ADAPTIVE DETERMINISTIC ENSEMBLE GAUSSIAN MIXTURE FILTERING FOR CISLUNAR TRACKING

Jakub Matoušek[*], Dalton Durant[†], and Renato Zanetti[‡]

The state estimation for navigation and tracking remains a central research topic in astrodynamics. The Ensemble Gaussian Mixture Filter has demonstrated strong performance, particularly in scenarios with long measurement outages, while retaining low computational complexity for on-board implementation. However, the standard EnGMF formulation is stochastic, and existing deterministic variants incur prohibitively high computational costs. Furthermore, an adaptive mechanism for selecting the number of ensembles is required to avoid unnecessary computational burden when additional samples provide diminishing improvements in estimation accuracy. The objective of this paper is therefore to develop a computationally efficient deterministic EnGMF with an adaptively selected ensemble size.

Keywords: ensemble Gaussian mixture filter, deterministic sampling, tracking

INTRODUCTION

State estimation of discrete-time stochastic dynamic systems from noisy measurements has been a cornerstone of navigation and tracking algorithms. Following the Bayesian approach, a general solution to the state estimation problem is given by the Bayesian recursive relations (BRRs), which compute the probability density functions (PDFs) of the state conditioned on the available measurements. These conditional PDFs provide a full statistical description of the state, which is not directly measurable.

However, the BRRs are analytically tractable only for a restricted class of models-typically those with linear dynamics and measurement relations. Classical example of an analytically derived filter is the Kalman filter.[1] For general nonlinear or non-Gaussian models, approximate solutions to the BRRs must be employed. These approximate estimation (or filtering) methods are commonly categorised according to the validity region of their underlying assumptions into local and global filters.[2]

Local (Gaussian) filters (LFs) provide efficient Gaussian-assumed estimates in the form of the conditional mean and covariance matrix only. Due to the underlying approximations, LFs can become unstable for models with higher non-linearity. The main representatives of this class are the extended and unscented Kalman filters.[3]

[*]Postdoctoral Fellow, Oden Institute for Computational Engineering and Sciences, The University of Texas at Austin, TX 78712, USA. Email: jakub.matousek@austin.utexas.edu.

[†]Postdoctoral Fellow, Department of Aerospace Engineering and Engineering Mechanics, The University of Texas at Austin, TX 78712, USA. Email: ddurant@utexas.edu

[‡]Associate Professor, Department of Aerospace Engineering and Engineering Mechanics, The University of Texas at Austin, TX, 78712. Email: renato@utexas.edu

Global filters (GFs), such as particle filters,[4] grid-based filters,[5] or the Gaussian sum filter,[6] provide estimates in the form of full conditional PDFs without imposing restrictive assumptions on the distribution family. These estimators can handle strongly nonlinear or non-Gaussian systems, however at the cost of increased computational demands.

This paper focuses on the Ensemble Gaussian Mixture Filter (EnGMF).[7,8] The EnGMF alternates between a Dirac-mixture PDF representation (as in a particle filter) and a Gaussian-mixture representation (as in the Gaussian sum filter). This hybrid structure preserves the robustness and accuracy associated with global methods while requiring significantly fewer samples/points, thus achieving a computational complexity closer to that of a local method. The filter is also notably robust in scenarios with long measurement outages-situations frequently encountered in e.g. cislunar tracking.

The main drawbacks of the standard EnGMF in this context are its stochastic nature, inherited from particle filters, and its fixed number of ensembles. In this paper, we address both challenges by developing an EnGMF variant that is fully deterministic and adaptively selects both the number of ensembles and the number of Gaussian-mixture components based on the current information content of the estimated state. The resulting method is deterministic, accurate, sufficiently fast for real-world applications and amenable to certification due to its deterministic structure.

For completeness, we summarise the state of the art related to these two issues. Existing efforts to enforce determinism within the EnGMF framework include:

- Sampling a Fibonacci grid for each Gaussian component, followed by an optimal-transport-based reduction of the ensemble size.[9]
- Employing projected cumulative distributions that are compared using the Cramér-von Mises distance.[10]

However, both methods suffer from high computational complexity and uncertain efficiency and accuracy when applied to high-dimensional problems. A way to make these methods readily compatible with an adaptive approach is also not apparent.

Regarding adaptivity, the following approaches have been proposed:

- Performing pruning, merging, and capping of Gaussian mixture components, after which the same number of Dirac ensembles are sampled as the number of remaining components.[11]
- Conducting a statistical test over multiple time steps that compares the predictive measurement PDF $p(\mathbf{z}_k \mid \mathbf{z}^{k-1})$ with the actual measurement and adjusts the number of ensembles accordingly.[7]

Building on the first line of work listed above, which applied pruning, merging, and capping as a separate mixture-editing stage, the present paper instead integrates these operations directly into the resampling procedure, resulting in a more coherent and theoretically grounded update. The second work adjusts the ensemble size only across multi-step windows; if the outcome is unsatisfactory, then the algorithm reprocesses the entire window, which reduces responsiveness and increases computational cost.

In contrast to the aforementioned state-of-the-art methods, the approach proposed here is fast, deterministic, and theoretically grounded. It introduces an adaptive mechanism that incrementally

adds samples within each iteration and evaluates a stopping condition to determine when sufficient convergence has been achieved.

MODEL DEFINITION AND STATE ESTIMATION

The following discrete-time state-space model of a nonlinear stochastic dynamic system with additive noises

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

$$\mathbf{z}_k = \mathbf{h}_k(\mathbf{x}_k) + \mathbf{v}_k, \tag{2}$$

is considered, where the vectors $\mathbf{x}_k \in \mathbb{R}^{n_x}$, and $\mathbf{z}_k \in \mathbb{R}^{n_z}$ represent the *unknown* state of the model and measurement at time instant k, respectively. The state function $\mathbf{f}_k : \mathbb{R}^{n_x} \to \mathbb{R}^{n_x}$, and measurement function $\mathbf{h}_k : \mathbb{R}^{n_x} \to \mathbb{R}^{n_z}$ are supposed to be *known*. Particular realisations of the state and measurement noises $\mathbf{w}_k$ and $\mathbf{v}_k$ are *unknown*, but their PDFs, i.e., the state noise PDF $p_{\mathbf{w}_k}(\mathbf{w}_k)$ and the measurement noise PDF $p_{\mathbf{v}_k}(\mathbf{v}_k)$, are supposed to be *known*, as well as the initial state PDF $p_{\mathbf{x}_0}(\mathbf{x}_0)$. The noises and initial state are independent.

Following the state-space formulation, the filtering task, can be formalised as an estimation of the state $\mathbf{x}_k$ based on the available measurements $\mathbf{z}^k = [\mathbf{z}_0, \mathbf{z}_1, \ldots, \mathbf{z}_k]$, and the state-space model (1), (2). The *Bayesian* estimation infers the PDF of the state conditioned on available measurements.

The *prior* PDF $p(\mathbf{x}_k|\mathbf{z}^{k-1})$ represents the predicted state distribution before processing the current measurement and reflects all information available up to time k–1. The *posterior* PDF $p(\mathbf{x}_k|\mathbf{z}^k)$ represents the updated state distribution after incorporating the measurement at time k.

STANDARD ENSEMBLE-GAUSSIAN MIXTURE FILTER

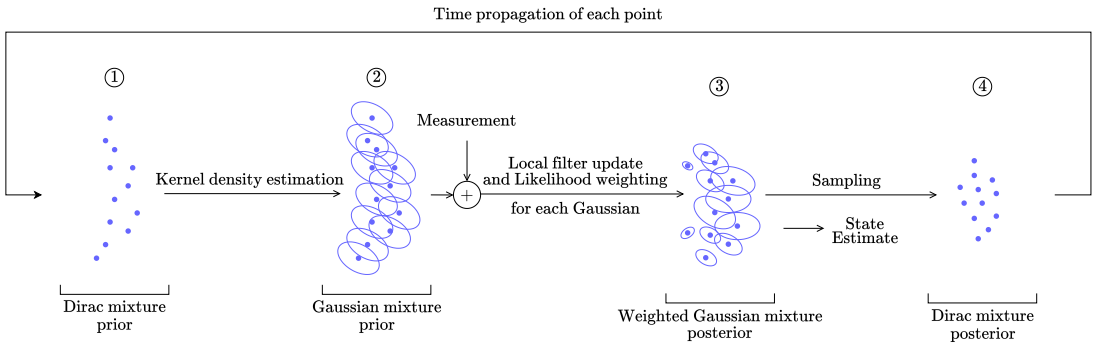

Figure 1. Diagram of standard EnGMF.

The diagram of steps that are generally performed in EnGMF can be seen in Fig. 1. The steps are further mathematically described with references to this Fig. such as ①.

The standard Ensemble Gaussian mixture filter (EnGMF) combines the features of Particle Filter (PF) and Gaussian sum filter (GSF). The measurement update starts with an equally weights prior

(or initial condition) Dirac mixture (DM) ①

$$\overline{p}(\mathbf{x}_{k+1}|\mathbf{z}^k) = \sum_{i=1}^{N} \overline{\omega}_{k+1|k}^{(i)} \, \delta\left(\mathbf{x} - \boldsymbol{\xi}_{k+1|k}^{(i)}\right), \tag{3}$$
$$\overline{\omega}_{k+1|k}^{(i)} = \tfrac{1}{N}, \; \forall i,$$

where $\overline{\omega}_{k+1|k}^{(i)}$ are constant weights for normalisation purposes, N is the number of and $\boldsymbol{\xi}_{k+1|k}^{(i)}$ are the positions of Dirac deltas.

However, to achieve an accurate Bayesian update with improved smoothness and reduced sampling noise, the EnGMF replaces the discrete Dirac mixture (DM) update of the PF with a continuous, smooth Gaussian mixture (GM) based update. Each ensemble member generates a Gaussian which then undergoes a measurement update, producing a posterior GM approximation that is smoother and more reliable than that of a particle filter, which typically requires far more samples to achieve comparable accuracy.

To obtain a Gaussian-mixture (GM) prior from a Dirac-mixture (DM) prior, kernel density estimation techniques are employed. The resulting GM typically consists of Gaussian components with a constant covariance defined by Silvermans rule of thumb.[12] Silvermans rule specifies the kernel bandwidth as a function of the sample covariance $\mathbf{P}_{k+1|k}$ and the number of samples N, minimising the mean integrated squared error for Gaussian data.[13] The kernel bandwidth is defined as

$$\mathbf{P}_s = \alpha \left(\frac{4}{(n_x+2)N} \right)^{2/(n_x+4)} \mathbf{P}_{k+1|k} + \mathbf{Q}_k, \tag{4}$$

where α is scaling parameter, and $\mathbf{P}_{k+1|k}$ is the sample covariance of prior DM (the ensembles are propagated without process noise realisations to keep deterministic behaviour). The prior GM ② is then given by

$$\widetilde{p}(x_{k+1}|z^k) = \sum_{i=1}^{N} \widetilde{\omega}_{k+1|k}^{(i)} \, \mathcal{N}(\mathbf{x}; \boldsymbol{\xi}_{k+1|k}^{(i)}, \mathbf{P}_s), \tag{5}$$
$$\widetilde{\omega}_{k+1|k}^{(i)} = \overline{\omega}_{k+1|k}^{(i)}.$$

Each component is then updated using a locally linearised measurement model. In this paper, linearisation is adopted for notational simplicity; however, a point-based approximation such as the unscented transform[14] can also be employed for higher accuracy. The measurement prediction and Jacobian are

$$\widehat{\mathbf{z}}_{k+1|k}^{(i)} = \mathbf{h}_{k+1}(\boldsymbol{\xi}_{k+1|k}^{(i)}), \tag{6}$$

$$\mathbf{H}_{k+1}^{(i)} = \left. \frac{\partial \mathbf{h}_{k+1}}{\partial \mathbf{x}} \right|_{\mathbf{x}=\boldsymbol{\xi}_{k+1|k}^{(i)}}. \tag{7}$$

Based on them, an innovation covariance, Kalman gain, and posterior mean and covariance of

each Gaussian component are calculated as (following the Gaussian sum filter (GSF) update[6])

$$\mathbf{S}_{k+1}^{(i)} = \mathbf{H}_{k+1}^{(i)} \mathbf{P}_{k+1|k}^{(i)} \mathbf{H}_{k+1}^{(i)^{\mathrm{T}}} + \mathbf{R}_{k+1}, \tag{8}$$

$$\mathbf{K}_{k+1}^{(i)} = \mathbf{P}_{k+1|k}^{(i)} \mathbf{H}_{k+1}^{(i)^{\mathrm{T}}} (\mathbf{S}_{k+1}^{(i)})^{-1}, \tag{9}$$

$$\mathbf{m}_{k+1|k+1}^{(i)} = \boldsymbol{\xi}_{k+1|k}^{(i)} + \mathbf{K}_{k+1}^{(i)} \left(\mathbf{z}_{k+1} - \widehat{\mathbf{z}}_{k+1|k}^{(i)} \right), \tag{10}$$

$$\begin{aligned} \mathbf{P}_{k+1|k+1}^{(i)} = & \left(\mathbf{I} - \mathbf{K}_{k+1}^{(i)} \mathbf{H}_{k+1}^{(i)} \right) \mathbf{P}_{k+1|k}^{(i)} \left(\mathbf{I} - \mathbf{K}_{k+1}^{(i)} \mathbf{H}_{k+1}^{(i)} \right)^{\mathrm{T}} \\ & + \mathbf{K}_{k+1}^{(i)} \mathbf{R}_{k+1} \mathbf{K}_{k+1}^{(i)^{\mathrm{T}}}. \end{aligned} \tag{11}$$

Then the mixture weights are updated proportionally to the component likelihoods

$$\widehat{\omega}_{k+1|k+1}^{(i)} = \widetilde{\omega}_{k+1|k}^{(i)} \mathcal{N}\left(\mathbf{z}_{k+1}; \widehat{\mathbf{z}}_{k+1|k}^{(i)}, \mathbf{S}_{k+1}^{(i)} \right), \tag{12}$$

$$\widetilde{\omega}_{k+1|k+1}^{(i)} = \frac{\widehat{\omega}_{k+1|k+1}^{(i)}}{\sum_{j=1}^{N} \widehat{\omega}_{k+1|k+1}^{(j)}}, \ \forall i.$$

Finally, the posterior Gaussian mixture (GM) ③ is given by

$$\begin{aligned} & \widetilde{p}(\mathbf{x}_{k+1}|\mathbf{z}^{k+1}) = \\ & \sum_{i=1}^{N} \widetilde{\omega}_{k+1|k+1}^{(i)} \mathcal{N}\left(\mathbf{x}_{k+1}; \mathbf{m}_{k+1|k+1}^{(i)}, \mathbf{P}_{k+1|k+1}^{(i)} \right). \end{aligned} \tag{13}$$

From this posterior GM the first two moments can be calculated as

$$\hat{\mathbf{x}}_{k+1|k+1}^{\mathrm{E}} = \sum_{i=1}^{N} \widetilde{\omega}_{k+1|k+1}^{(i)} \mathbf{m}_{k+1|k+1}^{(i)}, \tag{14}$$

$$\begin{aligned} \mathbf{P}_{k+1|k+1}^{\mathrm{E}} = & \sum_{i=1}^{N} \widetilde{\omega}_{k+1|k+1}^{(i)} \\ & \left[\mathbf{P}_{k+1|k+1}^{(i)} + (\mathbf{m}_{k+1|k+1}^{(i)} - \hat{\mathbf{x}}_{k+1|k+1})(\cdot)^{\mathrm{T}} \right]. \end{aligned} \tag{15}$$

These two moments are then reported as the estimate.

However, instead of maintaining explicit continuous Gaussian components, the EnGMF then samples equally weighted DM ④ out of the GM posterior (13), leading to

$$\overline{p}(\mathbf{x}_{k+1}|\mathbf{z}^{k+1}) = \sum_{i=1}^{N} \overline{\omega}_{k+1|k+1}^{(i)} \delta\left(\mathbf{x} - \boldsymbol{\xi}_{k+1|k+1}^{(i)} \right) \tag{16}$$

$$\overline{\omega}_{k+1|k+1}^{(i)} = \tfrac{1}{N}, \ \forall i,$$

where $\boldsymbol{\xi}_{k+1|k+1}^{(i)}$ are samples from the GM posterior (13). This replacement of continuous Gaussian components with Dirac delta samples greatly reduces computational complexity and avoids the exponential growth of mixture components typical in the GMF, while still capturing non-Gaussian and multimodal features of the posterior.

There are many ways to sample from a Gaussian mixture (GM). Some approaches are based on optimisation.[15,16] These methods, however, are typically computationally expensive, difficult to scale to higher dimensions, and raise the practical question of whether generating a large number of lower-quality samples would not be preferable to producing only a few high-quality ones at comparable computational cost. Other approaches rely on stochastic sampling,[17] which introduces challenges related to verifiability and reproducibility. There also exist deterministic methods that match only the first two moments of all Gaussian components jointly,[18] but these do not preserve the full structure of the mixture.

Lastly, the time-update is performed simply by propagation of the Dirac delta points (16) using model dynamics (1) leading to a new DM prior for next time step $k+2$ (back at ① in recursion)

$$\begin{aligned} \overline{p}(x_{k+2}|z^{k+1}) &= \sum_{i=1}^{N} \overline{\omega}^{(i)}_{k+2|k+1}\, \delta\Big(x - \mathbf{f}(\boldsymbol{\xi}^{(i)}_{k+1|k+1})\Big) \\ \overline{\omega}^{(i)}_{k+2|k+1} &= \frac{1}{N},\ \forall i. \end{aligned} \tag{17}$$

PROPOSED DETERMINISTIC ADAPTIVE ENGMF

In this section we first introduce the deterministic sampling method for EnGMF based on a Sobol sequence, and subsequently discuss stopping criteria used within the adaptive sampling procedure. After that the whole algorithm is presented.

Deterministic Sobol Sampling

Deterministic Quasi-Monte Carlo sampling based on Sobol sequences (SoS) has already been employed in particle filters. Its main advantages are determinism and the improved convergence rate resulting from the more uniform distribution of the generated samples. An implementation for particle filters, together with instructions on generating Sobol sequences, is provided in.[19] Sobol generators are available in many programming languages; in this work we rely on the *sobolset* object in MATLAB. The focus of this section is therefore to demonstrate how such Sobol sets can be used to deterministically sample from a Gaussian mixture in EnGMF algorithm.

Let us consider an (n_x+1)-dimensional SoS, filling the hypercube $[0,1]^{n_x+1}$, with elements

$$\begin{bmatrix} s_1^{(i)} \\ \mathbf{s}^{(i)} \end{bmatrix} \in [0,1]^{n_x+1}, \qquad i = 1, \ldots, N_s. \tag{18}$$

The first dimension $s_1^{(i)}$ is used to sample the index of the Gaussian mixture (13) component as

$$j^{(i)} = \min\big\{ j \in \{1, \ldots, N_g\} : s_1^{(i)} \leq c^{(j)} \big\}, \tag{19}$$

where N_g is the number of modes of the GM, c is the cumulative sum of the posterior GM weights $\widetilde{\omega}^{(j)}_{k+1|k+1}$,

$$c^{(j)} = \sum_{\ell=1}^{j} \widetilde{\omega}^{(\ell)}_{k+1|k+1}, \qquad j = 1, \ldots, N_g. \tag{20}$$

The state-space coordinates $\mathbf{s}^{(i)}$ are next mapped to standard normal variables via the componentwise inverse CDF,

$$\widehat{\mathbf{s}}^{(i)} = \Phi^{-1}\big(\mathbf{s}^{(i)}\big),\ \forall i., \tag{21}$$

where Φ^{-1} is the inverse of a Gaussian CDF. The sequence $\widehat{\mathbf{s}}^{(i)}$ is then whitened: we subtract its empirical mean and left-multiply by the inverse square root of its empirical covariance,

$$\overline{\mathbf{s}}^{(i)} = \sqrt{\mathbf{P}_{\widehat{\mathbf{s}}}}^{-1}\big(\widehat{\mathbf{s}}^{(i)} - \mu_{\widehat{\mathbf{s}}}\big), \tag{22}$$

where $\sqrt{\mathbf{P}_{\widehat{\mathbf{s}}}}$ denotes any matrix square root (e.g., its Cholesky factor). The resulting samples have zero mean and identity covariance by construction.

Finally, we construct a Dirac mixture approximation of the posterior Gaussian mixture by transforming each whitened Sobol sample according to its deterministically selected mixture index (19)

$$\boldsymbol{\xi}^{(i)}_{k+1|k+1} = \sqrt{\mathbf{P}^{(j^{(i)})}_{k+1|k+1}}\ \overline{\mathbf{s}}^{(i)} + \mathbf{m}^{(j^{(i)})}_{k+1|k+1}. \tag{23}$$

This yields the Dirac mixture (16) that is further used for simple time propagation.

Implementation tips

- The Sobol sequence can be generated very efficiently for each time step k, but it may also be precomputed offline as it does not depend on the model until the final moment-based transformation (23).
- For computational efficiency, MATLAB's `erfinv` may be used instead of the inverse CDF in (21).
- Some mixture components may not be selected at all, when others may be selected multiple times. Accounting for this while calculating the matrix square root $\sqrt{\mathbf{P}^{(j^{(i)})}_{k+1|k+1}}$ improves computational efficiency.

Adapting Gaussian Mixture Component Number

Generally, the goal of state estimation is to obtain an accurate and statistically consistent state estimate (14)(15), i.e. a high-quality posterior Gaussian mixture (13)③. Increasing the number of Gaussian components in the GM representation, or the number of Dirac samples in the DM representation, generally improves accuracy. However, the computational burden must remain low enough for real-time on-board execution, and moreover, beyond a certain point further increasing the number of GM/DM components yields only diminishing returns. For this reason, the objective is to construct the posterior GM (13)③ incrementally - Gaussian by Gaussian - until the mixture stabilises.

Ideally, the most principled indicator of whether the information content of the posterior GM is still changing (with addition of further modes) would be its differential entropy

$$\begin{aligned} H = \int \left(\sum_{i=1}^{N_g} \widetilde{\omega}^{(i)}_{k+1|k+1}, \mathcal{N}(\mathbf{x}_{k+1}; \mathbf{m}^{(i)}_{k+1|k+1}, \mathbf{P}^{(i)}_{k+1|k+1}) \right) \\ \log \left(\sum_{i=1}^{K} \widetilde{\omega}^{(i)}_{k+1|k+1}, \mathcal{N}_{\}}(\mathbf{x}_{k+1}; \mathbf{m}^{(i)}_{k+1|k+1}, \mathbf{P}^{(i)}_{k+1|k+1}) \right) \mathrm{d}\mathbf{x}_{k+1}, \end{aligned} \tag{24}$$

where N_g denotes the current number of Gaussian modes. However, the entropy of a Gaussian mixture (GM) (24) cannot be evaluated in a closed form and must therefore be approximated. The most direct-yet computationally most demanding-approach is Monte-Carlo integration; other approximations also exist, such as Taylor-series expansions.[20] Using such an entropy-based stopping condition would be computationally so expensive that one might just as well select a conservatively large number of mixture components.

Our goal is to demonstrate that, for the purposes of the proposed adaptive EnGMF, using the exact GM entropy as a stopping criterion is unnecessarily complicated. Instead, a lighter alternative criterion can be employed without significantly compromising the quality of the estimation.

A computationally light surrogate is the differential entropy of a single Gaussian with the same mean and covariance as the current posterior GM,

$$H = \frac{1}{2}\Big(n_x \log(2\pi e) + \log|\mathbf{P}^{1:N_g}_{k+1|k+1}|\Big), \tag{25}$$

where $\mathbf{P}^{1:N_g}_{k+1|k+1}$ denotes the covariance of the posterior GM constructed from the first N_g modes. This surrogate entropy is inexpensive to evaluate and provides a consistent indicator for stabilisation of the GM. To further reduce computational effort, one may compute only the diagonal of $\mathbf{P}^{1:N_g}_{k+1|k+1}$ and use it in the surrogate entropy formula.

Adapting Dirac Mixture Ensemble Number

Apart from adaptively selecting the number of posterior Gaussian mixture components ③, the number of sampled posterior Dirac ensembles ④ should also be adapted to ensure that the local filter updates remain valid, i.e., that the measurement function is approximately linear within each Gaussian mixture prior component. Ideally, one would adaptively draw Dirac deltas, compute the kernel bandwidth, and verify the near-linearity condition for every GM component. In practice, however, this is not feasible: the Silverman bandwidth cannot be evaluated before the sample covariance is available, and performing linearity checks for all components is computationally prohibitive.

We explored several theoretically motivated nonlinearity checks, but they were either too expensive or did not lead to considerable improvements. Ultimately, the only approach that consistently performed well across different parameter settings was a simple empirical rule: draw $\lceil r_d N_{\text{ESS}} \rceil$ samples, where r_d is a user-defined parameter and N_{ESS} is the effective sample size of the posterior mixture weights we are sampling from (13), computed as

$$N_{\text{ESS}} = \frac{1}{\sum_{i=1}^{N} \left(\omega^{(i)}_{k+1|k+1}\right)^2}. \tag{26}$$

Algorithm

The diagram of the proposed adaptive deterministic EnGMF is shown in Fig. 2. Owing to the algorithmic structure, the diagram now begins on the left with the GM posterior; however, the numbering is retained for consistency with the previous diagram and to facilitate comparison. An overview of the proposed estimator follows.

Adaptive Deterministic EnGMF

Parameters: Entropy stopping condition s_H, Silverman scale parameter α, number of ensembles for initialisation N_1, parameter for number of ensembles r_d, adaptation step for GM Δ_{GM}
Inputs: Initial condition mean $\hat{\mathbf{x}}_{0|0}$ and covariance matrix $\mathbf{P}_{0|0}$, measurements $\mathbf{z}^k$

Steps:

0. *Initialisation:* Set time step $k = 1$. Draw N_1 posterior DM (16) samples $\boldsymbol{\xi}_{k|k}$ from the initial condition to form DM posterior. Go to step 2.

1. *Ensemble Sampling:* Sample $\lceil r_d N_{\text{ESS}} \rceil$ (26) samples $\boldsymbol{\xi}_{k|k}$ from posterior GM distribution (13) using sobol sequence to obtain posterior DM (16).

2. *Ensemble propagation:* Deterministically propagate the ensembles $\boldsymbol{\xi}_{k|k}$ of posterior DM (16) through the model dynamics to form prior DM with ensembles $\boldsymbol{\xi}_{k+1|k}$ (17).

3. *Kernel density estimation:* Perform KDE with parameter α on the prior DM ensembles $\boldsymbol{\xi}_{k+1|k}$ to form prior GM (5).

4. *Recursive Local Updates:* Perform the local filter updates for next Δ_{GM} components of the prior GM (5). And save entropy of the resulting GM. Repeat until change of the entropy is less than s_H. Then set the resulting PDF as posterior GM (13).

5. *Record Estimates*: Calculate first two moments of resulting posterior GM as the estimate. Set $k = k + 1$ and go to step 1.

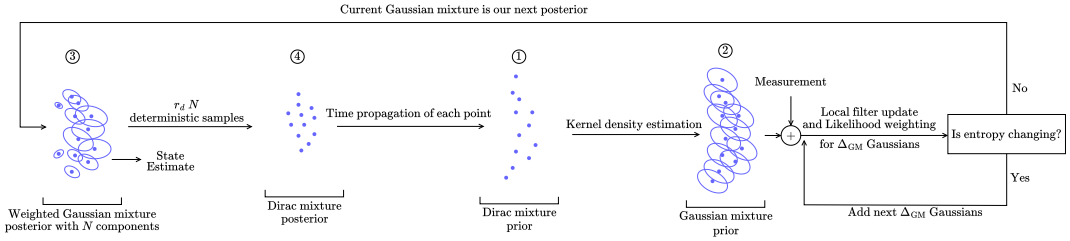

Figure 2. Diagram of proposed adaptive deterministic EnGMF.

CISLUNAR TRACKING

The estimators are evaluated in a cislunar relative-navigation scenario identical in configuration to,[11] but with modified parameter settings. A Target and a Chaser spacecraft fly on nearby Near Rectilinear Halo Orbits (NRHOs) in the Earth–Moon system. The Chaser is equipped with a vision-based sensor (e.g., a camera) that provides noisy line-of-sight angular measurements of the Target. Assuming perfect knowledge of its own inertial position and velocity, the estimation objective is the relative position and velocity of the Target with respect to the Chaser.

NRHOs are periodic trajectories in the Earth–Moon system that exhibit favorable long-term stability properties. At the same time, cislunar motion near NRHOs is known to be sensitive to initial

conditions and perturbations, which makes accurate estimation under sparse measurements challenging.[21,22]

Dynamics

A 9:2 synodic resonant NRHO is considered, meaning that the Target and Chaser complete nine revolutions around the Moon over two lunar synodic months. The orbit has an approximate 6.5-day period, a perilune radius of about $3{,}250\,\mathrm{km}$, and an apolune radius of approximately $71{,}000\,\mathrm{km}$. This configuration is commonly used as a baseline for Gateway-like cislunar missions.[23]

The translational dynamics are modeled using the Circular Restricted Three-Body Problem (CR3BP) for the Earth–Moon system in a synodic barycentric frame. The Earth and Moon are treated as point-mass primaries with gravitational parameters

$$\mu_\oplus = Gm_\oplus, \qquad \mu_{\leftmoon} = Gm_{\leftmoon},$$

where G is the gravitational constant and $m_\oplus$ and $m_{\leftmoon}$ denote the masses of the Earth and Moon, respectively. The dimensionless CR3BP mass parameter is defined as

$$\mu \triangleq \frac{\mu_{\leftmoon}}{\mu_\oplus + \mu_{\leftmoon}}.$$

Each spacecraft is described by the six-dimensional state

$$\mathbf{x} = [\mathbf{r}^\top, \mathbf{v}^\top]^\top \in \mathbb{R}^6,$$

where $\mathbf{r} \in \mathbb{R}^3$ and $\mathbf{v} \in \mathbb{R}^3$ are the Cartesian position and velocity in the synodic frame. In this frame, the Earth and Moon are fixed on the $\mathbf{r}(1)$-axis at the locations $[-\mu, 0, 0]^\top$ and $[1-\mu, 0, 0]^\top$, respectively. The distances to the primaries are defined as

$$r_\oplus = \sqrt{(\mathbf{r}(1)+\mu)^2 + \mathbf{r}(2)^2 + \mathbf{r}(3)^2}, \tag{27}$$

$$r_{\leftmoon} = \sqrt{(\mathbf{r}(1)-1+\mu)^2 + \mathbf{r}(2)^2 + \mathbf{r}(3)^2}. \tag{28}$$

The CR3BP equations of motion are written compactly as

$$\begin{aligned} \dot{\mathbf{r}} &= \mathbf{v}, \\ \dot{\mathbf{v}} &= \begin{bmatrix} \mathbf{r}(1) + 2\mathbf{v}(2) \\ \mathbf{r}(2) - 2\mathbf{v}(1) \\ 0 \end{bmatrix} - (1-\mu)\frac{\mathbf{r} + [\mu,0,0]^\top}{r_\oplus^3} - \mu\frac{\mathbf{r} - [1-\mu,0,0]^\top}{r_{\leftmoon}^3}. \end{aligned} \tag{29}$$

Distances and time are nondimensionalized using the length unit and time unit

$$\mathrm{LU} = 384{,}400\,\mathrm{km}, \qquad \mathrm{TU} = \sqrt{\frac{\mathrm{LU}^3}{\mu_\oplus + \mu_{\leftmoon}}}.$$

The continuous-time dynamics are numerically integrated using an embedded Runge–Kutta 8(7) method.[24] True initial conditions of the Target and Chaser, as well as the initial relative-state distribution used by the filters, are adopted directly from.[22]

To account for small unmodeled perturbations, such as solar radiation pressure and modeling errors, a continuous white-noise acceleration (CWNA) model is included in the filter. Let $w_a(t)$ denote an isotropic zero-mean Gaussian acceleration process with standard deviation σ_a. For the position–velocity state $\mathbf{x} = [\mathbf{r}^\top, \mathbf{v}^\top]^\top$ and filter time step Δt (in normalized time units), the induced discrete-time process-noise covariance is

$$\mathbf{Q} = \sigma_a^2 \begin{bmatrix} \frac{\Delta t^3}{3}\mathbf{I}_3 & \frac{\Delta t^2}{2}\mathbf{I}_3 \\ \frac{\Delta t^2}{2}\mathbf{I}_3 & \Delta t\,\mathbf{I}_3 \end{bmatrix}. \tag{30}$$

Measurements

The simulation spans five orbital periods and employs temporally sparse measurements, resulting in long intervals without observations. At the beginning of the simulation, the Chaser-mounted camera is oriented toward the Target; the camera attitude is held fixed, and the Target eventually exits the field of view. The filter is executed with a sampling interval of $\mathrm{dt} = 100\,\mathrm{min}$.

The measurement vector is

$$\mathbf{y} = [\eta, \varepsilon]^\top,$$

where η and ε denote the azimuth and elevation of the Target as seen from the Chaser. Let

$$\Delta\mathbf{r} \triangleq \mathbf{r}_{\text{Targ.}} - \mathbf{r}_{\text{Chas.}}$$

be the relative position, and let $\|\cdot\|$ denote the Euclidean norm. The noiseless measurement equations are

$$\eta = \tan^{-1}\left(\frac{\Delta\mathbf{r}(2)}{\Delta\mathbf{r}(1)}\right), \tag{31}$$

$$\varepsilon = \sin^{-1}\left(\frac{\Delta\mathbf{r}(3)}{\|\Delta\mathbf{r}\|}\right). \tag{32}$$

The measurements are corrupted by additive zero-mean Gaussian white noise with identical 1σ uncertainty ϕ arcseconds for both angles. Light-travel-time delay and measurement biases are not considered. Two noise configurations are used:

1. lower-noise: $\phi = 0.9$, $\sigma_a = 0$,
2. higher-noise: $\phi = 5$, $\sigma_a = 10^{-6}$.

ENTROPY STOPPING CONDITION ANALYSIS

To assess the behaviour of different proposed adaptive sampling stopping criterions, we ran the cislunar relative navigation for both parameter sets presented hereinbefore. Representative results are shown in Figs. 3(a) - 4(b). The full-entropy and surrogate-entropy stopping conditions exhibit nearly identical performance across all metrics (except computational cost). The diagonal-only approximation selects slightly more components and yields more conservative estimates, yet it remains the least computationally demanding option while producing results that are close to the other approaches. Therefore, the diagonal-only surrogate is adopted for the remaining evaluations.

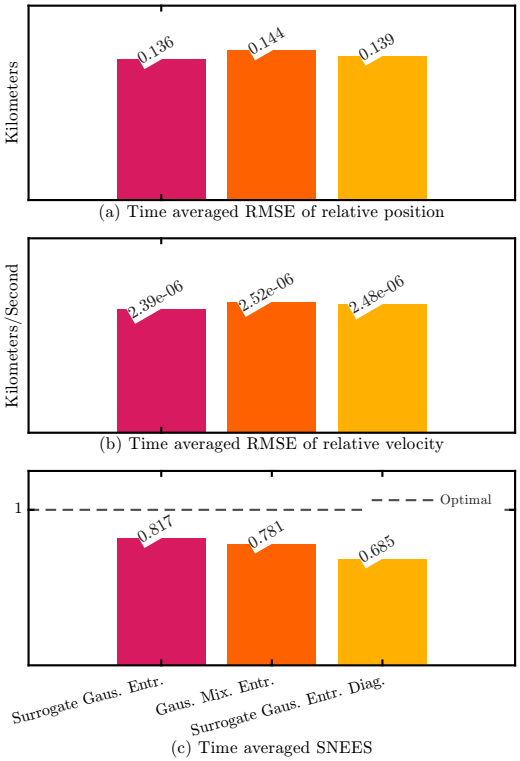

(a) RMSE comparison of stopping conditions for 100 MC (parameter set 1).

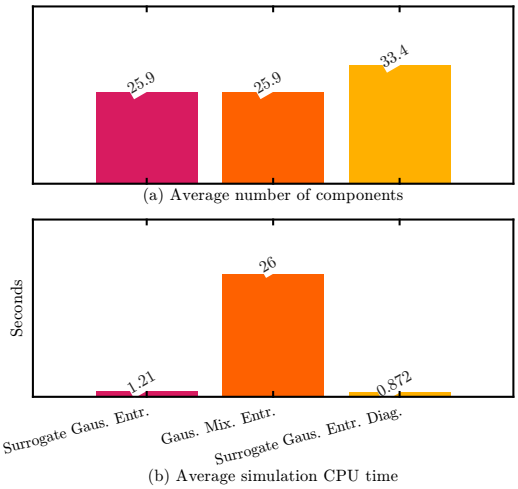

(b) Computation time comparison of stopping conditions for 100 MC (parameter set 1).

Figure 3. Comparison of stopping conditions for parameter set 1 using 100 Monte Carlo runs.

RESULTS

In this section, the proposed algorithms:

- **Deterministic** - deterministic EnGMF based on Sobol sequences,
- **Deterministic adaptive** - deterministic EnGMF with additional entropy-based adaptivity,

with $s_H = 5 \times 10^{-3}$, $\alpha = 1$, $N_1 = 200$, $r_d = 4$, $\Delta_{\text{GM}} = 10$ are tested and compared against:

- **EnGMF** - the standard stochastic EnGMF with systematic resampling,[8]
- **AEnGMF** - the stochastic adaptive EnGMF employing pruning, merging, and capping,[11]
- **ADEnGMF** - the deterministic adaptive EnGMF using Fibonacci-grid resampling based on minimizing the Cramér-von Mises distance, together with pruning, merging, and capping adaptivity.[11]

Extended and Unscented Kalman Filters are not suitable for the considered model, as shown in.[11] Global filters such as particle filters[25] or point-mass filters[5] are impractical due to their prohibitive computational complexity.

The performance of the filters is compared utilizing three metrics

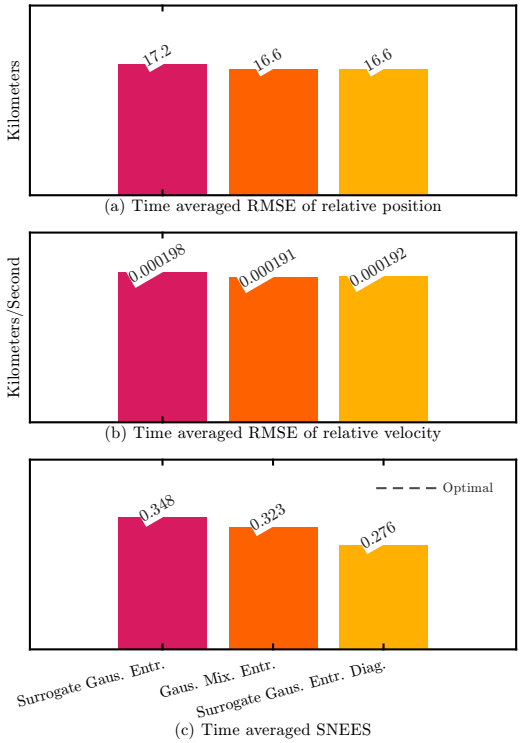

(a) RMSE comparison of stopping conditions for 100 MC (parameter set 2).

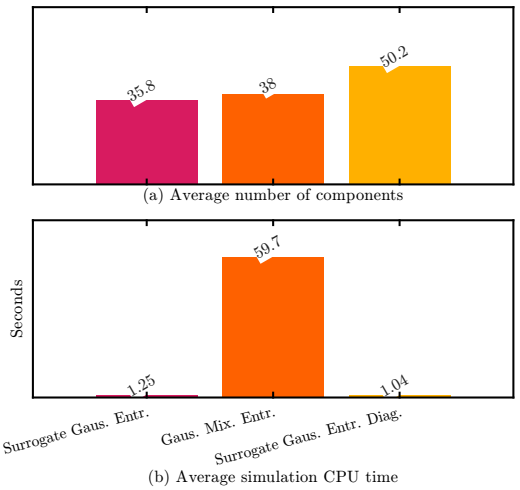

(b) Computation time comparison of stopping conditions for 100 MC (parameter set 2).

Figure 4. Comparison of stopping conditions for parameter set 2 using 100 Monte Carlo runs.

- Root mean square error - RMSE =

$$\sqrt{\frac{1}{M(T+1)}\sum_{m=1}^{M}\sum_{k=0}^{T}\left((\mathbf{x}_k)^{[m]}-(\hat{\mathbf{x}}^{\mathrm{E}}_{k|k})^{[m]}\right)^2}, \tag{33}$$

- Scaled Normalized Estimation Error Squared - SNEES =

$$\frac{1}{n_x M(T+1)}\sum_{m=1}^{M}\sum_{k=0}^{T}((\mathbf{x}_k)^{[m]}-(\hat{\mathbf{x}}^{\mathrm{E}}_{k|k})^{[m]})^{\top} ((\mathbf{P}^{\mathrm{E}}_{k|k})^{[m]})^{-1}((\mathbf{x}_k)^{[m]}-(\hat{\mathbf{x}}^{\mathrm{E}}_{k|k})^{[m]}), \tag{34}$$

- Average CPU time of whole simulation in seconds,

using $M = 500$ Monte Carlo (MC) simulations, where $(\mathbf{x}_k)^{[m]}$ is the the true state at time step k and m-th MC simulation, $(\hat{\mathbf{x}}^{\mathrm{E}}_{k|k})^{[m]} = \mathsf{E}[\mathbf{x}_k|\mathbf{z}^k]$ its filtering estimate, and $(\mathbf{P}^{\mathrm{E}}_{k|k})^{[m]}$ the corresponding filtering covariance.

The results for parameter set 1 are shown in Figs. 5(a) and 5(b), and for parameter set 2 in Figs. 6(a) and 6(b). It can be seen that the proposed deterministic algorithm achieves qualitatively the same accuracy, consistency, and simulation CPU time as the stochastic version, while

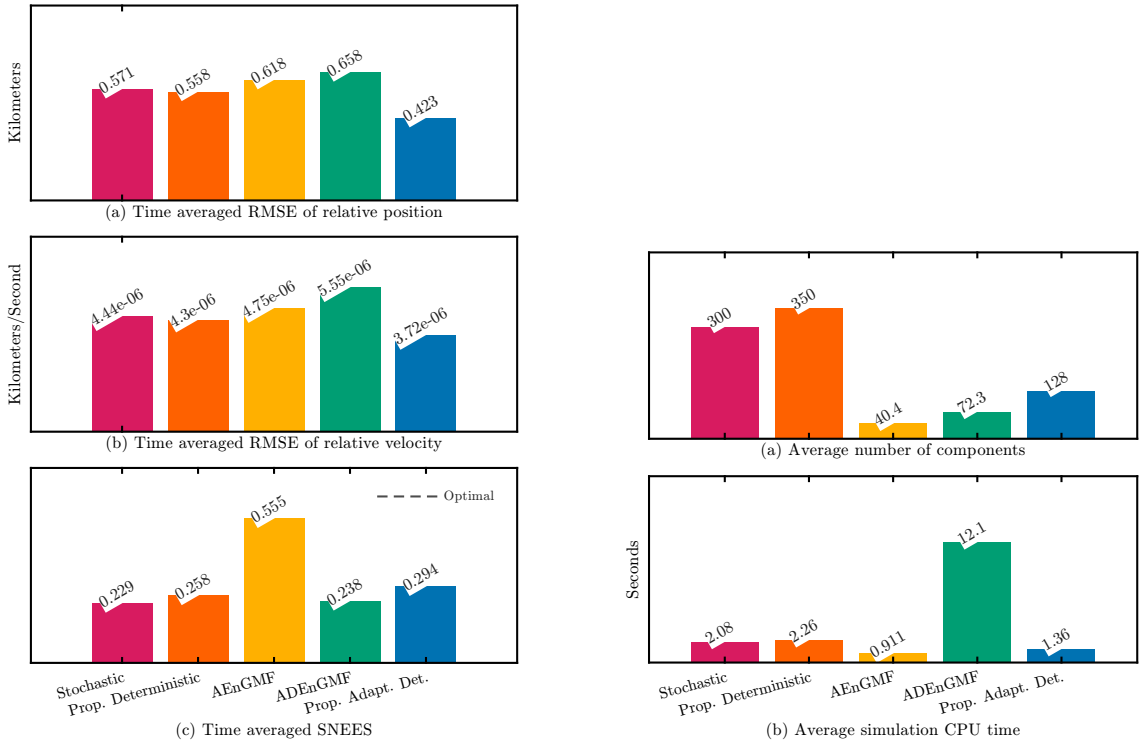

(a) Accuracy and consistency results for configuration 1.

(b) Computation time and number of components for configuration 1.

Figure 5. Performance results for configuration 1.

being fully deterministic. Moreover, incorporating the proposed adaptive component based on an entropy-based stopping condition leads to both higher accuracy and lower computational complexity for both parameter configurations. In contrast, the existing deterministic adaptive approach (ADEnGMF) results in a considerable increase in computational complexity and deteriorated accuracy and consistency for parameter configuration 2. The existing stochastic adaptive approach (AEnGMF) is computationally very efficient, but it is non-deterministic and also performs poorly for parameter configuration 2. In contrast, the proposed algorithms are deterministic and perform consistently across both scenarios.

CONCLUSION

A deterministic adaptive EnGMF was proposed based on sobol sequence based resampling and with adaptive number of modes given by entropy change. For the cislunar tracking it was shown that the proposed algorithm offers better results in a sense of estimation accuracy and complexity with similar consistency to standard EnGMF and even better robustness compared to state of the art approaches.

ACKNOWLEDGEMENTS

This material is based on research sponsored by the Air Force Office of Scientific Research (AFOSR) under agreement number FA9550-23-1-0646, Create the Future Independent Research Effort (CFIRE).

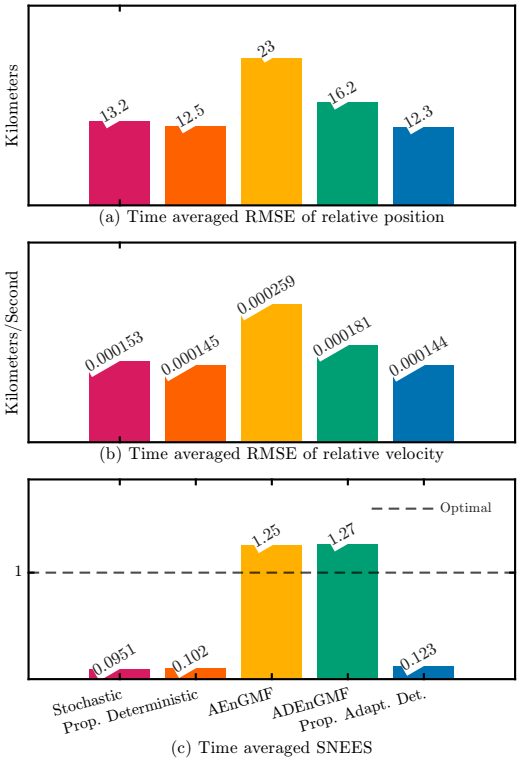

(a) Time averaged RMSE of relative position

(b) Time averaged RMSE of relative velocity

(c) Time averaged SNEES

(a) Accuracy and consistency results for configuration 2.

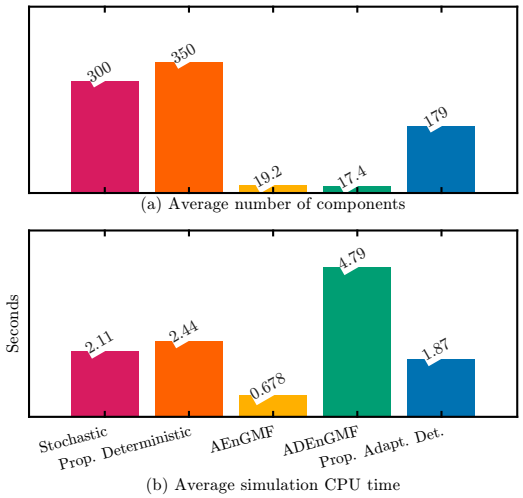

(a) Average number of components

(b) Average simulation CPU time

(b) Computation time and number of components for configuration 2.

Figure 6. Performance results for configuration 2.

REFERENCES

[1] R. E. Kalman, "A new approach to linear filtering and prediction problems," *Journal of Basic Engineering*, vol. 82, no. 1, pp. 35–45, 1960.

[2] J. Duník, S. K. Biswas, A. G. Dempster, T. Pany, and P. Closas, "State estimation methods in navigation: Overview and application," *IEEE Aerospace and Electronic Systems Magazine*, vol. 35, no. 12, pp. 16–31, 2020.

[3] S. Julier and J. Uhlmann, "A new extension of the kalman filter to nonlinear systems, proceedings of the," in *Signal Processing, Sensor Fusion, and Target Recognition VI Conference*, 1997, pp. 182–193.

[4] F. Gustafsson, F. Gunnarsson, N. Bergman, U. Forssell, J. Jansson, R. Karlsson, and P.-J. Nordlund, "Particle filters for positioning, navigation, and tracking," *IEEE Transactions on Signal Processing*, vol. 50, no. 2, pp. 425–437, 2002.

[5] J. Matoušek, J. Duník, and O. Straka, "Lagrangian grid-based filters with application to terrain-aided navigation," *IEEE Signal Processing Magazine*, 2025.

[6] H. W. Sorenson and D. L. Alspach, "Recursive Bayesian estimation using Gaussian sums," *Automatica*, vol. 7, pp. 465–479, 1971.

[7] 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.

[8] J. L. Anderson and S. L. Anderson, "A monte carlo implementation of the nonlinear filtering problem to produce ensemble assimilations and forecasts," *Monthly Weather Review*, vol. 127, no. 12, pp. 2741–2758, 1999.

[9] F. Giraldo-Grueso, A. A. Popov, U. D. Hanebeck, and R. Zanetti, "Optimal transport as a reduction technique for deterministic nonlinear filtering," in *2025 28th International Conference on Information Fusion (FUSION)*, Rio de Janeiro, Brazil, 2025.

[10] U. D. Hanebeck, D. Prossel, A. A. Popov, F. Giraldo-Grueso, and R. Zanetti, "Ensemble gaussian mixture filter based on projected cramér-von mises distance," in *2025 IEEE International Conference*

on Multisensor Fusion and Integration for Intelligent Systems (MFI). College Station, TX, USA: IEEE, September 2025.

[11] D. Durant, F. Giraldo-Grueso, and R. Zanetti, "An adaptive and deterministic ensemble gaussian mixture filter for cislunar relative navigation," in *2025 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI)*. College Station, TX, USA: IEEE, September 2025.

[12] B. W. Silverman, *Density Estimation for Statistics and Data Analysis*. Routledge, 2018.

[13] 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, 2025.

[14] E. Wan and R. Van Der Merwe, "The unscented kalman filter for nonlinear estimation," in *Proceedings of the IEEE 2000 Adaptive Systems for Signal Processing, Communications, and Control Symposium (Cat. No.00EX373)*, 2000, pp. 153–158.

[15] I. Gilitschenski, J. Steinbring, U. D. Hanebeck, and M. Simandl, "Deterministic dirac mixture approximation of gaussian mixtures," in *Proceedings of the 17th International Conference on Information Fusion (FUSION)*. IEEE, 2014, pp. 1–7.

[16] U. D. Hanebeck, D. Prossel, A. A. Popov, F. Giraldo-Grueso, and R. Zanetti, "Ensemble gaussian mixture filter based on projected cramér–von mises distance," in *Proceedings of the 2025 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI)*. IEEE, 2025, *Exact page numbers not yet available*.

[17] S. Yun and R. Zanetti, "Sequential monte carlo filtering with gaussian mixture sampling," *Journal of Guidance, Control, and Dynamics*, vol. 42, no. 9, pp. 2069–2077, 2019.

[18] B. Liu, B. Ait-El-Fquih, and I. Hoteit, "Efficient kernel-based ensemble gaussian mixture filtering," *Monthly Weather Review*, vol. 144, no. 2, pp. 781–800, 2016.

[19] T. Zhang, W. Zhang, Q. Zhao, J. Chen, and Y. Zhang, "Adaptive particle filter with randomized quasi-monte carlo sampling for unbalanced distribution system state estimation," *IEEE Transactions on Instrumentation and Measurement*, vol. 72, pp. 1–13, 2023.

[20] M. F. Huber, T. Bailey, H. Durrant-Whyte, and U. D. Hanebeck, "On entropy approximation for gaussian mixture random vectors," in *2008 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems*, 2008, pp. 181–188.

[21] A. F. Haapala and K. C. Howell, "A framework for constructing transfers linking periodic libration point orbits in the spatial circular restricted three-body problem," *International Journal of Bifurcation and Chaos*, vol. 26, no. 05, p. 1630013, 2016.

[22] E. M. Z. Spreen, "Trajectory design and targeting for applications to the exploration program in cislunar space," Ph.D. dissertation, Purdue University, 2021.

[23] D. E. Lee, "White Paper: Gateway destination orbit model: a continuous 15 year NRHO reference trajectory," NASA Johnson Space Center, Houston, TX, Tech. Rep. Document ID: 20190030294, 2019.

[24] J. R. Dormand and P. J. Prince, "A family of embedded Runge-Kutta formulae," *Journal of Computational and Applied Mathematics*, vol. 6, no. 1, pp. 19–26, 1980.

[25] A. Doucet, S. Godsill, and C. Andrieu, "On sequential monte carlo sampling methods for bayesian filtering," *Statistics and Computing*, vol. 10, no. 3, pp. 197–208, 2000.