The Journal of the Astronautical Sciences (2026) 73:67
https://doi.org/10.1007/s40295-026-00617-8

RESEARCH

Square Root Filters for Cislunar Angles-Only Relative Orbit Estimation

Andrea López[1] · Jay W. McMahon[1] · Hanspeter Schaub[1]

Received: 28 January 2026 / Accepted: 13 July 2026 / Published online: 30 July 2026

Abstract

Precise and robust relative orbit determination is critical for autonomous spacecraft operations in cislunar environments where nonlinear dynamics, sparse measurements, and limited ground support challenge traditional estimation approaches. This work investigates Square Root Unscented Kalman Filters (SRUKF) and their Gaussian Mixture extension (GMSRUKF) for angles-only relative orbit estimation in Near Rectilinear Halo Orbits within the Circular Restricted Three-Body Problem. While sigma-point filters like the SRUKF provide a capable baseline under nominal conditions, their single-Gaussian structure proves insufficient under degraded conditions likely in operational cislunar scenarios, including large initial uncertainty, biased initialization, and sparse measurements. To assess the robustness benefits of mixture-based filtering, this study implements a GMSRUKF with five components and no adaptive mechanisms, and shows that even this minimal extension substantially reduces failure rates and improves consistency compared to the SRUKF. The mixture-based representation enables recovery from poor initialization and intermittent measurements by capturing non-Gaussian uncertainty distributions that arise from three-body dynamics. A consider covariance formulation is also integrated to account for observer state uncertainty without joint estimation. These results demonstrate that the Gaussian Mixture square root filtering framework offers a viable path toward reliable autonomous relative navigation in cislunar space.

Keywords Relative orbit estimation · Angles-only · Gaussian mixture model · Square root unscented Kalman filter · Consider covariance · Cislunar space

1 Introduction

Accurate relative orbit estimation is a foundational capability for many space operations, including formation flying, autonomous rendezvous and docking, and on-orbit

✉ Andrea López
andrea.lopez@colorado.edu

[1] Ann and H.J. Smead Aerospace Engineering Sciences, University of Colorado Boulder, Boulder, CO 80309, USA

servicing. In recent years, angles-only relative orbit estimation has gained attention as a cost-effective and sensor-light alternative to traditional range-based navigation techniques. By relying solely on line-of-sight (LOS) angle measurements obtained from onboard optical sensors, angles-only relative orbit estimation reduces mission dependence on ground-based tracking or active ranging systems like GPS and radar, which may be unavailable, impractical, or too costly for deep-space or distributed missions [1, 2].

This approach is particularly well-suited for cislunar environments, where Earth-based tracking is sparse and maintaining continuous line-of-sight to ground stations is often infeasible [3, 4]. Optical sensors, already widely deployed for attitude determination and proximity operations, offer a passive, power-efficient, and compact means of acquiring relative navigation information. As missions push toward increased autonomy and coordination beyond Earth orbit, the ability to perform accurate, onboard relative navigation using passive sensors will be increasingly critical [5, 6].

However, angles-only relative orbit estimation presents significant challenges due to the inherent limitations of LOS measurements, which are purely angular and lack direct range information. This absence of range data results in a loss of depth perception, limiting observability particularly in the along-track direction [7, 8]. These difficulties are amplified in complex dynamical environments like the Earth-Moon three-body system, where nonlinear dynamics and chaotic behaviors can lead to rapid uncertainty growth and non-Gaussian state distributions. For example, Fig. 1 shows how uncertainty quickly becomes distorted and highly non-Gaussian in a Gateway-like Near Rectilinear Halo Orbit (NRHO), underscoring the need for advanced filtering techniques that can better represent and propagate uncertainty.

Traditional state estimation techniques, such as the Extended Kalman Filter (EKF), are widely used due to their simplicity and computational efficiency. However, their reliance on first-order linearizations makes them ill-suited for highly nonlinear problems like angles-only relative orbit determination (ROD), particularly in dynamical regimes where observability is already limited. As a result, EKF-based estimates often become inconsistent or diverge under large initial uncertainties, extended propagation times, or unmodeled nonlinear effects.

These same nonlinearities that make uncertainty propagation more challenging can also improve observability when appropriately leveraged. Unlike linearized formulations, sigma-point-based filtering techniques such as the Unscented Kalman Filter (UKF) can exploit the inherent nonlinear structure in both the dynamics and measurement models [9]. This enables disambiguation of the state and improved performance by extracting information that would otherwise be lost in a linearized setting [10, 11]. However, sigma-point filters still represent the posterior as a single Gaussian, an assumption that breaks down when the true distribution is multimodal or heavily skewed, as illustrated in Fig. 1.

To capture these non-Gaussian posteriors, recent research has focused on advanced probabilistic estimation techniques that maintain richer representations of uncertainty, such as Sequential Monte Carlo (SMC) methods [12], Particle Filters (PFs) [13], Ensemble Kalman Filters (EnKFs) [14], and Gaussian Mixture Models (GMMs) [15–17], which offer robust solutions for nonlinear and stochastic environments. By maintaining multiple weighted components or samples, these approaches can repre-

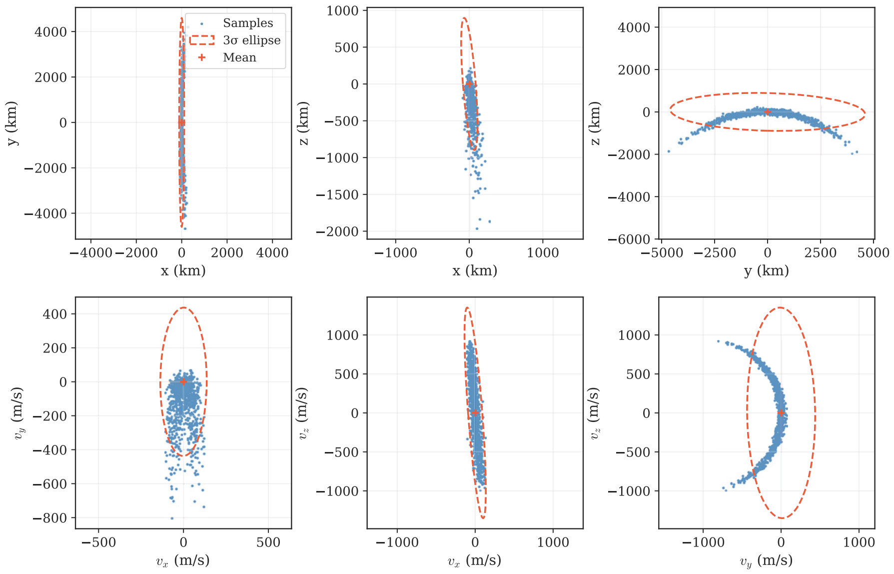

Fig. 1 Monte Carlo simulation showing the growth of uncertainty in a Gateway-like Near Rectilinear Halo Orbit (NRHO) after half an orbit. The simulation uses 1000 samples, initially drawn from a Gaussian distribution with a standard deviation of 1 km in position and 1 m/s in velocity. Over time, the uncertainty becomes significantly non-Gaussian, illustrating the challenges of propagating uncertainties in complex dynamic environments

sent complex posterior distributions and adapt to nonlinearity in both the dynamics and measurement models. Among them, Gaussian Mixture filters offer a compelling balance between expressiveness and computational tractability. They model uncertainty as a weighted sum of Gaussians, allowing for a flexible approximation of nonlinear belief evolution while retaining compatibility with Kalman-based updates [18].

Prior work on GMM-based angles-only filtering has addressed initial orbit determination [16, 17], including initial relative orbit determination in Earth orbit [15]. Square root filter formulations for high-uncertainty navigation have been developed separately [19–21]. This paper builds on these ideas by applying a Gaussian Mixture Square Root UKF to sequential angles-only relative orbit estimation in a Near Rectilinear Halo Orbit modeled with the Circular Restricted Three-Body Problem (CR3BP), which captures the essential nonlinear dynamical features that make this estimation problem challenging. The central finding of this study is that a single-Gaussian SRUKF, despite providing a capable baseline under nominal conditions, exhibits systematic failures under degraded conditions representative of operational cislunar scenarios. A fixed-mixand GMSRUKF with five components and no adaptive mechanisms demonstrates that the mixture representation itself, rather than algorithmic complexity or careful tuning, provides substantial improvements in robustness over the SRUKF and clear advantages over the EKF. A modified square root covariance update is introduced, incorporating noise and consider covariance terms after the rank-one Cholesky update [19] (Sect. 3.1). The consider covariance technique is also integrated into the square root framework to account for observer state uncertainty without full joint estimation.

The remainder of the paper is organized as follows. Section 2 describes the dynamical and measurement models and the simulation scenario, establishing the challenges that motivate the filtering approaches. Section 3 presents the SRUKF, GMSRUKF, and consider covariance formulations. Section 4 reports simulation results across baseline and degraded-condition test cases. Section 5 discusses implications, limitations, and directions for future work.

2 Estimation Problem Setup

This section outlines the assumptions, models, and test scenarios used to evaluate the performance of the filtering approaches. The dynamics and measurement models are presented first, followed by descriptions of the relative orbit geometries considered.

2.1 Dynamic and Measurement Models

Angles-only relative orbit estimation is explored using a dynamical environment based on the Earth-Moon circular restricted three-body problem (CR3BP), with observations generated from line-of-sight geometry. The models described below serve as the foundation for all simulation scenarios presented in this study.

2.1.1 Dynamic Model

The dynamic model used in this study is based on the Earth-Moon Circular Restricted Three-Body Problem (CR3BP), which models the motion of a massless particle in the Earth-Moon system, with the Earth and Moon following a mutual circular orbit. The dynamics of the particle are described in the synodic coordinate frame $\mathcal{R} : \{\hat{\boldsymbol{e}}_x, \hat{\boldsymbol{e}}_y, \hat{\boldsymbol{e}}_z\}$, centered at the barycenter of the system and rotating with the primaries. The $\hat{\boldsymbol{e}}_x$ vector points from the barycenter to the Moon, the $\hat{\boldsymbol{e}}_z$ vector points in the direction of the E-M angular momentum, and the $\hat{\boldsymbol{e}}_y$ vector completes the right-handed orthogonal triad. The mass parameter $\mu = m_M/(m_E + m_M) \approx 1.21506 \times 10^{-2}$ is the ratio of the Moon's mass to the total system mass, so that in normalized units the Earth and Moon contribute mass fractions $(1-\mu)$ and μ, respectively. The length unit LU is the distance between the primaries LU $= 389\,703$ km, and the time unit TU is the time it takes for the Moon to orbit the Earth and is TU $= 382\,981$ s. With these canonical considerations, the position of the Earth is fixed at $\boldsymbol{r}_E = (-\mu, 0, 0)^\top$ and the position of the Moon is fixed at $\boldsymbol{r}_M = (1-\mu, 0, 0)^\top$. The equations of motion in the synodic frame $\mathcal{R}$ for a particle with $\boldsymbol{r} = (x, y, z)^\top$ are given by [22–24]:

$$\begin{bmatrix} \dot{\boldsymbol{r}} \\ \ddot{\boldsymbol{r}} \end{bmatrix} = \boldsymbol{f}_{\mathrm{CR3BP}}(\boldsymbol{x}) = \begin{bmatrix} \dot{\boldsymbol{r}} \\ -\mu \frac{\boldsymbol{r}-\boldsymbol{r}_E}{\|\boldsymbol{r}-\boldsymbol{r}_E\|^3} - (1-\mu)\frac{\boldsymbol{r}-\boldsymbol{r}_M}{\|\boldsymbol{r}-\boldsymbol{r}_M\|^3} + 2\dot{\boldsymbol{r}} \times \hat{\boldsymbol{e}}_z \end{bmatrix} \tag{1}$$

where $\boldsymbol{x} = [\boldsymbol{r}^{\top} \dot{\boldsymbol{r}}^{\top}]^{\top}$ is the state vector expressed in the synodic frame, $\dot{\boldsymbol{r}}$ is the inertial velocity vector, and $\times$ denotes the cross product. The first term in the acceleration represents the gravitational attraction of the Earth, the second term represents the gravitational attraction of the Moon, and the third term represents the Coriolis force due to the rotation of the synodic frame.

The truth relative state of the target spacecraft with respect to the observer spacecraft for these simulations is simply the difference between their absolute states in the synodic frame:

$$\boldsymbol{x}_{\text{rel}} = \boldsymbol{x}_{\text{tar}} - \boldsymbol{x}_{\text{obs}} \tag{2}$$

where $\boldsymbol{x}_{\text{tar}}$ is the absolute state of the target spacecraft and $\boldsymbol{x}_{\text{obs}}$ is the absolute state of the observer spacecraft, both in the rotating (synodic) frame.

2.1.2 Measurement Model

The measurement model for angles-only relative orbit determination is based on the line-of-sight (LOS) angles between the observer spacecraft and the target spacecraft. This type of measurement model is common in space missions, where optical sensors are widely available and can provide accurate angular measurements. However, the LOS measurements do not provide direct range information, leading to challenges in estimating the full relative state vector. The LOS angles are defined as the angles between the line connecting the observer and target spacecraft and the camera frame. The measurement model is given by:

$$z = \boldsymbol{h}(\boldsymbol{x}) = \begin{bmatrix} \theta \\ \phi \end{bmatrix} = \begin{bmatrix} \arctan\left(\frac{\rho_T}{\rho_R}\right) \\ \arcsin\left(\frac{\rho_N}{\sqrt{\rho_R^2+\rho_T^2+\rho_N^2}}\right) \end{bmatrix} \tag{3}$$

where z is the measurement set containing θ, the azimuth angle, and ϕ, the elevation angle.

In this scenario, the camera frame is assumed to be aligned with the observer's RTN frame $\mathcal{O} : \{\hat{\boldsymbol{e}}_R, \hat{\boldsymbol{e}}_T, \hat{\boldsymbol{e}}_N\}$ that is defined using the state of the observer spacecraft relative to the Moon in the synodic frame. This coordinate frame is defined such that $\hat{\boldsymbol{e}}_R$ points radially outward from the Moon to the observer, $\hat{\boldsymbol{e}}_N$ points in the $\boldsymbol{r} \times \dot{\boldsymbol{r}}$ direction, and $\hat{\boldsymbol{e}}_T$ completes the right-handed orthogonal triad. Thus, the relative position vector $\boldsymbol{\rho} = (\rho_R, \rho_T, \rho_N)^{\top}$ between the observer and target spacecraft is expressed in the RTN frame as:

$$^{\mathcal{O}}\boldsymbol{\rho} = {}^{\mathcal{O}}\begin{bmatrix} \rho_R \\ \rho_T \\ \rho_N \end{bmatrix} = [OR]^{\mathcal{R}}\boldsymbol{r}_{\text{rel}} \tag{4}$$

where $[OR]$ is the orthogonal rotation matrix from the synodic frame to the RTN frame, and $^{\mathcal{R}}\boldsymbol{r}_{\text{rel}}$ is the relative position vector expressed in the synodic frame $\mathcal{R}$, the first three components from Eq. (2).

Table 1 Initial conditions for 9:2 NRHO relative navigation scenario (nondimensional rotating frame)

Component	Observer (x_{obs})	Relative (x_{rel})
x	1.0184925637968554e+00	− 5.4436659044920788e−06
y	− 1.3213113195312823e−08	− 8.9711501681069756e−04
z	− 1.7954585719053781e−01	2.2797627880627447e−05
$\dot{x}$	− 1.7056398720481548e−08	− 1.1581229295988791e−03
$\dot{y}$	− 9.5438750184004911e−02	1.8789496313231790e−05
$\dot{z}$	7.1429061824995221e−08	4.8502487579814343e−03

The measurement model noise is assumed to be Gaussian, with a zero mean and a covariance matrix given by:

$$\boldsymbol{R} = \begin{bmatrix} \sigma_\theta^2 & 0 \\ 0 & \sigma_\phi^2 \end{bmatrix} \tag{5}$$

with $\sigma_\theta = \sigma_\phi = 1 \times 10^{-4}$ rad $\approx$ 20 arcsec. This noise level is representative of the performance of typical optical sensors used in space missions, such as star trackers or cameras, and is consistent with the noise levels used in previous studies [7, 10, 15].

2.2 Relative Orbit Description

This section describes the simulation environment and initialization parameters used to evaluate the performance of the filtering approaches presented in this study. The scenario considered is a lead-follower configuration in a 9:2 Near Rectilinear Halo Orbit (NRHO) around the Moon, which is representative of the planned Gateway orbit for cislunar missions [5]. The initial condition was obtained from the JPL periodic orbits database and refined using a single-shooting corrector to improve periodicity within the simulation environment. This NRHO serves as an illustrative trajectory for cislunar missions with strong dynamical structure and observability variation.

The target spacecraft is initialized by propagating the observer's initial state forward by one hour. This setup provides a meaningful yet controlled relative geometry that is suitable for filter performance benchmarking under nonlinear dynamics. The resulting separation varies substantially over the orbit, ranging from 350 km near apolune to 6054 km near perilune (mean 1090 km), with the right-skewed distribution reflecting the longer time spent near apolune on an NRHO.

The initial conditions for this scenario are included in Table 1. The resulting relative trajectory is shown in Fig. 2a, and the corresponding bearing angles over time are shown in Fig. 2b.

3 Filtering Approaches

The estimation scenario described in the previous section poses several challenges that inform the choice of filtering approach. The combination of nonlinear CR3BP

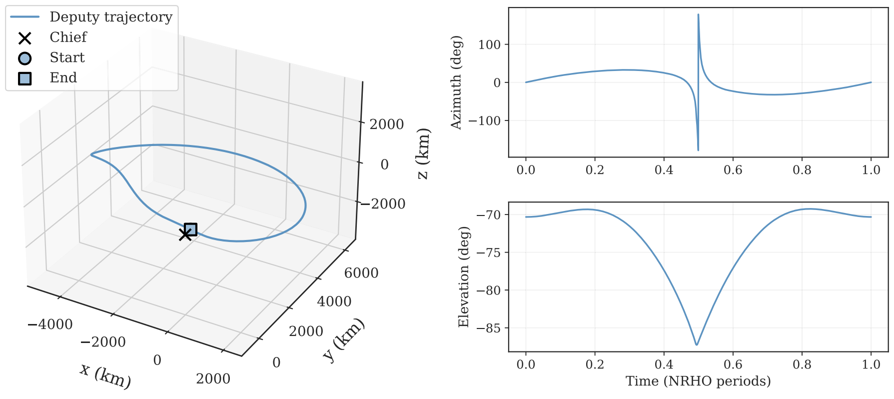

(a) Relative Orbit Trajectory (b) Bearing Angles over Time

Fig. 2 Relative orbit geometry in rotating frame and bearing angles for the lead-follower configuration in a 9:2 NRHO around the Moon. The target spacecraft is initialized by propagating the observer's initial state forward by one hour, resulting in a meaningful relative geometry for filter performance evaluation. In (a), the start marker (circle) coincides with the end marker (square) location and is not separately visible; the initial and final positions are nearly identical due to the periodic orbit geometry

dynamics, angles-only measurements with inherently limited observability, and the rapid uncertainty growth illustrated in Fig. 1 means that linearization-based filters are unlikely to maintain consistency over extended propagation horizons. At the same time, the non-Gaussian character of the propagated uncertainty suggests that a single Gaussian representation may be insufficient to capture the true posterior distribution, particularly under degraded conditions such as large initialization errors or sparse measurements. These considerations motivate the use of sigma-point filtering to handle nonlinearity without linearization, and Gaussian Mixture extensions to represent non-Gaussian posteriors. This study investigates the Square Root Unscented Kalman Filter (SRUKF) and its Gaussian Mixture extension (GMSRUKF) as the primary filtering architectures.

3.1 Square Root Unscented Kalman Filter (SRUKF)

The SRUKF is a numerically stable variant of the UKF that propagates the Cholesky factor of the covariance matrix instead of the full covariance [25]. This approach guarantees positive semi-definiteness of the covariance at each step and avoids direct computation of covariance square roots during every generation of sigma points, reducing numerical errors and improving robustness, particularly over long-duration propagations or under large uncertainty growth. In this implementation, process and measurement noise are treated as additive and incorporated directly into the covariance update. The SRUKF algorithm is summarized in Algorithm 1. Further algorithmic details and derivations are found in [19, 25].

Working directly with Cholesky factors throughout Algorithm 1 offers computational advantages that are less obvious from the formulation alone. The Kalman

Algorithm 1 Square Root Unscented Kalman Filter (SRUKF)

Require: Prior state $\boldsymbol{x}_{k-1}^{+}$, square root covariance $\boldsymbol{S}_{k-1}^{+}$ with $\boldsymbol{P}_{k-1}^{+} = \boldsymbol{S}_{k-1}^{+}\boldsymbol{S}_{k-1}^{+\top}$, measurement z_k, UKF weights $\{W_i^m, W_i^c\}_{i=0}^{2n}$, optional additive factors $[\boldsymbol{S}_{\text{add}}^{-}]$: prediction noise; $[\boldsymbol{S}_{\text{add}}^{+}]$: update noise

{——————- Prediction Step ——————-}

1: Generate sigma points $\{\boldsymbol{\chi}_{i,k-1}^{+}\}_{i=0}^{2n}$ from $\boldsymbol{x}_{k-1}^{+}$ and $\boldsymbol{S}_{k-1}^{+}$
2: **for all** $i = 0, \ldots, 2n$ **do**
3: Propagate: $\boldsymbol{\chi}_{i,k}^{-} \leftarrow \boldsymbol{f}(\boldsymbol{\chi}_{i,k-1}^{+})$
4: **end for**
5: Predicted state: $\boldsymbol{x}_k^{-} \leftarrow \sum_{i=0}^{2n} W_i^m \boldsymbol{\chi}_{i,k}^{-}$
6: Predicted square root covariance:
7: $\boldsymbol{S}_k^{-} \leftarrow \text{ComputeSqrtCov}(\{\boldsymbol{\chi}_{i,k}^{-}\},\ \boldsymbol{x}_k^{-},\ \{W_i^c\},\ [\boldsymbol{S}_{\text{add}}^{-}])$

{——————- Update Step ——————-}

8: Generate sigma points $\{\boldsymbol{\chi}_{i,k}^{-}\}_{i=0}^{2n}$ from $\boldsymbol{x}_k^{-}$ and $\boldsymbol{S}_k^{-}$
9: **for all** $i = 0, \ldots, 2n$ **do**
10: Predicted measurement: $\boldsymbol{\zeta}_i \leftarrow \boldsymbol{h}(\boldsymbol{\chi}_{i,k}^{-})$
11: **end for**
12: Measurement mean: $z_k^{-} \leftarrow \sum_{i=0}^{2n} W_i^m \boldsymbol{\zeta}_i$
13: Innovation square root covariance:
14: $\boldsymbol{S}_z \leftarrow \text{ComputeSqrtCov}(\{\boldsymbol{\zeta}_i\},\ z_k^{-},\ \{W_i^c\},\ [\boldsymbol{S}_{\text{add}}^{+}])$
15: Cross-covariance: $\boldsymbol{P}_{xz} \leftarrow \sum_{i=0}^{2n} W_i^c (\boldsymbol{\chi}_{i,k}^{-} - \boldsymbol{x}_k^{-})(\boldsymbol{\zeta}_i - z_k^{-})^{\top}$
16: Kalman gain: $\boldsymbol{K}_k \leftarrow \boldsymbol{P}_{xz}(\boldsymbol{S}_z\boldsymbol{S}_z^{\top})^{-1}$
17: Updated state: $\boldsymbol{x}_k^{+} \leftarrow \boldsymbol{x}_k^{-} + \boldsymbol{K}_k(z_k - z_k^{-})$
18: Updated square root covariance:
19: $\boldsymbol{U} \leftarrow \boldsymbol{K}_k \boldsymbol{S}_z$
20: $\boldsymbol{S}_k^{+} \leftarrow \text{CholUpdate}(\boldsymbol{S}_k^{-},\ \boldsymbol{U},\ -1)$

gain $\boldsymbol{K}_k$ at line 16 is computed via two triangular solves rather than explicit matrix inversion, which is more numerically stable and efficient. The innovation covariance $\boldsymbol{P}_{zz} = \boldsymbol{S}_z\boldsymbol{S}_z^{\top}$ does not need to be formed explicitly: $\boldsymbol{S}_z$ is computed directly via `ComputeSqrtCov` at line 14. The posterior covariance update at lines 18–20 uses a Cholesky downdate, maintaining the square root representation throughout.

The optional additive factors $\boldsymbol{S}_{\text{add}}^{-}$ and $\boldsymbol{S}_{\text{add}}^{+}$ in Algorithm 1 represent square root matrices that may be included for covariance inflation during prediction and update, respectively. These can represent process noise ($\boldsymbol{S}_Q$), measurement noise ($\boldsymbol{S}_R$), consider covariance contributions ($\boldsymbol{G}_f\boldsymbol{S}_{\text{obs}}$ or $\boldsymbol{G}_h\boldsymbol{S}_{\text{obs}}$), or combinations thereof. When omitted, the algorithm reduces to pure uncertainty propagation through the nonlinear dynamics and measurement model.

The `ComputeSqrtCov` function referenced in Algorithm 1 is detailed in Algorithm 2. This subroutine computes the square root (Cholesky factor) of a weighted covariance matrix directly from sigma points, without forming the full covariance. The approach is more numerically stable than computing the covariance $\boldsymbol{P}$ explicitly and then factoring it, particularly when negative weights are present (as occurs in some UKF parameterizations). The separate handling of the central sigma point via rank-one Cholesky update allows for proper downdating when $W_0^c < 0$, avoiding potential loss of positive definiteness.

Van der Merwe and Wan [19] perform the QR decomposition on the sigma-point deviations (points 1 through $2n$) augmented with the additive noise columns $\boldsymbol{S}_{\text{add}}$ in a single step, then apply the rank-one Cholesky update for the zeroth sigma point to the resulting triangular factor. When the additive noise and deviation columns differ substantially in magnitude, this combined QR produces a poorly conditioned triangular factor. The Givens rotations in the subsequent rank-one update are sensitive to column scaling, and the discrepancy introduces numerical errors in practice. Algorithm 2 instead performs QR on the sigma-point deviations alone, applies the rank-one update to the resulting factor, and incorporates $\boldsymbol{S}_{\text{add}}$ in a separate QR augmentation step afterward. This reordering is a minor change to the original algorithm, but it had a measurable effect on filter consistency in the high-uncertainty scenarios considered here, and the mismatch was nontrivial to identify as the source of numerical degradation.

Algorithm 2 ComputeSqrtCov: Numerically Stable Square Root Covariance Computation

Require: Sigma points $\{\boldsymbol{\xi}_i\}_{i=0}^{2n}$, mean $\bar{\boldsymbol{\xi}}$, weights $\{W_i^c\}_{i=0}^{2n}$, optional noise factor $[\boldsymbol{S}_{\text{add}}]$
Ensure: Square root factor $\boldsymbol{S}$ such that $\boldsymbol{P} = \boldsymbol{S}\boldsymbol{S}^\top = \sum_{i=0}^{2n} W_i^c(\boldsymbol{\xi}_i - \bar{\boldsymbol{\xi}})(\boldsymbol{\xi}_i - \bar{\boldsymbol{\xi}})^\top + [\boldsymbol{S}_{\text{add}}\boldsymbol{S}_{\text{add}}^\top]$
1: Compute weighted deviations:
2: $\boldsymbol{d}_0 = \sqrt{|W_0^c|}\,(\boldsymbol{\xi}_0 - \bar{\boldsymbol{\xi}})$
3: $\boldsymbol{d}_i = \sqrt{W_i^c}\,(\boldsymbol{\xi}_i - \bar{\boldsymbol{\xi}})$ for $i = 1, \ldots, 2n$
4: Perform QR decomposition: $\boldsymbol{S} \leftarrow \text{qr}\left([\boldsymbol{d}_1, \ldots, \boldsymbol{d}_{2n}]\right)$
5: Apply rank-one Cholesky update: $\boldsymbol{S} \leftarrow \text{CholUpdate}(\boldsymbol{S},\ \boldsymbol{d}_0,\ \text{sign}(W_0^c))$
6: **if** $\boldsymbol{S}_{\text{add}}$ provided **then**
7: Augment with noise: $\boldsymbol{S} \leftarrow \text{qr}\left([\boldsymbol{S}|\boldsymbol{S}_{\text{add}}]\right)$
8: **end if**
9: **return** $\boldsymbol{S}$

3.2 Gaussian Mixture Square Root Unscented Kalman Filter (GMSRUKF)

The GMSRUKF extends the SRUKF to a mixture of M parallel components (mixands), each carrying its own state estimate, square root covariance, and weight. Each mixand evolves independently through the SRUKF prediction and update steps, and weights are updated by likelihood-based reweighting. By maintaining multiple Gaussian components, the filter can represent non-Gaussian posteriors that a single SRUKF cannot capture. A summary of the algorithm is provided in Algorithm 3.

Within the square root formulation, likelihood evaluation follows naturally from the Cholesky factors already computed during the update step. The inverse of the innovation covariance for each mixand is obtained via triangular solves, consistent with the Kalman gain computation, and its determinant is recovered directly from the diagonal of $\boldsymbol{S}_z^{(j)}$ without forming the full covariance matrix explicitly.

3.3 Consider Covariance Implementation

In angles-only relative navigation, the observer spacecraft's own state uncertainty affects both the propagation of the relative state and the interpretation of measurements.

Algorithm 3 Gaussian Mixture Square Root Unscented Kalman Filter (GMSRUKF)

Require: Prior mixands $\{\boldsymbol{x}_{k-1}^{+(j)}, \boldsymbol{S}_{k-1}^{+(j)}, w_{k-1}^{(j)}\}_{j=1}^{M}$, measurement z_k

{———— Prediction Step ————}

1: **for all** mixands $j = 1$ to M **do**

2: Apply SRUKF Predict Step to $(\boldsymbol{x}_{k-1}^{+(j)}, \boldsymbol{S}_{k-1}^{+(j)})$ to obtain $(\boldsymbol{x}_k^{-(j)}, \boldsymbol{S}_k^{-(j)})$

3: **end for**

{———— Update Step ————}

4: **for all** mixands $j = 1$ to M **do**

5: Apply SRUKF Update Step to $(\boldsymbol{x}_k^{-(j)}, \boldsymbol{S}_k^{-(j)})$ with z_k to obtain $(\boldsymbol{x}_k^{+(j)}, \boldsymbol{S}_k^{+(j)})$

6: Compute likelihood $\ell^{(j)} = \mathcal{N}(z_k;\ z_k^{-(j)},\ \boldsymbol{S}_z^{(j)} \boldsymbol{S}_z^{(j)\top})$

7: **end for**

8: Normalize weights:

9: $w_k^{(j)} \leftarrow \dfrac{w_{k-1}^{(j)} \ell^{(j)}}{\sum_{i=1}^{M} w_{k-1}^{(i)} \ell^{(i)}}$

10: (Optional) Compute overall mixture mean and covariance:

11: $\boldsymbol{x}_k^{+} = \sum_j w_k^{(j)} \boldsymbol{x}_k^{+(j)}$

12: $\boldsymbol{P}_k^{+} = \sum_j w_k^{(j)} \left(\boldsymbol{S}_k^{+(j)} \boldsymbol{S}_k^{+(j)\top} + \boldsymbol{x}_k^{+(j)} \boldsymbol{x}_k^{+(j)\top}\right) - \boldsymbol{x}_k^{+} \boldsymbol{x}_k^{+\top}$

13: (Optional) Compute mixture square root: $\boldsymbol{S}_k^{+} \leftarrow \mathrm{chol}(\boldsymbol{P}_k^{+})$

When this is not accounted for, the reported covariance reflects only errors in the estimated relative state, omitting the additional uncertainty contributed by an imperfectly known observer. Jointly estimating observer and target states would address this, but at the cost of a significantly larger state and the need to solve a separate navigation problem for the observer.

The consider covariance technique [25] provides a way to account for this without joint estimation. The observer state is treated as a parameter with known uncertainty statistics. Rather than augmenting the state vector with observer state variables, the influence of the observer covariance $\boldsymbol{P}_{\mathrm{obs}}$ on the relative state covariance is captured through first-order sensitivity matrices.

During the time update, the propagated relative covariance is supplemented with the observer contribution as:

$$\boldsymbol{P} \leftarrow \boldsymbol{P} + \boldsymbol{G}_f \boldsymbol{P}_{\mathrm{obs}} \boldsymbol{G}_f^{\top} \tag{6}$$

where $\boldsymbol{G}_f = \partial \boldsymbol{f}_{\mathrm{rel}} / \partial \boldsymbol{x}_{\mathrm{obs}}$ is the sensitivity of the relative dynamics to the observer state. In the CR3BP, since both spacecraft are modeled as massless test particles, the target dynamics are independent of the observer state and $\boldsymbol{G}_f$ reduces to the negative Jacobian of the observer's absolute dynamics. The derivation and explicit form are given in Appendix 1.

Similarly, the innovation covariance is modified as:

$$\boldsymbol{P}_{zz} \leftarrow \boldsymbol{P}_{zz} + \boldsymbol{G}_h \boldsymbol{P}_{\mathrm{obs}} \boldsymbol{G}_h^{\top} \tag{7}$$

where $\boldsymbol{G}_h = \partial \boldsymbol{h} / \partial \boldsymbol{x}_{\mathrm{obs}}$ captures the dependence of the measurement function on the observer state through the projection into the camera frame (see Appendix 1). The CC terms are structurally distinct from process noise: process noise models uncharacterized forces acting on the target (dynamics model mismatch, unmodeled perturbations), whereas the CC terms model observer navigation uncertainty projected into the relative state and measurement spaces through the sensitivity matrices $\boldsymbol{G}_f$ and $\boldsymbol{G}_h$.

In the square root implementation, these contributions enter `ComputeSqrtCov` as additive factors: $\boldsymbol{S}_{\text{add}}^{-} = [\boldsymbol{G}_f \boldsymbol{S}_{\text{obs}} \mid \boldsymbol{S}_Q]$ during prediction and $\boldsymbol{S}_{\text{add}}^{+} = [\boldsymbol{G}_h \boldsymbol{S}_{\text{obs}} \mid \boldsymbol{S}_R]$ during the measurement update, where $\boldsymbol{S}_{\text{obs}}$ is the Cholesky factor of $\boldsymbol{P}_{\text{obs}}$. In scenarios without consider covariance, the additive factors reduce to $\boldsymbol{S}_{\text{add}}^{-} = \boldsymbol{S}_Q$ (when applicable), and $\boldsymbol{S}_{\text{add}}^{+} = \boldsymbol{S}_R$. For the Gaussian Mixture variants, $\boldsymbol{G}_f$ and $\boldsymbol{G}_h$ are shared across all mixands since all components reference the same observer state, which reduces computational overhead.

3.4 Initialization and Assumptions

As a baseline, all filters are initialized with a prior state $\boldsymbol{x}_0$ sampled from a multivariate normal distribution $\boldsymbol{x}_0 \sim \mathcal{N}(\boldsymbol{x}_{0,\text{truth}}, \boldsymbol{P}_0)$, where $\boldsymbol{x}_{0,\text{truth}}$ is the true initial state and $\boldsymbol{P}_0$ is the initial covariance. A complete summary of filter initialization parameters, measurement configuration, and simulation settings is provided in Table 2.

For the Gaussian Mixture filter, $M = 5$ mixands are initialized for each Monte Carlo trial by independently sampling from the prior distribution $\boldsymbol{x}_0^{(j)} \sim \mathcal{N}(\boldsymbol{x}_{0,\text{truth}}, \boldsymbol{P}_0)$. Each mixand evolves independently according to the SRUKF update rules, and weights are updated based on measurement likelihoods. No mixand merging or pruning is performed in this study, representing a minimal implementation of the Gaussian Mixture approach.

Although the filters operate on the relative state, sigma point generation in the prediction step is performed in absolute coordinates. In the normalized CR3BP, relative states in the NRHO lead-follower configuration are small relative to the absolute state magnitudes. Propagating these small relative quantities directly through the full CR3BP acceleration risks significant loss of floating-point precision. Instead, sigma points are constructed by first adding the current estimate of the observer's absolute state to the relative state estimate to recover absolute target states. These absolute sigma points are then propagated through the nonlinear CR3BP dynamics, and the observer's absolute state is subtracted afterward to recover the relative state. This keeps propagation in a regime where the dynamics and state magnitudes are consistent, avoiding precision loss in the sigma point deviations.

All simulation scenarios are evaluated using 100 Monte Carlo trials, each with an independently sampled initial state drawn from the scenario-specific prior distribution. This approach ensures statistical validity of the performance metrics and captures the variability in filter behavior across different realizations of the stochastic initialization process.

Table 2 Filter initialization and simulation parameters

Parameter	Value/Description
Filter Configuration	
Process noise covariance ($\boldsymbol{Q}$)	$\mathbf{0}$ (all filters)
Sigma Point Parameters (Van der Merwe [19])	
Spread parameter	$\alpha = 10^{-3}$
Distribution parameter	$\beta = 2$ (optimal for Gaussian)
Secondary spread parameter	$\kappa = 0$
GMSRUKF Configuration	
Number of mixands (M)	5
Weight update	Measurement likelihood
Splitting/merging/pruning	Not applied
Filter Initialization (Baseline)	
Initial State Uncertainty	
Position uncertainty	± 1 km per axis (1-σ)
Velocity uncertainty	± 1 m/s per axis (1-σ)
Covariance structure	Diagonal ($\boldsymbol{P}_0$)
SRUKF	
Prior mean ($\boldsymbol{x}_0$)	Sampled from $\mathcal{N}(\boldsymbol{x}_{0,\text{truth}}, \boldsymbol{P}_0)$
Prior covariance ($\boldsymbol{P}_0$)	As above
GMSRUKF	
Mixand prior mean ($\boldsymbol{x}_0^{(j)}$)	Sampled from $\mathcal{N}(\boldsymbol{x}_{0,\text{truth}}, \boldsymbol{P}_0)$
Mixand covariance ($\boldsymbol{P}_0^{(j)}$)	Initialized to $\boldsymbol{P}_0$
Initial weights ($w_0^{(j)}$)	$1/M$ (uniform)
Measurement Configuration	
Measurement interval (Δt)	60 s
Angular noise ($\sigma_\theta, \sigma_\phi$)	$\pm$ 1e-4 rad (1-σ)
Noise covariance ($\boldsymbol{R}$)	Diagonal, as per Eq. (5)
Field of view	No constraints (continuous visibility)
Simulation Parameters	
Simulation horizon	2 orbits ($\approx$ 13 days)
Reference frame	CR3BP rotating (synodic)
Monte Carlo trials per scenario	100
Occultation/illumination	Not modeled

4 Results

This section presents simulation results evaluating filter performance across a range of test cases in the cislunar relative navigation context. The analysis focuses on position and velocity root mean square (RMS) errors, failure statistics, and worst-case normalized errors. Results include comparisons between single-Gaussian and Gaussian Mixture filter variants. Three recurring figure types are used throughout.

In RMS error plots, solid lines show the mean RMS error across 100 Monte Carlo trials, shaded regions indicate the standard deviation of those RMS errors across trials, and dashed lines show the mean 3σ bounds computed as three times the square root of the corresponding diagonal element of the estimated state covariance $\boldsymbol{P}$, averaged across trials. The comparison of 3σ bounds against RMS error levels serves as a proxy for filter consistency.

Inconsistency histograms show, for each trial, the percentage of simulation time during which the RMS error exceeds the 3σ bound. Trials below 10% are classified as successful and those above 50% as failures. The shape of the histogram reveals the population-level distribution of filter reliability across trials.

Normalized error scatter plots show the 10 worst-performing trials (by maximum normalized position error) from each 100-trial Monte Carlo set. Each trial is represented by two points connected by a line: the maximum normalized position error (triangle) and the 90th percentile normalized position error (circle) for that same trial. A red dashed line marks the 3σ threshold. The line between them conveys the spread between worst-case and typical error within each trial.

4.1 Baseline Filter Comparison

Under nominal conditions, the EKF, UKF, and SRUKF are compared. All filters share the same initialization: prior uncertainty $\boldsymbol{P}_0$, measurement noise covariance $\boldsymbol{R}$, and no process noise ($\boldsymbol{Q} = \boldsymbol{0}$), as specified in Table 2. The intent is not to tune any filter to this specific scenario, but to evaluate each architecture's out-of-the-box performance under identical conditions.

Figure 3 shows the RMS position and velocity errors over two orbital periods. In this scenario and without process noise tuning, the EKF becomes inconsistent across all trials: its covariance collapses within one orbital period, leaving 3-sigma bounds that no longer capture the true error distribution. The UKF and SRUKF maintain consistent estimates within their uncertainty bounds, with nearly identical performance. State estimates converge to meter and millimeter-per-second levels after initial transients. The periodic error spikes at approximately 0.5 and 1.5 orbital periods correspond to perilune passages, where the CR3BP dynamics evolve most rapidly. These transient increases in error are expected and do not indicate filter divergence. The square root formulation matches the UKF in estimation quality while improving numerical stability.

With the SRUKF established as the single-Gaussian baseline, the effect of the number of Gaussian Mixture components is tested. The GMSRUKF is run with $M = 1, 3, 5$, and 10 mixands under the same nominal conditions. The $M = 1$ case reduces to a single SRUKF.

Figure 4 shows the RMS errors for each configuration. All perform comparably: the Gaussian Mixture extension introduces no penalty under favorable conditions, with slightly larger initial uncertainty bounds for higher M, but nearly indistinguishable performance after the first perilune passage.

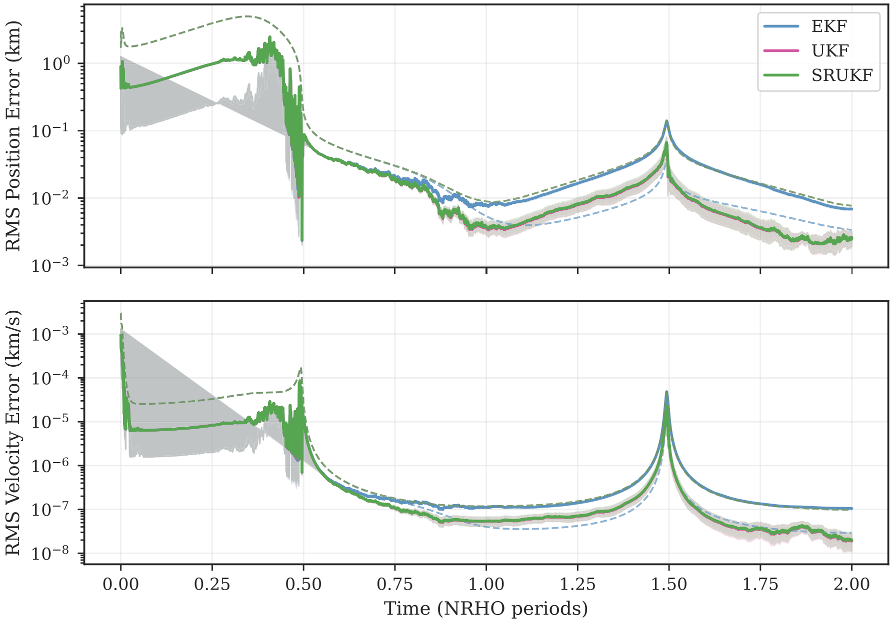

Fig. 3 Comparison of EKF, UKF, and SRUKF under nominal conditions with $\boldsymbol{Q} = \mathbf{0}$. The EKF becomes inconsistent rapidly, while the sigma-point filters (UKF, SRUKF) remain consistent

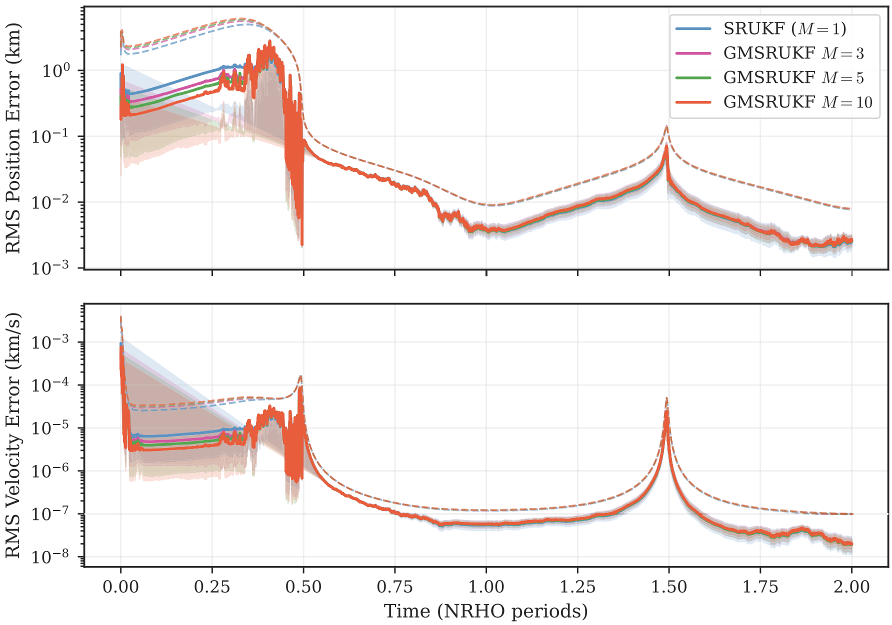

Fig. 4 Sensitivity of GMSRUKF performance to the number of mixands ($M = 1, 3, 5, 10$) under nominal conditions. All configurations achieve comparable accuracy, confirming that the Gaussian Mixture extension does not degrade baseline performance

As a balance between representational expressiveness and computational tractability, $M = 5$ is selected for the remainder of the study. The advantage of using multiple mixands over a single-Gaussian filter becomes apparent in the degraded-condition scenarios that follow, where larger initialization errors and sparser measurements stress the single-Gaussian approximation.

4.2 Robustness to Initial Uncertainty

Cislunar missions often face scenarios where target spacecraft states are only coarsely known due to limited ground tracking or delayed handoff from external navigation sources. To evaluate filter robustness under such conditions, simulations are conducted with varying levels of initial uncertainty and bias in the prior state estimate. Two levels of prior uncertainty are tested by scaling the nominal initial covariance by factors of $\alpha \in \{5, 10\}$:

$$\boldsymbol{P}_0' = \alpha^2 \boldsymbol{P}_0, \qquad \boldsymbol{x}_0' \sim \mathcal{N}(\boldsymbol{x}_{0,\text{truth}},\ \boldsymbol{P}_0') \tag{8}$$

Additionally, biased initializations are tested by shifting the prior mean by $2\alpha\sigma_{\text{pos}}$ along the in-track direction:

$$\boldsymbol{x}_0'' \sim \mathcal{N}(\boldsymbol{x}_{0,\text{truth}} + \boldsymbol{b},\ \boldsymbol{P}_0'), \qquad \boldsymbol{b} = 2\alpha\sigma_{\text{pos}}\,[OR]^\top \left[0\ 1\ 0\ 0\ 0\ 0\right]^\top \tag{9}$$

The in-track direction is chosen because it is the direction most susceptible to LOS alignment in this configuration and is thus poorly observable [7], making it a representative worst-case for initial bias recovery.

Figure 5 shows the distribution of initial condition errors quantified by the normalized distance $|\boldsymbol{e}|_2 = |(\boldsymbol{x}0 - \boldsymbol{x}0,\text{ truth})/\boldsymbol{\sigma}_0|$, which expresses the initial guess error in units of the filter's internal uncertainty. This gives a sense of how challenging each scenario is: unbiased cases have expected distance 2.35σ (chi distribution with $k = 6$; since $\sqrt{\cdot}$ is concave, $E[|\boldsymbol{e}|] < \sqrt{k}$, with a second-order Taylor correction yielding 2.35) and biased cases 3.05σ (noncentral chi with $k = 6$, $\lambda = 4$ from the 2σ in-track shift; same correction applied to $\sqrt{k + \lambda}$), meaning the biased initial guess is, on average, about three times the 1σ uncertainty from the truth. Observed means are consistent with these predictions, confirming correct sampling. Because the metric normalizes by $\boldsymbol{\sigma}_0$, both α levels collapse onto the same distribution. The GMSRUKF distances appear more concentrated than the SRUKF because the reported state is the mixture mean, which averages over independently sampled mixands and reduces apparent spread.

Figures 6 and 7 show the RMS position and velocity errors over two orbital periods for the SRUKF and GMSRUKF under varying levels of prior uncertainty, without and with bias in the initial estimate, respectively. For the unbiased scenario, both filters are shown to perform well at the moderate uncertainty level ($\alpha = 5$). However, as the initial uncertainty increases to $\alpha = 10$, the SRUKF's performance degrades significantly, with RMS errors exceeding the 3σ bounds for extended periods, indicating filter divergence or overconfidence. In contrast, the GMSRUKF maintains better consistency, with RMS errors remaining within bounds for most of the simulation.

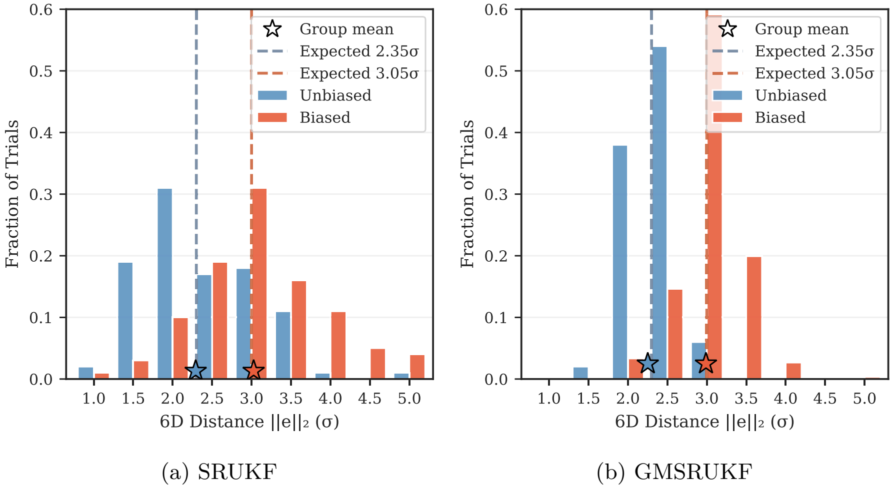

Fig. 5 Distribution of initial condition errors using the normalized distance metric $\|\boldsymbol{e}\|_2$. Dashed lines show the corrected theoretical expected values: 2.35σ for the unbiased case and 3.05σ for the biased case. Stars show observed group means. The GMSRUKF distribution appears more concentrated because the mixture mean is reported

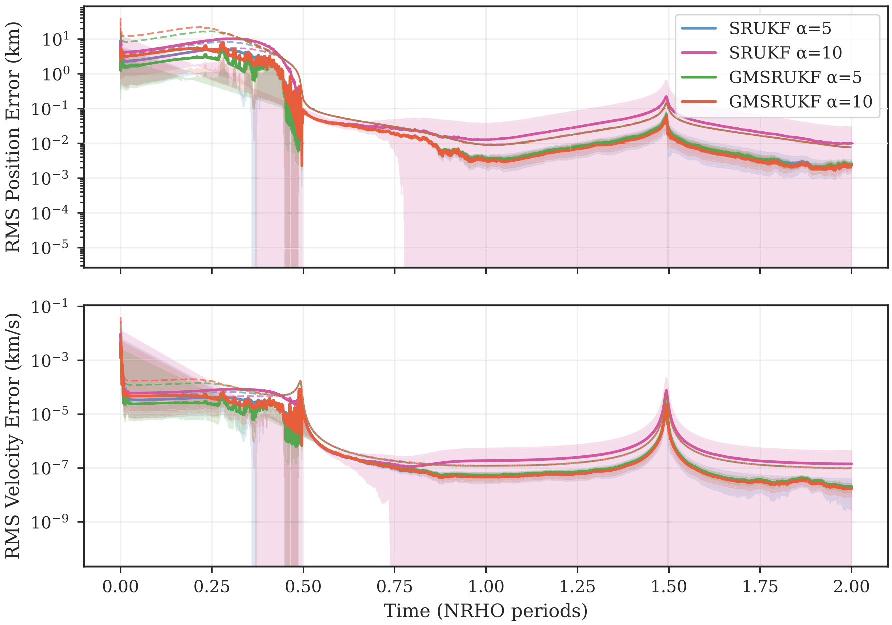

Fig. 6 RMS position and velocity errors for SRUKF and GMSRUKF under varying levels of prior uncertainty with unbiased initial guesses

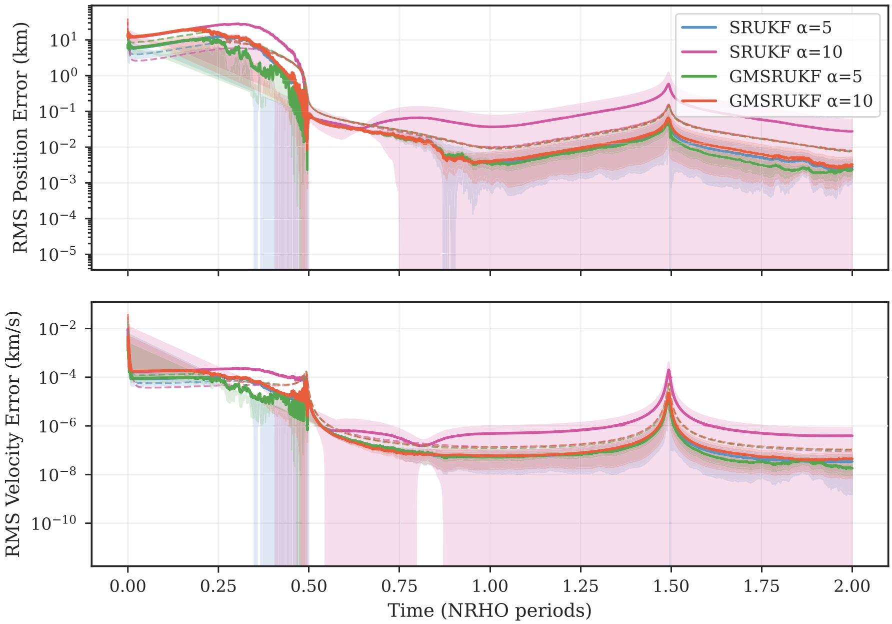

Fig. 7 RMS position and velocity errors for SRUKF and GMSRUKF under varying levels of prior uncertainty with biased initial guesses

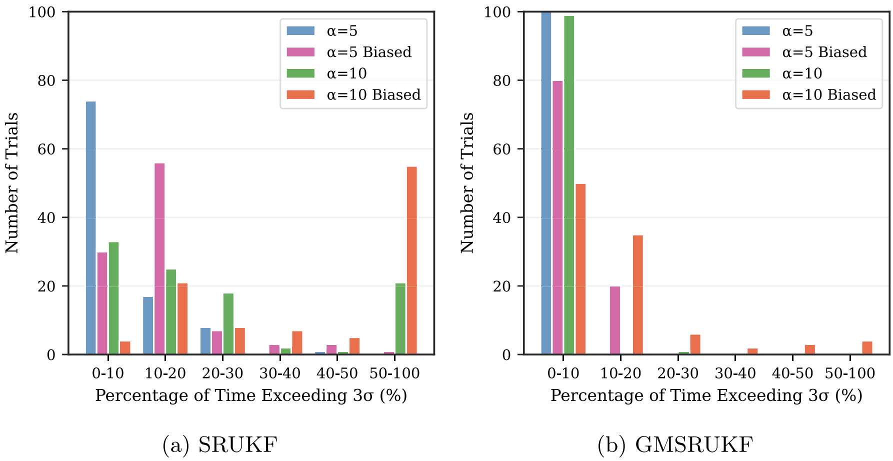

Fig. 8 Inconsistency histograms for SRUKF and GMSRUKF under varying levels of prior uncertainty and bias

When a bias is introduced in the initial estimate, the performance degradation of the SRUKF becomes more pronounced, particularly at higher uncertainty levels. The filter struggles to recover from the biased initial condition, leading to persistent errors outside the 3σ bounds. The GMSRUKF, however, demonstrates greater resilience to the bias, successfully converging in a significant portion of trials even at $\alpha = 10$.

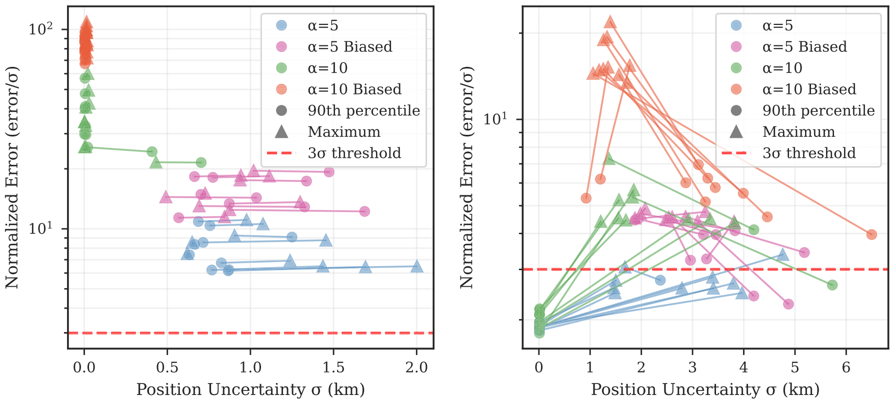

(a) SRUKF (b) GMSRUKF

Fig. 9 Normalized position errors for the SRUKF and GMSRUKF under varying levels of prior uncertainty and bias. Note that the y-axis scales differ between (a) and (b), reflecting the substantially larger normalized errors for the SRUKF under high-uncertainty conditions

Figure 8 shows inconsistency statistics for both filters. The SRUKF struggles even under unbiased, moderate-uncertainty conditions, with a limited number of successful trials across the board. The GMSRUKF, by contrast, succeeds in nearly all unbiased trials and maintains a high success rate under bias as well, with nearly all trials staying below 20% inconsistency. Even at the highest uncertainty and bias levels tested, fewer than 5 failures occur, underscoring the robustness of the mixture representation to poor initialization.

Figure 9 shows normalized position errors for SRUKF (left) and GMSRUKF (right) across varying initial uncertainty levels. Trials that remain near or below the 3-sigma threshold (red dashed line) indicate well-calibrated uncertainty estimates, while those significantly above indicate filter overconfidence. The GMSRUKF maintains errors closer to the 3-sigma bound across all uncertainty levels, particularly for the $\alpha = 10$ biased case, where SRUKF shows substantially larger normalized errors (note the scale difference between Fig. 9a and b). The SRUKF worst trials (particularly for $\alpha = 10$ with bias) cluster near zero position uncertainty, indicating covariance collapse where the filter becomes overconfident despite large errors. In contrast, GMSRUKF maintains reasonable uncertainty estimates (larger σ values) even in failed trials, with normalized errors staying closer to the 3-σ threshold.

4.3 Robustness to Measurement Dropout

Operational cislunar missions may experience intermittent measurement availability due to sensor field-of-view constraints, resource limitations or LOS obstructions. To evaluate filter robustness degraded measurement conditions, simulations are conducted with varying levels of randomly dropped observations. Measurement loss rates of $p_{\text{loss}} \in \{0.9, 0.99\}$ are tested, where p_{loss} represents the probability that any given

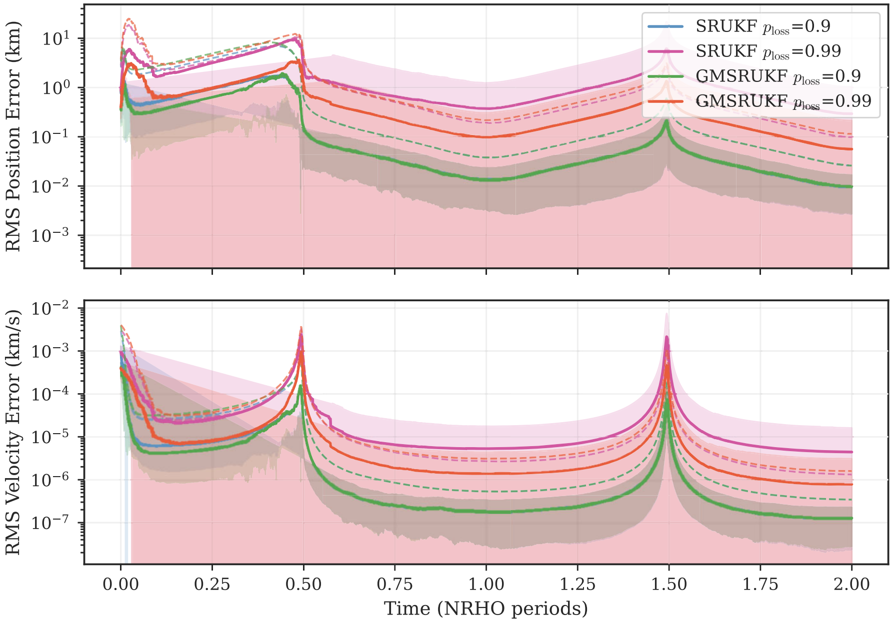

Fig. 10 RMS errors for the filters under varying levels of randomly dropped measurements

measurement is omitted at each time step. The filters are run under nominal initialization ($\boldsymbol{P}_0$) with no added prior bias or increased uncertainty.

One notable finding from these simulations is that both filters exhibit considerable robustness to measurement loss. With each LOS observation independently dropped with probability $p_{\text{loss}} = 0.9$ at each time step, the SRUKF and GMSRUKF maintain bounded estimation errors and perform nearly identically, as shown in Fig. 10. This robustness is notable given that the data loss is unstructured (not periodic subsampling), which presents a more challenging scenario for maintaining observability. The random dropout creates highly variable measurement gaps, averaging 10 min for $p_{\text{loss}} = 0.9$ and 100 min for $p_{\text{loss}} = 0.99$, during which the spacecraft state must be propagated through the highly nonlinear CR3BP dynamics with no corrective updates. These extended gaps can coincide with sensitive orbital phases, such as perilune passages, where dynamics evolve rapidly and small errors can grow quickly, making the measurement dropout scenarios particularly demanding for any filter.

Figure 10 reveals that while both filters generally succeed at $p_{\text{loss}} = 0.9$, a handful of trials experience significant inconsistency, likely when random dropout coincides with critical orbital phases. With each measurement independently dropped with probability $p_{\text{loss}} = 0.99$, performance differences become pronounced: the SRUKF exhibits substantial degradation with RMS errors exceeding 3σ bounds for extended periods in many trials, while the GMSRUKF maintains better consistency with fewer divergence events. This suggests that the GMSRUKF mixture-based approach provides enhanced robustness to sparse measurements by maintaining multiple mixands that can independently track different regions of the state space, allowing it to better adapt when

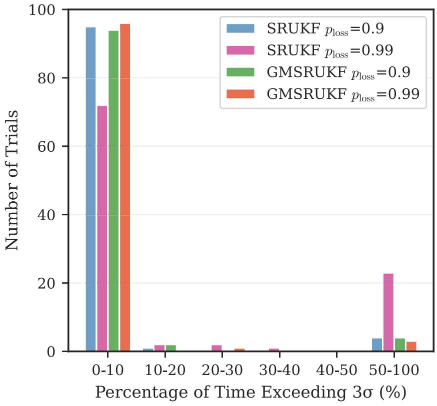

(a) Inconsistency histogram

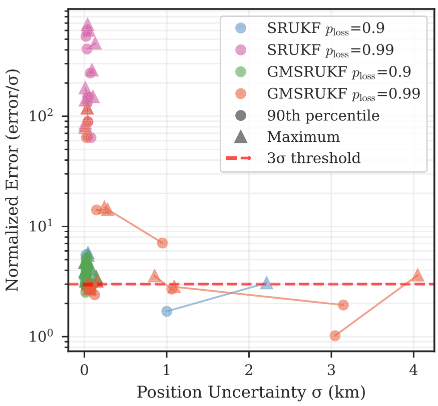

(b) Sigma scatter plot

Fig. 11 **a** Inconsistency histograms for SRUKF and GMSRUKF under varying levels of measurement loss. **b** Normalized position errors for the 10 worst-performing trials under varying levels of measurement loss

observations are infrequent. This is further supported by Fig. 11, where the inconsistency histogram and sigma scatter plot show that the GMSRUKF has fewer extreme outliers (error above 100-σ) compared to the SRUKF at high measurement loss levels.

4.4 Robustness to Observer State Uncertainty

In practical scenarios, the observer spacecraft state is not perfectly known, introducing additional uncertainty into the relative estimation problem. When the observer trajectory is perfectly known, the relative estimation problem simplifies considerably: the dynamics and measurement predictions depend only on the target state uncertainty, making the problem structurally similar to absolute orbit determination. However, when observer state uncertainty is present, both the predicted relative state and its uncertainty must account for errors in both the observer and target trajectories, fundamentally increasing the complexity of the estimation problem. Even when using the consider covariance approach (which avoids full joint estimation of observer and target states), the observer uncertainty contributes additional uncertainty to the relative state estimates through sensitivity to observer errors in both the dynamics and measurement models. This section evaluates filter performance when observer state uncertainty is incorporated via the consider covariance (CC) technique described in Sect. 3.3.

The observer state uncertainty is modeled with a constant covariance $\boldsymbol{P}_{\text{obs}}$ scaled as a fraction of the nominal initial uncertainty, which is a simplifying assumption that acts as a floor on the observer uncertainty, representative of a navigation solution with finite accuracy. In practice, $\boldsymbol{P}_{\text{obs}}$ could vary over time as the observer's onboard navigation solution evolves, and the CC formulation accommodates this naturally. Importantly, even under a constant $\boldsymbol{P}_{\text{obs}}$, the influence on the filter is not constant. The sensitivity matrices $\boldsymbol{G}_f$ and $\boldsymbol{G}_h$ vary along the orbit, so the CC terms $\boldsymbol{G}_f \boldsymbol{P}_{\text{obs}} \boldsymbol{G}_f^\top$

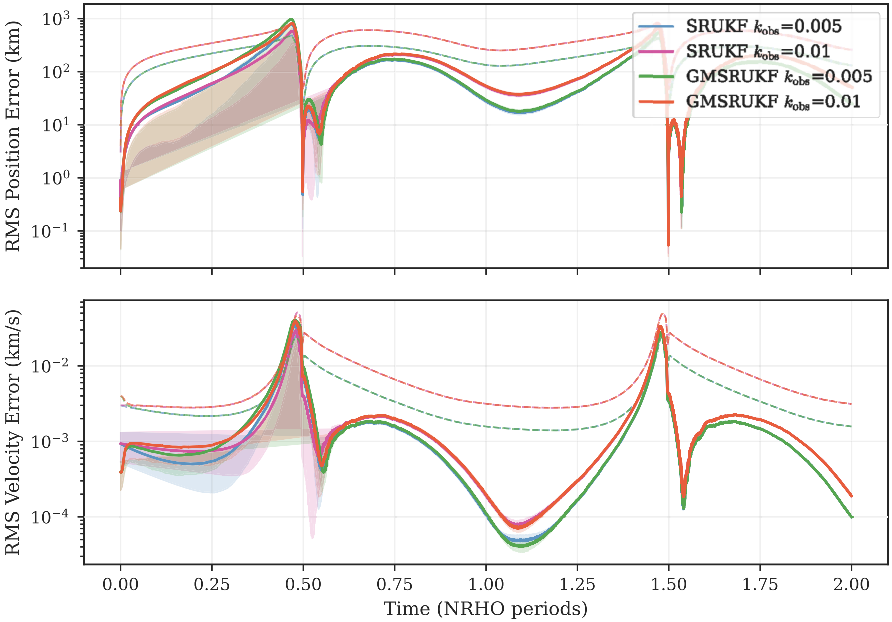

Fig. 12 RMS errors for the filters under varying levels of observer state uncertainty

and $\boldsymbol{G}_h \boldsymbol{P}_{\text{obs}} \boldsymbol{G}_h^\top$ capture how observer errors project into the relative dynamics and measurements differently at each point along the trajectory.

Specifically, two levels of observer uncertainty are tested: $k_{\text{obs}} \in \{0.005, 0.01\}$, such that $\boldsymbol{P}_{\text{obs}} = k_{\text{obs}}^2 \boldsymbol{P}_0$, giving standard deviations of 5 m and 10 m in position and 5 mm s^{-1} and 10 mm s^{-1} in velocity, respectively. These uncertainty levels span a range suitable for evaluating the sensitivity of the consider covariance approach to observer state errors.

To simulate observer trajectory deviations, additive white Gaussian noise is injected into the observer state that is input to the filter at each time step, with samples drawn from $\mathcal{N}(\mathbf{0}, \boldsymbol{P}_{\text{obs}})$. The filter is informed of the observer uncertainty magnitude through the CC formulation, allowing it to account for this uncertainty during prediction and update steps without joint estimation.

Figure 12 reveals a significant and interesting difference between the perfect observer baseline (Fig. 3) and scenarios with observer uncertainty. Both position and velocity RMS errors increase by approximately three orders of magnitude compared to the nominal case. More significantly, the temporal structure of the error evolution changes fundamentally. The perfect observer case exhibits sharp error spikes at perilune passages (0.5, 1.5 periods) with relatively uniform, low errors between these events. In contrast, the observer uncertainty cases show a distinct dome-shaped pattern: errors peak sharply at perilune, then exhibit a secondary maximum around 0.7 and 1.7 periods before decreasing to a minimum around 1.1 and 2.1 periods (post-apolune), and subsequently rising again toward the next perilune passage.

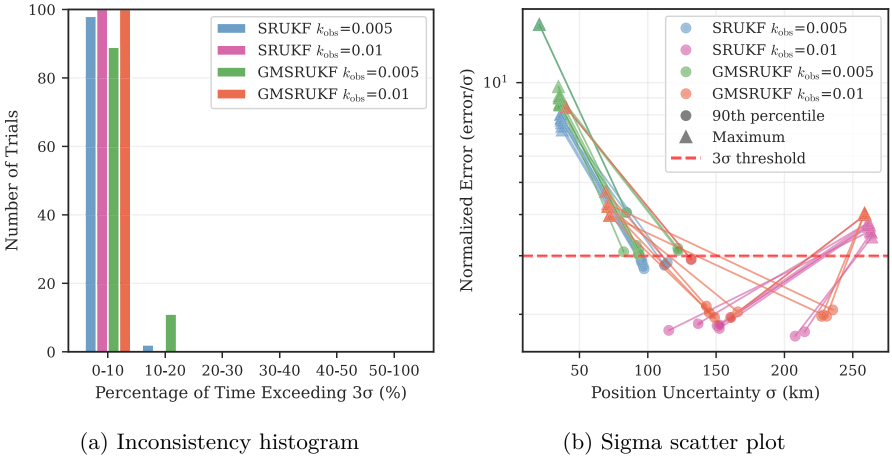

Fig. 13 **a** Inconsistency histograms for SRUKF and GMSRUKF under varying levels of observer state uncertainty. **b** Normalized position errors for the 10 worst-performing trials under varying levels of observer state uncertainty

This behavior indicates that sensitivity to observer state errors depends on the complex interplay between measurement geometry, orbital dynamics, and the linearization assumptions underlying the consider covariance approach. The quasi-steady measurement geometry during the post-apolune phase appears more favorable for separating observer errors from target state uncertainty than the rapidly evolving geometry in the near-perilune phase, despite the reduced information content at larger distances from the Moon. This orbit-phase dependence highlights a key challenge for operational deployment: the fidelity of the CC linearization varies substantially around the orbit, and is an important consideration for operational implementation of CC-based approaches in cislunar relative navigation.

Both SRUKF and GMSRUKF maintain bounded estimates with appropriately widened uncertainty bounds compared to the perfect observer baseline, demonstrating that the CC mechanism successfully captures the propagation of observer uncertainty. However, performance is nearly identical between the two filters in these scenarios. This is expected: observer uncertainty remains approximately Gaussian, so the Gaussian Mixture extension provides no representational advantage. The primary benefits of GMSRUKF (demonstrated in Sections Robustness to Initial Uncertainty and Robustness to Measurement Dropout) arise from handling non-Gaussian target state distributions.

Figure 13a quantifies the consistency across 100 Monte Carlo trials. Despite the elevated error magnitudes, both filters remain under 3-sigma bounds for more than 80% of the simulation duration in all trials, with small inconsistency regions observed near perilune passages in some cases. The maximum normalized errors for worst-case trials stay within 2–10 sigma, as shown in Fig. 13b. This indicates that the CC approach provides appropriately conservative uncertainty estimates that reflect the true difficulty of the estimation problem under observer state uncertainty.

The CC approach is fundamentally a linear approximation of how observer state uncertainty influences the relative state and measurement predictions through the sensitivity Jacobians $\boldsymbol{G}_f$ and $\boldsymbol{G}_h$ In the highly nonlinear CR3BP environment with bearing-angle measurements, the validity of this linearization varies substantially across the orbit. Preliminary investigations revealed that performance exhibits sensitivity to the relative magnitudes of observer uncertainty and linearization errors in the CC Jacobians. Comprehensive characterization of optimal tuning strategies across wider ranges of observer uncertainty, orbit geometries, and measurement cadences is beyond the scope of this paper but represents an important direction for future work. These results demonstrate that the square root formulation naturally accommodates the CC extension and maintains filter consistency when properly configured, extending applicability to realistic scenarios where observer navigation errors are non-negligible.

5 Discussion

The results presented in the previous section tell a consistent story about the role of uncertainty representation in cislunar relative navigation. Under nominal conditions (accurate initialization, full measurement availability, perfectly known observer state) the single-Gaussian SRUKF performs well, and one might conclude that sigma-point filtering alone is sufficient for this application. However, the degraded-condition scenarios reveal that this nominal performance is misleading. As initial uncertainty grows, as bias is introduced, or as measurements become sparse, the SRUKF's Gaussian assumption becomes increasingly inadequate for representing the true posterior distribution, leading to overconfidence, covariance collapse, and eventual divergence.

The GMSRUKF addresses this by maintaining multiple independent mixture components that collectively provide a richer representation of the state distribution. By sampling mixands across the prior distribution, the filter can better capture the true state even when the initial guess is poor or misaligned with the actual target state. Similarly, under sparse measurements, the independent evolution of mixands provides a form of implicit parallel exploration of the state space that helps prevent premature convergence to incorrect estimates. The measurement likelihood weighting adjusts each mixand's influence based on consistency with the observed data, down-weighting components that diverge from the measurement sequence.

These robustness improvements arise from the mixture structure itself rather than from algorithmic sophistication. The fixed-mixand implementation, with no splitting, merging, or pruning, directly isolates the contribution of mixture-based representation, and performance under nominal conditions is insensitive to the number of mixands (Sect. 4.1), confirming no penalty under favorable conditions. However, the computational cost scales linearly with M, amounting to roughly M times the cost of the SRUKF.

The consider covariance results revealed strong orbit-phase-dependent sensitivity to observer state errors. Unlike the other degraded scenarios, which show perilune-concentrated error spikes with otherwise low errors, the CC cases exhibit a distinct dome-shaped temporal pattern that reflects how the influence of observer state uncertainty on the relative dynamics and measurements varies continuously around the orbit, rather than concentrating at a single critical event. Under observer uncertainty, both filters performed nearly identically, as expected, since observer errors remain approximately Gaussian under the constant covariance model, providing no opportunity for the mixture-based representation to offer advantage. The robustness benefits of GMSRUKF arise specifically from capturing non-Gaussian target distributions that emerge from nonlinear propagation and ambiguous initialization.

The CR3BP dynamics are appropriate for evaluating the filtering architecture, as they capture the nonlinear features that make the problem challenging. Perturbations omitted by the model (e.g. SRP, higher-order lunar gravity) act on both spacecraft similarly, so their differential contribution to relative motion is considerably smaller than in the absolute case. The absence of process noise isolates the filter's handling of nonlinear uncertainty propagation. In practice, some level would be needed to account for residual unmodeled effects. The goal of this work is not to produce a mission-ready filter but to establish the viability of the Gaussian Mixture square root framework for cislunar relative navigation and identify where its advantages are most consequential.

Several specific limitations and directions for future work merit discussion.

- **Extreme initialization errors.** Under biased initialization with large uncertainty ($\alpha \geq 10$), both filters struggle to recover, indicating that even the mixture representation might not be able to compensate for an initial guess that is simultaneously imprecise and systematically offset. Some level of coarse initialization, e.g., from ground-based tracking or cooperative ranging, remains necessary for autonomous angles-only navigation.
- **No adaptive mechanisms.** The GMSRUKF uses a fixed number of mixands with no splitting, merging, or pruning. While this simplicity isolates the contribution of the mixture structure, adaptive mechanisms could improve both efficiency and representational fidelity, concentrating expressivity where it is actually needed [26–29].
- **Limited geometry testing.** This study focused on a single NRHO lead-follower configuration. Other relative orbit geometries, periodic orbit families, and formations with time-varying line-of-sight characteristics may exhibit different filter performance profiles and warrant systematic evaluation.
- **Pure uncertainty propagation.** The filters use no process noise, relying on the nonlinear dynamics and measurement model and consider covariance terms (when applicable) to propagate uncertainty. While this avoids artificial tuning parameters, it assumes perfect CR3BP dynamics. In practice, some process noise would be needed to account for the relative effect of unmodeled perturbations, such as solar radiation pressure or higher-order gravitational effects acting differently on each spacecraft.

6 Conclusion

This study evaluates Gaussian Mixture Square Root Unscented Kalman Filters for angles-only relative orbit determination in cislunar space. The baseline comparison shows that out-of-the-box linearization-based filters (EKF) become inconsistent for this problem, and that while the single-Gaussian SRUKF provides a capable baseline under nominal conditions, it exhibits systematic failures under degraded conditions likely in operational cislunar scenarios.

A minimal GMSRUKF implementation with five randomly initialized mixands, fixed throughout the simulation, with no adaptive mechanisms, demonstrates substantial improvements across all degraded scenarios. Under large initial uncertainty ($\alpha = 10$) and no bias, GMSRUKF maintained good consistency in 98% of trials compared to 32% for SRUKF. Under 99% measurement dropout ($p_{\text{loss}} = 0.99$), GMSRUKF failure rates remained below 5% versus 30% for SRUKF. These gains arise from the mixture representation's ability to capture non-Gaussian uncertainty distributions and self-correct via likelihood-based reweighting, at $M\times$ computational overhead.

Observer state uncertainty is incorporated via consider covariance, extending the framework to scenarios where observer navigation errors are non-negligible. The sensitivity Jacobians vary along the orbit, producing state-dependent covariance inflation that reflects the geometric relationship between observer uncertainty and relative state estimation. Because CC relies on a linear approximation of how observer state uncertainty influences the relative dynamics and measurements, its fidelity is tied to how well that linearization holds at each point in the orbit, a dependence that becomes visible in the error pattern observed in the results.

Taken together, these results show that a minimal Gaussian Mixture extension of the SRUKF substantially expands the operating envelope of cislunar angles-only relative navigation without requiring algorithmic complexity or mission-specific tuning. Its modular structure makes it a natural starting point for more capable architectures: adaptive mixture management, higher-fidelity dynamics, and tighter integration with onboard navigation solutions can each be incorporated independently, without redesigning the core filter.

Appendix 1 Sensitivity Matrix Expressions

This appendix provides explicit analytic expressions for the cross-sensitivity matrices $\boldsymbol{G}_f \in \mathbb{R}^{6\times 6}$ and $\boldsymbol{G}_h \in \mathbb{R}^{2\times 6}$ used in the consider covariance implementation of Sect. 3.3. All quantities are expressed in the non-dimensional synodic frame $\mathcal{R}$ and evaluated at the observer (chief) state $\boldsymbol{x}_{\text{obs}} = [\boldsymbol{r}^\top, \dot{\boldsymbol{r}}^\top]^\top$, $\boldsymbol{r} = (x, y, z)^\top$.

Appendix 1.1 Dynamics Cross-Sensitivity $\boldsymbol{G}_f$

The dynamics cross-sensitivity is the negated Jacobian of the CR3BP equations of motion evaluated at the observer state, $\boldsymbol{G}_f = -\boldsymbol{A}(\boldsymbol{x}_{\text{obs}})$. The Jacobian has the 6×6 block form

$$\boldsymbol{A}(\boldsymbol{x}_{\text{obs}}) = \begin{bmatrix} \boldsymbol{0}_{3\times 3} & \boldsymbol{I}_{3\times 3} \\ \boldsymbol{A}_{rr} & \boldsymbol{\Omega} \end{bmatrix}, \tag{10}$$

Table 3 Explicit entries of $\boldsymbol{A}_{rr}$ (symmetric: $[\boldsymbol{A}_{rr}]_{ij} = [\boldsymbol{A}_{rr}]_{ji}$). Here $c_i/r_i^2 = (1-\mu)/r_1^5$ for $i = 1$ and μ/r_2^5 for $i = 2$

Entry	Expression
$[\boldsymbol{A}_{rr}]_{11}$	$1 - c_1 - c_2 + \dfrac{3c_1(x+\mu)^2}{r_1^2} + \dfrac{3c_2(x-1+\mu)^2}{r_2^2}$
$[\boldsymbol{A}_{rr}]_{22}$	$1 - c_1 - c_2 + \dfrac{3c_1\, y^2}{r_1^2} + \dfrac{3c_2\, y^2}{r_2^2}$
$[\boldsymbol{A}_{rr}]_{33}$	$-(c_1 + c_2) + \dfrac{3c_1\, z^2}{r_1^2} + \dfrac{3c_2\, z^2}{r_2^2}$
$[\boldsymbol{A}_{rr}]_{12}$	$\dfrac{3c_1(x+\mu)\, y}{r_1^2} + \dfrac{3c_2(x-1+\mu)\, y}{r_2^2}$
$[\boldsymbol{A}_{rr}]_{13}$	$\dfrac{3c_1(x+\mu)\, z}{r_1^2} + \dfrac{3c_2(x-1+\mu)\, z}{r_2^2}$
$[\boldsymbol{A}_{rr}]_{23}$	$\dfrac{3c_1\, yz}{r_1^2} + \dfrac{3c_2\, yz}{r_2^2}$

where the Coriolis coupling block is

$$\boldsymbol{\Omega} = \begin{bmatrix} 0 & 2 & 0 \\ -2 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \tag{11}$$

and $\boldsymbol{A}_{rr}$ is the symmetric gravity-gradient (position–acceleration) Jacobian. Define the displacement vectors from each primary to the observer,

$$\boldsymbol{d}_1 = \begin{pmatrix} x+\mu \\ y \\ z \end{pmatrix}, \qquad \boldsymbol{d}_2 = \begin{pmatrix} x-1+\mu \\ y \\ z \end{pmatrix}, \tag{12}$$

with magnitudes $r_1 = \|\boldsymbol{d}_1\|$, $r_2 = \|\boldsymbol{d}_2\|$, and scalar gravity coefficients $c_1 = (1-\mu)/r_1^3$, $c_2 = \mu/r_2^3$. The position–acceleration Jacobian then admits the compact outer-product decomposition

$$\boldsymbol{A}_{rr} = \underbrace{\mathrm{diag}(1 - c_1 - c_2,\ 1 - c_1 - c_2,\ -(c_1 + c_2))}_{\text{centrifugal + gravity monopole}} + \frac{3c_1}{r_1^2}\boldsymbol{d}_1\boldsymbol{d}_1^\top + \frac{3c_2}{r_2^2}\boldsymbol{d}_2\boldsymbol{d}_2^\top. \tag{13}$$

Equation (13) is recognized as the Hessian of the CR3BP effective potential $U = \frac{1}{2}(x^2 + y^2) + (1-\mu)/r_1 + \mu/r_2$, confirming the symmetry $\boldsymbol{A}_{rr} = \boldsymbol{A}_{rr}^\top$. The six independent scalar entries are listed explicitly in Table 3 for implementation reference.

Appendix 1.2 Measurement Cross-Sensitivity G_h

The measurement cross-sensitivity $\boldsymbol{G}_h = \partial \boldsymbol{z} / \partial \boldsymbol{x}_{\text{obs}} \in \mathbb{R}^{2\times 6}$ is computed by the chain rule

$$\boldsymbol{G}_h = \underbrace{\frac{\partial \boldsymbol{z}}{\partial\, {}^{\mathcal{O}}\boldsymbol{\rho}}}_{2\times 3} \cdot \underbrace{\frac{\partial\, {}^{\mathcal{O}}\boldsymbol{\rho}}{\partial \boldsymbol{x}_{\text{obs}}}}_{3\times 6}, \tag{14}$$

holding the synodic-frame relative position ${}^{\mathcal{R}}\boldsymbol{\rho}$ fixed. Each factor is derived in turn below.

Factor 1: bearing-angle Jacobian

Writing ${}^{\mathcal{O}}\boldsymbol{\rho} = (\rho_R, \rho_T, \rho_N)^\top$, $\rho = \|{}^{\mathcal{O}}\boldsymbol{\rho}\|$, and $\rho_\perp = \sqrt{\rho_R^2 + \rho_T^2}$, the measurement model of Eq. (3) gives

$$\frac{\partial \boldsymbol{z}}{\partial\, {}^{\mathcal{O}}\boldsymbol{\rho}} = \begin{bmatrix} \dfrac{-\rho_T}{\rho_\perp^2} & \dfrac{\rho_R}{\rho_\perp^2} & 0 \\ \dfrac{-\rho_N\, \rho_R}{\rho^2\, \rho_\perp} & \dfrac{-\rho_N\, \rho_T}{\rho^2\, \rho_\perp} & \dfrac{\rho_\perp}{\rho^2} \end{bmatrix}. \tag{15}$$

Row 1 is the azimuth gradient $(\partial\theta/\partial\, {}^{\mathcal{O}}\boldsymbol{\rho})$ and Row 2 is the elevation gradient $(\partial\phi/\partial\, {}^{\mathcal{O}}\boldsymbol{\rho})$, obtained by differentiating $\theta = \arctan(\rho_T/\rho_R)$ and $\phi = \arcsin(\rho_N/\rho)$, respectively.

Factor 2: frame-orientation Jacobian

Because ${}^{\mathcal{O}}\boldsymbol{\rho} = [OR]\, {}^{\mathcal{R}}\boldsymbol{\rho}$, any change in the observer state rotates the camera frame and therefore reprojects the (fixed) synodic-frame vector into a different direction. Three intermediate vectors must be differentiated. Define the observer position relative to the Moon (P_2), the inertially-resolved velocity of that same vector expressed in $\mathcal{R}$, and the associated angular-momentum direction:

$$\boldsymbol{r}_2 = \boldsymbol{r}_{\text{obs}} - (1-\mu,\ 0,\ 0)^\top, \tag{16}$$

$$\dot{\boldsymbol{r}}_{2,\mathcal{N}} = \dot{\boldsymbol{r}}_{\text{obs}} + \boldsymbol{\omega}_{\mathcal{R}\mathcal{N}} \times \boldsymbol{r}_2, \quad \boldsymbol{\omega}_{\mathcal{R}\mathcal{N}} = (0, 0, 1)^\top, \tag{17}$$

$$\boldsymbol{h}_2 = \boldsymbol{r}_2 \times \dot{\boldsymbol{r}}_{2,\mathcal{N}}. \tag{18}$$

Their Jacobians with respect to $\boldsymbol{x}_{\text{obs}}$ follow directly:

$$\frac{\partial \boldsymbol{r}_2}{\partial \boldsymbol{x}_{\text{obs}}} = \left[\boldsymbol{I}_3\ \boldsymbol{0}_3\right], \tag{19}$$

$$\frac{\partial \dot{\boldsymbol{r}}_{2,\mathcal{N}}}{\partial \boldsymbol{x}_{\text{obs}}} = \left[\boldsymbol{\Delta}_\omega\ \boldsymbol{I}_3\right], \qquad \boldsymbol{\Delta}_\omega = [\boldsymbol{\omega}_{\mathcal{R}\mathcal{N}}]_\times = \begin{bmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \tag{20}$$

$$\frac{\partial \boldsymbol{h}_2}{\partial \boldsymbol{x}_{\text{obs}}} = \left[-[\dot{\boldsymbol{r}}_{2,\mathcal{N}}]_\times + [\boldsymbol{r}_2]_\times\, \boldsymbol{\Delta}_\omega\, [\boldsymbol{r}_2]_\times\right], \tag{21}$$

where $[\boldsymbol{v}]_\times$ denotes the 3×3 skew-symmetric matrix of $\boldsymbol{v}$ (i.e. $[\boldsymbol{v}]_\times \boldsymbol{u} = \boldsymbol{v} \times \boldsymbol{u}$). Equation (21) uses the identity $\partial(\boldsymbol{a} \times \boldsymbol{b})/\partial \boldsymbol{q} = -[\boldsymbol{b}]_\times(\partial \boldsymbol{a}/\partial \boldsymbol{q}) + [\boldsymbol{a}]_\times(\partial \boldsymbol{b}/\partial \boldsymbol{q})$.

The unit-vector Jacobians then follow from the standard projection formula $\partial \hat{\boldsymbol{v}}/\partial \boldsymbol{q} = \boldsymbol{\Pi}_{\hat{v}}/\|\boldsymbol{v}\| \cdot \partial \boldsymbol{v}/\partial \boldsymbol{q}$, where $\boldsymbol{\Pi}_{\hat{v}} = \boldsymbol{I}_3 - \hat{\boldsymbol{v}}\hat{\boldsymbol{v}}^\top$:

$$\frac{\partial \hat{\boldsymbol{r}}_2}{\partial \boldsymbol{x}_{\text{obs}}} = \frac{\boldsymbol{\Pi}_{\hat{r}_2}}{r_2}\left[\boldsymbol{I}_3\ \boldsymbol{0}_3\right], \tag{22}$$

$$\frac{\partial \hat{\boldsymbol{h}}_2}{\partial \boldsymbol{x}_{\text{obs}}} = \frac{\boldsymbol{\Pi}_{\hat{h}_2}}{h_2}\frac{\partial \boldsymbol{h}_2}{\partial \boldsymbol{x}_{\text{obs}}}, \tag{23}$$

$$\frac{\partial \hat{\boldsymbol{\theta}}_2}{\partial \boldsymbol{x}_{\text{obs}}} = -[\hat{\boldsymbol{r}}_2]_\times \frac{\partial \hat{\boldsymbol{h}}_2}{\partial \boldsymbol{x}_{\text{obs}}} + [\hat{\boldsymbol{h}}_2]_\times \frac{\partial \hat{\boldsymbol{r}}_2}{\partial \boldsymbol{x}_{\text{obs}}}, \tag{24}$$

with $r_2 = \|\boldsymbol{r}_2\|$, $h_2 = \|\boldsymbol{h}_2\|$. Equation (24) applies the product rule to $\hat{\boldsymbol{\theta}}_2 = \hat{\boldsymbol{h}}_2 \times \hat{\boldsymbol{r}}_2$.

Assembling Eqs. (22)–(24), the 3×6 frame-orientation Jacobian is

$$\frac{\partial\, {}^{\mathcal{O}}\boldsymbol{\rho}}{\partial \boldsymbol{x}_{\text{obs}}} = \begin{bmatrix} {}^{\mathcal{R}}\boldsymbol{\rho}^\top \left(\dfrac{\partial \hat{\boldsymbol{r}}_2}{\partial \boldsymbol{x}_{\text{obs}}}\right)^\top \\ {}^{\mathcal{R}}\boldsymbol{\rho}^\top \left(\dfrac{\partial \hat{\boldsymbol{\theta}}_2}{\partial \boldsymbol{x}_{\text{obs}}}\right)^\top \\ {}^{\mathcal{R}}\boldsymbol{\rho}^\top \left(\dfrac{\partial \hat{\boldsymbol{h}}_2}{\partial \boldsymbol{x}_{\text{obs}}}\right)^\top \end{bmatrix} \in \mathbb{R}^{3\times 6}, \tag{25}$$

where each row is the gradient of one Hill-frame component of ${}^{\mathcal{O}}\boldsymbol{\rho}$ with respect to the observer state. Substituting Eqs. (15) and (25) into Eq. (14) yields the complete 2×6 matrix $\boldsymbol{G}_h$.

Author Contributions A.L. developed the methodology, performed all simulations and analysis, prepared all figures, and wrote the original draft. J.W.M. and H.S. contributed to conceptualization, provided guidance throughout the research, and reviewed and edited the manuscript. All authors have read and approved the final manuscript.

Data Availability The simulation results presented in this study were generated using custom Python code implementing the filters described in the manuscript. The initial conditions and all filter parameters necessary to reproduce the results are provided in the manuscript. The source code and data files generated during this study are available from the corresponding author upon reasonable request.

Declarations

Conflict of interest The authors declare no conflict of interest.

References

1. Kruger, J., D'Amico, S., Roscoe, C., Westphal, J.: Angles-only tracking and navigation for approach and rendezvous in geosynchronous orbits. In: 33rd AAS/AIAA Space Flight Mechanics Meeting (2023)
2. Ardaens, J.-S., Gaias, G.: Angles-only relative orbit determination in low earth orbit. Adv. Space Res. **61**(11), 2740–2760 (2018). https://doi.org/10.1016/j.asr.2018.03.016
3. Sullivan, J., Lovell, T.A., D'Amico, S.: Angles-only navigation for autonomous on-orbit space situational awareness applications. In: AAS/AIAA Astrodynamics Specialist Conference (2018)
4. LeGrand, K.A., Khilnani, A.V., Iannamorelli, J.L.: Bayesian angles-only cislunar space object tracking. In: 33rd AAS/AIAA Space Flight Mechanics Meeting (2023)
5. Frueh, C., Howell, K., DeMars, K.J., Bhadauria, S.: Cislunar space situational awareness. In: 31st AAS/AIAA Space Flight Mechanics Meeting (2021)
6. Bhatia, S., Shah, K., Kamdar, J.: Review of strategies for cis-lunar space traffic management. In: 73rd International Astronautical Congress (IAC) (2022)
7. Woffinden, D.C., Geller, D.K.: Observability criteria for angles-only navigation. IEEE Trans. Aerosp. Electron. Syst. **45**(3), 1194–1208 (2009). https://doi.org/10.1109/TAES.2009.5259193
8. Kaufman, E., Lovell, T.A., Lee, T.: Nonlinear observability for relative orbit determination with angles-only measurements. J. Astronaut. Sci. **63**(1), 60–80 (2016). https://doi.org/10.1007/s40295-015-0082-9
9. Julier, S.J., Uhlmann, J.K.: Unscented filtering and nonlinear estimation. Proc. IEEE **92**(3), 401–422 (2004). https://doi.org/10.1109/JPROC.2003.823141
10. Sullivan, J., D'Amico, S.: Nonlinear Kalman filtering for improved angles-only navigation using relative orbital elements. J. Guid. Control. Dyn. **40**(9), 2183–2200 (2017). https://doi.org/10.2514/1.G002719
11. Butcher, E.A., Wang, J., Lovell, T.A.: On Kalman filtering and observability in nonlinear sequential relative orbit estimation. J. Guid. Control. Dyn. **40**(9), 2167–2182 (2017). https://doi.org/10.2514/1.G002702
12. Schlenker, L., Sinclair, A.J., Linares, R.: Angles-only orbit determination using Hamiltonian Monte Carlo. In: 28th AAS/AIAA Space Flight Mechanics Meeting (2018). https://doi.org/10.2514/6.2018-1975
13. McCabe, J.S., DeMars, K.J.: Particle filter methods for space object tracking. In: AIAA/AAS Astrodynamics Specialist Conference (2014). https://doi.org/10.2514/6.2014-4308
14. Yun, S., Zanetti, R., Jones, B.A.: Kernel-based ensemble Gaussian mixture filtering for orbit determination with sparse data. Adv. Space Res. **69**(12), 4179–4197 (2022). https://doi.org/10.1016/j.asr.2022.03.013
15. LeGrand, K.A., DeMars, K.J., Pernicka, H.J.: Bearings-only initial relative orbit determination. J. Guid. Control. Dyn. **38**(9), 1699–1713 (2015). https://doi.org/10.2514/1.G001003
16. Psiaki, M.L.: Gaussian-mixture Kalman filter for orbit determination using angles-only data. J. Guid. Control. Dyn. **40**(9), 2341–2347 (2017). https://doi.org/10.2514/1.G002812
17. DeMars, K.J., Jah, M.K.: Probabilistic initial orbit determination using Gaussian mixture models. J. Guid. Control. Dyn. **36**(5), 1324–1335 (2013). https://doi.org/10.2514/1.59844
18. Alspach, D.L., Sorenson, H.W.: Nonlinear Bayesian estimation using Gaussian sum approximations. IEEE Trans. Autom. Control **17**(4), 439–448 (1972). https://doi.org/10.1109/TAC.1972.1100034
19. Van Der Merwe, R., Wan, E.A.: The square-root unscented Kalman filter for state and parameter-estimation. In: 2001 IEEE International Conference on Acoustics, Speech, and Signal Process, Vol. 6, pp. 3461–3464 (2001). https://doi.org/10.1109/ICASSP.2001.940586
20. Geeraert, J.L., McMahon, J.W.: Square-root unscented Schmidt-Kalman filter. J. Guid. Control. Dyn. **41**(1), 280–287 (2018). https://doi.org/10.2514/1.G002921
21. Yang, Y.: Square-root higher-order unscented estimators for robust orbit determination. IEEE Trans. Aerosp. Electron. Syst. **60**(6), 7820–7837 (2024). https://doi.org/10.1109/TAES.2024.3454956

22. Szebehely, V.: Theory of Orbits: The Restricted Problem of Three Bodies. Academic Press, New York (1967). https://doi.org/10.1016/B978-0-12-395732-0.X5001-6
23. Koon, W.S., Lo, M.W., Marsden, J.E., Ross, S.D.: Dynamical systems, the three-body problem and space mission design. In: Fiedler, B., Gröger, K., Sprekels, J. (eds.) Equadiff 99: Proceedings of the International Conference on Differential Equations, pp. 1167–1181. World Scientific, Singapore (2000). https://doi.org/10.1142/9789812792617_0222
24. Roy, A.E.: Orbital Motion, vol. 4. CRC Press, Boca Raton (2020). https://doi.org/10.1201/9780367806620
25. Schutz, B., Tapley, B., Born, G.H.: Statistical Orbit Determination. Elsevier, Amsterdam (2004)
26. DeMars, K.J., Bishop, R.H., Jah, M.K.: Entropy-based approach for uncertainty propagation of nonlinear dynamical systems. J. Guid. Control. Dyn. **36**(4), 1047–1057 (2013). https://doi.org/10.2514/1.58987
27. Kulik, J., LeGrand, K.A.: Nonlinearity and Uncertainty Informed Moment-Matching Gaussian Mixture Splitting (2024). https://doi.org/10.48550/arXiv.2412.00343
28. Zanetti, R., Tuggle, K.: A novel Gaussian mixture approximation for nonlinear estimation. In: 21st International Conference on Information Fusion (FUSION), pp. 1100–1106 (2018). https://doi.org/10.23919/ICIF.2018.8455485
29. Runnalls, A.R.: Kullback-Leibler approach to Gaussian mixture reduction. IEEE Trans. Aerosp. Electron. Syst. **43**(3), 989–999 (2007). https://doi.org/10.1109/TAES.2007.4383588