Optimal Transport as a Reduction Technique for Deterministic Nonlinear Filtering

Felipe Giraldo-Grueso*, Andrey A. Popov†, Uwe D. Hanebeck‡, and Renato Zanetti*
*Department of Aerospace Engineering and Engineering Mechanics, The University of Texas at Austin, Austin, Texas 78712
Email: {fgiraldo, renato}@utexas.edu
†Department of Information and Computer Sciences, The University of Hawai'i at Mānoa, Honolulu, HI 96822
Email: apopov@hawaii.edu
‡Department of Computer Science, Karlsruhe Institute of Technology, Germany
Email: uwe.hanebeck@kit.edu

***Abstract*—The solution to the state estimation problem is given by the Bayesian recursive relations (BRRs). Recently, ensemble Gaussian mixture filters have shown to be an accurate and consistent solution to the state estimation problem. In this type of filters, the BRRs are solved by approximating the state probability density function (PDF) via Gaussian mixtures (GMs) and point masses (PMs). Throughout the propagation and measurement update steps, the approximated state PDF is constantly switching between GMs and PMs. Therefore, a key step for this solution involves optimally sampling PMs from GMs. For onboard applications, verifiable and computationally inexpensive sampling techniques are crucial. In previous work, a deterministic sampling technique was developed by minimizing a distance metric known as the modified Cramér-von Mises distance (MCVMD), yielding a verifiable solution. However, the computationally feasibility of this solution for onboard use was not considered. This work introduces a new sampling strategy that is both deterministic and computationally inexpensive compared to MCVMD approach. By solving the approximate optimal transport problem via an iterative Sinkhorn-Knopp algorithm, this new technique is able to sub-optimally sample from a GM, providing a computationally inexpensive filter.**

***Index Terms*—Nonlinear Estimation, Sequential Filtering, Optimal Transport, Sinkhorn-Knopp**

I. Introduction

In state estimation, quantities of interest are inferred from dynamic models and measurements. The Bayesian recursive relations (BRRs) [1] provide a solution to the state estimation problem. If the system is linear and Gaussian, it is known that the optimal estimator (in the minimum mean squared error sense) is the Kalman filter [1]. Most dynamics and measurement models are nonlinear, making the solution to the BRRs intractable. These scenarios require approximations and techniques to solve the BRRs numerically. Different filtering techniques approximate the solution of the BRRs by describing the state probability density function (PDF) in various functional forms.

These filtering techniques are generally expected to be accurate, consistent, and in some applications, computationally feasible and verifiable. Accuracy quantifies the error between the state estimate and the true state. Consistency refers to the estimated error distribution and/or covariance matrix being *consistent* with the true values [1]. A computationally feasible solution may have varying interpretations, but a fast filtering strategy is typically desirable [2]. Verifiability refers to producing expected and reproducible outputs from a known set of inputs [3]. Some of these requirements might not be met depending on the filtering strategy used to approximate the BRRs. Therefore, a filter that can optimally balance these requirements, based on the required application, is essential for state estimation.

Onboard navigation is a filtering scenario where there is a need to balance these four requirements. If a filter is to be used in real-time onboard a vehicle, computational resources may be constrained [2], and verifiability may be crucial [3]. In addition, low-margin of error onboard applications, such as atmospheric entry [4], [5] or autonomous landing [6], may also need very high accuracy with consistent error bounds. Thus, finding an optimal balance between accuracy, consistency, feasibility, and verifiability is important for such scenarios.

The point mass filter (PMF) is a filtering strategy used to solve the BRRs by using a deterministic grid of point masses (PMs) to describe the state PDF [7]–[10]. In recent work, the Silverman mass filter (SMF), a variation of the standard PMF, was introduced [11]. The SMF was shown to be an accurate solution to the BRRs, by placing the grid on the approximated posterior PDF obtained via a Gaussian sum filter (GSF) update [12], [13]. The SMF was also shown to yield consistent results by using kernel density estimation (KDE) techniques [14] to better approximate the prior PDF before performing the GSF update.

Recently, this filter was improved with a deterministic optimal sampling technique [15], making it equivalent to an ensemble Gaussian mixture filter (EnGMF) [16]–[19] with deterministic resampling. The deterministic nature of this filter makes it a good candidate for verifiability [3]. However, the optimal sampling technique was developed without considering the computational feasibility requirement. Therefore, in its current state, the SMF meets three out of the four requirements discussed, accuracy, consistency, and verifiability.

This work presents a solution to the computational feasibility requirement. The sampling technique, used to improve the

This material is based on research sponsored by the Air Force Office of Scientific Research (AFOSR) under award number: FA9550-23-1-0646.

SMF, optimally samples PMs from the approximated posterior PDF obtained via the GSF update, represented by a Gaussian mixture (GM). To obtain the PMs, deterministic samples are first drawn from each component of the GM, forming an initial PM approximation of the posterior PDF. This initial approximation is then reduced to the desired number of PMs by minimizing the modified Cramér-von Mises distance (MCVMD). The optimality of this reduction technique is, therefore, defined with respect to the MCVMD.

In this work, the initial PM approximation is reduced by using optimal transport (OT) techniques, similar to the strategies used in the ensemble transform particle filter (ETPF) [20]–[23], and in the pineapple filter [3]. However, the OT solution is approximated by an iterative use of the Sinkhorn-Knopp (SK) [24] algorithm, providing a faster runtime compared to previous work. This new reduction technique is shown to perform similar to the optimal sampling, but with a lower computational cost, thus satisfying the four filtering requirements discussed.

The remainder of this paper is organized as follows: first, a background of sequential filtering using PMs and GMs is provided in Section II. The optimal sampling strategy that minimizes the MCVMD is presented in Section III, and the use of OT as a reduction strategy is presented in Section IV. The performance of both techniques is tested on bivariate problem, and a sequential filtering problem in Section V. Section VI provides conclusions and future work.

II. Sequential Filtering using Point Masses and Gaussian Mixtures

Let $\boldsymbol{x}_k \in \mathbb{R}^{n_x}$ be the state, driven by discrete dynamics $\boldsymbol{f}_k : \mathbb{R}^{n_x} \to \mathbb{R}^{n_x}$ and measurements $\boldsymbol{h}_k : \mathbb{R}^{n_x} \to \mathbb{R}^{n_y}$ at time step k, such that

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

$$\boldsymbol{y}_k = \boldsymbol{h}_k\left(\boldsymbol{x}_k\right) + \boldsymbol{\eta}_k, \tag{2}$$

where $\boldsymbol{q}_k \in \mathbb{R}^{n_x}$ and $\boldsymbol{\eta}_k \in \mathbb{R}^{n_y}$ are process and measurement noise, respectively, and $\boldsymbol{y}_k \in \mathbb{R}^{n_y}$ is the measurement. To solve the BRRs, a hybrid PM and GM approach can be used by discretizing the state space with a grid of PMs.

In previous work, a PMF that places the grid centered on the approximated posterior PDF (referred to as the SMF) [11], [25], rather than the prior [8]–[10], has been developed. This filter begins by formulating the initial posterior state PDF as a GM with infinitesimal covariance,

$$p\left(\boldsymbol{x}_k|\boldsymbol{y}_k\right) \approx \sum_{i=1}^{N} w_{k|k}^{(i)}\, \delta\left(\boldsymbol{x}_k - \boldsymbol{\mathcal{X}}_{k|k}^{(i)}\right), \tag{3}$$

$$= \sum_{i=1}^{N} w_{k|k}^{(i)} \lim_{\boldsymbol{P}\to 0} \mathcal{N}\left(\boldsymbol{x}_k; \boldsymbol{\mathcal{X}}_{k|k}^{(i)}, \boldsymbol{P}\right), \tag{4}$$

where $\boldsymbol{\mathcal{X}}_{k|k} \in \mathbb{R}^{n_x \times N}$ is the collection of grid points, $\boldsymbol{\mathcal{X}}_{k|k}^{(i)} \in \mathbb{R}^{n_x}$ is the i-th grid point, $w_{k|k} \in \mathbb{R}_+^N$ are the probability masses, $w_{k|k}^{(i)} \in \mathbb{R}_+$ is the probability mass of the i-th grid point, and N is the total number of points.

Assuming that the process noise is Gaussian and zero-mean, the initial posterior PDF can be propagated by using the GSF algorithm [12], [13], and KDE techniques [14]. This results in an approximate prior distribution given by:

$$p\left(\boldsymbol{x}_{k+1}|\boldsymbol{y}_k\right) \approx \sum_{i=1}^{N} w_{k+1|k}^{(i)} \mathcal{N}\left(\boldsymbol{x}_{k+1}; \boldsymbol{f}_k\left(\boldsymbol{\mathcal{X}}_{k|k}^{(i)}\right), \boldsymbol{B}_k\right), \tag{5}$$

with $w_{k+1|k}^{(i)} = w_{k|k}^{(i)}$, and,

$$\boldsymbol{B}_k = \beta^2 \hat{\boldsymbol{P}}_{k+1|k} + \boldsymbol{Q}_k, \tag{6}$$

where $\hat{\boldsymbol{P}}_{k+1|k} \in \mathbb{R}^{n_x \times n_x}$ is the ensemble covariance of the propagated points, $\boldsymbol{Q}_k \in \mathbb{R}^{n_x \times n_x}$ is the process noise covariance, and $\beta^2 \in \mathbb{R}_+$ is the bandwidth parameter, given by Silverman's rule of thumb [14]. Note that in Eq. (6), the ensemble covariance and the process noise covariance balance each other. If $\boldsymbol{Q}_k$ is large, β^2 can be scaled by a factor $0 < \alpha \leq 1$, so that the filter is more consistent. On the contrary, if $\boldsymbol{Q}_k$ is small, β^2 does not need to be scaled [11].

A GSF update [12], [13] is performed when a measurement is obtained ($k \leftarrow k+1$) to approximate a posterior distribution of the form:

$$p\left(\boldsymbol{x}_k|\boldsymbol{y}_k\right) \approx \sum_{i=1}^{N} \widetilde{w}_{k|k}^{(i)} \mathcal{N}\left(\boldsymbol{x}_k; \widetilde{\boldsymbol{\mathcal{X}}}_{k|k}^{(i)}, \widetilde{\boldsymbol{\mathcal{P}}}_{k|k}^{(i)}\right), \tag{7}$$

where each mean, covariance, and weight of the GM are given by (assuming Gaussian and zero-mean measurement noise):

$$\boldsymbol{\mathcal{X}}_{k|k-1}^{(i)} = \boldsymbol{f}_k\left(\boldsymbol{\mathcal{X}}_{k|k}^{(i)}\right) \tag{8}$$

$$\boldsymbol{\nu}_k^{(i)} = \boldsymbol{y}_k - \boldsymbol{h}_k\left(\boldsymbol{\mathcal{X}}_{k|k-1}^{(i)}\right) \tag{9}$$

$$\widetilde{\boldsymbol{\mathcal{X}}}_{k|k}^{(i)} = \boldsymbol{\mathcal{X}}_{k|k-1}^{(i)} + \boldsymbol{K}_k^{(i)} \boldsymbol{\nu}_k^{(i)}, \tag{10}$$

$$\widetilde{\boldsymbol{\mathcal{P}}}_{k|k}^{(i)} = \boldsymbol{B}_{k-1} - \boldsymbol{K}_k^{(i)} \boldsymbol{W}_k^{(i)} \boldsymbol{K}_k^{(i)\mathrm{T}}, \tag{11}$$

$$\widetilde{w}_{k|k}^{(i)} \propto w_{k|k-1}^{(i)} \mathcal{N}\left(\boldsymbol{y}_k; \boldsymbol{h}_k\left(\boldsymbol{\mathcal{X}}_{k|k-1}^{(i)}\right), \boldsymbol{W}_k^{(i)}\right). \tag{12}$$

with $\boldsymbol{K}_k^{(i)} \in \mathbb{R}^{n_x \times n_y}$ and $\boldsymbol{W}_k^{(i)} \in \mathbb{R}^{n_y \times n_y}$ representing the Kalman gain and innovation covariance for each point, respectively.

After the GSF update, optimally representing the GSF-approximated posterior distribution (which could be multimodal) in Eq. (7), as a Dirac mixture (DM) approximation for the next iteration, becomes the main challenge,

$$p\left(\boldsymbol{x}_k|\boldsymbol{y}_k\right) \approx \sum_{i=1}^{N} w_{k|k}^{(i)} \delta\left(\boldsymbol{x}_k - \boldsymbol{\mathcal{X}}_{k|k}^{(i)}\right). \tag{13}$$

In other words, this filtering strategy requires a solution to optimally find the location of the new grid points $\boldsymbol{\mathcal{X}}_{k|k}$, and their associated weights, $w_{k|k}$, given the information from the GSF update. Figure 1 shows the key steps in the SMF solution.

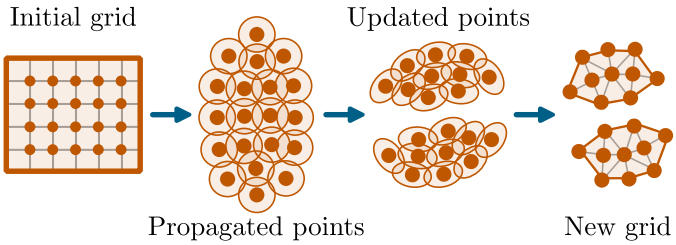

Fig. 1. Key steps of the SMF. First, the initial grid is propagated as a GM, then the points are updated, and a new grid is constructed.

III. An Optimal Reduction Solution

Recently, a solution based on optimal sampling of GMs has been proposed [15]. In this solution, the GM in Eq. (7) is deterministically sampled to revert to a PM approximation by minimizing the MCVMD. Thus, this technique is said to be optimal with respect to the MCVMD. To sample the GM, deterministic samples are first drawn from each component of the GM, forming an initial PM approximation of the posterior PDF. This initial approximation is then optimally reduced to the desired number of PMs.

The optimal sampling is summarized as follows. First, D samples are obtained from each Gaussian component[1],

$$\mathcal{N}\left(\boldsymbol{x}_k; \tilde{\boldsymbol{\mathcal{X}}}_{k|k}^{(i)}, \tilde{\boldsymbol{\mathcal{P}}}_{k|k}^{(i)}\right) \approx \frac{1}{D}\sum_{j=1}^{D}\delta\left(\boldsymbol{x}_k - \boldsymbol{\mathcal{D}}_{k|k,i}^{(j)}\right), \tag{14}$$

where $\boldsymbol{\mathcal{D}}_{k|k,i}^{(j)} \in \mathbb{R}^{n_x}$ is the j-th deterministic sample. Note that these samples are equally weighted. The initial PM approximation, therefore, consists of $N \cdot D$ PMs,

$$p\left(\boldsymbol{x}_k|\boldsymbol{y}_k\right) \approx \sum_{i=1}^{N}\frac{\tilde{w}_{k|k}^{(i)}}{D}\sum_{j=1}^{D}\delta\left(\boldsymbol{x}_k - \boldsymbol{\mathcal{D}}_{k|k,i}^{(j)}\right). \tag{15}$$

The second step is to optimally reduce the DM to the desired number of equally weighted points, in this case N,

$$\sum_{i=1}^{N}\frac{\tilde{w}_{k|k}^{(i)}}{D}\sum_{j=1}^{D}\delta\left(\boldsymbol{x}_k - \boldsymbol{\mathcal{D}}_{k|k,i}^{(j)}\right) \approx \frac{1}{N}\sum_{i=1}^{N}\delta\left(\boldsymbol{x}_k - \boldsymbol{\mathcal{X}}_{k|k}^{(i)}\right). \tag{16}$$

Note that this reduction step, where all the new PMs are equally weighted, makes this version of the SMF equivalent to an EnGMF [16]–[19] with deterministic resampling.

For this work, let $M = N \cdot D$ and let the points on the left-hand side of Eq. (16) be denoted as $\boldsymbol{\mathcal{Y}}_{k|k} \in \mathbb{R}^{n_x \times M}$ with weights $w_{k|k}^{\mathcal{Y}} \in \mathbb{R}_+^M$, that is:

$$\sum_{i=1}^{M} w_{k|k}^{\mathcal{Y}(i)}\delta\left(\boldsymbol{x}_k - \boldsymbol{\mathcal{Y}}_{k|k}^{(i)}\right) \approx \frac{1}{N}\sum_{i=1}^{N}\delta\left(\boldsymbol{x}_k - \boldsymbol{\mathcal{X}}_{k|k}^{(i)}\right). \tag{17}$$

[1]Deterministic samples can be obtained in various ways, this work uses Fibonacci grids. [26]

Given the two DMs in Eq. (17), the optimal points, $\boldsymbol{\mathcal{X}}_{k|k}$, are found by minimizing the MCVMD between the two DMs with respect to the points' location, given by [15]:

$$\begin{aligned} D_{\text{MCVM}}(\boldsymbol{\mathcal{X}}) = &\left(w_{k|k}^{\mathcal{Y}}\right)^{\mathrm{T}}\boldsymbol{M}_{\boldsymbol{\mathcal{Y}}_{k|k}\boldsymbol{\mathcal{Y}}_{k|k}}w_{k|k}^{\mathcal{Y}} \\ &- \frac{2}{N}\left(w_{k|k}^{\mathcal{Y}}\right)^{\mathrm{T}}\boldsymbol{M}_{\boldsymbol{\mathcal{Y}}_{k|k}\boldsymbol{\mathcal{X}}}\mathbf{1}_N \\ &+ \frac{1}{N^2}\left(\mathbf{1}_N\right)^{\mathrm{T}}\boldsymbol{M}_{\boldsymbol{\mathcal{X}}\boldsymbol{\mathcal{X}}}\mathbf{1}_N \\ &+ K\left|\left|\boldsymbol{\mathcal{Y}}_{k|k}w_{k|k}^{\mathcal{Y}} - \boldsymbol{\mathcal{X}}\frac{\mathbf{1}_N}{N}\right|\right|_2^2, \end{aligned} \tag{18}$$

where $K > 0$ is a constant, $\mathbf{1}_L \in \mathbb{R}^L$ is the one vector, and,

$$\boldsymbol{M}_{\boldsymbol{\mathcal{X}}_i\boldsymbol{\mathcal{X}}_j} = \operatorname{xlog}\left(\boldsymbol{D}_{\boldsymbol{\mathcal{X}}_i\boldsymbol{\mathcal{X}}_j}\right), \tag{19}$$

with $\boldsymbol{D}_{\boldsymbol{\mathcal{X}}_i\boldsymbol{\mathcal{X}}_j}$ denoting the Euclidean distance matrix between $\boldsymbol{\mathcal{X}}_i$ and $\boldsymbol{\mathcal{X}}_j$, and $\operatorname{xlog}(z) = z \cdot \log(z)$ element-wise.

Thus, the location of the new grid points is given by solving

$$\boldsymbol{\mathcal{X}}_{k|k} = \arg\min_{\boldsymbol{\mathcal{X}}} D_{\text{MCVM}}\left(\boldsymbol{\mathcal{X}}\right). \tag{20}$$

For this minimization problem, the location of the points can be initialized by either using the first N points of $\boldsymbol{\mathcal{Y}}_{k|k}$, or by using the means of the GSF-approximated posterior, $\tilde{\boldsymbol{\mathcal{X}}}_{k|k}$. Assuming $M \gg N$, the main complexity for calculating the MCVMD is $\mathcal{O}(M \cdot N)$ [27].

This solution has been shown to obtain high-quality points, improving the estimation performance of the SMF [15]. The deterministic nature of this solution is desirable for onboard applications, but solving the optimization problem in Eq. (20) can be a computational bottleneck. Therefore, this work explores a new technique by using OT to reduce the computational runtime of the SMF while preserving estimation performance.

IV. Optimal Transport as a Reduction Strategy

In OT, as formulated in Refs. [3], [21], [23], an approximate mapping between $\boldsymbol{\mathcal{X}}_{k|k}$ and $\boldsymbol{\mathcal{Y}}_{k|k}$ can be found, such that

$$\boldsymbol{\mathcal{X}}_{k|k} = \boldsymbol{\mathcal{Y}}_{k|k}\boldsymbol{T}_k^{\star}, \tag{21}$$

by minimizing the total transport known as the Monge-Kantorovich problem [20],

$$\boldsymbol{T}_k^{\star} = \arg\min_{\boldsymbol{T}_k}\sum_{i=1}^{M}\sum_{j=1}^{N}\boldsymbol{T}_k^{(i,j)}\left|\left|\boldsymbol{\mathcal{Y}}_{k|k}^{(i)} - \boldsymbol{\mathcal{X}}^{(j)}\right|\right|_2^2, \tag{22}$$

subject to

$$\sum_{i=1}^{M}\boldsymbol{T}_k^{\star(i,j)} = 1, \quad \sum_{j=1}^{N}\boldsymbol{T}_k^{\star(i,j)} = N w_{k|k}^{\mathcal{Y}(i)}, \quad \boldsymbol{T}_k^{\star(i,j)} \geq 0, \tag{23}$$

where $\boldsymbol{\mathcal{X}}$ is usually chosen to be the first N points of $\boldsymbol{\mathcal{Y}}_{k|k}$ or the means of the GSF-approximated posterior, $\tilde{\boldsymbol{\mathcal{X}}}_{k|k}$ [3]. Note that this strategy uses a different target function compared to the optimal sampling technique, where the MCVMD is minimized.

Figure 2 shows the use of OT as a reduction strategy. In essence, the initial weighted points $\boldsymbol{\mathcal{Y}}_{k|k}$ are *transported* with the operator $\boldsymbol{T}_k^\star \in \mathbb{R}^{M\times N}$ to equally weighted reduced points $\boldsymbol{\mathcal{X}}_{k|k}$. However, since finding $\boldsymbol{T}_k^\star$ is a linear programming problem, the computational worst-case complexity of all known solvers is exponential. This makes the OT approach unsuitable for reducing the complexity of the SMF.

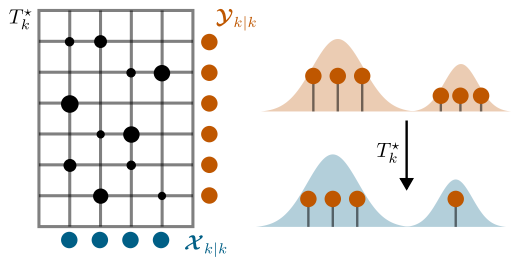

Fig. 2. Optimal transport as a reduction strategy.

A. Sinkhorn-Knopp Algorithm

What is interesting about this linear programming problem is that it can be approximately solved by using the SK algorithm [22], [24], [28]. This work proposes the use of the SK algorithm to approximately solve the OT problem as a reduction strategy.

The SK algorithm approximates $\boldsymbol{T}_k^\star$ by adding a regularization term to the cost function in Eq. (22),

$$\boldsymbol{T}_k^{\star,\mathrm{SK}} = \arg\min_{\boldsymbol{T}_k} \sum_{i=1}^{M}\sum_{j=1}^{N}\left[\boldsymbol{T}_k^{(i,j)}\|\boldsymbol{\mathcal{Y}}_{k|k}^{(i)} - \boldsymbol{\mathcal{X}}^{(j)}\|_2^2 + \frac{1}{\lambda}\boldsymbol{T}_k^{(i,j)}\log\frac{\boldsymbol{T}_k^{(i,j)}}{w_{k|k}^{\mathcal{Y}(i)}}\right], \tag{24}$$

subject to the same constraints, where $\lambda > 0$ is the regularization weight. For this setup, the SK algorithm has a guaranteed time complexity of $\mathcal{O}(M\cdot N)$ [22]. Note that this is the same complexity as evaluating the MCVMD once. Although the SK algorithm is an approximation of the OT problem, it is guaranteed to converge (in the fixed-point sense). This guarantee is essential for bounding computational time, which is crucial when running filtering algorithms in limited computational resources.

To assess convergence, the following criteria can be used. Note that, following the OT formulation, the SK solution has to satisfy

$$\sum_{i=1}^{M}\boldsymbol{T}_k^{\star,SK(i,j)} = 1. \tag{25}$$

Therefore, the algorithm is said to converge when

$$\left\|\left(\boldsymbol{T}_k^{\star,SK}\right)^{\mathrm{T}}\mathbf{1}_M - \mathbf{1}_N\right\|^2 < \varepsilon^2, \tag{26}$$

where $\varepsilon^2 > 0$ represents a small constant.

B. Reformulation of the Sinkhorn Knopp Algorithm

In addition to using the SK algorithm to solve the OT problem, this work proposes a simple yet significant improvement to the SK solution. Note that in Eq. (22) and Eq. (24), the choice of $\boldsymbol{\mathcal{X}}$ is fixed. Alternatively, by iterating a set amount of times on $\boldsymbol{\mathcal{X}}$, such that

$$\boldsymbol{\mathcal{X}}_{k|k}, \boldsymbol{T}_k^{\star,\mathrm{SK}} = \arg\min_{\boldsymbol{T}_k,\boldsymbol{\mathcal{X}}} \sum_{i=1}^{M}\sum_{j=1}^{L}\left[\boldsymbol{T}_k^{(i,j)}\|\boldsymbol{\mathcal{Y}}_{k|k}^{(i)} - \boldsymbol{\mathcal{X}}^{(j)}\|_2^2 + \frac{1}{\lambda}\boldsymbol{T}_k^{(i,j)}\log\frac{\boldsymbol{T}_k^{(i,j)}}{w_{k|k}^{\mathcal{Y}(i)}}\right], \tag{27}$$

the resulting operator $\boldsymbol{T}_k^{\star,\mathrm{SK}}$ and location of the new points $\boldsymbol{\mathcal{X}}_{k|k}$ can be significantly improved. In this case, the SK algorithm is run for a fixed number of iterations N_i, where in each iteration, $\boldsymbol{\mathcal{X}}$ is updated using the points found in the previous iteration. Instead of computing the Euclidean distance once, this reformulation repeatedly alternates between computing $\boldsymbol{T}_k^{\star,\mathrm{SK}}$ and updating $\boldsymbol{\mathcal{X}}$ over the course of the N_i iterations.

V. Numerical Examples

This section compares the accuracy, consistency, and computational runtime of the new SK approximation technique, the optimal reduction (OR) that minimizes the MCVMD [15], and solving the OT problem directly through linear programming [3]. Throughout these examples, the OR technique will be used as a benchmark, representing the optimal PM approximation. That is, the performance metric used for evaluating the location of the reduced points will be the MCVMD.

A. Clover Distribution

Figure 3 shows a two-dimensional example that compares the three reduction methods discussed. The sampling distribution is a GM with four equally weighted components in the form of a clover. The first two components are given by

$$p(\boldsymbol{x}_1) \sim \mathcal{N}\left(\boldsymbol{x}_{1,2}, P_{1,2}\right), \tag{28}$$

$$p(\boldsymbol{x}_2) \sim \mathcal{N}\left(-\boldsymbol{x}_{1,2}, P_{1,2}\right), \tag{29}$$

with

$$\boldsymbol{x}_{1,2} = \begin{bmatrix}2\\2\end{bmatrix}, \quad P_{1,2} = \begin{bmatrix}1 & 0.5\\0.5 & 1\end{bmatrix}. \tag{30}$$

The second two components are given by

$$p(\boldsymbol{x}_3) \sim \mathcal{N}\left(\boldsymbol{x}_{3,4}, P_{3,4}\right), \tag{31}$$

$$p(\boldsymbol{x}_4) \sim \mathcal{N}\left(-\boldsymbol{x}_{3,4}, P_{3,4}\right), \tag{32}$$

with

$$\boldsymbol{x}_{3,4} = \begin{bmatrix}-2\\2\end{bmatrix}, \quad P_{1,2} = \begin{bmatrix}1 & -0.5\\-0.5 & 1\end{bmatrix}. \tag{33}$$

From each component, 100 deterministic points are sampled using Fibonacci grids ($M = 400$). These points, referred to as $\boldsymbol{\mathcal{Y}}_{k|k}$ in the previous sections, are shown in gold. The union

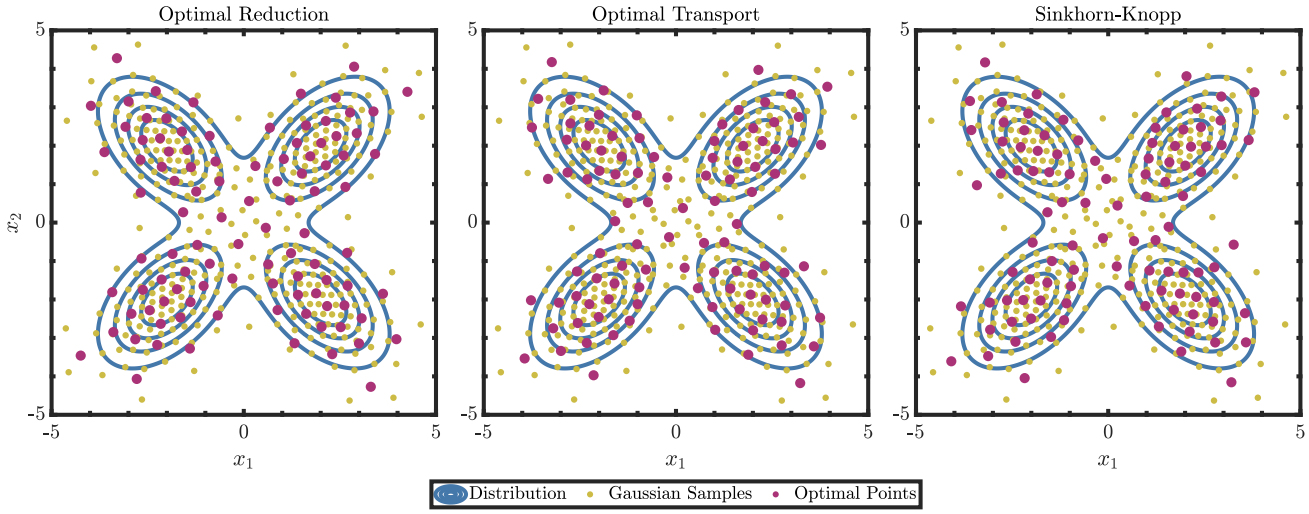

Fig. 3. Comparison of the OR (left), OT (center), and the SK approximation (right) for $N = 100$. The gold points show the first PM approximation of the GM plotted in blue. The purple points show the reduced PM approximation.

of the samples represent the initial PM approximation of the GM.

The initial PM approximation is then reduced to $N = 100$ equally weighted points, referred to as $\boldsymbol{\mathcal{X}}_{k|k}$ in the previous sections. This new points are shown in purple. The left panel shows the OR technique by minimizing the MCVMD using Matlab's `fminunc` function with default settings, as in Ref. [15]. The center panel shows the reduction technique that directly solves the OT problem, as in Ref. [3], using Matlab's `linprog` function with default settings. The right panel shows the SK approximation with $\varepsilon^2 = 1 \times 10^{-6}$, and $\lambda = 1 \times 10^3$. For all three approaches, $\boldsymbol{\mathcal{X}}$ is initialized at the first N points in $\boldsymbol{\mathcal{Y}}_{k|k}$. From the figure, it can be seen that all three approaches achieve similar reductions. In particular, the OT solution and the SK approximation yield nearly identical points.

However, if the number of desired points is reduced to $N = 50$, the OT solution and the SK approximation perform poorly compared to minimizing the MCVMD. This can be seen in the center panel of Fig. 4. The OT solution is not shown as it yields very similar points to the SK approximation. In this case, the reformulated approach, where the SK algorithm is run for multiple iterations, updating $\boldsymbol{\mathcal{X}}$ at the beginning of each one, can lead to significantly improved results. If $N_i = 10$ iterations are performed between $\boldsymbol{T}_k^{\star,\mathrm{SK}}$ and $\boldsymbol{\mathcal{X}}$, the resulting points represent the underlying distribution better, as shown in the right panel of Fig. 4. Note, however, that these points are still sub-optimal with respect to the MCVMD solution.

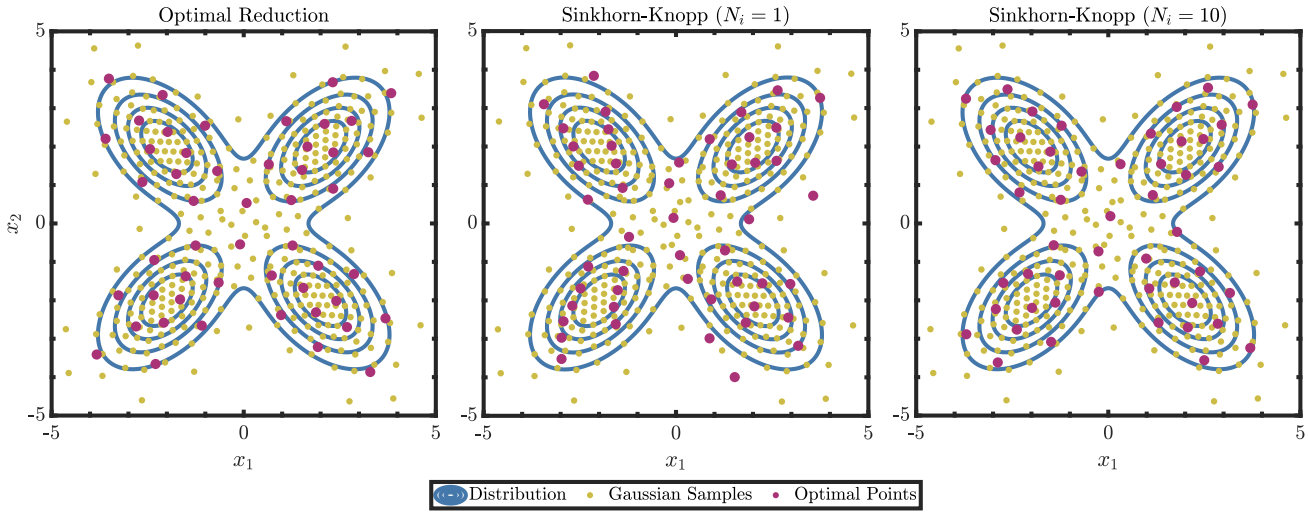

Fig. 4. Comparison of the OR (left), the SK approximation (center), and the reformulated SK approximation (right) for $N = 50$. The gold points show the first PM approximation of the GM plotted in blue. The purple points show the reduced PM approximation.

Performing a parameter sweep over N_i can show the benefits of the reformulated SK approximation over the standard SK and the OT solution. Figure 5 shows a comparison of the three reduction techniques as a function of N_i. The top panel shows the ratio, d_r, between the MCVMD, calculated with respect to the gold ($\boldsymbol{\mathcal{Y}}_{k|k}$) and purple points ($\boldsymbol{\mathcal{X}}_{k|k}$), for the OR technique and each of the other techniques,

$$d_r = \frac{D_{\text{MCVM}}\left(\boldsymbol{\mathcal{X}}_{k|k}^{(\text{OT or SK})}\right)}{D_{\text{MCVM}}\left(\boldsymbol{\mathcal{X}}_{k|k}^{(\text{OR})}\right)}. \tag{34}$$

That is, $d_r = 1$ represents the best reduction, in terms of minimizing the MCVMD, as given by the OR strategy. The bottom panel of Fig. 5 shows the runtime needed for each reduction. All strategies were implemented in MATLAB R2024a and were run on a MacBook Pro with an Apple M1 Pro chip and 16 GB of RAM.

From the figure, it can be seen that performing more iterations in the SK algorithm moves the reduced points towards the OR solution. By only performing two iterations, the SK solution achieves a PM representation that more closely aligns with the OR results when compared to the OT solution. In addition, it can be seen that the SK solution is less computationally expensive than both the OR and the OT techniques, even when performing multiple iterations of the algorithm. Even though the OT solution is faster than the OR technique, the reformulated SK approximation yields better quality points (in the sense of approaching the OR solution) in a faster runtime.

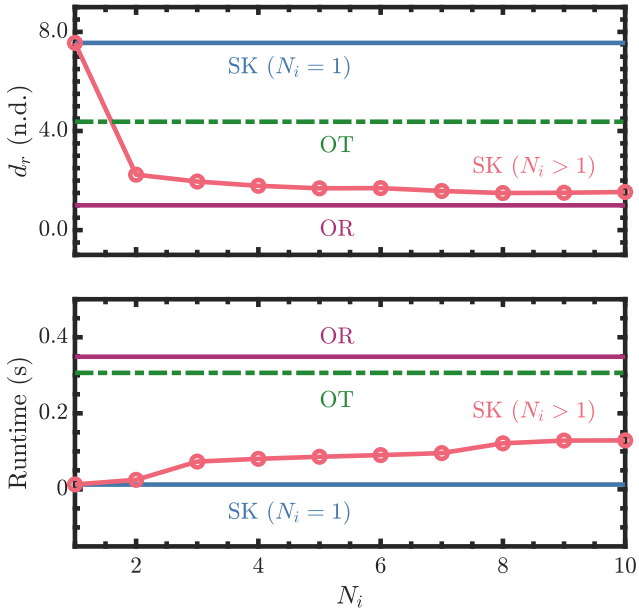

Fig. 5. Comparison between the three reduction techniques. The top panel shows the MCVMD ratio between the OR, and the other reduction techniques. The bottom panel shows the runtime required to perform the reduction.

B. Ikeda Map

The clover distribution example illustrates that the SK approximation can produce comparable points to the OR strategy, especially when iterating over both $\boldsymbol{T}_k^{\star,\text{SK}}$ and $\boldsymbol{\mathcal{X}}$. The SK framework may be more desirable from an onboard perspective as it avoids optimization schemes and guarantees convergence in the fixed-point sense. The remaining question is whether the points obtained with the SK approximation are representative enough of the underlying distribution. To test whether the SK approximation can yield results similar to the OR technique, the Ikeda map can be used to run a sequential filtering test.

The Ikeda map represents a discrete-time dynamical system, where [29], [30]

$$\boldsymbol{x}_{k+1}^{(1)} = 1 + u\left(\boldsymbol{x}_k^{(1)}\cos t_k - \boldsymbol{x}_k^{(2)}\sin t_k\right), \tag{35}$$

$$\boldsymbol{x}_{k+1}^{(2)} = u\left(\boldsymbol{x}_k^{(1)}\sin t_k + \boldsymbol{x}_k^{(2)}\cos t_k\right), \tag{36}$$

with t_k defined as

$$t_k = 0.4 - \frac{6}{1 + \left(\boldsymbol{x}_k^{(1)}\right)^2 + \left(\boldsymbol{x}_k^{(2)}\right)^2}. \tag{37}$$

For this example, $u = 0.9$ to have chaotic behavior, and additive white Gaussian process noise with covariance matrix $\boldsymbol{Q}_k = 1\times10^{-2} I_{2\times2}$ is assumed. Every time step, it is assumed that a nonlinear measurement is available,

$$y_k = \sqrt{\left(\boldsymbol{x}_k^{(1)}\right)^2 + \left(\boldsymbol{x}_k^{(2)}\right)^2} + \eta_k, \tag{38}$$

where η_k is white Gaussian measurement noise with scalar covariance matrix $\boldsymbol{R}_k = 1$. Figure 6 shows the final noiseless range measurement for different values of u. Each point represents a different trajectory simulated from an initial condition sampled from $\boldsymbol{x}_0 \sim \mathcal{N}(\boldsymbol{0}_{2\times1}, I_{2\times2})$ and propagated for 50 time steps. From the figure, it can be seen that for $u \geq 0.6$, the Ikeda map is a chaotic system.

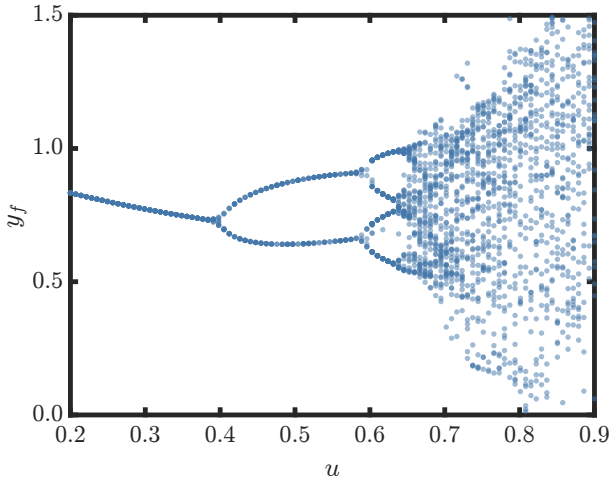

Fig. 6. Final range as a function of u for the Ikeda map.

For this sequential filtering example, 1000 Monte Carlo (MC) runs are propagated for 50 time steps, starting from an initial state $\boldsymbol{x}_0 \sim \mathcal{N}(\mathbf{0}_{2\times 1}, I_{2\times 2})$. Each reduction technique is tested within the SMF framework. The different filters used are summarized below.

1) **SMF-OR:** This filter uses the optimal reduction approach, sampling deterministic points from the GSF-updated points and minimizing the MCVMD to reduce them.
2) **SMF-SK ($N_i = 1$):** This filter uses the SK algorithm to reduce the deterministic points sampled from the GSF-updated points.
3) **SMF-SK ($N_i = i$):** This filter also uses the SK algorithm to reduce the deterministic points sampled from the GSF-updated points, but iterates i times between $\boldsymbol{T}_k^{\star,\mathrm{SK}}$ and $\boldsymbol{\mathcal{X}}$.

Note that the OT solution is not presented, as from the previous example, it was shown that the SK approximation can yield a better PM representation with a faster runtime. For all filters, $M = 125$ and $N = 25$. For the SMF-OR, the MCVMD is minimized using Matlab's `fminunc` function with default settings. In the case of both versions of the SMF-SK, $\varepsilon^2 = 1 \times 10^{-2}$ and $\lambda = 500$. For all reduction techniques, $\boldsymbol{\mathcal{X}}$ is initialized at the means of the GSF-approximated posterior, $\widetilde{\boldsymbol{\mathcal{X}}}_{k|k}$, and Silverman's rule of thumb is scaled by $\alpha = 0.4$.

In this example, two metrics are used to evaluate the performance of the filters. The root mean square error (RMSE) measures the accuracy of filters. The RMSE for this work is calculated as

$$\mathrm{RMSE}(k) = \sum_{j=1}^{n_{\mathrm{MC}}} \frac{1}{n_{\mathrm{MC}}} \sqrt{\frac{1}{n_x} \sum_{i=1}^{n_x} \left(\boldsymbol{x}_{k,j}^{(i)} - \hat{\boldsymbol{x}}_{k|k,j}^{(i)} \right)^2}, \quad (39)$$

where n_{MC} is the number of MC runs, n_x is the state dimension, $\boldsymbol{x}_{k,j}^{(i)}$ is the true state, and $\hat{\boldsymbol{x}}_{k|k,j}^{(i)}$ is the estimated state. The scaled normalized estimation error squared (SNEES) measures the consistency of the filter, obtained by

$$\mathrm{SNEES}(k) = \sum_{j=1}^{n_{\mathrm{MC}}} \frac{1}{n_{\mathrm{MC}}} \boldsymbol{e}_{k,j}^{\mathrm{T}} \left(P_{k|k,j} \right)^{-1} \boldsymbol{e}_{k,j}, \quad (40)$$

where

$$\boldsymbol{e}_{k,j} = \boldsymbol{x}_{k,j} - \hat{\boldsymbol{x}}_{k|k,j}. \quad (41)$$

A SNEES of one indicates a consistent filter. If the SNEES is lower than one, the filter is conservative. Alternatively, if the SNEES is higher than one, the filter is overconfident. [1]. As the SNEES is highly sensitive to outliers, any value greater than 1×10^3 was disregarded as a numerical instability for this example.

Figure 7 shows the RMSE vs. time steps for the SMF-OR and the two versions of the SMF-SK. The green line shows the approximate best performance in the current example, achieved with a regularized particle filter using 1×10^4 particles. As it can be seen, not iterating over $\boldsymbol{T}_k^{\star,\mathrm{SK}}$ and $\boldsymbol{\mathcal{X}}$, that is, $N_i = 1$, produces a slightly higher error in the SMF-SK compared to the SMF-OR. If $N_i = 5$, the SK approximation closely resembles the SMF-OR, with both solutions converging almost to Bayesian inference.

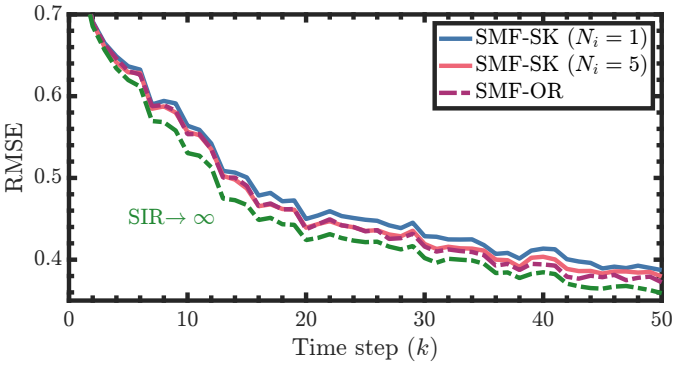

Fig. 7. RMSE vs. time steps for the SMF-OR and the SMF-SK ($N_i = i$). The dashed green line shows the RMSE obtained with a regularized particle filter with 1×10^4 particles.

A similar conclusion can be drawn from Fig. 8. The green dashed line in this figure shows an SNEES of one. The SMF-OR results show that minimizing the MCVMD yields the best overall consistency. However, performing more iterations on $\boldsymbol{T}_k^{\star,\mathrm{SK}}$ and $\boldsymbol{\mathcal{X}}$ improves consistency for the SMF-SK. Table I shows the time-averaged RMSE, SNEES, and runtime for each strategy. The runtime presented in this table represents the time required for a single propagation and update. Additionally, the runtime values are normalized with respect to the SMF-SK ($N_i = 1$) runtime, as this filter runs the fastest.

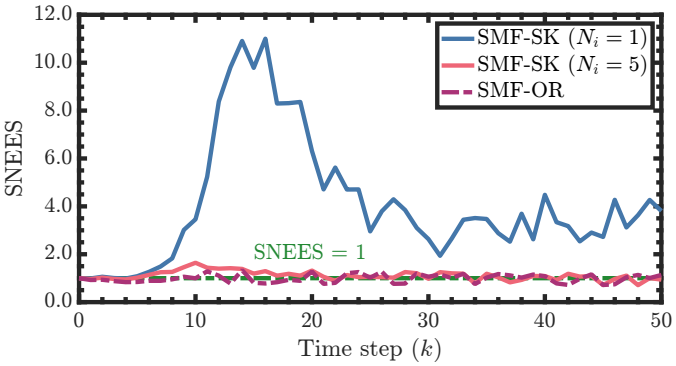

Fig. 8. SNEES vs. time steps for the SMF-OR and the SMF-SK ($N_i = i$). The green dashed lines shows a SNEES of one.

The SMF-SK ($N_i = 5$) is only slightly slower than the SMF-SK ($N_i = 1$), yet it achieves similar accuracy and consistency values to those of the SMF-OR. For this scenario, the SMF-OR is the most accurate and consistent filter. However, it runs the slowest, as it has to minimize the MCVMD for each reduction step. It is important to note that there are various techniques to speed up the runtime of the SMF-OR, but as presented, the SMF-SK ($N_i = 5$) provides the best balance between accuracy, consistency, and runtime.

TABLE I
TIME AVERAGED RMSE, SNEES, AND RUNTIME FOR EACH STRATEGY

Filter	RMSE	SNEES	Runtime
SMF-SK ($N_i = 1$)	0.4862	4.0909	1.0000
SMF-SK ($N_i = 5$)	0.4774	1.1074	3.3440
SMF-OR	0.4751	0.9889	54.2780

VI. CONCLUSIONS

State estimation for nonlinear systems requires filters that numerically approximate the solution to the Bayesian recursive relations. These filtering strategies are expected to be accurate and consistent, and in some applications, verifiable and computationally feasible. Filters that parametrize the state probability density function as Gaussian mixtures and point masses, such as the ensemble Gaussian mixture filter, have shown to be an accurate and consistent solution to the state estimation problem. However, they often rely on stochastic resampling methods, which makes the filter not verifiable.

In recent work, an improved version of the Silverman mass filter, equivalent to a deterministic ensemble Gaussian mixture filter, was developed [15]. In this new filter, the resampling step was made deterministic with the sampling of point masses from the posterior Gaussian mixture by minimizing a distance measure. Although this new technique improved the filtering outcomes by providing a verifiable solution, the computational cost for the sampling step was not considered.

This work presents a new strategy for deterministically sampling point masses from a Gaussian mixture, based on an approximate solution to the optimal transport problem. By solving the approximate optimal transport problem via an iterative Sinkhorn-Knopp algorithm, the sampling step can be performed faster, while preserving accuracy and consistency. This new strategy was tested on a bivariate static example and a sequential filtering problem involving the Ikeda map. In both cases, using the iterative Sinkhorn-Knopp algorithm was shown to be less computationally expensive than previous work, while still achieving comparable results.

This new sampling technique presents a filtering strategy that is accurate, consistent, verifiable, and computationally feasible. Future work will implement this new strategy in onboard scenarios, such as spacecraft navigation, to validate its use in higher dimensional and realistic scenarios.

REFERENCES

[1] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan, *Estimation with Applications to Tracking and Navigation*. John Wiley & Sons, Inc, 2001.

[2] S. K. Biswas, "Computationally efficient non-linear kalman filters for on-board space vehicle navigation," Ph.D. dissertation, UNSW, 2017.

[3] A. A. Popov and R. Zanetti, "Deterministic optimal transport-based gaussian mixture particle filtering for verifiable applications," *arXiv preprint arXiv:2501.17302*, 2025.

[4] F. Giraldo-Grueso, A. A. Popov, and R. Zanetti, "Precision mars entry navigation with atmospheric density adaptation via neural networks," *Journal of Aerospace Information Systems*, vol. 21, no. 12, 2024.

[5] ——, "Adaptive mars entry guidance with atmospheric density estimation," in *Proceedings of the AAS/AIAA Astrodynamics Specialist Conference*, Broomfield, Colorado, Aug. 2024.

[6] R. Zanetti, "Advanced navigation algorithms for precision landing," Ph.D. dissertation, The University of Texas at Austin, 2007.

[7] R. Bucy and K. Senne., "Digital synthesis of non-linear filters," *Automatica*, vol. 7, no. 3, pp. 287–298, 1971.

[8] N. Bergman, "Recursive bayesian estimation: Navigation and tracking applications," Ph.D. dissertation, Linköping University, 1999.

[9] M. Šimandl, J. Královec, 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.

[10] ——, "Advanced point-mass method for nonlinear state estimation," *Automatica*, vol. 42, no. 7, pp. 1133–1145, 2006.

[11] F. Giraldo-Grueso, A. A. Popov, and R. Zanetti, "Gaussian mixture-based point mass filtering with applications to terrain relative navigation," *IEEE Transactions on Aerospace and Electronic Systems*, pp. 1–15, 2024.

[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] B. W. Silverman, *Density estimation for statistics and data analysis*. Routledge, 2018.

[15] F. Giraldo-Grueso, A. A. Popov, U. D. Hanebeck, and R. Zanetti, "Optimal grid point sampling for point mass filtering," in *Proceedings of the AAS/AIAA Spaceflight Mechanics Meeting*, Kaua'i, Hawaii, Jan. 2025.

[16] 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.

[17] A. A. Popov and R. Zanetti, "An adaptive covariance parameterization technique for the ensemble gaussian mixture filter," *arXiv preprint arXiv:2212.10323*, 2022.

[18] ——, "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.

[19] ——, "Ensemble-localized kernel density estimation with applications to the ensemble gaussian mixture filter," *arXiv preprint arXiv:2308.14143*, 2023.

[20] V. I. Bogachev and A. V. Kolesnikov, "The Monge-Kantorovich problem: achievements, connections, and perspectives," *Russian Mathematical Surveys*, vol. 67, no. 5, pp. A2013–A2024, 2012.

[21] S. Reich, "A nonparametric ensemble transform method for bayesian inference," *SIAM Journal on Scientific Computing*, vol. 35, no. 4, pp. A2013–A2024, 2013.

[22] W. Acevedo, J. De Wiljes, and S. Reich, "Second-order accurate ensemble transform particle filters," *SIAM Journal on Scientific Computing*, vol. 39, no. 5, pp. A1834–A1850, 2017.

[23] A. A. Popov, A. N. Subrahmanya, and A. Sandu, "A stochastic covariance shrinkage approach to particle rejuvenation in the ensemble transform particle filter," *Nonlinear Processes in Geophysics*, vol. 29, no. 2, pp. 241–253, 2022.

[24] M. Cuturi, "Sinkhorn distances: Lightspeed computation of optimal transportation distances," in *Advances in Neural Information Processing Systems*, 2013.

[25] F. Giraldo-Grueso, A. A. Popov, and R. Zanetti, "Gaussian mixture-based point mass filtering," in *Proceedings of the 2024 27th International Conference on Information Fusion (FUSION)*, Venice, Italy, Jul. 2024.

[26] D. Frisch and U. D. Hanebeck, "Deterministic gaussian sampling with generalized fibonacci grids," in *Proceedings of the 2021 IEEE International Conference on Information Fusion (FUSION)*, Sun City, South Africa, Nov. 2021.

[27] U. D. Hanebeck, "Optimal reduction of multivariate dirac mixture densities," *arXiv preprint arXiv:1411.4586*, 2014.

[28] A. Corenflos, J. Thornton, G. Deligiannidis, and A. Doucet, "Differentiable particle filtering via entropy-regularized optimal transport," *arXiv preprint arXiv:2102.07850*, 2021.

[29] 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.

[30] 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.