An Adaptive and Deterministic Ensemble Gaussian Mixture Filter for Cislunar Relative Navigation*

Dalton Durant[1], Felipe Giraldo-Grueso[2], and Renato Zanetti[3]

***Abstract*— The objective of this paper is to develop an onboard nonlinear filter for cislunar relative navigation. An effective onboard filter must be accurate, consistent, efficient, and verifiable. Recent work has shown that the Ensemble Gaussian Mixture Filter (EnGMF) is an accurate and consistent nonlinear filter capable of handling non-Gaussian probability distributions. However, its fixed number of Gaussian components restricts efficiency, and its random resampling step prevents verifiability. To address these issues, the EnGMF is refined such that it adaptively reduces the number of Gaussian components, and it deterministically resamples using Fibonacci grids with optimal reduction. The presented cislunar scenario relies solely on noisy azimuth and elevation angle measurements obtained from vision-based sensors (*e.g.*, cameras), demonstrating both the need for nonlinear filtering and its potential for deep-space operations. Simulation results indicate that the refined EnGMF approaches maintain the accuracy and consistency of the traditional EnGMF, and also improve its efficiency and ensure verifiability.**

I. Introduction

To be an effective onboard filter, it must be accurate, consistent, efficient, and verifiable. To date, most filters onboard spacecraft orbiting the Earth are linear filters which assume Gaussian probability distributions like the Extended Kalman Filter (EKF) [1] and the Unscented Kalman Filter (UKF) [2], [3]—linear filters being ones that incorporate measurement information linearly in their update step. They are extensively used because of their tried-and-true lineage. However, as missions shift to the cislunar environment, nonlinear filters will need to be used to account for non-Gaussian probability distributions.

The Ensemble Gaussian Mixture Filter (EnGMF) [4], [5] has been shown to be an accurate and consistent nonlinear filter for Low Earth Orbit (LEO) [6]–[8], Geostationary Orbit (GEO) [9], and cislunar orbit determination problems [10], [11]. However, its intended usefulness is limited to ground tracking. This is because the EnGMF traditionally requires a fixed number of Gaussian components, which restricts its efficiency, and it randomly samples components after its update step, meaning it is not suitable as a verifiable filter either. Therefore, in its current form, the EnGMF will never be used as an onboard filter for spacecraft because it is neither efficient nor verifiable.

It is the purpose of this work to make the EnGMF both efficient and verifiable by adaptively reducing the number of Gaussian components needed, and deterministically resampling using Fibonacci grids with optimal reduction. This work tests the filters against a simulated cislunar relative navigation scenario, relying solely on noisy azimuth and elevation angle measurements obtained from a vision-based sensor (*e.g.*, camera), demonstrating both the need for nonlinear filters, like the EnGMF, and their potential for deep-space operations. The simulation results indicate that the proposed EnGMF refinements maintain the accuracy and consistency of the traditional EnGMF, and also improve its efficiency and ensure verifiability.

In Section II, this work gives background to what makes a filter accurate, consistent, efficient, and verifiable. Then, in Section III, this work presents the methodology for the EnGMF, including its efficient and verifiable versions. Section IV compares the results of the different EnGMFs in a simulated cislunar relative navigation scenario, and also compares them against the EKF and UKF for completeness. Finally, Section V gives a brief discussion with concluding remarks.

II. Background

A. Notation

This work models the state dynamics and measurements as

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

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

where $\mathbf{x} \in \mathbb{R}^{n_\mathbf{x}}$ is the true state, $f\colon \mathbb{R}^{n_\mathbf{x}} \to \mathbb{R}^{n_\mathbf{x}}$ is the known discrete dynamics propagation function that propagates the state from time step $k-1$ to time step k, and $\boldsymbol{\upsilon}_k \in \mathbb{R}^{n_\mathbf{x}}$ is zero-mean Gaussian white process noise with covariance $\mathbf{Q}$. Additionally, $\mathbf{y} \in \mathbb{R}^{n_\mathbf{y}}$ is the measurement, $h\colon \mathbb{R}^{n_\mathbf{x}} \to \mathbb{R}^{n_\mathbf{y}}$ is the known measurement function, and $\boldsymbol{\eta}_k \in \mathbb{R}^{n_\mathbf{y}}$ is zero-mean Gaussian white measurement noise with covariance $\mathbf{R}$, uncorrelated to $\boldsymbol{\upsilon}_k$.

In this work, the hat notation $\hat{(\cdot)}$ is used to indicate variables estimated by the filter. For example, $\hat{\mathbf{x}}$ and $\hat{\mathbf{P}}$

*This material is partially based on research sponsored by Air Force Research Laboratory (AFRL) under agreement number FA9453-21-2-0064. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of Air Force Research Laboratory (AFRL) and or the U.S. Government.

[1]Dalton Durant is a Ph.D Candidate for the Dept. of Aerospace Engineering and Engineering Mechanics at The University of Texas at Austin, Austin, TX, USA ddurant@utexas.edu

[2]Felipe Giraldo-Grueso is a Ph.D Candidate for the Dept. of Aerospace Engineering and Engineering Mechanics at The University of Texas at Austin, Austin, TX, USA fgiraldo@utexas.edu

[3]Renato Zanetti is an Associate Professor for the Dept. of Aerospace Engineering and Engineering Mechanics at The University of Texas at Austin, Austin, TX, USA renato@utexas.edu

Case Number: AFRL-2025-2642. The material was assigned a clearance of CLEARED on 22 May 2025.

represent the filter-estimated state and covariance, respectively. Additionally, the subscript notation $(\cdot)_{k|k-1}$ denotes a predicted estimate from $k-1$ to k. For example, $\hat{\mathbf{x}}_{k|k-1}$ and $\hat{\mathbf{P}}_{k|k-1}$ represent the predicted state and covariance. Similarly, the notation $(\cdot)_{k|k}$ is used to indicate an updated estimate at time step k, such as $\hat{\mathbf{x}}_{k|k}$ and $\hat{\mathbf{P}}_{k|k}$.

B. Accuracy

The following is the general definition of estimation error:

$$\tilde{\mathbf{x}}_{k|k} \stackrel{\text{def}}{=} \mathbf{x}_k - \hat{\mathbf{x}}_{k|k}, \tag{3}$$

where a filter is considered *accurate* if this error is sufficiently small.

For multi-dimensional problems, the Root Mean Square Error (RMSE) is typically used, and is computed by

$$\text{RMSE} = \sqrt{\frac{1}{n_{\mathbf{x}}}(\tilde{\mathbf{x}}_{k|k})^{\top}\tilde{\mathbf{x}}_{k|k}}, \tag{4}$$

where $(\cdot)^{\top}$ is notation for the transpose of a vector or matrix. A lower RMSE indicates a more accurate filter, *i.e.*, the estimate more closely aligns with the truth.

C. Consistency

A filter is considered *consistent* if its estimation errors are unbiased and their covariance matches the filter-calculated covariance. A more in-depth review of this topic can be found in Ref. [12]. However, in brief, under multi-dimensional linear Gaussian assumptions, the expectation of the squared norm of the estimation error is equal to its dimension:

$$\mathbb{E}[(\tilde{\mathbf{x}}_{k|k})^{\top}(\hat{\mathbf{P}}_{k|k})^{-1}\tilde{\mathbf{x}}_{k|k}] = n_{\mathbf{x}}. \tag{5}$$

This results in the Scaled Normalized (state) Estimation Error Squared (SNEES):

$$\text{SNEES} = \frac{1}{n_{\mathbf{x}}}(\tilde{\mathbf{x}}_{k|k})^{\top}(\hat{\mathbf{P}}_{k|k})^{-1}\tilde{\mathbf{x}}_{k|k}. \tag{6}$$

The SNEES is an $n_{\mathbf{x}}$-scaled version of the NEES from Ref. [12], which sets the optimal filter consistency to a value of 1 for linear Gaussian systems. Therefore, a value less than 1 indicates that the filter is conservative, while a value greater than 1 suggests that the filter is overly confident. Although real-world systems are typically not linear or Gaussian, the SNEES remains a valid measure of relative consistency between filters, *i.e.*, how consistent one filter is compared to another.

D. Efficiency

The concept of filter efficiency considered in this work differs from the "efficiency check" described in Ref. [12]. Instead, this work focuses on *computational* filter efficiency, *i.e.*, the time required for the filter to process data and generate estimates.

Efficiency is measured using wall-clock time, which represents the total elapsed time for a simulation to run from start to finish, as if it were observed by an external observer. This differs from CPU time, which only accounts for the time the processor was actively working. Wall-clock time includes both active computation and any waiting periods, making it a more comprehensive measure of real-world performance.

Achieving a filter that is accurate, consistent, and computationally efficient is highly desirable.

E. Verifiability

Verifiable filters avoid (pseudo) random subroutines, ensuring consistent behavior across multiple runs. Common particle and ensemble-based filters rely on stochastic resampling procedures to avoid particle or ensemble collapse. The stochasticity of these strategies makes these filters not verifiable, and therefore, not a very realistic choice for mission-critical applications. [13]

In hybrid filters, where the state probability density function (pdf) is parametrized by Gaussian mixtures and point masses, optimally sampling particles from a Gaussian mixture is a crucial step. This step can be performed deterministically by using optimal transport strategies [13], or by minimizing a distance metric known as the modified Cramér-von Mises distance [14]–[16]. These strategies make ensemble-based filters deterministic and thus verifiable.

III. Methodology

A. The Ensemble Gaussian Mixture Filter

The following is a brief introduction to the EnGMF, with detailed algorithmic descriptions available in Ref. [5]. Recent studies have demonstrated its effectiveness in tracking targets even under sparse observation conditions [5]–[11], [17]–[19].

Fig. 1 provides an illustrative diagram of the EnGMF steps. Before step ①, the EnGMF assumes a Dirac delta empirical measure, approximating the prior distribution $p(\mathbf{x}_{k-1})$ using N independent and identically distributed (i.i.d.) particles $\{\boldsymbol{\mathcal{X}}_{k-1}^{(i)}\}_{i=1}^{N}$:

$$p(\mathbf{x}_{k-1}) \approx \sum_{i=1}^{N}\frac{1}{N}\delta(\mathbf{x}_{k-1}-\boldsymbol{\mathcal{X}}_{k-1}^{(i)}), \tag{7}$$

where $\delta(\cdot)$ represents the Dirac delta distribution, which can be interpreted as a normal distribution with covariance approaching the zero matrix in the limit.

In step ① of Fig. 1, the particles are propagated to time step k, and in step ②, they are converted into a Gaussian mixture using the Kernel Density Estimation (KDE) method from Ref. [20]. This KDE method transforms each particle into a Gaussian component with nonzero covariance and equal weight:

$$p(\mathbf{x}_k) \approx \textstyle\sum_{i=1}^{N}\frac{1}{N}\mathcal{N}\big(\mathbf{x}_k;\hat{\mathbf{x}}_{k|k-1}^{(i)},\beta_{\text{Silv.}}\text{Cov}(\{\hat{\mathbf{x}}_{k|k-1}^{(i)}\}_{i=1}^{N})\big), \tag{8}$$

where $\hat{\mathbf{x}}_{k|k-1}^{(i)}=f(\boldsymbol{\mathcal{X}}_{k-1}^{(i)})$ is the predicted mean of the i-th particle, $\text{Cov}(\{\hat{\mathbf{x}}_{k|k-1}^{(i)}\}_{i=1}^{N})$ is the sample covariance matrix of all the propagated particles, and $\beta_{\text{Silv.}}$ is the bandwidth parameter. To reduce computational costs, Silverman's Rule of Thumb [20] is used to determine the bandwidth parameter for each component:

$$\beta_{\text{Silv.}} = \left(\frac{4}{n_{\mathbf{x}}+2}\right)^{\frac{2}{n_{\mathbf{x}}+4}} N^{-\frac{2}{n_{\mathbf{x}}+4}}. \tag{9}$$

Case Number: AFRL-2025-2642. The material was assigned a clearance of CLEARED on 22 May 2025.

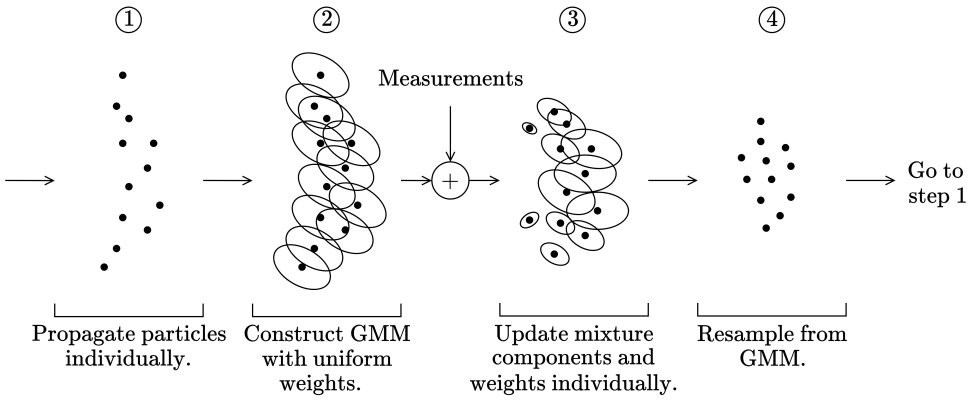

Fig. 1: Diagram of the Ensemble Gaussian Mixture Filter (EnGMF) steps.

Next, measurements are incorporated, and in step ③ of Fig. 1, the posterior distribution is approximated as a weighted Gaussian mixture:

$$p(\mathbf{x}_k|\mathbf{y}_k) \approx \sum_{i=1}^{N} w_k^{(i)} \mathcal{N}(\mathbf{x}_k; \hat{\mathbf{x}}_{k|k}^{(i)}, \hat{\mathbf{P}}_{k|k}^{(i)}), \tag{10}$$

where each Gaussian component's mean and covariance are updated individually using a linear filter, such as the EKF or UKF. The mixture weights are computed as

$$w_k^{(i)} = \frac{\frac{1}{N}\mathcal{N}\left(\mathbf{y}_k; h(\hat{\mathbf{x}}_{k|k-1}^{(i)}), \mathbf{H}_k^{(i)}\hat{\mathbf{P}}_{k|k-1}^{(i)}(\mathbf{H}_k^{(i)})^\top + \mathbf{R}\right)}{\sum_{j=1}^{N}\frac{1}{N}\mathcal{N}\left(\mathbf{y}_k; h(\hat{\mathbf{x}}_{k|k-1}^{(j)}), \mathbf{H}_k^{(j)}\hat{\mathbf{P}}_{k|k-1}^{(j)}(\mathbf{H}_k^{(j)})^\top + \mathbf{R}\right)}, \tag{11}$$

where all weights are strictly positive and sum to 1. Here, $\mathbf{H}_k^{(i)} = \frac{\partial h(\mathbf{x})}{\partial \mathbf{x}}\big|_{\mathbf{x}=\hat{\mathbf{x}}_{k|k-1}^{(i)}}$ is the Jacobian of the measurement mapping function evaluated at the prior estimate.

For analysis purposes (not affecting the filter), the extracted mean and covariance are:

$$\bar{\hat{\mathbf{x}}}_k = \sum_{i=1}^{N} w_k^{(i)} \hat{\mathbf{x}}_{k|k}^{(i)}, \tag{12}$$

$$\bar{\hat{\mathbf{P}}}_k = \sum_{i=1}^{N} w_k^{(i)} \left[\hat{\mathbf{P}}_{k|k}^{(i)} + (\hat{\mathbf{x}}_{k|k}^{(i)} - \bar{\hat{\mathbf{x}}}_k)(\hat{\mathbf{x}}_{k|k}^{(i)} - \bar{\hat{\mathbf{x}}}_k)^\top\right]. \tag{13}$$

Finally, in step ④ of Fig. 1, the EnGMF generates N i.i.d. new samples from the Gaussian mixture approximation of the posterior distribution using the resampling method described in Ref. [21]. These newly drawn samples then serve as the initial particles for the next iteration of the filter.

B. Adaptively Changing the Number of EnGMF Components

The goal here is to decrease the number of Gaussian components carried by the EnGMF, and in turn increase its computational efficiency. In its current form, the EnGMF has a fixed number of components, and to that regard it is different from other nonlinear filters like the Gaussian Sum Filter (GSF) [22], whose number of Gaussian components continue to increase combinatorially. The GSF mitigates its ever growing number of components by often using a pruning, merging, and capping scheme. Although the EnGMF retains the same number of Gaussian components at every time step, this same GSF pruning, merging, and capping scheme can be used to *prune* (discard) components with small weights, *merge* (combine) components that are similar, and *cap* (limit) the total number of components altogether.

Pruning and capping are computationally cheap, whereas merging is more expensive, requiring a comminatory search of all components to find the merging pairs. It is common to merge pairs based on a distance criteria, which is typically the Mahalanobis distance and is used in the merging schemes of [23]–[25]. This work uses the pruning, merging, and capping algorithm from Ref. [24] and is provided in Algorithm I.

This pruning, merging, and capping algorithm is well and good, but it presents a new problem for the EnGMF—*degeneracy*. By using the algorithm, there could come a

TABLE I: Pruning, merging, and capping steps.

Given: $\{w_k^{(i)}, \hat{\mathbf{x}}_{k|k}^{(i)}, \hat{\mathbf{P}}_{k|k}^{(i)}\}_{i=1}^{N}$, pruning threshold T, merging threshold U, and capping limit $N_{\max}$.

Step 1: *Pruning*
set $n = 0$, and $I = \{i = 1, \ldots, N | w_k^{(i)} \geqslant T\}$
renormalize weights

Step 2: *Merging*
repeat
$n \leftarrow n + 1$
$j = \arg_{i \in I}\max\ w_k^{(i)}$
$L = \{i \in I | (\hat{\mathbf{x}}_{k|k}^{(i)} - \hat{\mathbf{x}}_{k|k}^{(j)})^\top (\hat{\mathbf{P}}_{k|k}^{(i)})^{-1} (\hat{\mathbf{x}}_{k|k}^{(i)} - \hat{\mathbf{x}}_{k|k}^{(j)}) \leqslant U\}$
$\bar{w}_k^{(n)} = \sum_{i \in L} w_k^{(i)}$
$\bar{\mathbf{x}}_{k|k}^{(n)} = \frac{1}{\bar{w}_k^{(n)}} \sum_{i \in L} w_k^{(i)} \hat{\mathbf{x}}_{k|k}^{(i)}$
$\bar{\mathbf{P}}_{k|k}^{(n)} = \frac{1}{\bar{w}_k^{(n)}} \sum_{i \in L} w_k^{(i)} (\hat{\mathbf{P}}_{k|k}^{(i)} + (\hat{\mathbf{x}}_{k|k}^{(i)} - \bar{\mathbf{x}}_{k|k}^{(n)})(\hat{\mathbf{x}}_{k|k}^{(i)} - \bar{\mathbf{x}}_{k|k}^{(n)})^\top)$
$I \leftarrow I \backslash L$
until $I = \varnothing$
renormalize weights

Step 3: *Capping*
if $n > N_{\max}$, then keep $N_{\max}$ with largest weights in $\{\bar{w}_k^{(i)}, \bar{\mathbf{x}}_{k|k}^{(i)}, \bar{\mathbf{P}}_{k|k}^{(i)}\}_{i=1}^{n}$
renormalize weights

Output: $\{\bar{w}_k^{(i)}, \bar{\mathbf{x}}_{k|k}^{(i)}, \bar{\mathbf{P}}_{k|k}^{(i)}\}_{i=1}^{n}$

point in which the number of components collapses to a single Gaussian component and never recovers. To prevent this, there needs to be injected some diversity. This work does this by requesting a minimum number of particles, N_{min}, to be added onto the resampling step. Fig. 2 shows the same process of the EnGMF from Fig. 1, except now with a pruning, merging, and capping procedure occurring at step ④, after which there are N components, and in step ⑤ it resamples $N + N_{\text{min}}$ particles to prevent degeneracy.

C. Deterministic Resampling of the EnGMF Using Fibonacci Grids and Optimal Reduction

As seen in Fig. 1, in the EnGMF, particles must be resampled from the posterior GMM. In order to make the filter verifiable, this work uses the deterministic resampling technique presented in Refs. [15], [16]. This technique optimally samples the posterior GMM in two steps, summarized as follows.

The first step is to samples deterministic points from each Gaussian component of the GMM,

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

where D are the total number of points to be sampled from each component and $\{\boldsymbol{\mathcal{D}}_{k|k,i}^{(j)}\}_{j=1}^{D}$ are the deterministic samples. For this work, Fibonacci grids [26] are used to sample the deterministic points. With this first step, the posterior GMM is approximated as a Dirac mixture with $N \cdot D$ points,

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

The second step is to reduce the Dirac mixture back to N equally weighted points, such that

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

where $\{\boldsymbol{\mathcal{X}}_k^{(i)}\}_{i=1}^{N}$ are the newly sampled particles. That is, this step finds the optimal location of the new particles by minimizing the modified Cramér-von Mises distance [14] between the two Dirac mixtures.

For this step, the modified Cramér-von Mises distance between the two Dirac mixtures is minimized using Matlab's `fminunc` function, as in Refs. [15], [16]. For the minimization, the reduced points are initialized with the means of the posterior GMM.

IV. Numerical Experiment

The following scenario involves a Target and a Chaser in a similar Near Rectilinear Halo Orbit (NRHO). The Target is being tracked by the Chaser using a vision-based sensor (e.g. camera), which produces noisy azimuth and elevation measurements. Assuming knowledge of its own inertial position and velocity perfectly, the goal of the Chaser is to estimate its relative position and velocity with respect to the Target.

Cislunar NRHOs are a class of trajectories that exist in the vicinity of the Moon and are typically positioned close to Lagrange points associated with the Earth-Moon system. NRHOs offer a degree of long-term stability and remain relatively stable over extended periods. This stability is valuable for missions requiring long-duration observations or operations.

However, they are also known for their chaotic behavior and sensitivity to initial conditions, as reported by [27] and [28]. Minor errors in the initial state or measurements can lead to significant divergence in estimated states, making accurate estimation challenging.

A. Dynamics

This work uses a 9:2 synodic resonant NRHO, which means the Target and Chaser make nine revolutions around the Moon for every two lunar synodic months. This specific type of NRHO has a 6.5-day period, a perilune radius of about 3,250 km, and an apolune radius of approximately 71,000 km. It is the lowest-altitude NRHO with a useful resonance, and serves as the baseline orbit for NASA's Lunar Gateway mission [29].

This work models cislunar NRHO dynamics using the Circular Restricted Three Body Problem (CR3BP) for the Earth-Moon system with a 6-dimensional state space represented by $\mathbf{x} = [\mathbf{r}(1), \mathbf{r}(2), \mathbf{r}(3), \mathbf{v}(1), \mathbf{v}(2), \mathbf{v}(3)]^\top$:

$$\begin{aligned}
&\dot{\mathbf{r}}(1) = \mathbf{v}(1), \qquad \dot{\mathbf{r}}(2) = \mathbf{v}(2), \qquad \dot{\mathbf{r}}(3) = \mathbf{v}(3), \\
&\dot{\mathbf{v}}(1) = \mathbf{r}(1) + 2\mathbf{v}(2) - \frac{(1-\mu)(\mathbf{r}(1)+\mu)}{r_{\oplus}^3} - \\
&\qquad \frac{\mu(\mathbf{r}(1)-1+\mu)}{r_{\mathbb{C}}^3}, \\
&\dot{\mathbf{v}}(2) = \mathbf{r}(2) - 2\mathbf{v}(1) - \frac{(1-\mu)\mathbf{r}(2)}{r_{\oplus}^3} - \frac{\mu\mathbf{r}(2)}{r_{\mathbb{C}}^3}, \\
&\dot{\mathbf{v}}(3) = -\frac{(1-\mu)\mathbf{r}(3)}{r_{\oplus}^3} - \frac{\mu\mathbf{r}(3)}{r_{\mathbb{C}}^3},
\end{aligned} \tag{17}$$

where $\mathbf{r}(1), \mathbf{r}(2), \mathbf{r}(3)$ and $\mathbf{v}(1), \mathbf{v}(2), \mathbf{v}(3)$ represent the scaled Cartesian positions and velocities of the space objects with respect to the Barycenter origin, μ is the scaled Moon geocentric gravitational constant, and $r_{\oplus}$ and $r_{\mathbb{C}}$ are the distances of the space objects with respect to the Earth and Moon in the Barycenter reference frame:

$$\mu = \frac{\mu_{\mathbb{C}}}{\mu_{\oplus} + \mu_{\mathbb{C}}}, \tag{18}$$

$$r_{\oplus} = \sqrt{(\mathbf{r}(1)+\mu)^2 + \mathbf{r}(2)^2 + \mathbf{r}(3)^2}, \tag{19}$$

$$r_{\mathbb{C}} = \sqrt{(\mathbf{r}(1)-1+\mu)^2 + \mathbf{r}(2)^2 + \mathbf{r}(3)^2}. \tag{20}$$

In this work, $\mu_{\oplus}{=}G \cdot m_{\oplus}$ and $\mu_{\mathbb{C}}{=}G \cdot m_{\mathbb{C}}$. The gravitational constant is $G{=}6.6743 \times 10^{-11}\ \text{m}^3\,\text{s}^{-2}\,\text{kg}^{-1}$, the mass of the Earth is $m_{\oplus}{=}5.972 \times 10^{24}$ kg, and the mass of the Moon is $m_{\mathbb{C}}{=}7.342 \times 10^{22}$ kg.

The units for distance and time are normalized by length units LU$=384\,400 \times 10^3$ m and time units TU$=\sqrt{\text{LU}^3/(\mu_{\oplus}+\mu_{\mathbb{C}})}$ s. The system dynamic equations are numerically integrated with an embedded Runge-Kutta 8(7) method [30].

Case Number: AFRL-2025-2642. The material was assigned a clearance of CLEARED on 22 May 2025.

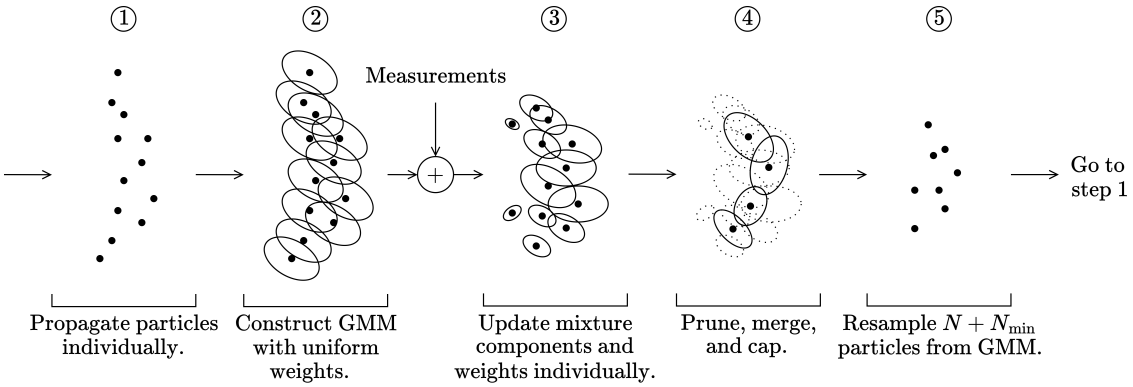

Fig. 2: Diagram of the Adaptive Ensemble Gaussian Mixture Filter (AEnGMF) steps.

The true states of the Chaser start at the dimensionless coordinates $\mathbf{x}_{0,\text{Chas.}}$ from Ref. [28]:

$$\mathbf{x}_{0,\text{Chas.}} = [1.0110350588,\ 0,\ -0.1731500000,\ 0,\ -0.0780141199,\ 0]^\top, \tag{21}$$

and the Chaser has a dimensionless period of roughly 1.3632096570 [28]. Whereas, the true states of the Target start along $1/5$-th of the Chaser's orbit:

$$\mathbf{x}_{0,\text{Targ.}} = [1.0072370971,\ -0.0199527269,\ -0.1523270340,\ -0.0278278626,\ -0.0630584615,\ 0.1565627028]^\top. \tag{22}$$

For the filters, each Monte Carlo simulation starts with an initial Gaussian distribution centered at the following dimensionless relative coordinates $\hat{\mathbf{x}}_{0|0}$, with covariance $\hat{\mathbf{P}}_{0|0}$:

$$\hat{\mathbf{x}}_{0|0} = [-0.0037979617,\ -0.0199527269,\ 0.0208229660,\ -0.0278278626,\ 0.0149556584,\ 0.1565627028]^\top, \tag{23}$$

$$\hat{\mathbf{P}}_{0|0} = \text{diag}([1\times10^{-4},\ \ 1\times10^{-4},\ \ 1\times10^{-4},\ \ 1\times10^{-6},\ \ 1\times10^{-6},\ \ 1\times10^{-6}]^2). \tag{24}$$

The different filters do not assume any additive discrete process noise: $\mathbf{Q}=\mathbf{0}_{n_\mathbf{x}\times n_\mathbf{x}}$.

B. Measurements

The full simulation lasts for 5 orbits. Measurements are sparse in time such that they are infrequent over long periods. The camera on the Chaser is pointed strictly in the positive Barycenter-inertial Y-direction. As the Target crosses the camera's field of view (FOV) of 90°, measurements are recorded. The filter operates at a rate of $\text{dt} = 100$ min.

In this example, the Target is assumed trackable such that it has sufficient size and reflectivity. This absolves any issues related to detection for this work, while studies such as [31] provide more details on the detectability of cislunar objects.

The measurement vector $\mathbf{y}=[\alpha,\varepsilon]^\top$ contains azimuth α and elevation ε of the Target mapped to the Chaser:

$$\alpha = \tan^{-1}\left(\frac{\mathbf{r}_{\text{Targ.}}(2) - \mathbf{r}_{\text{Chas.}}(2)}{\mathbf{r}_{\text{Targ.}}(1) - \mathbf{r}_{\text{Chas.}}(1)}\right), \tag{25}$$

$$\varepsilon = \sin^{-1}\left(\frac{\mathbf{r}_{\text{Targ.}}(3) - \mathbf{r}_{\text{Chas.}}(3)}{\|\mathbf{r}_{\text{Targ.}} - \mathbf{r}_{\text{Chas.}}\|}\right), \tag{26}$$

where $\|\cdot\|$ is the Euclidean 2-norm, $\mathbf{r}_{\text{Targ.}}=[\mathbf{r}_{\text{Targ.}}(1),\mathbf{r}_{\text{Targ.}}(2),\mathbf{r}_{\text{Targ.}}(3)]^\top$ is the position of the Target, and $\mathbf{r}_{\text{Chas.}}=[\mathbf{r}_{\text{Chas.}}(1),\mathbf{r}_{\text{Chas.}}(2),\mathbf{r}_{\text{Chas.}}(3)]^\top$ is the position of the Chaser. The measurements are corrupted by additive zero-mean Gaussian white noise having 1σ-uncertainties of 0.9 arcseconds for both azimuth and elevation angles. Light-travel time delay and measurement biases are not considered.

C. Results

This section tests the EKF, UKF, EnGMF, AEnGMF, and ADEnGMF in the presented cislunar relative navigation problem. The EKF is from [1] and approximates its state transition matrix (STM) using the matrix exponential of the dynamics Jacobian times dt, which provides a reasonable approximation for small time steps. The UKF is the 3-parameter formulation from [3], which uses α=1, β=2, and $\kappa=3-n_\mathbf{x}$ and re-Gaussianizes after every prediction step. The EnGMF is from [5] and is recalled in Section III-A. It uses N=250 components. The AEnGMF adaptively changes the number of Gaussian components as needed and is detailed in Section III-B. It starts with N=250 components, but also uses a $T=1\times10^{-3}$ pruning threshold, U=4 merging threshold, $N_{\text{max}}=N$ caping limit, and N_{min}=10 minimum number of resampled particles. Finally, the ADEnGMF is the AEnGMF that deterministically resamples using Fibonacci grids with optimal reduction and is detailed in Section III-C. It uses the same parameters as the AEnGMF, but with the addition of $N_{\text{Fib.}}=2n_\mathbf{x}+1$ deterministic Fibonacci points.

Fig. 3 shows the accuracy and consistency of the compared filters over the 5 orbits. The first thing to notice is that the EKF and UKF diverge because they assume a single Gaussian probability distribution. This assumption is insufficient for this scenario because the dynamics and measurements are nonlinear, and the state is point-wise-in-time unobservable from the angles-only measurements. The next thing to notice is that in Figs. 3 (a), (b), and (c), the EnGMF, AEnGMF, and ADEnGMF each experience the same peaks and valleys. This is because the Target is not always visible, and when

Case Number: AFRL-2025-2642. The material was assigned a clearance of CLEARED on 22 May 2025.

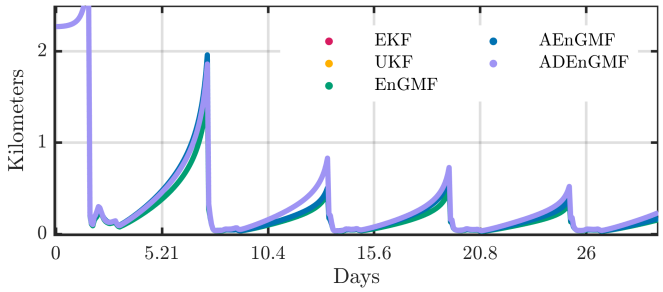

(a) RMSE of relative position vs. time

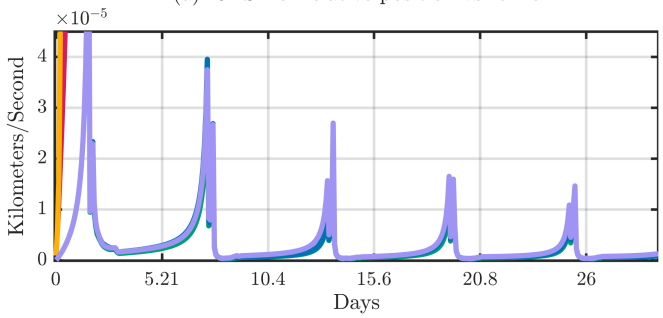

(b) RMSE of relative velocity vs. time

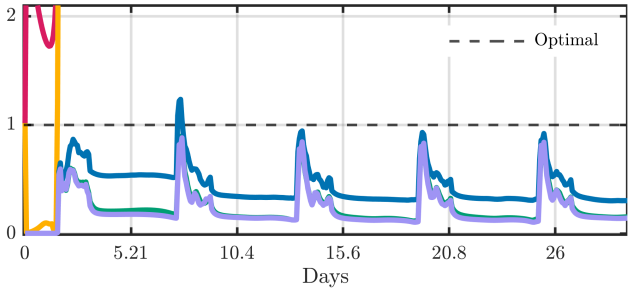

(d) SNEES vs. time

Fig. 3: This figure compares the accuracy (RMSE) and consistency (SNEES) of each filter vs. time. Results are averaged over 1,000 Monte Carlo simulations.

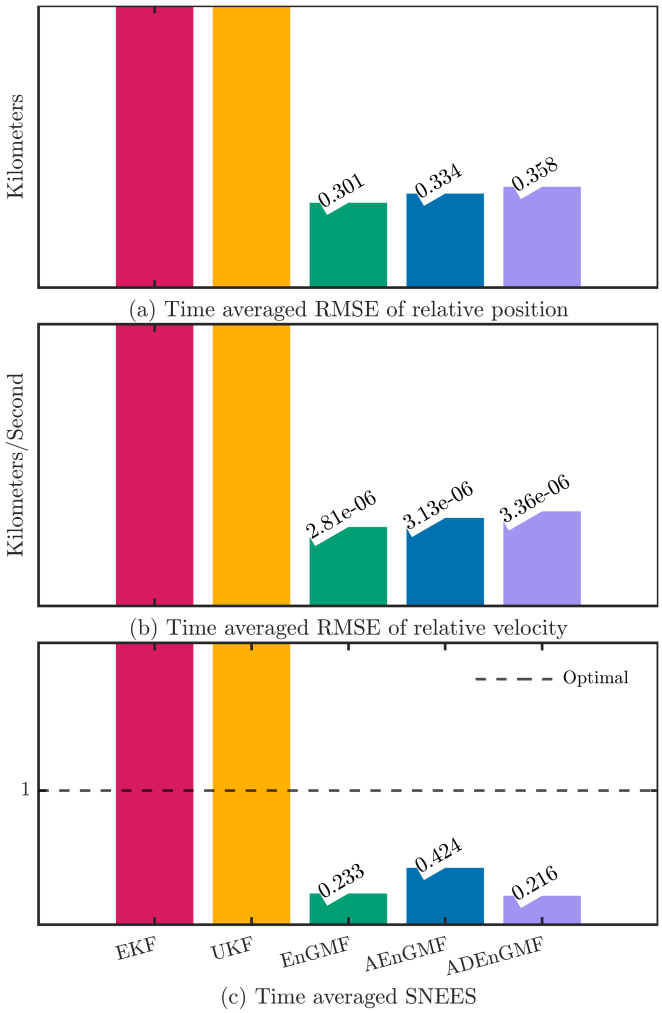

(a) Time averaged RMSE of relative position

(b) Time averaged RMSE of relative velocity

(c) Time averaged SNEES

Fig. 4: This figure compares the overall accuracy (RMSE) and consistency (SNEES) of each filter averaged over time. Results are further averaged over 1,000 Monte Carlo simulations.

this occurs, the uncertainty of the filter grows, causing the peaks. When the target becomes visible, measurements are given to the filter and it improves its estimate, causing the valleys.

In Figs. 3 (a) and (b), the EnGMF, AEnGMF, and ADEnGMF all overlap nicely, indicating similar accuracies even with the refinements made by AEnGMF and ADEnGMF. However, in Fig. 3 (c), the consistency of the AEnGMF is slightly more confident than the EnGMF and ADEnGMF, likely due to having fewer Gaussian components, but this is still acceptable performance.

Fig. 4 just reiterates what was shown in Fig. 3, except averaging over the 5 orbits to get a single data point. The EKF and UKF diverge, while the EnGMF, AEnGMF, and ADEnGMF each have similar accuracies and consistencies despite their differences.

Finally, Fig. 5 shows the efficiency of the filters broken down into two sections. Fig. 5 (a) shows the average number of components each filter has and Fig. 5 (b) shows their average Monte Carlo simulation wall-clock times. Notice that the AEnGMF improves the efficiency of the EnGMF considerably, almost catching up to the UKF. However, as soon as the deterministic sampling is introduced, the ADEnGMF's efficiency soars, hindering its marketability as an onboard filter. If the deterministic sampling methods can sample just as fast as random sampling methods, then the EnGMF can be a real contender as an effective onboard nonlinear filter. Regardless, the presented results are promising and indicate this work is heading in the right direction.

V. Conclusion

Recent work has shown that the Ensemble Gaussian Mixture Filter (EnGMF) is an accurate and consistent nonlinear filter with applications to Low Earth Orbit (LEO), Geostationary Orbit (GEO), and cislunar orbit determination problems. However, it is neither efficient nor verifiable, which limits its applicability and bars it from ever being used as an effective onboard navigation filter.

In this work, the EnGMF is refined to be both efficient and verifiable by adaptively reducing the number of Gaus-

Case Number: AFRL-2025-2642. The material was assigned a clearance of CLEARED on 22 May 2025.

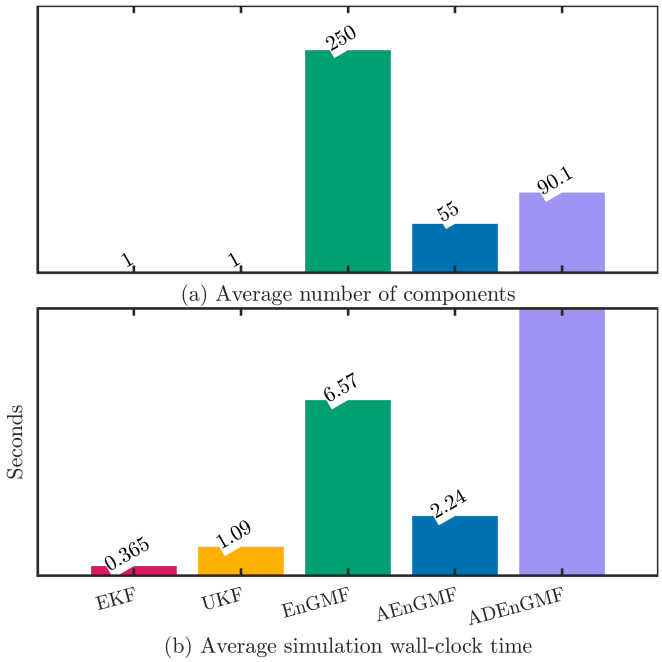

(a) Average number of components

(b) Average simulation wall-clock time

Fig. 5: This figure compares the efficiency of the filters and how their number of components relate to their wall-clock times. Results are averaged over 1,000 Monte Carlo simulations. Using an Intel Core i7 9700K CPU at a base speed of 3.00 GHz and with 16 GB of RAM.

sian components needed, and deterministically resampling using Fibonacci grids with optimal reduction. This work compares the proposed EnGMFs in a simulated cislunar relative navigation scenario, and also compares their performance against the Extended Kalman Filter (EKF) and Unscented Kalman Filter (UKF) for completeness. The presented cislunar scenario relies solely on noisy azimuth and elevation angle measurements obtained from vision-based sensors (*e.g.*, cameras), demonstrating both the need for nonlinear filters, like the EnGMF, and their potential for deep-space operations. Simulation results indicate that the proposed EnGMFs maintain the accuracy and consistency of the traditional EnGMF, and also improve its efficiency and ensure verifiability.

Although the results show a lot of promise, the current bottleneck lies in the speed of the deterministic sampling algorithms, *i.e.*, the Fibonacci grid sampling with optimal reduction is too slow. If this can be improved to achieve similar resampling times as a random sampler, then the EnGMF will be a considerable contender as an effective onboard filter.

References

[1] A. Gelb, *Applied Optimal Estimation*. MIT Press, 1974.

[2] S. J. Julier and J. K. Uhlmann, "New extension of the Kalman filter to nonlinear systems," in *Signal Processing, Sensor Fusion, and Target Recognition VI*, vol. 3068. Spie, 1997, pp. 182–193.

[3] R. van der Merwe, "Sigma-point Kalman filters for probabilistic inference in dynamic state-space models," Ph.D. dissertation, The OGI School of Science and Engineering at Oregon Health and Science University, 2004.

[4] 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.

[5] 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.

[6] B. L. Reifler, S. Yun, B. A. Jones, and R. Zanetti, "Multi-target ensemble Gaussian mixture tracking with sparse observations," in *AMOS Conf. Proc*, 2021.

[7] B. L. Reifler, A. A. Popov, B. A. Jones, and R. Zanetti, "Large-scale space object tracking in a proliferated LEO scenario," in *26th International Conference on Information Fusion (FUSION)*. IEEE, 2023, pp. 1–8.

[8] D. Durant, A. A. Popov, and R. Zanetti, "MCMC EnGMF for sparse data orbit determination," in *AAS/AIAA Astrodynamics Specialist Conference*, 2023, pp. 23–356.

[9] S. Yun, N. Ravago, B. L. Reifler, R. Zanetti, and B. A. Jones, "Generalized labeled multi-Bernoulli filter with kernel-based ensemble Gaussian mixture filtering for orbit determination with sparse data," in *AMOS Conf. Proc*, 2022.

[10] D. Durant, A. A. Popov, and R. Zanetti, "Processing angles-only tracklets for cislunar multi-target tracking," in *AAS/AIAA Astrodynamics Specialist Conference*, 2024, pp. 24–339.

[11] ——, "What are you weighting for? Improved weights for Gaussian mixture filtering with application to cislunar orbit determination," *arXiv preprint arXiv:2405.11081*, 2024.

[12] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan, *Estimation with Applications to Tracking and Navigation: Theory Algorithms and Software*. John Wiley & Sons, 2001.

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

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

[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] ——, "Optimal grid point sampling for point mass filtering with applications to cislunar orbit determination," *In preparation. Journal of Astronautical Sciences*, pp. 1–30, 2025.

[17] A. A. Popov and R. Zanetti, "Ensemble Gaussian mixture filtering with particle-localized covariances," in *26th International Conference on Information Fusion (FUSION)*. IEEE, 2023, pp. 1–7.

[18] ——, "An adaptive covariance parameterization technique for the ensemble Gaussian mixture filter," *SIAM Journal on Scientific Computing*, vol. 46, no. 3, pp. A1949–A1971, 2024.

[19] A. A. Popov, E. M. Zucchelli, and R. Zanetti, "Ensemble-regularized Kernel density estimation with applications to the ensemble Gaussian mixture filter," *Computational Geosciences*, vol. 29, no. 2, pp. 1–23, 2025.

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

[21] 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.

[22] H. W. Sorenson and D. L. Alspach, "Recursive Bayesian estimation using Gaussian sums," *Automatica*, vol. 7, no. 4, pp. 465–479, 1971.

[23] A. Rauh, K. Briechle, and U. D. Hanebeck, "Nonlinear measurement update and prediction: Prior density splitting mixture estimator," in *2009 IEEE Control Applications,(CCA) & Intelligent Control,(ISIC)*. IEEE, 2009, pp. 1421–1426.

[24] B.-N. Vo and W.-K. Ma, "The Gaussian mixture probability hypothesis density filter," *IEEE Transactions on Signal Processing*, vol. 54, no. 11, pp. 4091–4104, 2006.

[25] W. Wu, H. Sun, M. Zheng, and W. Huang, *Target Tracking with Random Finite Sets*. Springer, 2023.

[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] 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.
[28] E. M. Z. Spreen, "Trajectory design and targeting for applications to the exploration program in cislunar space," Ph.D. dissertation, Purdue University, 2021.
[29] 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.
[30] 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.
[31] C. Frueh, K. Howell, K. J. DeMars, and S. Bhadauria, "Cislunar space situational awareness," in *31st AIAA/AAS Space Flight Mechanics Meeting*, 2021, pp. 21–290.