

Gaussian Mixture-Based Point Mass Filtering

Felipe Giraldo-Grueso

*Dept. of Aerospace Engineering and
Engineering Mechanics*

The University of Texas at Austin

Austin, TX, USA

fgiraldo@utexas.edu

Andrey A. Popov

*Oden Institute for Computational
Engineering and Sciences*

The University of Texas at Austin

Austin, TX, USA

andrey.a.popov@utexas.edu

Renato Zanetti

*Dept. of Aerospace Engineering and
Engineering Mechanics*

The University of Texas at Austin

Austin, TX, USA

renato@utexas.edu

Abstract—The accuracy of the point mass filter (PMF) relies on the precise placement of grid points. Since the approximated probability distributions are evaluated only at these points, sub-optimal choices in grid placement can result in an inaccurate representation of the posterior distribution. This work addresses this issue by representing the propagated grid points as a Gaussian mixture, enabling a Gaussian sum filter (GSF) update before grid construction. The use of the GSF update enhances the accuracy of the mean and covariance estimates, from which a new grid can be constructed. This approach leads to improved grid placement and reduces the number of points required to achieve satisfactory results. A comparative analysis is conducted between this new approach, the traditional PMF, and a PMF variant that uses an unscented Kalman filter update before grid construction. Using a simple bivariate example, the new variant is shown to approximate the posterior distribution better than the other filters. Furthermore, the new approach is evaluated in two sequential filtering problems: the first involves the Ikeda map, and the second focuses on terrain-relative navigation for Martian exploration. The results show a more accurate, and more consistent filter compared to the other two PMF variants considered.

Index Terms—Nonlinear Estimation, Point Mass Filters, Gaussian Mixtures, Grid Design

I. INTRODUCTION

Estimating quantities of interest from dynamic models and noisy measurements poses a significant challenge, especially in the context of nonlinear systems. A key approach to addressing this challenge involves solving the Bayesian recursive relations (BRR) [1]. In these relations, an initial state probability density function (pdf) undergoes propagation using the Chapman-Kolmogorov equation (in the context of discrete-time stochastic dynamic systems). Subsequently, the pdf is updated upon obtaining a measurement through the application of Bayes' rule. However, these relations can become analytically intractable when dealing with highly nonlinear dynamics and measurement models [1]. Consequently, when faced with such complexity, approximations to the solution of the BRR become necessary.

Point mass filters (PMF) represent a common methodology for numerically solving the BRR [2]. These types of filters approximate the BRR solutions through the application of deterministic grid-based integration techniques. The

PMF specifically approximates pdfs only at the grid points, making the strategic placement of these points pivotal for its performance. Therefore, the majority of PMF research has concentrated on refining grid design. Most approaches have focused on optimizing the predictive grid, often neglecting the integration of measurements at the current time [3]–[5]. While this strategy has proven to improve the performance of the filter, incorporating the measurement to construct the predictive grid has the potential to yield a more accurate approximation of the posterior pdf.

Recent studies have advanced the idea of incorporating measurements as additional information when creating the predictive grid. Dunfk et al. [6] introduced an algorithm that leverages an auxiliary unscented Kalman filter (UKF) to estimate the first two moments of the prior and posterior distributions in order to increase grid resolution in the approximated support of these pdfs. In a similar fashion, Choe and Park [7] developed an algorithm using log-homotopy induced flow to derive a grid support that accurately represents the non-negligible region of the posterior distribution.

This work introduces a new approach for generating the predictive grid. By representing the propagated points from the initial filtering grid as a Gaussian mixture, a Gaussian sum filter (GSF) update is performed, leading to a more precise estimation of the mean and covariance of the posterior pdf, similar to the techniques used in the ensemble Gaussian mixture filter [8]–[11]. This information is then used to construct an improved predictive grid. The application of this new technique results in more tighter grids, requiring fewer grid points to achieve accurate estimation performance. To demonstrate the performance of this approach, a simple bivariate ‘banana’ example is used, demonstrating that the new technique outperforms other PMF variants in approximating the posterior distribution. Furthermore, this new filter is validated using two sequential filtering scenarios, one involving the Ikeda map and the other pertaining to terrain-relative navigation for Martian exploration.

The remainder of this paper is organized as follows: first, background of state estimation in discrete systems is provided in Section II. The PMF is introduced in Section III, and the new improvement to this filter is presented in Section IV. A graphic comparison between the PMF, a PMF using a UKF update before grid construction, and the new variant

is illustrated in Section V using a simple bivariate example. The performance of these filters is tested on a sequential filtering problem using the Ikeda map in Section VI. All filters are evaluated on a more practical example regarding terrain-relative navigation in Section VII. Finally, Section VIII provides conclusions and future work.

II. STATE ESTIMATION FOR DISCRETE SYSTEMS

In this section, brief background on state estimation for discrete systems is provided. Given the evolution of a state \mathbf{x}_k governed by discrete dynamics \mathbf{f}_k and measurements \mathbf{h}_k ,

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

$$\mathbf{y}_k = \mathbf{h}_k(\mathbf{x}_k, \boldsymbol{\eta}_k), \quad (2)$$

where \mathbf{y}_k is the associated measurement, \mathbf{q}_k is the process noise and $\boldsymbol{\eta}_k$ is measurement noise, state estimation deals with determining the time evolution of the pdf of the state given the dynamic knowledge and measurements obtained.

The solution to the state estimation problem is obtained by solving the BRR. Starting from an initial pdf, the Chapman-Kolmogorov equation is used to propagate this pdf over time [1],

$$p(\mathbf{x}_{k+1}|\mathbf{y}_k) = \int_{\mathcal{S}(\mathbf{x}_k)} p(\mathbf{x}_{k+1}|\mathbf{x}_k) p(\mathbf{x}_k|\mathbf{y}_k) d\mathbf{x}_k, \quad (3)$$

where $\mathcal{S}(\mathbf{x}_k)$ denotes the support of \mathbf{x}_k . Once a measurement is obtained ($k \leftarrow k+1$), the pdf is updated using Bayes' rule [1],

$$p(\mathbf{x}_k|\mathbf{y}_k) = \frac{p(\mathbf{y}_k|\mathbf{x}_k, \mathbf{y}_{k-1}) p(\mathbf{x}_k|\mathbf{y}_{k-1})}{p(\mathbf{y}_k|\mathbf{y}_{k-1})}. \quad (4)$$

The solution to these relations can become intractable when dealing with highly nonlinear dynamics or measurements. In such cases, state estimation involves implementing efficient and accurate approximations to iteratively solve these two equations until a desired time is reached.

III. POINT MASS FILTER

The PMF aims to solve the state estimation problem deterministically through a structured grid of point particles. Despite the increased computational overhead, this filter has shown superior estimation accuracy compared to standard particle filters, especially in cases involving distributions with heavy tails [5]. For this filtering strategy, instead of describing the state pdf with realizations of exchangeable samples (as done in particle filters), the state pdf is discretized by strategically placing a finite number of possible realizations at deterministic grid points, where each point is assigned a finite probability.

In the PMF, the initial posterior pdf can be approximated as a Dirac mixture, such that [2],

$$p(\mathbf{x}_k|\mathbf{y}_k) \approx \sum_{i=1}^N w_{k|k}^{(i)} \delta(\mathbf{x}_k - \boldsymbol{\chi}_{k|k}^{(i)}), \quad (5)$$

where $\boldsymbol{\chi}_{k|k}^{(i)}$ are the posterior discretization points, $w_{k|k}^{(i)}$ represent the probability of each point, and N is the total

number of points. Using the Chapman-Kolmogorov equation, the approximated pdf is then propagated,

$$p(\mathbf{x}_{k+1}|\mathbf{y}_k) \approx \int_{\mathcal{S}(\mathbf{x}_k)} p(\mathbf{x}_{k+1}|\mathbf{x}_k) \sum_{i=1}^N w_{k|k}^{(i)} \delta(\mathbf{x}_k - \boldsymbol{\chi}_{k|k}^{(i)}) d\mathbf{x}_k, \quad (6)$$

$$= \sum_{i=1}^N w_{k|k}^{(i)} \int_{\mathcal{S}(\mathbf{x}_k)} p(\mathbf{x}_{k+1}|\mathbf{x}_k) \delta(\mathbf{x}_k - \boldsymbol{\chi}_{k|k}^{(i)}) d\mathbf{x}_k, \quad (7)$$

$$= \sum_{i=1}^N w_{k|k}^{(i)} p(\mathbf{x}_{k+1}|\boldsymbol{\chi}_{k|k}^{(i)}). \quad (8)$$

Using a Dirac mixture approximation, the resulting predictive distribution can be expressed as,

$$p(\mathbf{x}_{k+1}|\mathbf{y}_k) \approx \sum_{j=1}^M w_{k+1|k}^{(j)} \delta(\mathbf{x}_{k+1} - \boldsymbol{\chi}_{k+1}^{(j)}), \quad (9)$$

where $\boldsymbol{\chi}_{k+1}^{(j)}$ are the M new discretization points, and the propagated weights are proportional to:

$$w_{k+1|k}^{(j)} \propto \sum_{i=1}^N w_{k|k}^{(i)} p(\boldsymbol{\chi}_{k+1}^{(j)}|\boldsymbol{\chi}_{k|k}^{(i)}). \quad (10)$$

It is important to note that the number of grid points, M , in the discretization of the support of \mathbf{x}_{k+1} , does not have to be the same as in the initial discretization, N , of the support of \mathbf{x}_k . To generate the discretization points for the predictive pdf, a new grid is constructed, referred to as the predictive grid. The center of the grid is obtained by calculating the mean of the propagated points, while the orientation and expanse of the grid is given by the propagated covariance. In the case of zero-mean additive process noise with covariance \mathbf{Q}_k , the mean and covariance of the propagated points are given by [2]:

$$\hat{\mathbf{x}}_{k+1|k} = \sum_{i=1}^N w_{k|k}^{(i)} \mathbf{f}_k(\boldsymbol{\chi}_{k|k}^{(i)}), \quad (11)$$

$$\begin{aligned} \mathbf{P}_{k+1|k} &= \mathbf{Q}_k - \hat{\mathbf{x}}_{k+1|k} \hat{\mathbf{x}}_{k+1|k}^T \\ &+ \sum_{i=1}^N w_{k|k}^{(i)} \left[\mathbf{f}_k(\boldsymbol{\chi}_{k|k}^{(i)}) \mathbf{f}_k(\boldsymbol{\chi}_{k|k}^{(i)})^T \right]. \end{aligned} \quad (12)$$

Once the new grid has been generated and a new measurement has been obtained ($k \leftarrow k+1$), the weights of the grid are updated to generate the filtering grid. Using Bayes' rule,

$$p(\mathbf{x}_k|\mathbf{y}_k) \propto p(\mathbf{y}_k|\mathbf{x}_k) p(\mathbf{x}_k|\mathbf{y}_{k-1}), \quad (13)$$

$$\approx \sum_{j=1}^M w_{k|k-1}^{(j)} p(\mathbf{y}_k|\mathbf{x}_k) \delta(\mathbf{x}_k - \boldsymbol{\chi}_k^{(j)}), \quad (14)$$

$$= \sum_{j=1}^M w_{k|k}^{(j)} \delta(\mathbf{x}_k - \boldsymbol{\chi}_k^{(j)}), \quad (15)$$

where the new weights are proportional to the measurement likelihood,

$$w_{k|k}^{(j)} \propto w_{k|k-1}^{(j)} p(\mathbf{y}_k|\boldsymbol{\chi}_k^{(j)}). \quad (16)$$

The new weighted points then become the posterior discretization points, such that,

$$\mathbf{x}_{k|k}^{(j)} = \mathbf{x}_k^{(j)}, \quad (17)$$

where the first two moments of the approximated posterior pdf are calculated as:

$$\hat{\mathbf{x}}_{k|k} = \sum_{j=1}^M w_{k|k}^{(j)} \mathbf{x}_{k|k}^{(j)}, \quad (18)$$

$$P_{k|k} = \sum_{j=1}^M w_{k|k}^{(j)} \left[\mathbf{x}_{k|k}^{(j)} \left(\mathbf{x}_{k|k}^{(j)} \right)^T \right] - \hat{\mathbf{x}}_{k|k} \hat{\mathbf{x}}_{k|k}^T. \quad (19)$$

These equations summarize the standard PMF, which achieves accurate and consistent results when the grid separation is smaller than the standard deviation of the process noise. In situations where there is low process noise, this filter can result in denser grids and higher computational costs. In this work, a different approach for constructing the predictive grid is introduced, aiming to achieve more accurate results while using a reduced number of grid points.

IV. FUSION MASS FILTER

This section introduces a new method for creating the predictive grid in the PMF, referred to as the fusion mass filter (FMF) since it leverages information from both dynamics and measurements to construct a more accurate grid. This approach begins with a similar approximation to the PMF, by discretizing the initial posterior state pdf as a Dirac mixture, however, it introduces a subtle distinction by reformulating the Dirac mixture as a Gaussian mixture with infinitesimal covariance,

$$p(\mathbf{x}_k | \mathbf{y}_k) \approx \sum_{i=1}^N \omega_{k|k}^{(i)} \delta(\mathbf{x}_k - \mathbf{x}_{k|k}^{(i)}), \quad (20)$$

$$= \sum_{i=1}^N \omega_{k|k}^{(i)} \lim_{P \rightarrow 0} \mathcal{N}(\mathbf{x}_k; \mathbf{x}_{k|k}^{(i)}, P), \quad (21)$$

Assuming zero-mean additive process noise with covariance Q_k and using the GSF algorithm [12], [13], the resulting Gaussian mixture can be propagated in time to approximate the predictive distribution as:

$$p(\mathbf{x}_{k+1} | \mathbf{y}_k) \approx \sum_{i=1}^N \omega_{k+1|k}^{(i)} \mathcal{N}(\mathbf{x}_{k+1}; \mathbf{f}_k(\mathbf{x}_{k|k}^{(i)}), Q_k), \quad (22)$$

where,

$$\omega_{k+1|k}^{(i)} = \omega_{k|k}^{(i)}. \quad (23)$$

Instead of constructing a predictive grid directly from the propagated points, this new approach uses a GSF update prior to assembling the predictive grid. This process has the potential to improve the estimation of the mean and covariance of the state pdf thus resulting in a more accurate representation of the posterior pdf.

Once a measurement arrives ($k \leftarrow k + 1$), a posterior distribution can be approximated from the prior Gaussian mixture [12], [13], such that,

$$p(\mathbf{x}_k | \mathbf{y}_k) \approx \sum_{i=1}^N \tilde{w}_{k|k}^{(i)} \mathcal{N}(\mathbf{x}_k; \tilde{\mathbf{x}}_{k|k}^{(i)}, \tilde{P}_{k|k}^{(i)}), \quad (24)$$

where the notation $\tilde{\cdot}$ signifies that these are auxiliary variables obtained with the GSF update. Each mean and covariance of the Gaussian mixture are given by:

$$\tilde{\mathbf{x}}_{k|k}^{(i)} = \mathbf{x}_{k|k-1}^{(i)} + K_k^{(i)} \boldsymbol{\nu}_k^{(i)}, \quad (25)$$

$$\tilde{P}_{k|k}^{(i)} = Q_{k-1} - K_k^{(i)} W_k^{(i)} K_k^{(i)T}, \quad (26)$$

with the intermediate variables defined as:

$$\mathbf{x}_{k|k-1}^{(i)} = \mathbf{f}_{k-1}(\mathbf{x}_{k-1|k-1}^{(i)}), \quad (27)$$

$$\boldsymbol{\nu}_k^{(i)} = \mathbf{y}_k - \mathbf{h}_k(\mathbf{x}_{k|k-1}^{(i)}), \quad (28)$$

$$K_k^{(i)} = Q_{k-1} H_k^{(i)T} (W_k^{(i)})^{-1}, \quad (29)$$

$$W_k^{(i)} = H_k^{(i)} Q_{k-1} H_k^{(i)T} + R_k, \quad (30)$$

$$H_k^{(i)} = \frac{\partial \mathbf{h}_k(\mathbf{x})}{\partial \mathbf{x}} \Big|_{\mathbf{x}=\mathbf{x}_{k|k-1}^{(i)}}. \quad (31)$$

The weights of this Gaussian mixture are defined as proportional to the probability of the obtained measurement,

$$\tilde{w}_{k|k}^{(i)} \propto w_{k|k-1}^{(i)} \mathcal{N}(\mathbf{y}_k; \mathbf{h}_k(\mathbf{x}_{k|k-1}^{(i)}), W_k^{(i)}). \quad (32)$$

As mentioned previously, rather than creating a predictive grid with the propagated mean and covariance estimates, this approach discretizes the support based on the updated mean and covariance estimates. The predictive grid is assembled after the GSF update, using both the updated mean and covariance estimates, given by:

$$\tilde{\mathbf{x}}_{k|k} = \sum_{i=1}^N \tilde{w}_{k|k}^{(i)} \tilde{\mathbf{x}}_{k|k}^{(i)}, \quad (33)$$

$$\tilde{P}_{k|k} = \sum_{i=1}^N \tilde{w}_{k|k}^{(i)} \left(\tilde{P}_{k|k}^{(i)} + \tilde{\mathbf{x}}_{k|k}^{(i)} \left(\tilde{\mathbf{x}}_{k|k}^{(i)} \right)^T - \tilde{\mathbf{x}}_{k|k} \tilde{\mathbf{x}}_{k|k}^T \right), \quad (34)$$

where the center of the grid is placed at the updated mean, and the orientation and expanse are set to align with the updated covariance. With the new grid, the posterior distribution is approximated as a Dirac mixture,

$$p(\mathbf{x}_k | \mathbf{y}_k) \approx \sum_{j=1}^M w_{k|k}^{(j)} \delta(\mathbf{x}_k - \mathbf{x}_k^{(j)}), \quad (35)$$

where $\mathbf{x}_k^{(j)}$ are the new grid points, and $w_{k|k}^{(j)}$ are the new weights of each point. To calculate these new weights and generate the filtering grid, the posterior pdf is expressed as,

$$p(\mathbf{x}_k | \mathbf{y}_k) = \int_{\mathcal{S}(\mathbf{x}_{k-1})} p(\mathbf{x}_k, \mathbf{x}_{k-1} | \mathbf{y}_k) d\mathbf{x}_{k-1}, \quad (36)$$




Fig. 1. Comparative analysis between the generation of the predictive grid in the PMF (top) and in the FMF (bottom). The FMF updates the propagated points via a GSF update before creating the new grid. Once the predictive grid has been assembled, the grid points are weighted to create the filtering grid.

where,

$$p(\mathbf{x}_k, \mathbf{x}_{k-1} | \mathbf{y}_k) \propto p(\mathbf{y}_k | \mathbf{x}_k) p(\mathbf{x}_k | \mathbf{x}_{k-1}) \cdot \dots \cdot p(\mathbf{x}_{k-1} | \mathbf{y}_{k-1}). \quad (37)$$

Which results in weights proportional to:

$$w_{k|k}^{(j)} \propto p(\mathbf{x}_k^{(j)} | \mathbf{y}_k), \quad (38)$$

$$\approx \sum_{i=1}^N p(\mathbf{x}_k^{(j)}, \mathbf{x}_{k-1|k-1}^{(i)} | \mathbf{y}_k), \quad (39)$$

$$= \sum_{i=1}^N p(\mathbf{y}_k | \mathbf{x}_k^{(j)}) p(\mathbf{x}_k^{(j)} | \mathbf{x}_{k-1|k-1}^{(i)}) w_{k-1|k-1}^{(i)}. \quad (40)$$

Just as with the PMF, the new weighted points become the posterior discretization points and the first two moments of the approximated posterior pdf are calculated as in (18) and (19). These equations summarize the new methodology for constructing a more precise predictive grid within the PMF framework, offering the potential to improve accuracy and consistency in filtering outcomes.

Figure 1 presents a comparative analysis between the generation of the predictive grid in the PMF and in the proposed FMF. The FMF describes each propagated point as a Gaussian distribution with an associated covariance, facilitating a GSF update. By updating the propagated points prior to grid creation, there is an improvement in the accuracy of mean and covariance estimates. This can result in a comparatively tighter grid in the FMF as opposed to the PMF. The difference in grid sizes may suggest a potential benefit in scenarios with low process noise or when seeking accurate results with a small number of grid points.

V. ILLUSTRATION WITH A SIMPLE BANANA EXAMPLE

To illustrate the potential improvement in grid construction of the FMF, a simple two-dimensional problem is presented [14]. The problem starts from a Gaussian posterior distribution,

$$p(\mathbf{x}_k | \mathbf{y}_k) \sim \mathcal{N} \left(\begin{bmatrix} -3.5 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \right). \quad (41)$$

From this posterior distribution, a uniform grid composed of 625 points, centered at the mean and scaled up to 3σ , is created. The dynamics are assumed to induce a correlation between the two states such that,

$$p(\mathbf{x}_{k+1} | \mathbf{y}_k) \sim \mathcal{N} \left(\begin{bmatrix} -3.5 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 & 0.5 \\ 0.5 & 1 \end{bmatrix} \right). \quad (42)$$

Using this predictive distribution, a uniform predictive grid is created for the PMF with the same number of points as the starting grid, but centered at the new mean and scaled by the new covariance. The new points of the predictive grid are convoluted with the starting grid using zero-mean additive process noise with a covariance matrix of $Q = 2 \times 10^{-1} I_{2 \times 2}$. The points on the predictive grid are weighted to generate the filtering grid, according to the following measurement,

$$y = h(\mathbf{x}) = \sqrt{x_1^2 + x_2^2}, \quad (43)$$

with a value of $y = 1$ and associated scalar measurement error covariance of $R = 0.1^2$.




Fig. 2. Comparison between the PMF (left), UMF (center) and FMF (right). The prior distribution is shown in blue, the measurement likelihood is plotted in red, and the true posterior distribution is depicted in yellow, resembling the shape of a banana. All filtering grids are shown in purple, with each grid point plotted based on its weight. Larger points indicate higher weights, while smaller points carry lower weights.

For the FMF, 625 different realizations are sampled from the prior distribution, from which a GSF update is performed by using the same measurement, measurement covariance and process noise covariance matrices as for the PMF. With the updated mean and covariance estimates obtained from the GSF update, a new uniform grid is created centered at the updated mean and scaled by the updated covariance. The new points of the grid are weighted as in (40), generating the filtering grid.

For comparison purposes, a third filter, incorporating the concepts presented by Duník et al. [6] that are compatible with the framework discussed in this work, is used to compare the FMF to related work. For this filter, instead of using a GSF update to estimate a posterior mean and covariance, a UKF update is applied before assembling the new grid. This PMF variant will be referred to as the unscented mass filter (UMF). Just as with the FMF, once the updated mean and covariance estimates are obtained, a new uniform grid is centered at the updated mean, scaled by the updated covariance and weighted as described in (40).

Figure 2 shows the resulting filtering grids for the three approaches. The left figure illustrates the grid obtained using the PMF, the center figure shows the grid obtained with the UMF, while the right figure presents the grid from the FMF. From this figure, it can be seen that the FMF results in a filtering grid that more accurately approximates the true posterior distribution compared to the PMF and UMF. By incorporating a GSF update before grid creation, the filtering grid in the FMF is closely centered on the true posterior, depicted in yellow and resembling the shape of a banana. Additionally, the resulting grid covers a smaller area, resulting in more points located near the true posterior. In contrast, the PMF generates a predictive grid that is too large, resulting in a filtering grid with only a few points in proximity to the true posterior distribution. Similarly, the UKF linear update in the UMF is not as effective as the GSF update in the FMF, making the placement of the grid less efficient.

VI. SEQUENTIAL FILTERING WITH THE IKEDA MAP

The following example aims to demonstrate the performance improvement of the FMF by conducting a grid study in a sequential filtering problem. To this end, the Ikeda map is used, representing a discrete-time dynamical system:

$$x_{k+1}^{(1)} = 1 + u \left(x_k^{(1)} \cos t_k - x_k^{(2)} \sin t_k \right), \quad (44)$$

$$x_{k+1}^{(2)} = u \left(x_k^{(1)} \sin t_k + x_k^{(2)} \cos t_k \right), \quad (45)$$

$$t_k = 0.4 - \frac{6}{1 + \left(x_k^{(1)} \right)^2 + \left(x_k^{(2)} \right)^2}, \quad (46)$$

This dynamical system is designed to model light circulating within a nonlinear optical resonator [15], [16]. To introduce chaos and create a more challenging sequential filtering problem, the parameter u is set to $u = 0.9$ and the dynamics are propagated with additive white Gaussian process noise with covariance matrix $Q = 1 \times 10^{-2} I_{2 \times 2}$. A nonlinear measurement model is assumed, with measurements given by:

$$y_k = \sqrt{\left(x_k^{(1)} \right)^2 + \left(x_k^{(2)} \right)^2} + \eta_k, \quad (47)$$

where η_k represents white Gaussian measurement noise with scalar covariance matrix $R = 1$. For every tested filter configuration, different grid sizes scaled up to 3σ are used to estimate the state in 1000 distinct trajectories simulated for 50 time steps, each with an initial true state defined as $x_0 \sim \mathcal{N}(0_{2 \times 1}, I_{2 \times 2})$.

To compare each filter, two metrics are used. The time-averaged root mean squared error (RMSE) is used to assess the accuracy of the filter. For this work, the RMSE is defined as,

$$\text{RMSE} = \sum_{k=1}^{N_t} \frac{1}{N_t} \sum_{j=1}^{N_m} \frac{1}{N_m} \sqrt{\frac{\sum_{i=1}^{N_s} \left(x_{k,j}^{(i)} - \hat{x}_{k|k,j}^{(i)} \right)^2}{N_s}}, \quad (48)$$

where N_t is total number of discrete steps in each simulation, N_m is the number of independent trajectories, N_s is the total number of states, $x_{k,j}^{(i)}$ is the true state and $\hat{x}_{k|k,j}^{(i)}$ is the estimated state. Figure 3 shows the RMSE for the PMF, UMF and FMF as a function of the total grid points. As expected, with an increment in grid points, all filters show improved accuracy. The RMSE curve for the FMF consistently remains below that of the PMF and UMF, showing improved accuracy with a reduced requirement of grid points.




Fig. 3. Root mean squared error for the PMF (blue with diamond markers), UMF (red with square markers) and FMF (purple with cross markers) as a function of total grid points. The cyan line represents the results obtained with a regularized particle filter with 10000 particles.

To evaluate the consistency of the filter, the time-averaged scaled normalized estimation error squared (SNEES) is used,

$$\text{SNEES} = \sum_{k=1}^{N_t} \frac{1}{N_t} \sum_{j=1}^{N_m} \frac{1}{N_m N_s} \sum_{i=1}^{N_s} e_{k,j}^T (P_{k|k,j})^{-1} e_{k,j}, \quad (49)$$

where,

$$e_{k,j} = x_{k,j} - \hat{x}_{k|k,j}, \quad (50)$$

is the estimation error, and $P_{k|k,j}$ is the estimated covariance. A consistent filter will obtain a SNEES close to one. If the SNEES value significantly exceeds one, it suggests overconfidence in the estimator. Conversely, if the value is considerably smaller than one, the estimator is overly conservative [1]. It is important to note that for this work, any SNEES value over 1×10^4 was disregarded as a numerical instability.

Figure 4 shows the SNEES for all filters as a function of total grid points. Similar to the RMSE figure, all filters exhibit increasing consistency with the use of more grid points. Both the UMF and FMF approach a SNEES close to one at around 300 grid points, while the PMF requires more than 400 grid points to achieve comparable values. Furthermore, the FMF consistently outperforms the UMF and PMF in this aspect, as the SNEES curve of the FMF resides below that of the UMF and PMF, approaching a value closer to one across all of the grid sizes tested.

Considering that both the UMF and FMF involve one additional step compared to the PMF, specifically the UKF




Fig. 4. Scaled normalized estimation error squared for the PMF (blue with diamond markers), UMF (red with square markers) and FMF (purple with cross markers) as a function of total grid points. The light blue dashed line shows a SNEES value of one.

or GSF update, it is important to quantify the computational run time of these three algorithms. Figure 5 illustrates the RMSE of each filter as a function of their required run time for a single propagation and update. All filters have been implemented to maximize their run time efficiency for fair comparisons. The extra step in the UMF and FMF is evident by the slight shift of the curve to the right compared to the PMF. Regardless, as it was shown before, less grid points are needed for the FMF to achieve more accurate results when compared to the PMF and UMF. The use of the GSF update allows the implementation of smaller grids, which translates into faster run time accompanied by a low RMSE. This benefit in run time for the FMF can be seen in Fig. 5, as the time curve consistently lies below that of the PMF and UMF.




Fig. 5. Root mean squared error for the PMF (blue with diamond markers), UMF (red with square markers) and FMF (purple with cross markers) as a function of required run time for one propagation and update.

VII. MARS TERRAIN-RELATIVE NAVIGATION

The PMF is a common approach for terrain-relative navigation [6], [17], [18]. In this type of problem, an onboard

elevation map is used to estimate the horizontal position of a spacecraft based on measurements from a barometric sensor and an altimeter and velocity estimates provided by an onboard inertial navigation system (INS) [6], [17]. To show the performance of the FMF in a more practical setup, a grid study is presented for a Mars terrain-relative navigation problem. A trajectory over the Nili Fossae, located north of Jezero Crater, is simulated. Figure 6 shows the digital elevation map¹, generated from the Mars Reconnaissance Orbiter (MRO) Context Camera (CTX) [19], and the simulated nominal trajectory. The nominal trajectory follows a constant-velocity and constant-altitude flight over a valley, reminiscent of possible reconnaissance flights on Mars by small helicopters, such as Ingenuity, but with a longer flying range.




Fig. 6. Digital elevation map of the Nili Fossae, located north of Jezero Crater. The simulated nominal trajectory is shown in cyan.

The true dynamics of the trajectory are given by:

$$x_{k+1}^{(1)} = x_k^{(1)} + \Delta t_k v^{(1)}, \quad (51)$$

$$x_{k+1}^{(2)} = x_k^{(2)} + \Delta t_k v^{(2)}, \quad (52)$$

where Δt_k is the time-step and v is velocity. The dynamics used by the filter are,

$$x_{k+1|k}^{(1)} = x_{k|k}^{(1)} + \Delta t_k \hat{v}_k^{(1)}, \quad (53)$$

$$x_{k+1|k}^{(2)} = x_{k|k}^{(2)} + \Delta t_k \hat{v}_k^{(2)}, \quad (54)$$

where \hat{v}_k represents the velocity estimated by the INS. In this work, the true dynamics are propagated with additive white Gaussian process noise with covariance matrix $Q = 0.5I_{2 \times 2}$ and the estimated velocity is sampled at each time step from $v_k \sim \mathcal{N}(10/\sqrt{2}I_{2 \times 1} \text{ m/s}, 0.1I_{2 \times 2} \text{ m}^2/\text{s}^2)$, suggesting accurate and uncorrelated INS estimates. In this case, the covariance matrix of the transitional pdf used for the convolution in the PMF, and to describe the predictive distribution in the UMF and FMF, has to account for the uncertainty in velocity, where,

$$Q_k = \Delta t_k^2 (0.1I_{2 \times 2}) + Q. \quad (55)$$

¹Elevation map obtained from The Pacific Regional Planetary Data Center hosted by the Hawai'i Institute of Geophysics and Planetology.

The measurement used for this problem is the terrain elevation given by the map at the current position, simulating possible measurements from a barometric sensor and an altimeter. As the map provides elevations at discrete positions, cubic spline interpolation is used to obtain the elevation at the query positions. Each measurement is corrupted with additive white Gaussian noise with scalar covariance matrix $R = 0.025^2$, based on the laser altimeter Garmin LIDAR-Lite V3 [20].

For every filter tested, all grids are constructed up to 5σ and 1000 distinct trajectories are simulated for 250 seconds with a time-step of $\Delta t = 1$ second, each with an initial true state defined as:

$$x_0 \sim \mathcal{N} \left(\begin{bmatrix} 2300 \\ 800 \end{bmatrix} \text{ m}, \begin{bmatrix} 50 & 0 \\ 0 & 50 \end{bmatrix} \text{ m}^2 \right). \quad (56)$$

Figure 7 shows the RMSE for the PMF, UMF and FMF as a function of the total grid points used. Just as with the Ikeda map, the curve for the FMF remains below the curve of the PMF and UMF, showing a clear improvement in filtering accuracy. Figure 8 shows the SNEES for the three filters as a function of total grid points. In addition to being more accurate, this figure shows that the FMF is a more consistent filter than the PMF and UMF, specifically for lower grid sizes, as the SNEES for this filter stays closer to one. The SNEES of the PMF shows that this filter struggles with lower number of grid points, as the results show high SNEES values up until around the use of 700 grid points.




Fig. 7. Root mean squared error for the PMF (blue with diamond markers), UMF (red with square markers) and FMF (purple with cross markers) as a function of total grid points.

The FMF also proves to be more efficient in terms of computational cost for this problem setup, as smaller grid sizes are needed to achieve accurate and consistent results. To calculate the efficiency gains of the FMF over the PMF and UMF, the run time required for the filters to achieve an RMSE of 1.3 meters was determined by interpolating the results shown in Fig. 7 and the recorded run time. This resulted in a speed-up of 4.174x for the FMF over the PMF and approximately 1.596x for the FMF over the UMF. It is important to note that the derivative of the measurement




Fig. 8. Scaled normalized estimation error squared for the PMF (blue with diamond markers), UMF (red with square markers) and FMF (purple with cross markers) as a function of total grid points. The light blue dashed line shows a SNEES value of one.

model, used to perform the GSF update, was obtained through numerical differentiation with cubic spline interpolation of the elevation map. More efficient methods for calculating this derivative could be used, representing a potentially significant improvement in efficiency.

VIII. CONCLUSIONS

In this work, a new variant of the point mass filter has been introduced. In this new variant, coined the fusion mass filter, a Gaussian sum filter update is performed before the creation of the predictive grid. The resulting updated mean and covariance estimates are then used to generate the grid. Through a simple bivariate example, it was shown that incorporating the update prior to grid creation yields a tighter grid. This tighter grid provides a more accurate description of the posterior distribution compared to both the standard point mass filter and a point mass filter with an unscented Kalman filter update performed before creating the grid. Furthermore, by applying this new technique to two different sequential filtering examples, it was observed that the fusion mass filter outperforms the point mass filter and the unscented Kalman filter variation in terms of accuracy and consistency.

While this paper has shown potential improvements in the point mass filter, it is essential to consider additional developments from both theoretical and practical perspectives. On the theoretical side, there is an opportunity to explore various expressions for the predictive distribution in (22), with the aim of potentially yielding more favorable outcomes during the Gaussian sum filter update. Furthermore, it is worth considering alternative methodologies for assigning weights to the new predictive grid, which, in turn, may contribute to the derivation of a more effective filtering grid.

From a practical standpoint, multi-density grids have been used in the point mass filter to improve its computational run time. A potential opportunity for future work lies in evaluating the performance of this new algorithm within the

context of more efficient grid choices, such as sparse or non-uniform grids. Moreover, instead of exclusively relying on a point mass approximation for the posterior distribution, exploring alternative methodologies, such as describing this starting distribution using uniform kernels, offers a possibility to potentially improve the effectiveness of this approach.

ACKNOWLEDGMENT

The authors would like to acknowledge the Ingenuity Mars helicopter for its inspiring, now completed, mission.

REFERENCES

- [1] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan, *Estimation with Applications to Tracking and Navigation*. John Wiley & Sons, Inc, 2001.
- [2] R. Bucy and K. Senne., "Digital synthesis of non-linear filters," *Automatica*, vol. 7, no. 3, pp. 287–298, 1971.
- [3] N. Bergman, "Recursive bayesian estimation: Navigation and tracking applications," Ph.D. dissertation, Linköping University, 1999.
- [4] M. Šimandl, J. Královce, and T. Söderström, "Anticipative grid design in point-mass approach to nonlinear state estimation," *IEEE Transactions on Automatic Control*, vol. 47, no. 4, pp. 699–702, 2002.
- [5] —, "Advanced point-mass method for nonlinear state estimation," *Automatica*, vol. 42, no. 7, pp. 1133–1145, 2006.
- [6] J. Duník, O. Straka, and J. Matoušek, "Conditional density driven grid design in point-mass filter," in *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 9180–9184.
- [7] Y. Choe and C. G. Park, "Point-mass filtering with boundary flow and its application to terrain referenced navigation," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 57, no. 6, pp. 3600–3613, 2021.
- [8] S. Yun, R. Zanetti, and B. A. Jones, "Kernel-based ensemble gaussian mixture filtering for orbit determination with sparse data," *Advances in Space Research*, vol. 69, no. 12, pp. 4179–4197, 2022.
- [9] A. A. Popov and R. Zanetti, "An adaptive covariance parameterization technique for the ensemble gaussian mixture filter," *arXiv preprint arXiv:2212.10323*, 2022.
- [10] —, "Ensemble gaussian mixture filtering with particle-localized covariances," in *Proceedings of the 2023 26th International Conference on Information Fusion (FUSION)*, Charleston, South Carolina, Jul. 2023.
- [11] —, "Ensemble-localized kernel density estimation with applications to the ensemble gaussian mixture filter," *arXiv preprint arXiv:2308.14143*, vol. 69, 2023.
- [12] H. Sorenson and D. Alspach, "Recursive bayesian estimation using gaussian sums," *Automatica*, vol. 7, no. 4, pp. 465–479, 1971.
- [13] D. Alspach and H. Sorenson, "Nonlinear bayesian estimation using gaussian sum approximations," *IEEE Transactions on Automatic Control*, vol. 17, no. 4, pp. 439–448, 1972.
- [14] K. Michaelson, A. A. Popov, and R. Zanetti, "Ensemble kalman filter with bayesian recursive update," in *Proceedings of the 2023 26th International Conference on Information Fusion (FUSION)*, Charleston, South Carolina, Jul. 2023.
- [15] K. Ikeda, "Multiple-valued stationary state and its instability of the transmitted light by a ring cavity system," *Optics Communications*, vol. 30, no. 2, pp. 257–261, 1979.
- [16] K. Ikeda, H. Daido, and O. Akimoto, "Optical turbulence: Chaotic behavior of transmitted light from a ring cavity," *Phys. Rev. Lett.*, vol. 45, pp. 709–712, 1980.
- [17] N. Bergman, "A bayesian approach to terrain-aided navigation," *IFAC Proceedings Volumes*, vol. 30, no. 11, pp. 1457–1462, 1997.
- [18] J. Duník, M. Soták, M. Veselý, O. Straka, and W. Hawkinson, "Design of rao-blackwellized point-mass filter with application in terrain aided navigation," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 55, no. 1, pp. 251–272, 2019.
- [19] M. C. Malin, J. F. Bell III, B. A. Cantor, M. A. Caplinger, W. M. Calvin, R. T. Clancy, K. S. Edgett, L. Edwards, R. M. Haberle, P. B. James, S. W. Lee, M. A. Ravine, P. C. Thomas, and M. J. Wolff, "Context camera investigation on board the mars reconnaissance orbiter," *Journal of Geophysical Research: Planets*, vol. 112, no. E5, 2007.
- [20] D. S. Bayard, D. T. Conway, R. Brockers, J. H. Delaune, L. H. Matthies, H. F. Grip, G. B. Merewether, T. L. Brown, and A. M. S. Martin, "Vision-based navigation for the nasa mars helicopter," 2019.