

Filtering on the Unit Sphere Using Spherical Harmonics

Florian Pfaff*, Gerhard Kurz*, and Uwe D. Hanebeck*

Abstract—For manifolds with topologies that strongly differ from the standard topology of \mathbb{R}^n , using common filters created for linear domains can yield misleading results. While there is a lot of ongoing research on estimation on the unit circle, higher-dimensional problems particularly pose a challenge. One important generalization of the unit circle is the unit hypersphere. In this paper, we propose a recursive Bayesian estimator for the unit sphere S^2 based on spherical harmonics for arbitrary likelihood functions and rotationally symmetric system noises. In our evaluation, the proposed filter outperforms the particle filter in a target tracking scenario on the sphere.

I. INTRODUCTION

In the rapidly evolving field of recursive Bayesian estimation, the linear topology of \mathbb{R}^n is often taken for granted. However, many frequently occurring phenomena feature an underlying periodic manifold and can only be modeled accurately by taking the periodicity into account. The periodic manifold that is easiest to handle using recursive Bayesian estimation is the unit circle, which has gained considerable attention and for which fast and versatile filters [1], [2], [3], [4] have been developed. Another manifold that is very important in practice is the unit sphere. Among many others, applications include tracking on a sphere [5], speaker tracking using microphone arrays [6], and estimating the orientation of rotationally symmetric objects.

For estimation on the sphere, the von Mises–Fisher filter [5] and the Bingham filter [7] have been proposed. These filters assume that the densities involved are distributed according to a von Mises–Fisher or Bingham density, respectively. While nonlinear variants have been proposed [8], the applicability of the filters is limited by the densities that can be represented accurately using the respective distribution. While a particle filter [9] can easily be adapted to spherical domains, its usual disadvantages apply. A large number of particles are required in general and due to particle degeneracy, the number of particles that are of practical use is far lower than the total number of particles employed. Furthermore, the particle filter is non-deterministic and thus only yields acceptable results on average, even if high numbers of particles are used.

For circular manifolds, we have proposed filters that do not require that the density can be represented well using one of the parametric densities. One approach is to use grid filters [4] that are based on a discretization of the state space. Another approach is to use Fourier filters [3]. In the Fourier

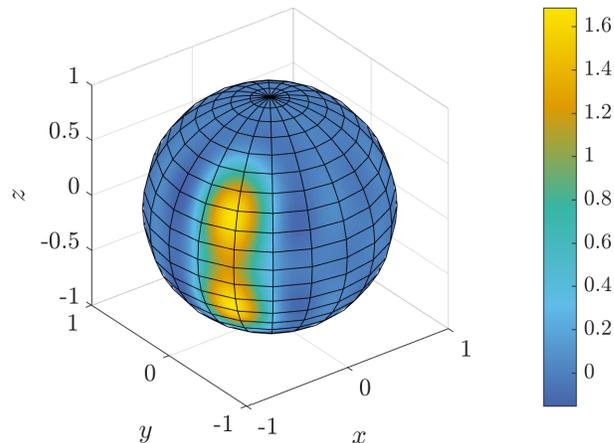


Fig. 1: An approximation of a bimodal, non-antipodally symmetric density function as obtained by the use of our filter with a maximum degree of 10.

filters, the density (or its square root) is approximated using a trigonometric polynomial [10].

In this paper, we adapt the idea of representing densities using a complete orthogonal sequence of functions to densities on the sphere by the use of spherical harmonics. Using spherical harmonics, we can model densities that are non-antipodally symmetric and multimodal (such as the one shown in Fig. 1), which pose a challenge to filters that only support specific densities. Spherical harmonics are well researched as they are useful for signal processing [11], weather modeling [12], physics [13], [14], chemistry [15], and computer graphics [16].

In the next section, we first give a review of the basics of spherical harmonics. In the third section, we introduce the prediction and update step of our spherical harmonics filter. Afterward, we evaluate our filter by comparing it with a particle filter on the sphere. Last, we provide a conclusion.

II. BASICS OF SPHERICAL HARMONICS

Spherical harmonics are a complete orthonormal sequence in the function space of square integrable functions on S^2 [11, Theorem 2.12]. When performing a spherical harmonics analysis, the function on the sphere is represented by a sum of spherical harmonics Y_l^m weighted by coefficients w_l^m according to [11, Sec. 2.9.4]

$$f(\theta, \phi) = \sum_{l=0}^L \sum_{m=-l}^l w_l^m Y_l^m(\theta, \phi),$$

in which m is called the order, l the degree, and L is the maximum degree used. To approximate arbitrary continuous

*Florian Pfaff, Gerhard Kurz, and Uwe D. Hanebeck are with the Intelligent Sensor-Actuator-Systems Laboratory (ISAS), Institute for Anthropomatics and Robotics, Karlsruhe Institute of Technology (KIT), Germany. florian.pfaff@kit.edu, gerhard.kurz@kit.edu, uwe.hanebeck@kit.edu

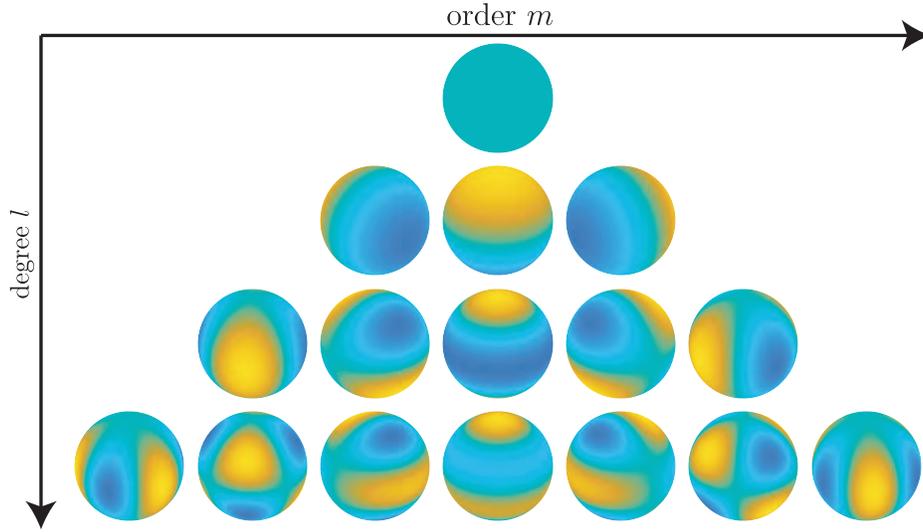


Fig. 2: Real spherical harmonics up to degree 3 (from high function values in yellow to low values in blue). The harmonics are listed in ascending order according to their degree (vertically) and order (horizontally). All spheres shown are unit spheres.

densities, $L = \infty$ is necessary, which results in an infinite series. For practical use, we only use coefficients up to a certain degree $L \in \mathbb{N}$.

Analogous to Fourier series, real and complex spherical harmonics exist. While the former can only be used to model real functions on the sphere, the latter can be used to represent both real and complex functions¹. For both representations, multiple conventions exist, especially due to the different fields of application in the individual communities.

In this paper, we use the complex spherical harmonics according to the convention [11, Ch 7.3.3, Eq. (7.19)]

$$Y_l^m(\theta, \phi) = N_l^m P_l^m(\cos(\theta)) e^{im\phi},$$

depending on the associated Legendre function P_l^m including the Condon–Shortley phase² defined by [11, Eq. (7.23-7.24)]

$$P_l^m(x) = \begin{cases} \frac{(-1)^m}{2^l l!} (1-x^2)^{m/2} \frac{d^{l+m}}{dx^{l+m}} (x^2-1)^l, & m \in \mathbb{Z}_0^+ \\ (-1)^m \frac{(l+m)!}{(l-m)!} P_l^{-m}(x), & m \in \mathbb{Z}^- \end{cases}$$

and the normalization constant [11, Eq. (7.20)]

$$N_l^m = \sqrt{\frac{2l+1}{4\pi} \frac{(l-m)!}{(l+m)!}}$$

that ensures that Y_l^m is normalized in regard to the L^2 norm.

To obtain the spherical harmonics coefficients of a square integrable function on the sphere $f(\theta, \phi)$ parametrized using

¹When representing a real function using complex spherical harmonics, the complex coefficients include some redundancy.

²The Condon–Shortley phase is an optional factor $(-1)^m$ that causes the signs of the spherical harmonics to alternate with increasing m for positive m [13, Ch. 15.5]. Some authors, e.g. [17], do not include it in the definition of the associated Legendre functions but rather explicitly include it in the definition of the spherical harmonics. The Condon–Shortley phase factor is not essential to the definition of the spherical harmonics but simplifies the calculations in certain applications.

spherical coordinates, we can calculate the integral

$$w_l^m = \int_0^{2\pi} \int_0^\pi f(\theta, \phi) \bar{Y}_l^m(\theta, \phi) \sin(\theta) d\theta d\phi,$$

with \bar{Y}_l^m being the complex conjugate of Y_l^m .

The unit sphere can be regarded as a two-dimensional manifold embedded in \mathbb{R}^3 . Many parametric densities are not given depending on spherical coordinates, but use a vector $\underline{x} \in \mathbb{R}^3$ instead. To be able to use such densities in the spherical harmonics framework, we use the formulae

$$\begin{aligned} x_1 = x &= \cos(\theta) \cos(\phi), \\ x_2 = y &= \cos(\theta) \sin(\phi), \\ x_3 = z &= \sin(\theta) \end{aligned}$$

to convert the spherical coordinates to Cartesian coordinates. Whenever we give a function parametrized by a vector in \mathbb{R}^3 , we implicitly assume that a coordinate transformation is performed, if necessary.

For each degree l , there are $2l+1$ spherical harmonics, which we index using the order $m \in \{-l, \dots, l\}$. We illustrate the increasing number of spherical harmonics for increasing degree by showing the real spherical harmonics up to degree 3 in Fig. 2. When using all spherical harmonics up to a maximum degree L , there are a total of $(L+1)^2$ spherical harmonics and spherical harmonics coefficients. As can be seen in Fig. 3, the approximation of the density improves with every additional degree. Especially densities that are highly concentrated and multimodal (and are thus particularly “peaky”) can benefit from using higher degrees.

It should be noted that even if a valid density is approximated, the approximation can attain negative function values unless a transformation such as the square root as applied in [3], [18], [19] is used. Examples for approximations that have negative function values can be seen in Fig. 1 and Fig. 3. Approximations of densities with regions close to 0 and approximations using only few coefficients are more

susceptible to this effect. We disregard this effect as we observed in [19] that it can also make sense to invest the additional computational power required for the square root version of the filter into using more coefficients instead.

For the Fisher–Bingham distribution, formulae involving infinite sums to calculate the spherical harmonics coefficients have been proposed [20]. For arbitrary distributions, calculating the two-dimensional integral to obtain the coefficients may be too costly. Analogous to the FFT for trigonometric polynomials, fast algorithms to calculate spherical harmonics coefficients from function values on a grid have been proposed. Important insights were obtained in [21] and a fast algorithm with a computational complexity of $O(L^2 \log L)$ for a maximum degree L was presented in [22]. In our implementation, we use an algorithm implemented in Matlab as used in [23], which we have tweaked for higher performance.

III. THE SPHERICAL HARMONICS FILTER

For a recursive Bayesian estimator, we have to implement the two required operations—namely, the update and the prediction step—that we lay out in the following subsections.

A. Update Step

Given the measurement likelihood, the update step at time step t can be formulated based on Bayes’ law as

$$f_t^e(\underline{x}_t | \underline{z}_1, \dots, \underline{z}_t) \propto f_t^{\text{LH}}(\underline{z}_t | \underline{x}_t) f_t^p(\underline{x}_t | \underline{z}_1, \dots, \underline{z}_{t-1}),$$

depending on the state vector \underline{x}_t and the measurements \underline{z} at the respective time steps, the likelihood function f_t^{LH} , the prior density f_t^p , and the posterior density f_t^e .

A useful identity regarding the multiplication of functions in a spherical harmonics representation is the equation [21, Theorem 2]

$$Y_{l_1}^{m_1} \cdot Y_{l_2}^{m_2} = \sum_{l_c=|l_1-l_2|}^{l_1+l_2} \sqrt{\frac{(2l_1+1)(2l_2+1)}{4\pi(2l_c+1)}} \cdot C_{0,0,0}^{l_1,l_2,l_c} \cdot C_{m_1,m_2,m_1+m_2}^{l_1,l_2,l_c} \cdot Y_{l_c}^{m_1+m_2} \quad (1)$$

depending on the Wigner symbol C (using the convention of [14, Eq. (3.165)]). This formula allows the insight that the result of the multiplication of two spherical harmonics can be represented using a weighted sum of multiple spherical harmonics. By regarding each pair of spherical harmonics separately, this formula can be used to determine (aside from numerical imprecision) the exact result of the multiplication of two functions represented by their spherical harmonics coefficients. As can be seen from the formula, even if both spherical harmonics are of degree l , spherical harmonics of degree $2l$ are required in general to represent the multiplication result.

To multiply two functions in their spherical harmonics representations, we can first expand the product of the two sums with $O(L^2)$ terms each. This leads to a sum of $O(L^2 \cdot L^2)$ multiplied pairs of spherical harmonics from the first and the second function with their corresponding coefficients. We can then apply (1) to represent the multiplication result of each pair multiplied using weighted spherical harmonics.

For each pair, we obtain $O(L)$ terms to be added to the coefficients of the spherical harmonics representation of the multiplication result. Thus, the total complexity of the multiplication performed like this is in $O(L^5)$. Due to the multiplication of pairs of spherical harmonics of degree L , we may have nonzero entries for degrees up to $2L$ in the multiplication result.

Thus, if we use this approach to derive the exact result, the degree of the approximation of the resulting density will increase with every update step, inducing the need for a truncation to prevent an ever-increasing number of coefficients. Moreover, while the formula can be used to calculate the result directly using the spherical harmonics coefficients, this approach is not cheap.

To avoid such a costly operation, we use the fast approximations of the spherical harmonics analysis and synthesis operations used by the authors of [23] in our implementation instead. We first use the fast spherical harmonics synthesis to calculate function values of f_t^p on a grid of points on the sphere from the spherical harmonics coefficients. We then proceed similarly to a grid filter and use the likelihood function directly to obtain the (unnormalized) function values of f_t^e on the same grid. For the next prediction step, we calculate the new spherical harmonics coefficients using the fast analysis operation. In total, the effort never exceeds $O(L^2 \log L)$ as required by the spherical harmonics analysis and synthesis operations. Because the synthesis and analysis operations only yield approximations and because no higher degrees than L are used to approximate the multiplication result, approximation errors are generally inevitable.

The normalization after the multiplication that is required for the update step can easily be performed. The integral of a real function given by its spherical harmonics coefficients over the entire unit sphere only depends on the coefficient of degree zero w_0^0 as all harmonics except Y_0^0 cancel out in the integral. As $Y_0^0(\theta, \phi) = 1/\sqrt{4\pi}$ [11, Example 7.9], the integral of $w_0^0 Y_0^0(\theta, \phi)$ over the 4π surface of the sphere is $\sqrt{4\pi} w_0^0$. Thus, we can normalize the function by dividing all coefficients by $\sqrt{4\pi} w_0^0$ to ensure that the coefficient of degree and order zero of the resulting function is $1/\sqrt{4\pi}$, resulting in an integral of 1.

B. Prediction Step

For the prediction step, the predicted density f_{t+1}^p fulfilling the Chapman–Kolmogorov equation

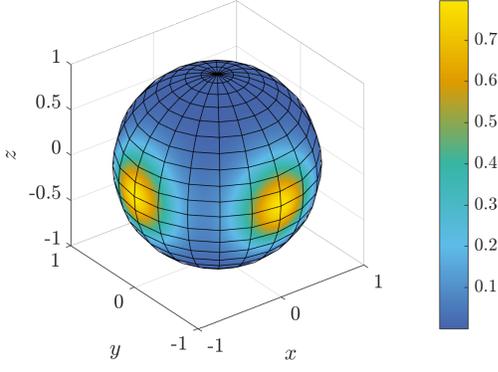
$$f_{t+1}^p(\underline{x}_{t+1} | \underline{z}_1, \dots, \underline{z}_t) = \int_{\Omega_x} f_t^T(\underline{x}_{t+1} | \underline{x}_t) f_t^e(\underline{x}_t | \underline{z}_1, \dots, \underline{z}_t) d\underline{x}_t \quad (2)$$

is to be obtained. One very useful transition density is an analog to the identity system model with additive noise

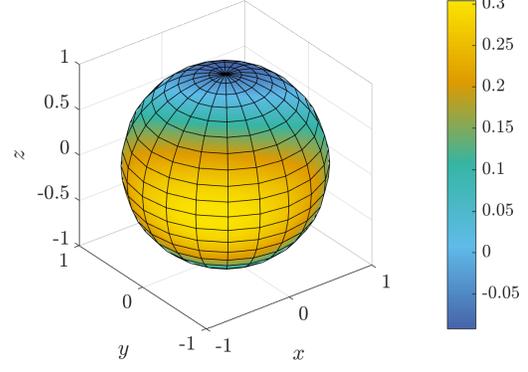
$$f_t^T(\underline{x}_{t+1} | \underline{x}_t) = \mathcal{VMF}(\underline{x}_{t+1}; \underline{\mu} = \mathbf{Q}_t \underline{x}_t, \kappa_t^w) \quad (3)$$

presented in [8, Eq. (6)]. It involves a rotation using a rotation matrix $\mathbf{Q} \in SO(3)$ and a von Mises–Fisher distributed [24] noise term parametrized by the concentration parameter κ_t^w .

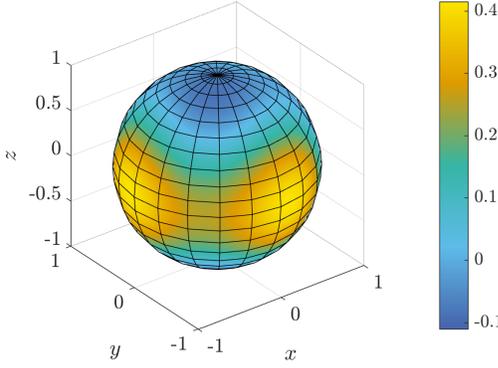
In the context of spherical harmonics, there are special functions that are called zonal. A function is called zonal



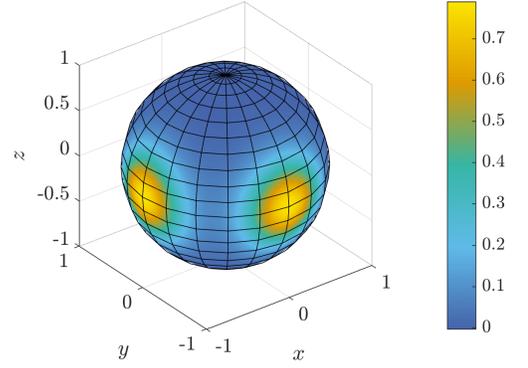
(a) The von Mises–Fisher mixture that is approximated in the subsequent subfigures. Both von Mises–Fisher components are equally weighted and have a dispersion according to $\kappa = 10$.



(b) Approximation using all spherical harmonics up to degree 2. The approximation is not yet able to capture the bimodality correctly and the modes are fused instead.



(c) Approximation using all spherical harmonics up to degree 3. The approximation has now become bimodal.



(d) Approximation using all spherical harmonics up to degree 9. The Hellinger distance to the ground truth is approximately 10^{-4} .

Fig. 3: Visualization of approximations of a von Mises–Fisher mixture with two modes. Since no square root transformation as used in [3] is employed, negative values are possible, especially when using few coefficients.

when it is rotationally symmetric regarding, e.g., the z -axis. The proposed filter works for arbitrary likelihoods but requires zonal noise terms for the prediction step. Due to the zonality of the von Mises–Fisher distribution, we can rewrite (3) as

$$f_t^T(\underline{x}_{t+1}|\underline{x}_t) = \mathcal{VMF}\left(\left[0, \sqrt{1 - (\underline{x}_{t+1}^T \mathbf{Q}_t \underline{x}_t)^2}, \underline{x}_{t+1}^T \mathbf{Q}_t \underline{x}_t\right]^T; \underline{\mu} = [0, 0, 1]^T, \kappa_t^w\right). \quad (4)$$

To get from (3) to (4), we regard the angle between $\mathbf{Q}_t \underline{x}_t$ and \underline{x}_{t+1} and calculate its cosine via $\underline{x}_{t+1}^T \mathbf{Q}_t \underline{x}_t$ and its sine via $\sqrt{1 - (\underline{x}_{t+1}^T \mathbf{Q}_t \underline{x}_t)^2}$. These values are then used to rotate $\underline{\mu}$ around the x -axis. Based on the representation (4), we say that our noise term is $f_t^v(\underline{v}_t) = \mathcal{VMF}(\underline{v}_t; \underline{\mu} = [0, 0, 1]^T, \kappa_t^w)$.

We only support zonal transition densities. For non-zonal transition densities, the rotation of the density in (3) around $\mathbf{Q}_t \underline{x}_t$ would have to be declared but there is no good basis to decide upon a specific rotation. In (4), the density is only evaluated along a great circle of the sphere containing $\underline{\mu}$. If the density was not zonal, the function values along the different great circles containing $\underline{\mu}$ would not be identical and it would be necessary to decide upon a specific great circle.

The convolution-like operation with a zonal density in the prediction step can be calculated efficiently when the spherical harmonics coefficients of f_t^e and f_t^v are given. If w_l^m are the spherical harmonics coefficients for f_t^e and v_l^m are the coefficients for the zonal noise density f_t^v , then the resulting coefficients c_l^m can be calculated according to [11, Sec. 9.3, Theorem 9.2]

$$c_l^m = \sqrt{\frac{4\pi}{2l+1}} w_l^m v_l^0. \quad (5)$$

When the spherical harmonics coefficients are given for both densities, this formula allows us to perform a prediction step with a complexity of $O(L^2)$ for a maximum degree L as only a fixed number of operations are required per coefficient in (5). If the spherical harmonics coefficients of the noise density have to be determined first, the complexity increases to $O(L^2 \log L)$. However, if the noise density does not change over time, the coefficients only have to be calculated once.

IV. EVALUATION

In our simulation-based evaluation, we regard the challenge of locating a moving target on a sphere using measurements

³As we only require the entries of order zero for the zonal density, we could reduce the effort to $O(L \log L)$.

that only give information about a single axis. The state is not observable using only a single measurement, which can lead to posterior densities that are not unimodal. As densities with multiple modes can usually not be approximated well using parametric densities, this scenario is unsuited for the use of filters that are based on parametric densities, such as the von Mises–Fisher filter. Therefore, we only compare our filter with a particle filter that we have adapted to the spherical domain.

A common distance measure between two points on the sphere is the orthodromic distance. The orthodromic distance describes the length of the shortest path along the sphere from one point to the other. To compare the approaches, we evaluate the orthodromic distance between the true position of the target and the estimated position as given by the mean resultant vector [8]

$$\underline{m} = \mathbb{E}(\underline{x}) = \int_{S^2} \underline{x} f(\underline{x}) d\underline{x} .$$

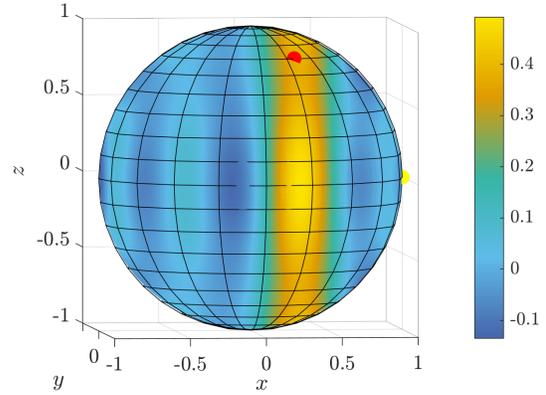
While our spherical harmonics filter can provide a probability density in each time step, we only evaluate the point estimate of the last position of the target. We do not regard the quality of the approximation of the posterior density as, to our knowledge, there is no established way to convert the particles of the particle filter to a continuous density on the sphere.

We compared the two filters regarding the estimation quality and run time on a laptop with an Intel Core i7-5500U processor, 12 GB of RAM, and Matlab 2016a on Windows 10. For the particle filter, different numbers of particles were used and for the spherical harmonics filter, we used coefficients up to different degrees. We always include all coefficients of the utilized degree so the number of coefficients scales quadratically with increasing degree.

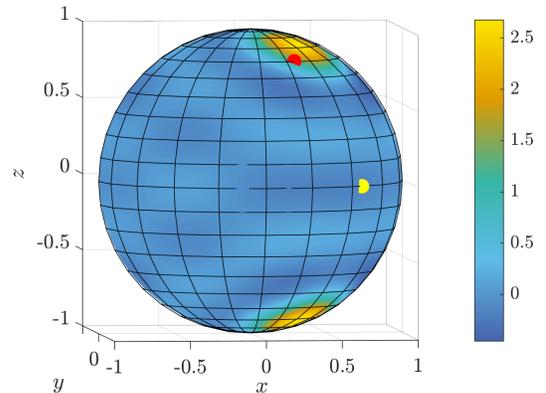
A. Scenario Description

The initial position of the target is set by sampling from a uniform distribution on the sphere. First, we obtain 5 measurements of the position along the x -axis, then 5 measurements of the y -axis position, and then 5 measurements of the position along the z -axis. Each measurement is perturbed by a (truncated) Gaussian noise with $\sigma = 0.3$ along the respective axis. The big challenge in this scenario is that the state is, in general, only partially observable using each of the measurements. Hence, the position of the target cannot be estimated to arbitrary accuracy using only measurements of the position along one axis⁴. This can clearly be seen in the intermediate results of the filters shown in Fig. 4 and Fig. 5. When only measurements along one axis have been obtained, the approximation of the posterior density given by the spherical harmonics filter depicted in Fig. 4a clearly shows that the posterior density has high probability density along a small circle of the sphere. This is also reflected in the weights of the particles of the particle filter shown in Fig. 5a.

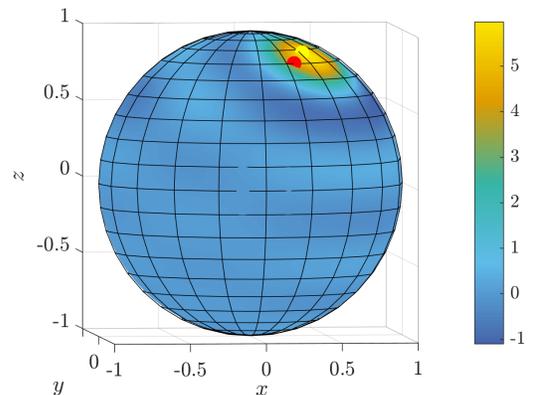
⁴An exception is, e.g., when the target is at -1 or 1 along the x -axis as there is only one point with this x coordinate.



(a) Result after the first 5 measurements of the position along the x -axis have been obtained. A whole region along a small circle of the sphere has high probability density. Because only part of the state was observable so far, the mean resultant vector cannot be used to reliably estimate the true position of the target.

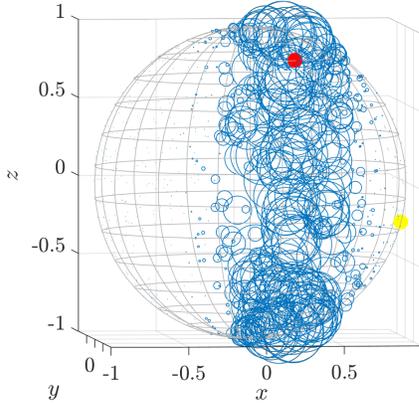


(b) Result after incorporating the 5 measurements of the position along the y -axis. The probability density is still bimodal and when using the mean resultant vector to derive an estimate, we obtain a point between the two modes.

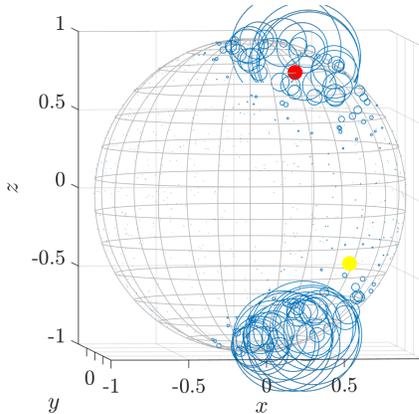


(c) Result after incorporating the 5 measurements of the position along the z -axis. The ambiguity has resolved and the estimate is now close to the true position.

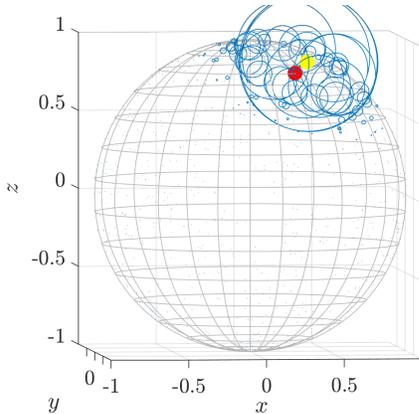
Fig. 4: Visualizations of the intermediate results for the first 15 measurements for the spherical harmonics filter after all measurements of the individual axis have been obtained. Coefficients up to degree 13 are used. The red circle indicates the true position of the target, whereas the yellow circle indicates the current estimate.



(a) Result after the 5 measurements of the position along the x -axis have been obtained. Particles along a whole small circle of the sphere have high weights.



(b) Result after incorporating the 5 measurements of the position along the y -axis. Particles at the two modes of the density (which can be seen in the result of the spherical harmonics filter) have high weights. Significant particle degeneracy can be observed.



(c) Result after incorporating the 5 measurements of the position along the z -axis. Only particles in the vicinity of the true position have height weights.

Fig. 5: Visualizations of the intermediate results for the first 15 measurements of the particle filter for 800 particles. The particles are shown as blue circles and the sizes indicate the particles' weights. The red circle indicates the true position of the target, whereas the yellow circle indicates the current estimate.

After these 15 measurements, the target moves according to the transition density given in (3) with $\mathbf{Q} = \mathbf{I}$ and $\kappa = 10$. For these parameters, the transition density induces a motion behavior resembling a random walk on the real plane. After this prediction step, 5 measurements of the position along each axis are obtained again and another prediction step is performed. Finally, 5 measurements of each axis are obtained sequentially again and the performance of the estimate at the last time step is determined (no prediction step is performed in the last time step). Thus, a total of 45 update steps and 2 prediction steps are performed.

B. Evaluation Results

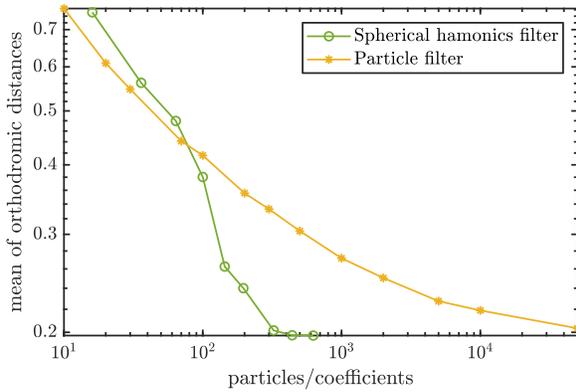
In the results of our evaluation, we provide the errors and run times based on an average of 5000 runs for all configurations. For the particle filter, we plot the results shown in Fig. 6 against the number of particles used, and for the spherical harmonics filter, we plot them against the total number of coefficients. For an identical number of parameters, the particle filter is superior in the range up until 121 particles. For more than 121 coefficients, the spherical harmonics filter performs better compared on a per coefficient basis. As shown in Fig. 6a, the spherical harmonics filter reaches an estimation quality close to its optimal quality using only 324 coefficients. The estimation quality achieved is superior to that of the particle filter, even when using 50 000 particles.

The run times shown in Fig. 6b are in accordance with the complexities of the two filters. The computation time of the particle filter increases linearly in the number of particles and the run time of the spherical harmonics filter rises approximately linearly with increasing number of coefficients. This is to be expected as the total number of spherical harmonics coefficients is in $O(L^2)$ for the maximum degree L and the complexity is in $O(L^2 \log L)$. For very few coefficients, the relative increase in the run time is very small for both filters as the constant overhead is predominant in that case.

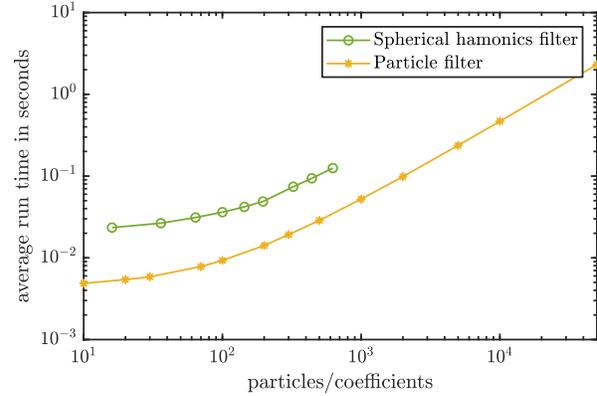
Compared on a run time basis, the particle filter is superior for low numbers of particles as it is very fast per particle and little constant overhead is involved. However, using 324 coefficients, the spherical harmonics filter achieves a quality that is not surpassed by the particle filter, even if a lot of computational power is invested and 50 000 particles are used. Since the spherical harmonics filter only requires about 74 ms for the entire evaluation scenario that involves 45 update steps and 2 prediction steps in this configuration, this is a reasonable configuration to use even in many real time applications.

V. CONCLUSION

The spherical harmonics filter is a versatile new filter for estimation problems on the unit sphere. Using fast and efficient algorithms in our implementation of the update and prediction steps, we are able to achieve good estimation quality and run time performance. In our evaluation, the filter outperforms the particle filter in a task of estimating the position of a target on the sphere when a high precision is desired. Furthermore, the filter provides an entire density



(a) Errors in form of the orthodromic distances for both filters shown using logarithmic scales.



(b) Run times of both filters shown using logarithmic scales.

Fig. 6: Errors and run times of both filters in the evaluated scenario.

on the sphere in every time step and is fully deterministic, allowing the user to generate reproducible results.

In future work, we intend to work on a square root version of the filter as done for the Fourier filters [3], [18], [19] and assess how this extension affects its performance and robustness. This is also of theoretical interest as negative function values can be completely prevented by the use of the square root transformation. Furthermore, using a different complete orthogonal sequence or regarding other manifolds may lead to new promising filters.

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 the grant HA 3789/13-1.

REFERENCES

- [1] 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 2009)*, 2009, pp. 882–887.
- [2] G. Kurz, I. Gilitschenski, and U. D. Hanebeck, "Recursive Bayesian Filtering in Circular State Spaces," *IEEE Aerospace and Electronic Systems Magazine*, vol. 31, no. 3, pp. 70–87, Mar. 2016.
- [3] 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.
- [4] G. Kurz, F. Pfaff, and U. D. Hanebeck, "Application of Discrete Recursive Bayesian Estimation on Intervals and the Unit Circle to Filtering on SE(2) (to appear)," *IEEE Transactions on Industrial Informatics*, 2017.
- [5] A. Chiuso and G. Picci, "Visual Tracking of Points as Estimation on the Unit Sphere," in *The Confluence of Vision and Control*. Springer, 1998, vol. 237, pp. 90–105.
- [6] 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 (MLSP)*, Sep. 2014.
- [7] G. Kurz, I. Gilitschenski, S. J. 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 2013)*, Istanbul, Turkey, Jul. 2013.
- [8] G. Kurz, I. Gilitschenski, and U. D. Hanebeck, "Unscented von Mises–Fisher Filtering," *IEEE Signal Processing Letters*, vol. 23, no. 4, pp. 463–467, Apr. 2016.
- [9] 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.
- [10] A. Zygmund, *Trigonometric Series*, 3rd ed. Cambridge University Press, 2003, vol. 1 and 2.
- [11] R. A. Kennedy and P. Sadeghi, *Hilbert Space Methods in Signal Processing*. Cambridge University Press, 2013.
- [12] N. P. Wedi, M. Hamrud, and G. Mozdzyński, "A Fast Spherical Harmonics Transform for Global NWP and Climate Models," *Monthly Weather Review*, vol. 141, no. 10, pp. 3450–3461, 2013.
- [13] G. B. Arfken and H. J. Weber, *Mathematical Methods for Physicists*, 7th ed. Academic Press, 2012.
- [14] L. C. Biedenharn and J. D. Louck, *Angular Momentum in Quantum Physics*. Cambridge University Press, 1984.
- [15] C. D. H. Chisholm, *Group Theoretical Techniques in Quantum Chemistry*. London: Academic Press, 1976.
- [16] R. Ramamoorthi and P. Hanrahan, "An Efficient Representation for Irradiance Environment Maps," in *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*. ACM, 2001, pp. 497–500.
- [17] M. A. Blanco, M. Flórez, and M. Bermejo, "Evaluation of the Rotation Matrices in the Basis of Real Spherical Harmonics," *Journal of Molecular Structure (Theochem)*, vol. 419, pp. 19–27, 1997.
- [18] 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.
- [19] —, "Multivariate Angular Filtering Using Fourier Series," *Journal of Advances in Information Fusion*, vol. 11, no. 2, pp. 206–226, Dec. 2016.
- [20] Y. F. Alem, Z. Khalid, and R. A. Kennedy, "Spherical Harmonic Expansion of Fisher–Bingham Distribution and 3-D Spatial Fading Correlation for Multiple-Antenna Systems," *IEEE Transactions on Vehicular Technology*, vol. 65, no. 7, pp. 5695–5700, July 2016.
- [21] J. R. Driscoll and D. M. Healy, Jr., "Computing Fourier Transforms and Convolutions on the 2-Sphere," *Advances in Applied Mathematics*, vol. 15, no. 2, pp. 202–250, 1994.
- [22] R. Suda and M. Takami, "A Fast Spherical Harmonics Transform Algorithm," *Mathematics of computation*, vol. 71, no. 238, pp. 703–715, 2002.
- [23] F. J. Simons, F. Dahlen, and M. A. Wiecek, "Spatiospectral Concentration on a Sphere," *SIAM review*, vol. 48, no. 3, pp. 504–536, 2006.
- [24] R. Fisher, "Dispersion on a Sphere," *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences*, vol. 217, no. 1130, pp. 295–305, 1953.