

GAUSSIAN MIXTURE SQUARE ROOT FILTERS FOR CISLUNAR ANGLES-ONLY RELATIVE ORBIT DETERMINATION

Andrea López^{*}, Jay W. McMahon[†] and Hanspeter Schaub[‡]

Precise and robust relative orbit determination is critical for tasks such as autonomous spacecraft rendezvous, docking, and formation flying. Angles-only techniques, which typically rely on angular measurements from optical sensors, offer a practical alternative in scenarios where traditional range measurements are unavailable or impractical. However, challenges arise due to the nonlinear nature of line-of-sight measurements and complex dynamics in environments like the Earth-Moon three-body system. This work investigates the use of Square Root Unscented Kalman Filters (SRUKF) and their Gaussian Mixture extension (GM-SRUKF) for angles-only relative orbit estimation in Near Rectilinear Halo Orbit regimes. The filters are evaluated across cases with varying initial uncertainty, measurement dropout, and relative orbit geometry.

INTRODUCTION

Accurate relative orbit estimation is a foundational capability for many space operations, including formation flying, autonomous rendezvous and docking, and on-orbit 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 radar and GPS, 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 situational awareness, 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 (nearly aligned with LOS vector) [7, 8]. These difficulties are amplified in

^{*}PhD Student, Ann and H.J. Smead Department of Aerospace Engineering Sciences, University of Colorado, Boulder, Boulder, CO, 80309. AIAA Student Member.

[†]Associate Professor, Department of Aerospace Engineering Sciences, University of Colorado, 429 UCB, Colorado Center for Astrodynamics Research, Boulder, CO 80303

[‡]Distinguished Professor and Department Chair, Schaden Leadership Chair, Ann and H.J. Smead Department of Aerospace Engineering Sciences, University of Colorado, Boulder, 431 UCB, Colorado Center for Astrodynamics Research, Boulder, CO, 80309. AAS Fellow, AIAA Fellow.

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, Figure 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.

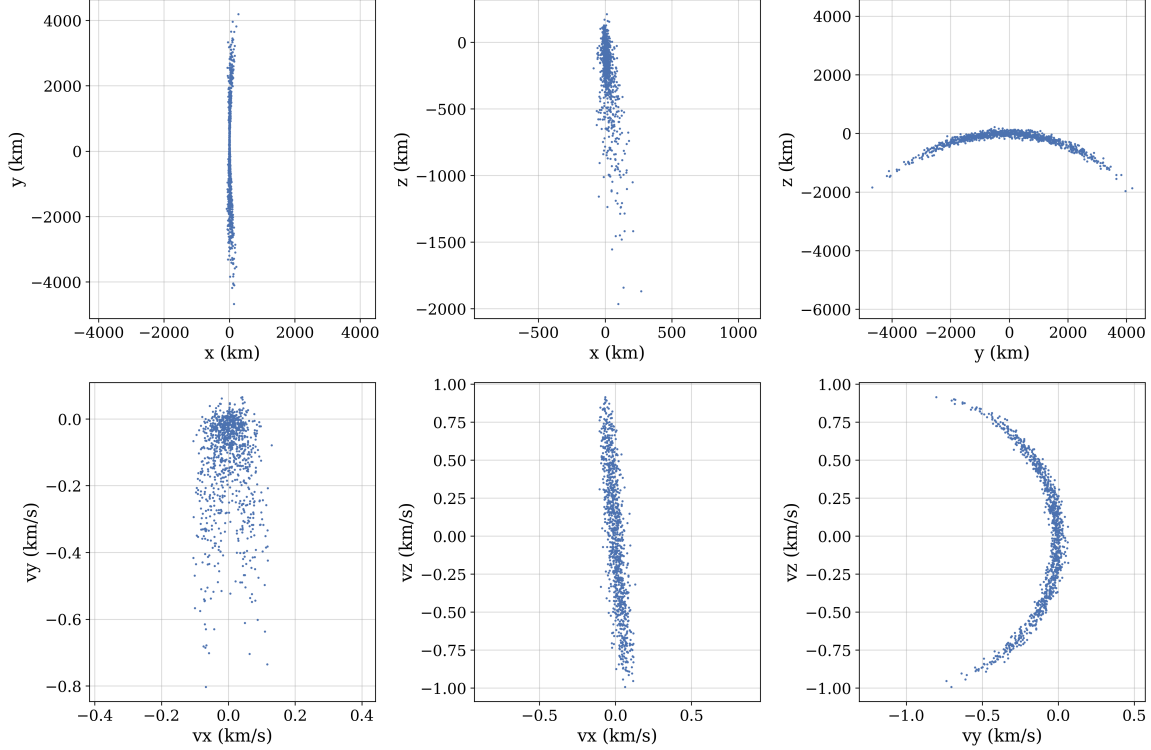


Figure 1: Monte Carlo simulation showing the growth of uncertainty in a Gateway-like Near Rectilinear Halo Orbit (NRHO) after half an orbit. The simulation uses 1000 samples, initially drawn from a Gaussian distribution with a standard deviation of 1 km in position and 1 m/s in velocity. Over time, the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

These same nonlinearities that make uncertainty propagation more challenging can also improve observability when appropriately leveraged. Unlike linearized formulations, nonlinear filtering techniques like 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]. In this sense, the complexity of cislunar dynamics presents not only a challenge but also an opportunity for more informative estimation when nonlinearities are fully utilized.

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 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.

In response to these challenges, recent research has focused on advanced probabilistic estimation techniques that allow for more meaningful 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, 16, 17], which offer robust solutions for nonlinear and stochastic environments. By maintaining multiple hypotheses or samples, these approaches can represent 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].

This paper builds on these ideas by exploring the use of Gaussian Mixture filters based on square root Kalman filter algorithms, specifically the Square Root Unscented Kalman Filter (SRUKF). Square root variants improve numerical stability and covariance accuracy, which are critical for maintaining filter consistency in long-duration or high-uncertainty missions [19, 20]. By combining square root filtering with the expressive power of Gaussian Mixtures, the resulting filters are well-positioned to handle the challenges of angles-only relative navigation in cislunar space, including nonlinear dynamics, sparse measurements, and non-Gaussian uncertainty growth.

In the following sections, we evaluate the performance of these filtering techniques in a series of simulated angles-only relative estimation scenarios. These simulations consider a range of initial uncertainties and measurement availability conditions, with particular attention to estimation performance in NRHO-based and quasi-periodic orbit regimes. The results demonstrate the potential of Gaussian Mixture Square Root filters to support accurate and robust relative navigation in challenging operational environments, laying the groundwork for more autonomous and resilient space mission architectures.

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.

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.

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 $\mathcal{R} : \{\hat{e}_x, \hat{e}_y, \hat{e}_z\}$, centered at the barycenter of the system and rotating with the primaries. The \hat{e}_x vector points from the barycenter to the Moon, the \hat{e}_z vector points in the direction of the E-M angular momentum, and the \hat{e}_y vector completes the right-handed orthogonal

triad. The mass of the system is normalized by the sum of the masses of the Earth and Moon, such that the mass of the Earth is $m_E = 1 - \mu$ and the mass of the Moon is $m_M = \mu$, where $\mu = \frac{m_M}{m_E + m_M} \approx 1.21506 \cdot 10^{-2}$. The unit of distance is the distance between the primaries, and the time unit is the time it takes for the Moon to orbit the Earth. With these considerations, the position of the Earth is fixed at $\mathbf{r}_E = (-\mu, 0, 0)^\top$ and the position of the Moon is fixed at $\mathbf{r}_M = (1 - \mu, 0, 0)^\top$. The equations of motion in the synodic frame \mathcal{R} for a particle with $\mathbf{r} = (x, y, z)^\top$ are given by:

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

where $\mathbf{x} = [\mathbf{r}^\top \dot{\mathbf{r}}^\top]^\top$ is the state vector expressed in the synodic frame, $\dot{\mathbf{r}}$ is the 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:

$$\mathbf{x}_{\text{rel}} = \mathbf{x}_T - \mathbf{x}_O, \quad (2)$$

where \mathbf{x}_T is the absolute state of the target spacecraft and \mathbf{x}_O is the absolute state of the observer spacecraft, both in the rotating (synodic) frame.

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:

$$\mathbf{z} = \mathbf{h}(\mathbf{x}) = \begin{bmatrix} \theta \\ \phi \end{bmatrix} = \begin{bmatrix} \arctan\left(\frac{\rho_\theta}{\rho_r}\right) \\ \arcsin\left(\frac{\rho_h}{\sqrt{\rho_r^2 + \rho_\theta^2 + \rho_h^2}}\right) \end{bmatrix}, \quad (3)$$

where \mathbf{z} is the measurement vector, θ is the azimuth angle, ϕ is the elevation angle.

In this scenario, the camera frame is assumed to be aligned with the observer's modified Hill frame $\mathcal{O} : \{\hat{\mathbf{e}}_r, \hat{\mathbf{e}}_\theta, \hat{\mathbf{e}}_h\}$ that is defined using the state of the chief spacecraft relative to P2 in the synodic frame. This coordinate frame is defined such that $\hat{\mathbf{e}}_r$ points radially outward from the Moon to the observer, $\hat{\mathbf{e}}_h$ points in the $\mathbf{r} \times \dot{\mathbf{r}}$ direction, and $\hat{\mathbf{e}}_\theta$ completes the right-handed orthogonal triad. Thus, the relative position vector $\boldsymbol{\rho} = (\rho_r, \rho_\theta, \rho_h)^\top$ between the observer and target spacecraft is expressed in the modified Hill frame as:

$${}^{\mathcal{O}}\boldsymbol{\rho} = \begin{bmatrix} \rho_r \\ \rho_\theta \\ \rho_h \end{bmatrix} = [OR]{}^{\mathcal{R}}\mathbf{r}_{\text{rel}}, \quad (4)$$

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

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

$$\mathbf{R} = \begin{bmatrix} \sigma_{\theta}^2 & 0 \\ 0 & \sigma_{\phi}^2 \end{bmatrix} \quad (5)$$

with $\sigma_{\theta} = \sigma_{\phi} = 1 \times 10^{-3} \text{rad} \approx 200 \text{arcsec}$. 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, 15, 10].

Relative Orbit Scenarios

This section describes the simulation environment and initialization parameters used to evaluate the performance of the filtering approaches presented in this study. Two scenarios are considered: a baseline lead-follower configuration and a more complex Near Rectilinear Halo Orbit to Quasi-Periodic Orbit (NRHO-QPO) setup, which introduces richer dynamics and line-of-sight variation. The NRHO-QPO results are included only in preliminary form due to ongoing investigation. Key simulation parameters include initial conditions, process and measurement noise assumptions, sensor characteristics, and filter initialization.

Lead-Follower Scenario

The chief spacecraft is initialized on a 9:2 Near Rectilinear Halo Orbit (NRHO) around the Moon, consistent with the planned Gateway orbit. 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 a representative trajectory for cislunar missions with strong dynamical structure and observability variation.

For the lead-follower scenario, the deputy spacecraft is initialized by propagating the chief's 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 over time greatly, despite the small initial offset.

The exact initial conditions are not included in this section due to the high numerical precision required for accurate propagation. However, the NRHO initial state is available in the JPL periodic orbit database, and the target state is generated by propagating this reference forward by one hour. The resulting relative trajectory is shown in Figure 2a.

NRHO-QPO Scenario

The second scenario explores a more nonlinear and dynamically rich regime. The chief spacecraft remains in a Near Rectilinear Halo Orbit (NRHO) around the Moon, representative of Gateway-like missions. The target spacecraft is initialized on a quasi-periodic orbit (QPO) about this NRHO, with the same period as the chief.

This configuration presents a challenging estimation environment due to time-varying line-of-sight geometry, complex orbital motion, and fluctuating observability across the orbit. The goal is to evaluate filter performance in regimes characterized by strong nonlinearities and varying sensitivity to different state directions.

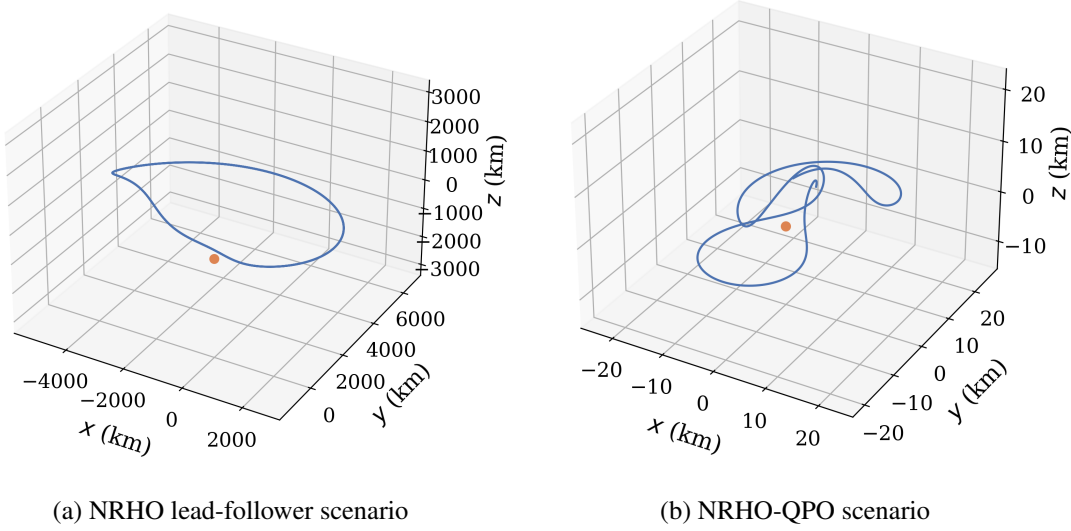


Figure 2: Relative orbits for NRHO lead-follower and NRHO-QPO scenarios propagated for two orbits. The yellow circle marks the position of the observer spacecraft. (a) shows the lead-follower configuration with the chief spacecraft in NRHO and the target spacecraft initialized one hour later. (b) illustrates the more complex NRHO-QPO scenario, where the target follows a quasi-periodic orbit around the chief’s NRHO.

The QPO initial condition is generated by computing the monodromy matrix of the NRHO, identifying the oscillatory mode direction, and applying a small perturbation along that direction. This perturbation creates a quasi-periodic trajectory that oscillates around the NRHO reference, allowing for rich dynamics while maintaining bounded relative motion.

The relative trajectory for the NRHO-QPO scenario is shown in Figure 2b.

FILTERING APPROACHES

This study focuses on the Square Root Unscented Kalman Filter (SRUKF) and its Gaussian Mixture extension (GMSRUKF) for angles-only relative orbit estimation. These filters were selected for their numerical stability, robustness in degraded conditions, and capacity to represent non-Gaussian uncertainty through mixture modeling.

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 [21]. This approach avoids direct computation of covariance square roots during each generation of sigma points, reducing numerical errors and improving robustness, particularly over long-duration propagations or under large uncertainty growth. Further algorithmic details and derivations can be found in [22, 21].

Algorithm 1 Square Root Unscented Kalman Filter (SRUKF)

Require: Prior state \mathbf{x}_{k-1}^+ , Cholesky factor \mathbf{S}_{k-1}^+ such that $\mathbf{P}_{k-1}^+ = \mathbf{S}_{k-1}^+ \mathbf{S}_{k-1}^{+\top}$, measurement \mathbf{z}_k

- 1: {————— **Prediction Step** —————}
- 2: Generate sigma points $\{\chi_{i,k-1}^+\}_{i=0}^{2n}$ using \mathbf{x}_{k-1}^+ , \mathbf{S}_{k-1}^+ , and weights $\{W_i^m, W_i^c\}$
- 3: **for all** i **do**
- 4: Propagate: $\chi_{i,k}^- = \mathbf{f}(\chi_{i,k-1}^+)$
- 5: **end for**
- 6: Compute predicted mean at time k : $\mathbf{x}_k^- = \sum_{i=0}^{2n} W_i^m \chi_{i,k}^-$
- 7: Compute square root of predicted covariance at k :
- 8: Deviations: $\mathbf{d}_0 = \sqrt{|\mathbf{W}_0^c|} (\chi_{0,k}^- - \mathbf{x}_k^-)$, $\mathbf{d}_i = \sqrt{W_i^c} (\chi_{i,k}^- - \mathbf{x}_k^-)$ $i = 1, \dots, 2n$
- 9: Solve: $\mathbf{S}_k^- \leftarrow \text{qr}([\mathbf{d}_1, \dots, \mathbf{d}_{2n} | \mathbf{S}_Q])$
- 10: Apply cholupdate: $\mathbf{S}_k^- \leftarrow \text{cholupdate}(\mathbf{S}_k^-, \mathbf{e}_0, \text{sign}(W_0^c))$
- 11: {————— **Update Step** —————}
- 12: Generate sigma points $\{\chi_{i,k}^-\}_{i=0}^{2n}$ using \mathbf{x}_k^- , \mathbf{S}_k^- , and weights $\{W_i^m, W_i^c\}$
- 13: **for all** i **do**
- 14: Predict measurement: $\mathbf{z}_i = \mathbf{h}(\chi_{i,k}^-)$
- 15: **end for**
- 16: Measurement mean at k : $\mathbf{z}_k^- = \sum_{i=0}^{2n} W_i^m \mathbf{z}_i$
- 17: Compute square root of innovation covariance:
- 18: Deviations: $\mathbf{e}_0 = \sqrt{|\mathbf{W}_0^c|} (\mathbf{z}_0 - \mathbf{z}_k^-)$, $\mathbf{e}_i = \sqrt{W_i^c} (\mathbf{z}_i - \mathbf{z}_k^-)$ $i = 1, \dots, 2n$
- 19: Solve: $\mathbf{S}_z \leftarrow \text{qr}([\mathbf{e}_1, \dots, \mathbf{e}_{2n} | \mathbf{S}_R])$
- 20: Apply cholupdate: $\mathbf{S}_z \leftarrow \text{cholupdate}(\mathbf{S}_z, \mathbf{e}_0, \text{sign}(W_0^c))$
- 21: Compute cross-covariance: $\mathbf{P}_{xz} = \sum_{i=0}^{2n} W_i^c (\chi_{i,k}^- - \mathbf{x}_k^-)(\mathbf{z}_i - \mathbf{z}_k^-)^\top$
- 22: Compute Kalman gain using back-substitution without explicit inversion:
- 23: Solve: $\mathbf{T} = \text{solve_triangular}(\mathbf{S}_z, \mathbf{P}_{xz}^\top, \text{lower} = \text{True})^\top$
- 24: Solve: $\mathbf{K}_k = \text{solve_triangular}(\mathbf{S}_z^\top, \mathbf{T}^\top, \text{lower} = \text{False})^\top$
- 25: Update state at k : $\mathbf{x}_k^+ = \mathbf{x}_k^- + \mathbf{K}_k(\mathbf{z}_k - \mathbf{z}_k^-)$
- 26: Update square root covariance via Cholesky downdate:
- 27: Compute: $\mathbf{U} = \mathbf{K}_k \mathbf{S}_z$
- 28: Update: $\mathbf{S}_k^+ \leftarrow \text{cholupdate}(\mathbf{S}_k^-, \mathbf{U}, -1)$

In this SRUKF implementation, the qr function computes the QR decomposition, and cholupdate performs a numerically stable rank-one Cholesky factor update. The matrices \mathbf{S}_Q and \mathbf{S}_R are the square root of the process noise covariance and measurement noise covariance, respectively. The Kalman gain is computed via two triangular solves corresponding to $\mathbf{P}_{xz}(\mathbf{S}_z \mathbf{S}_z^\top)^{-1}$, avoiding explicit matrix inversion.

Gaussian Mixture Square Root Unscented Kalman Filter (GMSRUKF)

To better represent non-Gaussian, complex uncertainties in angles-only relative navigation, this study implements a Gaussian Mixture extensions of the SRUKF. These filters maintain multiple parallel SRUKF components (called mixands), each representing a portion of the state distribution. The combined state estimate is computed by weighting the contribution of each mixand according to its likelihood and prior weight. Each mixand in the GMSRUKF independently propagates its own square root covariance, retaining the numerical stability and efficiency of the SRUKF. Unlike the non-square root variant, which requires repeated matrix inversions for each mixand during weight

updates, the GMSRUKF leverages Cholesky factors to compute likelihoods more efficiently and stably. Likelihood-based reweighting follows standard Gaussian Mixture filtering procedures. A summary of the GMSRUKF algorithm is provided in Algorithm 2.

Algorithm 2 Gaussian Mixture Square Root Unscented Kalman Filter (GMSRUKF)

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

- 1: {————— **Prediction Step** —————}
- 2: **for all** mixands $j = 1$ to M **do**
- 3: Apply SRUKF Predict Step to $(\mathbf{x}_{k-1}^{+(j)}, \mathbf{S}_{k-1}^{+(j)})$ to obtain $(\mathbf{x}_k^{-(j)}, \mathbf{S}_k^{-(j)})$
- 4: **end for**
- 5: {————— **Update Step** —————}
- 6: **for all** mixands $j = 1$ to M **do**
- 7: Apply SRUKF Update Step to $(\mathbf{x}_k^{-(j)}, \mathbf{S}_k^{-(j)})$ with \mathbf{z}_k to obtain $(\mathbf{x}_k^{+(j)}, \mathbf{S}_k^{+(j)})$
- 8: Compute likelihood $\ell^{(j)} = \mathcal{N}(\mathbf{z}_k; \mathbf{z}_k^{-(j)}, \mathbf{S}_z^{(j)} \mathbf{S}_z^{(j)\top})$
- 9: **end for**
- 10: Normalize weights:
- 11: $w_k^{(j)} \leftarrow \frac{w_{k-1}^{(j)} \ell^{(j)}}{\sum_{i=1}^M w_{k-1}^{(i)} \ell^{(i)}}$
- 12: (Optional) Compute overall mixture mean and covariance:
- 13: $\mathbf{x}_k^+ = \sum_j w_k^{(j)} \mathbf{x}_k^{+(j)}$
- 14: $\mathbf{P}_k^+ = \sum_j w_k^{(j)} \left(\mathbf{S}_k^{+(j)} \mathbf{S}_k^{+(j)\top} + \mathbf{x}_k^{+(j)} \mathbf{x}_k^{+(j)\top} \right) - \mathbf{x}_k^+ \mathbf{x}_k^{+\top}$
- 15: (Optional) Compute mixture square root: $\mathbf{S}_k^+ \leftarrow \text{chol}(\mathbf{P}_k^+)$

For the GMSRUKF, the calculation of the likelihoods is also more efficient than in the non-square root variant, as it avoids repeated matrix inversions. Instead, the inverse of the innovation covariance for each mixand is computed via two triangular solves, similarly to the Kalman gain computation. The square root of the determinant of the innovation covariance is simply the product of the diagonal elements of the Cholesky factor $\mathbf{S}_z^{(j)}$.

Initialization and Assumptions

As a baseline, all filters are initialized with a prior state \mathbf{x}_0 sampled from a multivariate normal distribution $\mathbf{x}_0 \sim \mathcal{N}(\bar{\mathbf{x}}_{0, \text{truth}}, \mathbf{P}_0)$, where $\bar{\mathbf{x}}_{0, \text{truth}}$ is the true initial state and \mathbf{P}_0 is the initial covariance. Unless otherwise noted, \mathbf{P}_0 is diagonal, with standard deviations of 10 km in position and 10 m s⁻¹ in velocity.

To prevent covariance collapse due to nonlinearities and maintain filter consistency, a small process noise term is added to the simulations. Specifically, a scaled identity matrix $\mathbf{Q} = \sigma_Q^2 \mathbf{I}$ with $\sigma_Q = 1 \times 10^{-8}$ is used in the discrete-time propagation model. This process noise does not reflect any specific physical model but instead serves as a regularization term to mitigate overconfidence in the filter estimates, particularly in scenarios with high nonlinearity or limited observability. Similar strategies have been employed in previous nonlinear filtering studies to ensure stable numerical behavior over long-duration propagations [23].

The measurement noise covariance \mathbf{R} that is input to the filters is the same one that is used to generate the synthetic measurements, defined in Equation 5.

Measurements are generated at a fixed cadence of $\Delta t = 60$ s, and the simulation horizon covers two full orbits in the synodic frame (about 13 days). No field-of-view (FOV) constraints are applied in this study, and it is assumed that the observer has continuous visibility of the target throughout the simulation. Future work may incorporate realistic camera pointing models, FOV limits, illumination considerations and occultation effects.

For the Gaussian Mixture filter, $N = 5$ mixands are initialized by sampling from a multivariate Gaussian distribution centered at \mathbf{x}_0 with covariance \mathbf{P}_0 . Each mixand evolves independently according to the SRUKF update rules, and weights are updated based on measurement likelihoods.

Although the filters operate on the relative state, sigma point generation in the prediction step and dynamic propagation are performed in absolute coordinates for numerical stability. Specifically, sigma points are first constructed by adding the current estimate of the chief’s absolute state to the relative state estimate. These absolute sigma points are then propagated through the nonlinear CR3BP dynamics. After propagation, the chief’s absolute state is again subtracted after reconstructing the sigma points to obtain the target state in the relative frame.

This approach prevents numerical issues associated with directly propagating small relative states in a highly nonlinear environment, and ensures consistency when computing Jacobians and measurement predictions. It is particularly important in scenarios like the NRHO-QPO configuration, where relative scales can vary significantly over time.

RESULTS AND DISCUSSION

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 robustness to degraded measurement conditions and uncertain initial knowledge. Results include comparisons between single-model and Gaussian Mixture filter variants. Each scenario is repeated for 50 Monte Carlo runs.

To establish a baseline, Figure 3 compares the behavior of several filter formulations (EKF, UKF, SRUKF, GMUKF, and GMSRUKF) under nominal conditions. This scenario uses the baseline initial uncertainty, no injected bias, and full measurement availability. As expected, the EKF fails across all trials, with estimation errors diverging rapidly due to its inability to capture the system’s nonlinearities, and 3-sigma bounds not capturing the true error distribution. In contrast, the UKF, SRUKF, and both Gaussian Mixture variants perform similarly well, producing consistent estimates that remain within their uncertainty bounds throughout the simulation. The comparable performance of these nonlinear filters under nominal conditions allows the rest of this study to focus on the SRUKF and GMSRUKF as representative single-model and mixture-based approaches, respectively, for more challenging scenarios.

Varying Knowledge of Initial Conditions

In this experiment, we evaluate filter robustness under increasingly uncertain or biased initial knowledge. The prior covariance in these experiments is $\mathbf{P}'_0 = \alpha \mathbf{P}_0$, with $\alpha \in \{1, 10^2, 10^4\}$. The true initial state \mathbf{x}'_0 is then sampled from a Gaussian distribution centered at the nominal initial state \mathbf{x}_0 , with covariance \mathbf{P}'_0 . Additionally, cases with biased initial guesses are tested by shifting the initial guess by 1-sigma such that $\mathbf{x}''_0 \sim \mathcal{N}(\mathbf{x}_0, \mathbf{P}'_0) + \sqrt{\text{diag}(\mathbf{P}'_0)}$, simulating scenarios where the target spacecraft’s initial state is poorly known or misestimated. This setup simulates scenarios where the target spacecraft’s initial state is only coarsely known, which is common in cislunar




Figure 3: Comparison of EKF, UKF, SRUKF, GMUKF, and GMSRUKF under nominal conditions. Solid lines indicate mean RMS errors across 50 Monte Carlo runs, while dashed lines show the mean 3-sigma bounds. The EKF diverges rapidly, while the UKF, SRUKF, and GMSRUKF remain consistent within their uncertainty bounds.

missions where ground-based tracking may be sparse, unavailable, or reserved for higher-priority assets. In such cases, autonomous estimation must contend with large uncertainties or potentially biased initial guesses, particularly for small or secondary spacecraft that may lack direct Earth communication or onboard navigation support.

As shown in Figures 4 and 5, both the SRUKF and GMSRUKF perform well under nominal conditions ($\alpha = 1$), even when a small bias is introduced in the initial estimate. The root mean square (RMS) errors remain within the 3σ consistency bounds for both position and velocity, indicating that the filters maintain reliable and well-calibrated uncertainty estimates when the initialization is accurate.

However, as the prior uncertainty increases to $\alpha = 10^2$, the SRUKF's performance deteriorates rapidly. In both the biased and unbiased cases, the RMS errors exceed the 3σ bounds for extended periods, suggesting that the filter is overconfident and unable to correctly track the target. In contrast, the GMSRUKF remains more consistent. For the unbiased case, its performance stays largely within the bounds, and even in the biased case, it succeeds in roughly half the trials. This suggests that the GMSRUKF is more robust to large initial uncertainty and bias due to its ability to represent multimodal or dispersed belief states.

When the initial uncertainty increases further to $\alpha = 10^4$, both filters fail to produce consistent estimates. In all cases, the RMS errors remain outside the 3σ bounds for most of the simulation,




Figure 4: RMS position and velocity errors over time for SRUKF under varying levels of prior uncertainty and biased initial guesses. Solid lines indicate mean RMS errors across 50 Monte Carlo runs, while dashed lines show the mean 3-sigma bounds across the same runs.

indicating divergence or failure to converge. In practical terms, this means the filters are underestimating their uncertainty: they predict tight confidence bounds while the true state lies well outside them. This breakdown illustrates the limitations of both filters under extreme initialization errors, though the GMSRUKF still shows relatively delayed degradation compared to the SRUKF.

Table 1: Failure Statistics Under Varying Prior Knowledge: Percentage of Runs Exceeding 3σ Bounds in Position / Velocity for 50 Monte Carlo trials

Filter	Threshold	$\alpha = 1$	$\alpha = 1$, biased	$\alpha = 10^2$	$\alpha = 10^2$, biased
SRUKF	> 10% time	8% / 8%	16% / 16%	90% / 90%	100% / 100%
	> 50% time	0% / 0%	0% / 0%	56% / 58%	80% / 82%
GMSRUKF	> 10% time	0% / 0%	6% / 6%	34% / 36%	68% / 68%
	> 50% time	0% / 0%	0% / 0%	6% / 10%	40% / 44%

Table 1 summarizes failure rates for the SRUKF and GMSRUKF under varying levels of initial uncertainty and bias. Failure is defined as exceeding the 3-sigma bound in position or velocity estimates either for more than 10% or 50% of the simulation duration. As expected, both filters perform well under nominal conditions, but their robustness diverges under degraded initialization. SRUKF shows steep degradation, with failure rates climbing to 90-100% even for moderate in-




Figure 5: RMS position and velocity errors for GMSRUKF under varying levels of prior uncertainty and biased initial guesses. Solid lines indicate mean RMS errors across 50 Monte Carlo runs, while dashed lines show the mean 3-sigma bounds.

creases in initial uncertainty, and further worsening when bias is added. In contrast, the GMSRUKF demonstrates significantly better resilience, particularly in the $\alpha = 10^2$ scenarios, where failure rates remain below 40-70% for the 10% threshold and below 10-44% for the 50% threshold. However, both filters break down completely under the most extreme $\alpha = 10^4$ conditions (omitted from table as failure rates are 100% for all cases). These results highlight the GMSRUKF's ability to accommodate greater uncertainty and model mismatch, albeit with practical performance limits.

A closer look at the post-fit measurement residuals, $\mathbf{r}^+\theta$ and $\mathbf{r}^+\phi$, provides further insight into the behavior of the filters. In Figure 6, the SRUKF residuals for the $\alpha = 10^4$ cases clearly diverge from the expected distribution. A similar trend is observed in the GMSRUKF results shown in Figure 7. For the $\alpha = 10^2$ cases, the SRUKF with a biased initial guess exhibits noticeable divergence after the first perilune passage and, to a lesser extent, during the second. The corresponding unbiased case follows the same trends, though the deviations are less pronounced. In contrast, the GMSRUKF shows more consistent behavior across the $\alpha = 10^2$ cases. The unbiased case aligns well with the expected distribution, while the biased case shows some divergence after the first perilune passage, similar to the unbiased SRUKF. The $\alpha = 1$ cases, included for completeness, show well-behaved residuals across both filters and do not warrant further discussion.

Focusing on the $\alpha = 10^2$ biased GMSRUKF, the fact that its residuals remain close to the expected distribution after the first perilune passage suggests that the filter is extracting as much information as possible from the measurements. However, Figure 5 shows that the RMS errors remain

outside the 3-sigma bounds. This may indicate the presence of non-unique solutions to the relative state estimation problem that are consistent with the available measurements under current noise levels, implying that the GMSRUKF cannot fully resolve the underlying ambiguity.

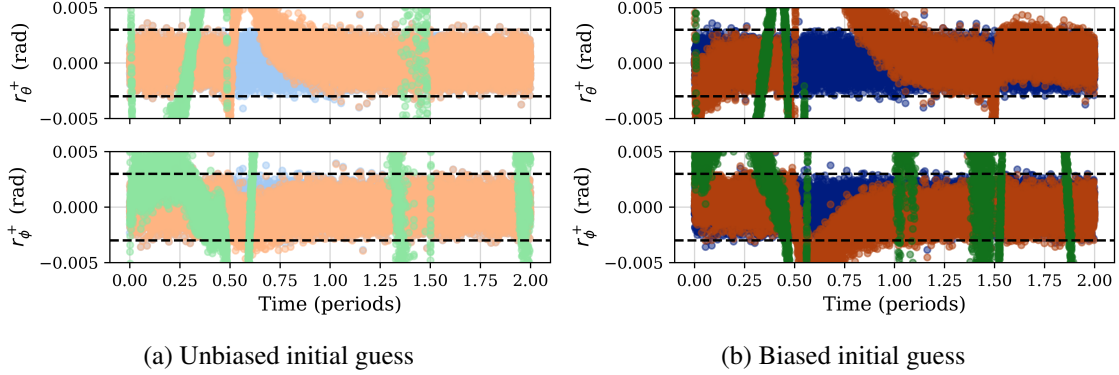


Figure 6: Post-fit measurement residuals for SRUKF with unbiased and biased initial guess, averaged over 50 Monte Carlo runs. The blue color corresponds to $\alpha = 1$, the orange color to $\alpha = 10^2$, and the green color to $\alpha = 10^4$. The dashed lines indicate the expected 3-sigma bounds.




Figure 7: Post-fit measurement residuals for GMSRUKF with unbiased and biased initial guess, averaged over 50 Monte Carlo runs. The blue color corresponds to $\alpha = 1$, the orange color to $\alpha = 10^2$, and the green color to $\alpha = 10^4$. The dashed lines indicate the expected 3-sigma bounds.

Impact of Missing Measurements

This experiment evaluates how each filter handles randomly dropped measurements. Measurement loss percentages of 25%, 50%, 75% and 90% are tested by randomly omitting LOS angle observations at each step with the corresponding probability. The filters are run under nominal initialization (P_0) with no added prior bias or observer uncertainty.

One notable result from this experiment is that both filters exhibit considerable robustness to measurement loss. Even with 50% of LOS observations randomly dropped, both the SRUKF and GMSRUKF are generally able to track the target effectively, maintaining bounded estimation errors. This is notable given that the data loss is unstructured, which is a much more challenging scenario for maintaining observability compared to regular subsampling.




Figure 8: RMS errors for SRUKF under varying levels of randomly dropped measurements. Solid lines indicate mean RMS errors across 50 Monte Carlo runs, while dashed lines show the mean 3-sigma bounds.

As shown in Figure 8, the SRUKF begins to struggle with convergence during the early portion as the loss rate increases beyond 50%, but in many cases recovers as more informative measurements are received (e.g., near perilune). The GMSRUKF, in contrast, maintains better initial stability and converges more quickly even with 75% loss, as shown in Figure 9. This is due to its ability to maintain multiple hypotheses and reweight them based on measurement likelihoods, allowing it to self-correct when measurements are sparse.

Table 2 further quantifies these trends. For the 10% failure threshold, the SRUKF begins to show sensitivity as early as 25% measurement loss, with failure rates increasing steadily through 75%. In contrast, the GMSRUKF maintains near-zero failure up to 50% loss and shows a much more gradual degradation thereafter. This reflects the GMSRUKF’s stronger ability to stay within uncertainty bounds even under sparse observations. For the 50% threshold, the two filters show more comparable performance, with both degrading noticeably at 90% loss. However, the GMSRUKF retains a slight advantage, especially in its ability to maintain consistency during the early orbit phases.

Impact of relative orbit geometry

To explore how different relative motion geometries might affect estimation performance, a preliminary analysis was conducted using a quasi-periodic orbit (QPO) configuration, where the target spacecraft oscillates around a Near Rectilinear Halo Orbit (NRHO) reference, while the observer




Figure 9: RMS errors for GMSRUKF under varying levels of randomly dropped measurements. Solid lines indicate mean RMS errors across 50 Monte Carlo runs, while dashed lines show the mean 3-sigma bounds.

spacecraft is in that reference orbit. This setup presents a particularly challenging estimation environment due to complex relative dynamics and potential observability limitations. Although only a single geometry is tested here, it serves as a first look into how relative orbit configuration might influence filter behavior.

In this scenario, all filters are initialized with an unbiased prior drawn from a multivariate normal distribution with a prior covariance with standard deviations of 1 km in position and 1 m s^{-1} in velocity. Process noise is omitted to isolate the impact of geometry in a setup that is highly sensitive to initial conditions.

Figure 10 shows RMS errors and 3-sigma bounds for both the SRUKF and GMSRUKF over 50 Monte Carlo trials. Although the error levels for both filters are comparable, the uncertainty bounds for the SRUKF decrease very quickly, indicating that the filter is becoming overconfident in its estimates. In contrast, the GMSRUKF maintains more consistent uncertainty bounds throughout the simulation, suggesting that it is better able to adapt to the challenging dynamics and maintain a more realistic estimate of its uncertainty.

As summarized in Table 3, the SRUKF either succeeds or fails decisively: 64% of the Monte Carlo trials exceed the 3σ error bounds for more than 10%, 50%, and even 90% of the simulation duration. This flat failure profile suggests that in scenarios where the SRUKF fails to converge early on, potentially due to a more challenging initial guess, it does not recover later. In contrast, the GMSRUKF shows a more nuanced performance: while 74% of trials exceed the 3σ bounds for

Table 2: Failure Statistics Under Measurement Loss: Percentage of Runs Exceeding 3σ Bounds in Position / Velocity for 50 Monte Carlo Runs

Filter	Threshold	25% loss	50% loss	75% loss	90% loss
SRUKF	> 10% time	12% / 12%	24% / 24%	36% / 36%	80% / 80%
	> 50% time	0% / 0%	0% / 0%	2% / 2%	58% / 58%
GMSRUKF	> 10% time	0% / 0%	0% / 0%	24% / 24%	60% / 60%
	> 50% time	0% / 0%	0% / 0%	4% / 4%	42% / 42%

Table 3: Failure Statistics for the NRHO-QPO Scenario: Percentage of Runs Exceeding 3σ Bounds in Position / Velocity for 50 Monte Carlo Runs

Filter	Threshold	Failure Rate (Position / Velocity)
SRUKF	> 10% time	64% / 64%
	> 50% time	64% / 64%
	> 90% time	64% / 64%
GMSRUKF	> 10% time	74% / 74%
	> 50% time	34% / 34%
	> 90% time	0% / 0%

more than 10% of the time, only 34% do so for more than 50% of the simulation duration, and none exceed the bounds for more than 90%. This indicates that while the GMSRUKF is also challenged by this geometry, it has a better chance of recovering from early estimation errors compared to the SRUKF. Inspection of the mean RMS error across trials in Figure 10 shows that the GMSRUKF tends to track the true state with error magnitudes comparable to those of the SRUKF, but with uncertainty bounds that more reliably contain those errors over time.

A closer look at individual GMSRUKF runs shows what might be causing this behavior. Figure 11 illustrates the RMS errors for each mixand in a representative Monte Carlo run. The RMS errors for each mixand vary significantly, with some performing much better than others, while the overall mixture estimate is worse than the best-performing mixand. This suggests that averaging across diverse hypotheses may dilute useful information. Figure 12 shows the post-fit measurement residuals for each mixand in the same run. Despite the spread in state estimates, the residuals are nearly identical across all mixands, indicating that multiple distinct state hypotheses are consistent with the available measurements. This is a sign of ambiguity in the estimation problem, similar to what was observed in the high-uncertainty initialization cases, where filters were unable to resolve the state uniquely despite producing reasonable residuals.

These results hint that relative geometry could play a significant role in relative estimation performance, even for mixture-based filters like the GMSRUKF. Further study will be needed to assess whether these observations generalize across other geometries, and whether more adaptive strategies could help mitigate the ambiguity in these settings.




Figure 10: RMS errors for SRUKF and GMSRUKF for the NRHO-QPO scenario. Solid lines indicate mean RMS errors across 50 Monte Carlo runs, while dashed lines show the mean 3-sigma bounds.

CONCLUSION

Traditional filtering approaches like the Extended Kalman Filter (EKF) fall short in cislunar angles-only scenarios due to their reliance on linear approximations. Robust estimation in these nonlinear, uncertainty-rich environments demands more expressive and numerically stable techniques.

This study evaluated the performance of the Gaussian Mixture Square Root Unscented Kalman Filter (GMSRUKF) for angles-only relative orbit determination in cislunar space. Even with a minimal implementation using only five mixands and no adaptive mechanisms, the GMSRUKF demonstrated improved robustness and estimation accuracy over the single-model Square Root UKF (SRUKF) across all tested scenarios.

Together, the experiments with uncertain initial conditions, measurement dropout, and complex NRHO-QPO relative geometries show that the GMSRUKF consistently outperforms the SRUKF in terms of maintaining uncertainty bounds and recovering from poor initialization. While both filters perform well under nominal conditions, the SRUKF quickly becomes overconfident when prior uncertainty or measurement loss increases, whereas the GMSRUKF maintains consistency for a broader range of degraded cases. In the NRHO-QPO scenario, where the line-of-sight geometry varies nonlinearly over time, the GMSRUKF is again better able to manage uncertainty, though it still faces challenges recovering from ambiguous initializations.




Figure 11: RMS errors for GMSRUKF mixands in the NRHO-QPO scenario. Solid lines indicate mean RMS errors of each mixand (and the mixture) for a single Monte Carlo run, while dashed lines show the mean 3-sigma bounds.

Interestingly, both the high-uncertainty initialization and NRHO-QPO experiments hint at deeper observability limitations. In cases where mixands exhibit substantially different estimated states, their post-fit residuals can remain nearly indistinguishable. This suggests the possibility of multiple distinct solutions producing similar measurement profiles at current measurement noise levels, raising questions about the uniqueness of the solution.

Notably, this performance was achieved without any adaptive mechanisms such as mixand splitting or merging, jitter addition, likelihood-based pruning, or resampling. Despite its advantages, this GMSRUKF implementation also carries limitations inherited from its simplicity. In the absence of adaptive mechanisms, mixands that diverge or become inconsistent are retained throughout the simulation, even when they no longer contribute meaningful information. This can result in weight collapse or inefficient use of computational resources, especially in highly nonlinear scenarios where hypotheses separate rapidly. More sophisticated approaches (such as mixand pruning, resampling, or adaptive reconfiguration) could help improve long-term performance and computational efficiency.

Potential extensions of this work include integrating such adaptive mechanisms, as well as addressing more complex mission scenarios involving alternative relative motion scenarios, maneuvering targets, field-of-view constraints, and illumination conditions. These directions would help evaluate the generalizability and operational readiness of Gaussian Mixture filters in real-world applications for autonomous navigation in deep space.




Figure 12: Post-fit measurement residuals for GMSRUKF mixands in the NRHO-QPO scenario, for the same Monte Carlo run as Figure 11. Each color corresponds to a different mixand. The dashed lines indicate the expected 3-sigma bounds.

REFERENCES

- [1] J. Kruger, S. D’Amico, C. Roscoe, and J. Westphal, “Angles-Only Tracking and Navigation for Approach and Rendezvous in Geosynchronous Orbits,” *33rd AAS/AIAA Space Flight Mechanics Meeting*, AAS Paper, 2023, pp. 23–309.
- [2] J.-S. Ardaens and G. Gaias, “Angles-only relative orbit determination in low earth orbit,” *Advances in Space Research*, Vol. 61, No. 11, 2018, pp. 2740–2760.
- [3] J. Sullivan, T. A. Lovell, and S. D’Amico, “Angles-only navigation for autonomous on-orbit space situational awareness applications,” *AAS/AIAA Astrodynamics Specialist Conference*, Snowbird, UT, 2018.
- [4] K. A. LeGrand, A. V. Khilnani, and J. L. Iannamorelli, “Bayesian angles-only cislunar space object tracking,” *33rd AAS/AIAA Space Flight Mechanics Meeting*, 2023.
- [5] C. Frueh, K. Howell, K. J. DeMars, and S. Bhadauria, “Cislunar space situational awareness,” *31st AIAA/AAS Space Flight Mechanics Meeting*, Vol. 144, 2021.
- [6] S. Bhatiaa, K. Shahb, and J. Kamdarc, “Review of strategies for cis-lunar space traffic management,” *73rd International Astronautical Congress (IAC)*, 2022.
- [7] D. C. Woffinden and D. K. Geller, “Observability criteria for angles-only navigation,” *IEEE Transactions on Aerospace and Electronic Systems*, Vol. 45, No. 3, 2009, pp. 1194–1208.
- [8] E. Kaufman, T. A. Lovell, and T. Lee, “Nonlinear Observability for Relative Orbit Determination with Angles-Only Measurements,” *The Journal of the Astronautical Sciences*, Vol. 63, Mar. 2016, pp. 60–80.
- [9] S. J. Julier and J. K. Uhlmann, “Unscented filtering and nonlinear estimation,” *Proceedings of the IEEE*, Vol. 92, No. 3, 2004, pp. 401–422.
- [10] J. Sullivan and S. D’Amico, “Nonlinear Kalman filtering for improved angles-only navigation using relative orbital elements,” *Journal of Guidance, Control, and Dynamics*, Vol. 40, No. 9, 2017, pp. 2183–2200.

- [11] E. A. Butcher, J. Wang, and T. A. Lovell, "On Kalman filtering and observability in nonlinear sequential relative orbit estimation," Journal of Guidance, Control, and Dynamics, Vol. 40, No. 9, 2017, pp. 2167–2182.
- [12] L. Schlenker, A. J. Sinclair, and R. Linares, "Angles-Only Orbit Determination Using Hamiltonian Monte Carlo," 2018 Space Flight Mechanics Meeting, 2018, p. 1975.
- [13] J. S. McCabe and K. J. DeMars, "Particle filter methods for space object tracking," AIAA/AAS Astrodynamics Specialist Conference, 2014, p. 4308.
- [14] 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, 2022, pp. 4179–4197.
- [15] K. A. LeGrand, K. J. DeMars, and H. J. Pernicka, "Bearings-only initial relative orbit determination," Journal of Guidance, Control, and Dynamics, Vol. 38, No. 9, 2015, pp. 1699–1713.
- [16] M. L. Psiaki, "Gaussian-mixture Kalman filter for orbit determination using angles-only data," Journal of Guidance, Control, and Dynamics, Vol. 40, No. 9, 2017, pp. 2341–2347.
- [17] K. J. DeMars and M. K. Jah, "Probabilistic initial orbit determination using Gaussian mixture models," Journal of Guidance, Control, and Dynamics, Vol. 36, No. 5, 2013, pp. 1324–1335.
- [18] D. L. Alspach and H. W. Sorenson, "Nonlinear Bayesian estimation using Gaussian sum approximations," IEEE Transactions on Automatic Control, Vol. 17, No. 4, 1972, pp. 439–448.
- [19] J. L. Geeraert and J. W. McMahon, "Square-Root Unscented Schmidt–Kalman Filter," Journal of Guidance, Control, and Dynamics, Vol. 41, No. 1, 2018, pp. 280–287.
- [20] Y. Yang, "Square-root higher-order unscented estimators for robust orbit determination," IEEE Transactions on Aerospace and Electronic Systems, Vol. 60, No. 6, 2024, pp. 7820–7837.
- [21] B. Schutz, B. Tapley, and G. H. Born, Statistical orbit determination. Elsevier, 2004.
- [22] R. Van Der Merwe and E. A. Wan, "The square-root unscented Kalman filter for state and parameter-estimation," 2001 IEEE International Conference on Acoustics, Speech, and Signal Processing. Proceedings (Cat. No. 01CH37221), Vol. 6, IEEE, 2001, pp. 3461–3464.
- [23] D. Simon, Optimal State Estimation: Kalman, H Infinity, and Nonlinear Approaches. John Wiley & Sons, 2006.