represented by the reference signal $\tilde{s}[k]$ introduced in (4).

The sensor trajectory remains the same over all Monte-Carlo runs. The position of the emitter is chosen uniformly at random from a given area of interest. The localization accuracy is also evaluated w.r.t. the signal-to-noise ratio (SNR). A sensor trajectory that is similar to the one of the field experiments (Section VII) is used for the simulations. Fig. 7 shows the trajectory as well as the area of interest in which possible emitters are located.

We use the following definition of SNR for the simulations:

$$\text{SNR}_{[\text{dB}]} = 10 \log_{10} \frac{P_s}{P_n} \quad (25)$$

with P_s being the mean signal power and P_n the mean noise power. A total of 250 Monte-Carlo runs were performed. Each Monte-Carlo run consists of the following:

- 1) An emitter position is chosen at random from the area of interest.
- 2) A random start drift of the broadcast signal is generated.
- 3) Signal noise for each sensor is generated.

- 4) For the given observer trajectory and emitter position and time of measurement, corresponding TOAs are calculated.
- 5) The broadcast signal is embedded into noise in accordance to the respective TOAs and scaled to meet given SNR value.
- 6) Localization results are calculated using both estimation methods
 - a) The initialization is done by evaluating a grid of the respective cost functions for the area of interest.
 - b) The position is estimated using Nelder Mead simplex optimization.
- 7) Points 5 to 6 are repeated for all SNR values in question.

B. Position Estimation

The position estimation for the S^4 TDOA method is divided into two main steps. In the first step, the received signal is correlated with the stored reference signal. The maximum of this correlation function yields the TOA of each measurement. TOAs of two observation steps form one TDOA measurement. In the second step, the emitter position is estimated based on a set of TDOA measurements.

For the DS^4 TDOA localization, from the two observation steps that form the TDOA measurement in the above described case, cross correlate the cross correlation functions of the respective received signals and the reference signal. Estimate the emitter position from a set of those cross correlation functions.

For both methods, the respective cost functions are minimized using Nelder/Mead simplex optimization. The initialization problem is solved by evaluating the cost functions of each method for a grid over the area of interest. For the simulations, the grid points were spaced by 100×100 m.

C. Results

Fig. 8 depicts the results of the simulations. For each SNR value, the RMSE of the position estimation over all 250 simulation runs is calculated. The red line with red dots shows the RMSE using the S^4 TDOA, DS^4 TDOA is plotted using red diamonds. The accuracy of the DS^4 TDOA localization approach outperforms the S^4 TDOA localization method.

VII. EXPERIMENTAL RESULTS

A. Experimental Setup

Field experiments were conducted to verify the presented method for real data. A GPS time-synchronized sensor node was used to gather data from a GSM mobile station. The sensors receiving antenna was mounted under the wing of an aircraft. The sensor itself and a PC for data processing were installed inside the aircraft.




Fig. 8. Simulation results: Localization RMSE over SNR for DS⁴TDOA (red diamonds) and S⁴TDOA (red dots).

Every five seconds, data from the broadcast channel of the GSM base station was recorded at a sample rate of $f_s = 1$ MHz.

Along with the signal data, corresponding timestamps and position information from the GPS receiver of the sensor are recorded. For each observation time step n , the received signals are filtered and the CCF is calculated. From this CCF the TOA $\hat{\tau}_n$ of the signal is estimated as described in Section IV. The CCF, the estimated TOA, the sensors position and time are used in the localization step. The localization estimates for both methods are calculated using the same initialization for the optimization algorithm.

Fig. 9 depicts the sensors trajectory, the position of the GSM base station as well as the localization results using the presented S⁴TDOA and DS⁴TDOA method. The presented localization approach is evaluated for different levels of signal strength. Here, a threshold P_t is applied to the measurements. If the mean received signal strength $P_{z_n[k]}$ is below the threshold, the measurement is not used in the localization step. The mean signal strength of a signal $z_n[k]$ is defined as

$$P_{z_n[k]} = \frac{z_n^*[k]z_n[k]}{K} \quad (26)$$

with K being the total number of samples.

B. Results

Fig. 10 to Fig. 13 show the localization cost functions evaluated for a grid of possible emitter positions. The black line indicates the flight trajectory where the black dots indicate the measurements that are taken into account in the localization step according to the received signal strength threshold. The true position of the emitter is marked by a green dot. The position estimate of the S⁴TDOA method is shown by a yellow x, the respective DS⁴TDOA estimate by a red circle. The achieved localization accuracy is given in Table I.




Fig. 9. Scenario of field experiments. Sensor trajectory and localization results.

TABLE I
Localization accuracy in [m] of field experiments data.

RSS	S ⁴ TDOA	DS ⁴ TDOA	RSS	S ⁴ TDOA	DS ⁴ TDOA
-80	4049	385	-69	1003	129
-79	889	257	-68	2026	82
-78	681	449	-67	1542	190
-77	842	223	-66	186	166
-76	1203	50	-65	161	179
-75	1424	451	-64	170	268
-74	1272	89	-63	312	141
-73	1374	145	-62	458	404
-72	1623	116	-61	258	66
-71	1546	112	-60	199	119
-70	901	332			

As can be seen in Fig. 10, the minimum of the cost function of the S⁴TDOA for a received signal strength threshold level of $P_t = -74$ dBm is not located at the true emitter position due to the choice of one or more faulty TOA values (maximum peaks of the CCF). This results in a larger localization error. Here, the advantage of the DS⁴TDOA approach can be seen. Fig. 11 depicts the cost function of the DS⁴TDOA method for the same scenario. As can be observed, the minimum of the cost function is located near the true emitter position and the localization result is more accurate. For this scenario with a received signal strength threshold of $P_t = -74$ dBm, the 3-D localization error of the S⁴TDOA is 1272 m. Using the DS⁴TDOA localization algorithm, the position estimation error is 89 m.

The cost function of the S⁴TDOA and a received signal strength threshold of $P_t = -60$ dBm is shown in Fig. 12. Less measurements are used to localize the emitter, but due to the higher signal level, the choice of the peak of the CCF as TOA value tends towards the correct peak. With more accurate TDOA estimation, the localization result becomes more accurate. The cost function using the direct localization method (Fig. 13) is




Fig. 10. Normalized cost function for S^4 TDOA (signal threshold -74 dBm).




Fig. 12. Normalized cost function for S^4 TDOA (signal threshold -60 dBm).




Fig. 11. Normalized cost function for DS^4 TDOA (signal threshold -74 dBm).




Fig. 13. Normalized cost function for DS^4 TDOA (signal threshold -60 dBm).

very similar to the afore mentioned, also the localization results are nearly the same.

The localization accuracy for the S^4 TDOA improves from 1272 m ($P_t = -74$ dBm) to 199 m ($P_t = -60$ dBm). For the DS^4 TDOA location estimation method, a slight degradation of accuracy from 89 m ($P_t = -74$ dBm) to 119 m ($P_t = -60$ dBm) is noticed.

Fig. 14 shows the comparison of the localization errors of both methods over different signal strength levels. It can be observed, that the DS^4 TDOA method is more robust to smaller received signal strength and outperforms the S^4 TDOA-based method. As the TOA estimation using the signal structure information relies on the amplitude of the signal, with lower SNR, the TOA estimation becomes more and more noisy up until peaks that do not correspond to the signal are chosen as TOA. Since the DS^4 TDOA method does not require choosing one peak of the CCF, the localization results remain more stable for lower signal level values.

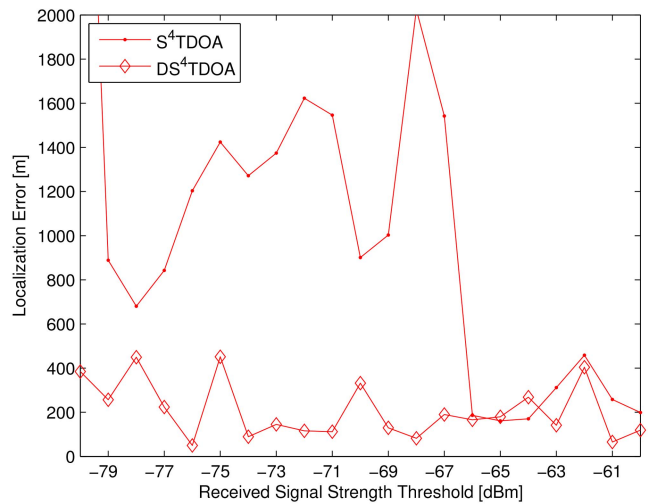


Fig. 14. Localization accuracy for different received signal strength thresholds of field experiment data.

C. Discussion

A large amount of localization error of a real world TDOA system can be caused by time and position inaccuracies of the sensors. In our experiments, we used GPS to determine the observers position during the flight. Especially the elevation estimation of a GPS receiver is known to be imprecise. Although we employed GPS disciplined oscillators, the time synchronization error might be the largest cause of localization error. In the case of stationary observers, a synchronization to UTC in the range of 25 ns is achievable. For in-flight use, the accuracy of the local clock can degrade up to 200 ns. Even though an exact time stamp is not necessary for the (D)S⁴TDOA^(*) methods, the employed experimental system allows only processing of one seconds of signal each five seconds. If continuous streaming of data is possible, a stable oscillator without exact time information is sufficient (the accuracy issue remains the same). Another real world error lies in the clock accuracy of the emitter which needs to be stable enough for all (D)S⁴TDOA^(*) methods to be applicable.

VIII. CONCLUSION

We evaluated four (direct) localization approaches for the use with a single moving sensor. The methods are based on the S⁴TDOA found in [13]. The direct localization solution DS⁴TDOA firstly introduced in [14] is derived in Section IV-B. The performance in means of emitter localization accuracy of S⁴TDOA and DS⁴TDOA are evaluated in simulations (Section VI). Field experiments dealing with the localization of GSM base stations using a single airborne sensor are presented in Section VII. Additionally, two methods (D)S⁴TDOA* that do not require the explicit representation of the signal structure are introduced in Section IV-D. All four approaches are evaluated in Monte-Carlo simulations and compared to the CRLB (Section V).

All presented methods allow emitter localization with a light weight and small sensor node. Only one reception channel combined with an omni-directional antenna is needed. The requirements on the communication channel bandwidth between sensor and situation display system are small. Even if the position estimate is not determined at the sensor node but is calculated at a control station on ground, for (D)S⁴TDOA only CCF and corresponding time and position information need to be transmitted. Classic TDOA approaches require the transmission of raw signal data to a reference sensor or control station, thus having higher demands on the communication channel.

If processing power and storage capacity is not a limiting factor, direct localization using DS⁴TDOA* is shown in simulations to give the best localization results.

The feasibility of determining the position of an emitter using (D)S⁴TDOA^(*) is shown. The proposed DS⁴TDOA^(*) direct localization is more robust to smaller SNR and outperforms the S⁴TDOA^(*) localization in both simulations and field experiments.

REFERENCES

- [1] A. Amar and A. J. Weiss
Direct geolocation of stationary wideband radio signal based on time delays and doppler shifts,
in *Proc. 15th Workshop on Statistical Signal Processing*, Cardiff, Wales, Sep. 2009, pp. 101–104.
- [2] K. Becker
An efficient method of passive emitter location,
IEEE Trans. Aerosp. Electron. Syst., vol. 28, pp. 1091–1104, Oct. 1992.
- [3] ———
Target Motion Analysis (TMA),
in *Advanced Signal Processing Handbook*, S. Stergioulas, Ed. New York, NY: CRC Press, 2001, ch. 9, pp. 284–301.
- [4] M. L. Fowler
Analysis of single-platform passive emitter location with terrain data,
IEEE Trans. Aerosp. Electron. Syst., vol. 37, pp. 495–507, Apr. 2001.
- [5] M. L. Fowler and X. Hu
Signal models for TDOA/FDOA estimation,
IEEE Trans. Aerosp. Electron. Syst., vol. 44, pp. 1543–1550, Oct. 2008.
- [6] R. Kaune, D. Musicki, and W. Koch
On passive emitter tracking in sensor networks,
in *Sensor Fusion and its Applications*, C. Thomas, Ed. InTech, 2010, ch. 13, pp. 293–318.
- [7] Z. Madadi, F. Quitin, and W. P. Tay
Periodic rf transmitter geolocation using a mobile receiver,
in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, South Brisbane, QLD, Australia, Apr. 2015, pp. 2584–2588.
- [8] M. Oispuu
Direct state determination of multiple sources with intermittent emission,
in *Proc. 17th European Signal Processing Conference*, Glasgow, Scotland, Aug. 2009, pp. 1948–1952.
- [9] M. Oispuu and U. Nickel
Direct detection and position determination of multiple sources with intermittent emission,
Signal Processing, vol. 90, pp. 3056–3064, Dec. 2010.
- [10] M. Pourhomayoun and M. L. Fowler
Distributed computation for direct position determination emitter location,
IEEE Trans. Aerosp. Electron. Syst., vol. 50, pp. 2878–2889, Oct. 2014.
- [11] F. Quitin, Z. Madadi, and W. P. Tay
Rf transmitter geolocation based on signal periodicity: Concept and implementation,
in *Proc. IEEE International Conference on Communications (ICC)*, London, England, Jun. 2015, pp. 4593–4598.
- [12] S. V. Schell and W. A. Gardner
High-resolution direction finding,
in *Handbook of Statistics, Signal Processing and its Applications*, N. K. Bose and C. R. Rao, Eds. Elsevier, 1993, vol. 10, ch. 18, pp. 755–817.
- [13] C. Steffes
Novel TDoA-based single sensor localization approach using signal structure information,
in *Proc. 12th Workshop on Positioning, Navigation and Communication (WPNC)*, Dresden, Germany, Mar. 2015.

- [14] C. Steffes and M. Oispuu
Direct Single Sensor TDOA Localization Using Signal Structure Information, in *18th International Conference On Information Fusion (FUSION)*, Washington, D.C., USA, July 2015.
- [15] S. Stein
Algorithms for ambiguity function processing, *IEEE Trans. Acoust., Speech, Signal Processing*, vol. 29, pp. 588–599, Jun. 1981.
- [16] D. J. Torrieri
Statistical theory of passive location system, *IEEE Trans. Aerosp. Electron. Syst.*, vol. 20, pp. 183–198, Mar. 1984.
- [17] E. Tzoref, B. Z. Bobrovsky, and A. J. Weiss
Single receiver emitter geolocation based on signal periodicity with oscillator instability, *IEEE Trans. Signal Processing*, vol. 62, pp. 1377–1385, Feb. 2014.
- [18] A. J. Weiss
Direct position determination of narrowband radio frequency transmitters, *IEEE Signal Processing Lett.*, vol. 11, pp. 513–516, May 2004.
- [19] A. J. Weiss and A. Amar
Direct position determination of multiple radio transmitters, in *Advances in Direction-of-Arrival Estimation*, S. Chandran, Ed. Norwood, MA: Artech House, 2006, ch. 11, pp. 213–239.
- [20] A. Yeredor and E. Angel
Joint tdoa and fdoa estimation: A conditional bound and its use for optimally weighted localization, *IEEE Trans. Signal Processing*, vol. 59, pp. 1612–1623, Apr. 2011.



Christian Steffes received his Diploma in Computer Science at University of Bonn, Germany. Since 2010 he works at Fraunhofer Institute for Communication, Information Processing and Ergonomics FKIE, Wachtberg, Germany, in the department Sensor Data and Information Fusion. His research is focused on passive emitter localization.



Marc Oispuu received the diploma degree in electrical engineering from the RWTH Aachen, Germany, in 2005. Currently he is a research assistant with the Department of Sensor Data and Information Fusion at Fraunhofer Institute for Communication, Information Processing and Ergonomics FKIE, Wachtberg, Germany. Additionally he achieved his Ph.D. degree at the Center for Sensor Systems, University of Siegen, in 2013. His research interests focus on passive emitter localization and advanced array signal processing.

INTERNATIONAL SOCIETY OF INFORMATION FUSION

ISIF Website: <http://www.isif.org>

2015 BOARD OF DIRECTORS*

2013–2015	2014–2016	2015–2017
Jean Dezert	Sten F. Andler	Joachim Biermann
Gee-Wah Ng	Murat Efe	Frederik Gustafsson
Anne-Laure Jousset	Lyudmila Mihaylova	Jesús García

*Board of Directors are elected by the members of ISIF for a three year term.

PAST PRESIDENTS

Darin Dunham, 2015	Elisa Shahbazian, 2009	Xiao-Rong Li, 2003
Darin Dunham, 2014	Darko Musicki, 2008	Yaakov Bar-Shalom, 2002
Wolfgang Koch, 2013	Erik Blasch, 2007	Pramod Varshney, 2001
Roy Streit, 2012	Pierre Valin, 2006	Yaakov Bar-Shalom, 2000
Joachim Biermann, 2011	W. Dale Blair, 2005	Jim Llinas, 1999
Stefano Coraluppi, 2010	Chee Chong, 2004	Jim Llinas, 1998

SOCIETY VISION

The International Society of Information Fusion (ISIF) is the premier professional society and global information resource for multidisciplinary approaches for theoretical and applied information fusion technologies.

SOCIETY MISSION

Advocate

To advance the profession of fusion technologies, propose approaches for solving real-world problems, recognize emerging technologies, and foster the transfer of information.

Serve

To serve its members and engineering, business, and scientific communities by providing high-quality information, educational products, and services.

Communicate

To create international communication forums and hold international conferences in countries that provide for interaction of members of fusion communities with each other, with those in other disciplines, and with those in industry and academia.

Educate

To promote undergraduate and graduate education related to information fusion technologies at universities around the world. Sponsor educational courses and tutorials at conferences.

Integrate

Integrate ideas from various approaches for information fusion, and look for common threads and themes—look for overall principles, rather than a multitude of point solutions. Serve as the central focus for coordinating the activities of world-wide information fusion related societies or organizations. Serve as a professional liaison to industry, academia, and government.

Disseminate

To propagate the ideas for integrated approaches to information fusion so that others can build on them in both industry and academia.

Call for Papers

The Journal of Advances in Information Fusion (JAIF) seeks original contributions in the technical areas of research related to information fusion. Authors are encouraged to submit their manuscripts for peer review <http://isif.org/journal>.

Call for Reviewers

The success of JAIF and its value to the research community is strongly dependent on the quality of its peer review process. Researchers in the technical areas related to information fusion are encouraged to register as a reviewer for JAIF at <http://jaif.msubmit.net>. Potential reviewers should notify via email the appropriate editors of their offer to serve as a reviewer.