

TRAJECTORY-INFORMED ADAPTIVE MULTI-FIDELITY PROPAGATION IN CISLUNAR SPACE

Cedric Petion, Benjamin L. Reifler, and Brandon A. Jones

The University of Texas at Austin, 2617 Wichita St, Austin, TX 78712, USA

Email: {cpetion, benjamin.reifler, brandon.jones}@utexas.edu

ABSTRACT

As the number of missions to cislunar space increases, the population of space objects in this region is expected to grow, making efficient uncertainty propagation essential for space situational awareness (SSA). This is complicated by the cislunar domain's vastness, chaotic dynamical environment, and limited availability of measurements. This paper presents an adaptive multi-fidelity uncertainty propagation method that dynamically adjusts the included perturbing forces based on position in cislunar space, minimizing computation time while maintaining a prescribed modeling accuracy. The proposed adaptive method is then integrated into a multi-target tracking framework to reduce the computational cost of track prediction without sacrificing accuracy, which is important for managing the growing number of objects in cislunar space. The effectiveness of the approach is demonstrated in simulated test cases relevant to upcoming cislunar missions and SSA applications, resulting in a significant reduction in computational cost compared to a non-adaptive approach while achieving equivalent or superior accuracy.

Keywords: Space situational awareness; cislunar; multi-fidelity; uncertainty propagation; multi-target tracking.

1. INTRODUCTION

The cislunar regime is increasingly receiving attention from governmental and commercial entities, with the growing number of planned missions expected to significantly increase the number of anthropogenic space objects (ASOs) in the region. This proliferation of ASOs, along with cislunar space's size and illumination conditions that contribute to sparse measurements and the chaotic multi-body dynamics that complicate uncertainty propagation, presents a need for advances in cislunar space situational awareness (SSA) [1, 2, 3, 4].

An essential capability within cislunar SSA is the ability to track a large number of ASOs, which may include both operational spacecraft and space debris. Multi-

target tracking (MTT) algorithms provide a framework to achieve this, but incur a computational cost for the prediction step proportional to the cost of propagating a single object's state probability density function (PDF) times the number of hypothesized tracks. Thus, to maintain computational tractability, efficient orbit uncertainty propagation methods must be employed. This paper presents an adaptive method for multi-fidelity orbit propagation in cislunar space. The proposed method is applied to an ensemble Gaussian mixture filter (EnGMF), a particle-based filter that is robust to sparse measurements [5, 6, 7]. The EnGMF is used in a generalized labeled multi-Bernoulli filter (GLMBF), which is a multiple hypothesis multi-target filter based on labeled random finite sets (RFSs) [8, 9, 10].

Multi-fidelity methods were first proposed for orbit uncertainty propagation by [11]. This approach propagates a particle-based representation of a state space PDF using a fast-to-compute but possibly inaccurate low-fidelity model, and then identifies a small subset of "important" samples to be repropagated with an accurate yet computationally expensive high-fidelity model. These low- and high-fidelity particles are then used in a stochastic collocation procedure to produce an approximate high-fidelity solution (designated the multi-fidelity solution) for each point in the full particle ensemble. The multi-fidelity solution will then have accuracy close to that of the high-fidelity model but runtime close to the low-fidelity model.

The accuracy and runtime of the multi-fidelity solution are contingent on the low- and high-fidelity models that are selected. For cislunar applications, the optimal model pair varies throughout cislunar space as the relative contributions of Earth gravity, Lunar gravity, and perturbing forces change throughout different domains. Thus, when propagating a PDF through cislunar space with the multi-fidelity method, the optimal model pair may change over time and space and may not be known a priori.

Our solution is an adaptive approach to model selection that varies the perturbations included in the high-fidelity model as a function of position in cislunar space. We pose the issue as an optimization problem, with the objective being to minimize runtime subject to an upper bound on the expected error in acceleration magnitude. The solution is obtained using a precomputed library of

DISTRIBUTION A: Approved for public release; distribution is unlimited. Public Affairs release approval #AFRL-2025-1167.

Proc. 9th European Conference on Space Debris, Bonn, Germany, 1–4 April 2025, published by the ESA Space Debris Office

Editors: S. Lemmens, T. Flohrer & F. Schmitz, (<http://conference.sdo.esoc.esa.int>, April 2025)

recommended gravity field terms from [12]. This work advances previous research from [13] that used a limited set of three models: one pair of models near the Earth, one near the Moon, and one in all other regimes, as well as the work of [14] that examined non-adaptive cislunar multi-fidelity propagation.

Background on the multi-fidelity method, the EnGMF, and the GLMBF are provided in Sections 2.1, 2.2, and 2.3, respectively. The adaptive multi-fidelity model selection method is outlined in Section 3. Lastly, numerical results for simulated test cases are presented in Section 4, which assesses the propagation accuracy, tracking accuracy, and runtime for cislunar SSA scenarios.

2. BACKGROUND

2.1. Multi-Fidelity Orbit Uncertainty Propagation

The multi-fidelity method employed in this paper was first proposed in [15] and later extended to orbit uncertainty propagation by [11]. The approach used in this paper is bi-fidelity, with a low-fidelity model that is intended to be fast to evaluate but may possess relatively poor accuracy, and a high-fidelity model with greater accuracy at the expense of increased computation time.

Reference [11] suggests using a general perturbations propagator for the low-fidelity model to provide analytic or semi-analytic solutions that are rapid to compute. However, in cislunar space, where the gravities of the Earth and the Moon are of similar magnitudes, general perturbations methods, which typically assume small third-body perturbations (if any at all), often fail to provide sufficiently accurate solutions. Therefore, in this work, the low-fidelity model is a special perturbations propagator that numerically integrates the differential equations for a limited, and thus relatively fast to evaluate, force model. The high-fidelity model is also a special perturbations propagator, but with a force model containing substantially more perturbations that are relatively slow to compute. See Section 3 for an in-depth discussion on the model selection.

Conceptually, the multi-fidelity method works by using a low-fidelity dynamics model to propagate a set of particles (such as Monte Carlo samples or sigma points), and then repropagates a limited subset that are dynamically important with a high-fidelity dynamics model and uses these “important” samples to adjust the low-fidelity samples via stochastic collocation.

The remainder of this section provides an abbreviated explanation of the multi-fidelity method in an orbit propagation context. For a complete derivation and discussion, see [15] and [11]. Begin by letting $\xi \in \mathbb{R}^d$ be a random input vector. In this work, it corresponds to a Monte Carlo sample at the initial time and may also be concatenated with a process noise acceleration vector. Let the

propagated low-fidelity and high-fidelity samples be denoted by $x^L(\xi)$ and $x^H(\xi)$, respectively. Then, for the set of random inputs $\Xi = \{\xi_i\}_{i=1}^m$, define the snapshot matrix of low-fidelity samples

$$\mathbf{X}^L(\Xi) = [x^L(\xi_1) \ \dots \ x^L(\xi_m)] \in \mathbb{R}^{n \times m}, \quad (1)$$

and the subspace

$$\mathbb{X}^L(\Xi) = \text{span}(\mathbf{X}^L(\Xi)). \quad (2)$$

We desire a multi-fidelity approximation provided by a stochastic collocation surrogate

$$x^H(\xi) \approx \hat{x}^H(\xi) = \sum_{l=1}^r c_l(\xi) x^H(\bar{\xi}_l), \quad (3)$$

where c_l are expansion coefficients, $\bar{\xi}_l$ are the important samples, and r is the rank of the surrogate with $r \ll m$.

The important samples are identified iteratively, with each iteration’s important sample being the point that maximizes the distance to the subspace of previously identified important samples. Thus, the k th important sample is

$$\bar{\xi}_k = \arg \max_{\xi \in \Xi} \text{dist}(x^L(\xi), \mathbb{X}^L(\Xi^{k-1})), \quad (4)$$

where

$$\text{dist}(x, \mathbb{X}) \equiv \inf_{y \in \mathbb{X}} \|x - y\|, \quad (5)$$

and

$$\Xi^k = \Xi^{k-1} \cup \{\bar{\xi}_k\}. \quad (6)$$

The maximum number of important samples is limited to the rank of the surrogate, with $r \leq n$. Since orbit state vectors typically have dimension $n = 6$, [11] suggests augmenting $x(\xi)$ with the state trajectory at multiple points in time to increase the rank, thus permitting surrogates with more important samples, which was found to improve the accuracy of the multi-fidelity solution.

To solve the discrete optimization problem in Eq. 4, [15] suggests a greedy approach via the pivoted Cholesky decomposition to produce an approximate solution given by

$$[\mathbf{X}^L]^\top \mathbf{X}^L = \mathbf{A}^\top \mathbf{L} \mathbf{L}^\top \mathbf{A}, \quad (7)$$

where $[\mathbf{X}^L]^\top \mathbf{X}^L$ is the Gramian matrix of low-fidelity samples, \mathbf{L} is lower triangular, and \mathbf{A} is a pivot matrix.

The expansion coefficients are solved such that the basis of low-fidelity important samples may be used to approximately reconstruct the snapshot matrix. The expansion coefficients c_l must then satisfy

$$x^L(\xi) \approx \hat{x}^L(\xi) = \sum_{l=1}^r c_l(\xi) x^L(\bar{\xi}_l), \quad (8)$$

which parallels Eq. 3. Reference [11] proposes an iterative approach to determine r by incrementing it until the reconstructed snapshot matrix produced by Eq. 8 is within some user-specified tolerance of the original snapshot matrix from Eq. 1.

Expressing Eq. 8 in matrix form and combining it with Eq. 7, we obtain

$$LL^T \mathbf{c} = \mathbf{X}(\bar{\Xi})^T \mathbf{X}(\Xi), \quad (9)$$

where $\bar{\Xi} = \{\bar{\xi}_i\}_{i=1}^r$ and \mathbf{c} is a matrix of the coefficients, c_l . The coefficients \mathbf{c} , which may be solved using forward / backward substitution, may then be used in Eq. 3 to obtain the multi-fidelity solution.

2.2. The Ensemble Gaussian Mixture Filter

The EnGMF enables accurate but efficient nonlinear estimation. The filter parameterizes the estimated state PDF as a set of particles and uses kernel density estimation (KDE) to avoid particle depletion, allowing it to function with fewer particles than a particle filter (PF) [5, 6]. The EnGMF is initialized by drawing N random samples from the initial PDF. The PDF is predicted by propagating each particle forward in time. To update the PDF, we first convert the particles to a Gaussian mixture model (GMM) via KDE using Silverman's rule [16, 6]: Each particle becomes the mean of a GMM component with weight N^{-1} and covariance

$$B_S = \left(\frac{4}{N(d+2)} \right)^{\frac{2}{d+4}} P, \quad (10)$$

where P is the particles' sample covariance. The GMM is then updated using an appropriate GMM-based filter and N new particles are sampled from the updated GMM. The KDE GMM is also used to extract the estimated mean and covariance, as opposed to using the sample mean and covariance directly.

In this paper, the GMM update is performed using the square root unscented Kalman filter (SR-UKF) update [17] to enable accurate nonlinear updates with fewer particles. This is intended to reduce the cost of prediction and high-fidelity correction, but does increase the computational cost of the measurement update per particle. The EnGMF generally performs better for orbit determination when its particles' states are parameterized by equinoctial orbital elements, due to improved linearity when applying Silverman's rule [18, 6]. However, because equinoctial elements are designed to represent two-body orbits, here we instead use Cartesian coordinates.

2.3. The Generalized Labeled Multi-Bernoulli Filter

The GLMBF is a labeled RFS-based multi-target filter. An RFS may be thought of as a set of random vectors

whose cardinality is also a random variable. Random finite sets enable Bayesian estimation of the states of unknown numbers of objects. The GLMBF is a closed-form solution to the Bayes multi-target filter recursion [8, 9]. This paper uses the joint predict-update formulation of the GLMBF [10]. This description of the GLMBF omits the parts of the filter that deal with the addition of new objects to the tracked population (spontaneous birth and spawning models), because the simulations in this paper do not use those capabilities.

2.3.1. Notation

In this section, lowercase letters denote vectors, capital letters denote sets, blackboard bold letters represent spaces, and calligraphic letters denote sets of sets. Bold symbols indicate the use of label-augmented vectors. For variables that change over time, the subscript k to indicate time is not written and the subscript $k+1$ to indicate the following timestep is abbreviated to a subscript plus sign. The multi-object exponential is defined as $f^X = \prod_{x \in X} f(x)$, the Kronecker delta is defined as

$$\delta_Y[X] = \begin{cases} 1, & X = Y, \\ 0, & \text{otherwise,} \end{cases} \quad (11)$$

the indicator function is defined as

$$1_Y(X) = \begin{cases} 1, & X \subseteq Y, \\ 0, & \text{otherwise,} \end{cases} \quad (12)$$

and $1_Y(x) = 1_Y(\{x\})$. The function \mathcal{F} denotes the set of all finite subsets of a space. A label-augmented RFS is defined on $\mathbb{X} \times \mathbb{L}$, where \mathbb{X} is the state space and \mathbb{L} is a discrete label space. This means that each element in a realization of a labeled RFS is of the form $\mathbf{x} = (x, l)$, where x is the state and l is the label. In this work, object labels are of the form $l = (i, k, j)$, where i is the index of the sensor that first detected the object, k is the timestep at which the object's initial track was created, and j is the object's unique index out of all objects first detected by sensor i at time k . A labeled RFS realization may not contain duplicate labels, and this is enforced using the distinct label indicator

$$\Delta(\mathbf{X}) = \delta_{|\mathbf{X}|} \left[|\text{lab}(\mathbf{X})| \right], \quad (13)$$

where $\text{lab}(x, l) = l$ denotes the projection of the label-augmented state space $\mathbb{X} \times \mathbb{L}$ onto its discrete label space \mathbb{L} and $\text{lab}(\mathbf{X}) = \{\text{lab}(\mathbf{x}_1), \dots, \text{lab}(\mathbf{x}_n)\}$.

2.3.2. Random Finite Set PDFs

The PDF for a δ -generalized labeled multi-Bernoulli (GLMB) RFS may be parameterized by components $(I, h) \in \mathcal{F}(\mathbb{L}) \times \mathbb{H}$, where \mathbb{H} is a discrete space, and associated weights $w^{(I, h)}$. For MTT, each component

typically represents a data association hypothesis, with I being a set of objects that may exist, h being their combined measurement association history, and $w^{(I,h)}$ being the estimated probability that the hypothesis is true. A δ -GLMB RFS density is of the form

$$\begin{aligned} \pi(\mathbf{X}) &= \Delta(\mathbf{X}) \sum_{(I,h) \in \mathcal{F}(\mathbb{L}) \times \mathbb{H}} w^{(I,h)} \delta_I[\text{lab}(\mathbf{X})] \left(p^{(h)} \right)^{\mathbf{X}}, \end{aligned} \quad (14)$$

where $p^{(h)}(\cdot, l)$ is an object's state-space PDF given association history h and label l .

2.3.3. Joint Predict–Update Recursion

Given the initial filtering density in Eq. 14 at timestep k , the predicted and updated density at timestep $k + 1$ is given by

$$\begin{aligned} \pi_+(\mathbf{X}_+) &\propto \Delta(\mathbf{X}_+) \sum_{I, h, I_+, \theta_+} w^{(I,h)} w_+^{(I,h,I_+, \theta_+)}(Z_+) \\ &\quad \times \delta_{I_+}[\text{lab}(\mathbf{X}_+)] \\ &\quad \times \left(p_+^{(h, \theta_+)}(\cdot | Z_+) \right)^{\mathbf{X}_+}, \end{aligned} \quad (15)$$

where $I \in \mathcal{F}(\mathbb{L})$, $h \in \mathbb{H}$, $I_+ \in \mathcal{F}(\mathbb{L}_+)$, and $\theta_+ \in \Theta_+$, where $\mathbb{L}_+ = \mathbb{L} \cup \mathbb{B}_+$, \mathbb{B}_+ is the space of object labels that could be born at this time, and Θ_+ is the set of maps $\theta_+ : \mathbb{L}_+ \rightarrow \{0 : |Z_+|\}$ assigning measurements in the current measurement set Z_+ to object labels, where $\theta_+(l) = 0$ implies that label l is not assigned a measurement. Additionally,

$$\begin{aligned} w_+^{(I,h,I_+, \theta_+)}(Z_+) &= (r_{B,+})^{\mathbb{B}_+ \cap I_+} (1 - r_{B,+})^{\mathbb{B}_+ - I_+} \left(\bar{p}_S^{(h)} \right)^{I \cap I_+} \\ &\quad \times \left(1 - \bar{p}_S^{(h)} \right)^{I - I_+} \left(\psi_+^{(h, \theta_+(\cdot))}(\cdot | Z_+) \right)^{I_+}, \end{aligned} \quad (16)$$

$$p_+^{(h, \theta_+)}(x_+, l | Z_+) = \frac{\bar{p}_+^{(h)}(x_+, l) \psi_+^{(\theta_+(l))}(x_+, l | Z_+)}{\bar{\psi}_+^{(h, \theta_+(l))}(l | Z_+)}, \quad (17)$$

$$\begin{aligned} \bar{p}_+^{(h)}(x_+, l) &= 1_{\mathbb{B}_+}(l) p_{B,+}(x_+, l) \\ &\quad + 1_{\mathbb{L}}(l) \frac{\left\langle p_S(\cdot, l) f_{S,+}(x_+ | \cdot, l), p^{(h)}(\cdot, l) \right\rangle}{\bar{p}_S^{(h)}(l)}, \end{aligned} \quad (18)$$

$$\bar{p}_S^{(h)}(l) = \left\langle p^{(h)}(\cdot, l), p_S(\cdot, l) \right\rangle, \quad (19)$$

$$\bar{\psi}_+^{(h,j)}(l | Z_+) = \left\langle \bar{p}_+^{(h)}(\cdot, l), \psi_+^{(j)}(\cdot, l | Z_+) \right\rangle, \quad (20)$$

$$\begin{aligned} \psi_+^{(j)}(x_+, l | Z_+) &= \delta_0[j] (1 - p_D(x_+, l)) \\ &\quad + (1 - \delta_0[j]) \frac{p_D(x_+, l) g(z_{+,j} | x_+, l)}{\kappa(z_{+,j})}, \end{aligned} \quad (21)$$

where $r_{B,+}(l)$ is the probability that object l is born, $p_{B,+}(\cdot, l)$ is the single-target PDF for the newborn object, p_S is the probability of survival from one step to the next, p_D is the probability of detection, $f_{S,+}(x_+ | \cdot, l)$ is the surviving object transition density, g is the measurement likelihood, and κ is the expected clutter intensity.

The number of possible hypotheses grows exponentially over time. To maintain computational tractability, the set of new hypotheses (I, h, I_+, θ_+) in Eq. 15 resulting from prior hypothesis (I, h) is truncated using a ranked assignment algorithm and the cost matrix shown in Eq. 22, where $c(l, j) = -\log(\eta^{(I,h)}(l, j))$ and

$$\eta^{(I,h)}(l, j) = \begin{cases} 1 - r_{B,+}^{(l)}, & j < 0 \wedge l \in \mathbb{B}_+, \\ r_{B,+}^{(l)} \bar{\psi}_+^{(h,j)}(l | Z_+), & j \geq 0 \wedge l \in \mathbb{B}_+, \\ 1 - \bar{p}_S^{(h)}(l), & j < 0 \wedge l \in I, \\ \bar{p}_S^{(h)}(l) \bar{\psi}_+^{(h,j)}(l | Z_+), & j \geq 0 \wedge l \in I. \end{cases} \quad (23)$$

If an object is assigned to the left block of the cost matrix, it does not exist, if it is assigned to the center block, it exists but was not detected at time $k + 1$, and if it is assigned to the right block, it exists and produced measurement $z_{+,j}$. The classic approach to solve this ranked assignment problem is to use Murty's algorithm [19, 8, 9], but for large cost matrices, a Gibbs sampler-based approach is more efficient [10].

3. ADAPTIVE MULTI-FIDELITY MODEL SELECTION

The accuracy and computational efficiency of the multi-fidelity solution depend on the choice of low- and high-fidelity models. We seek the optimal model pair that minimizes computation time while satisfying a user-defined requirement for accuracy in the multi-fidelity solution. Accordingly, the low- and high-fidelity force models must incorporate sufficient perturbing forces to maintain the desired accuracy while avoiding extraneous perturbations that would unnecessarily increase computational cost. Thus, for cislunar applications, the optimal model pair will vary throughout cislunar space as the relative contributions of Earth gravity, Lunar gravity, and perturbing forces change as a function of position.

For the filtering applications considered in this work, we would ideally define the accuracy requirement as an upper limit on the PDF divergence between the multi-fidelity and a full-fidelity solution in the prediction step

$$\left[\begin{array}{ccc|ccc|ccc} c(l_1, -1) & \cdots & \infty & c(l_1, 0) & \cdots & \infty & c(l_1, 1) & \cdots & c\left(l_1, \left|Z_+^{(i)}\right|\right) \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ \infty & \cdots & c(l_M, -1) & \infty & \cdots & c(l_M, 0) & c(l_M, 1) & \cdots & c\left(l_M, \left|Z_+^{(i)}\right|\right) \end{array} \right] \quad (22)$$

of the filter, as proposed by [13]. Formulated as an optimization problem, this becomes

$$\begin{aligned} \min_{(f^L, f^H) \in F \times F} \quad & t_R(f^L, f^H) \\ \text{subject to} \quad & \mathcal{D}(p, \hat{p}(f^L, f^H)) < \varepsilon_{\mathcal{D}}, \end{aligned} \quad (24)$$

where f^L and f^H are the low- and high-fidelity models, respectively, that solve the optimization problem and belong to a set of force models $F = \{f_1, f_2, \dots\}$ that vary in which perturbing forces are included. We seek a map

$$m : \mathbb{X} \rightarrow F \times F \quad \text{such that} \quad (f^L, f^H) = m(\hat{x}), \quad (25)$$

where \hat{x} is the mean state estimate at some time t_k . Also, $t_R(f^L, f^H)$ is the runtime of selecting and evaluating the model pair (f^L, f^H) , $\mathcal{D}(\cdot, \cdot)$ is a PDF divergence metric, p is the state-space PDF predicted from time t_k to t_{k+1} using a full-fidelity model, $\hat{p}(f^L, f^H)$ is the PDF predicted using the multi-fidelity algorithm with f^L and f^H , and $\varepsilon_{\mathcal{D}}$ is a user-defined value.

As posed, this optimization problem is intractable to compute online because evaluating $\mathcal{D}(p, \hat{p}(f^L, f^H))$ requires both a full-fidelity solution for p , which is prohibitively expensive to compute, and the multi-fidelity solution $\hat{p}(f^L, f^H)$, which is itself the unknown quantity we seek as the end result, making its direct computation in the optimization problem impractical. We therefore reformulate the optimization problem to have an accuracy requirement that the expected error in acceleration magnitude remain below some threshold. Since evaluating the acceleration at a particular state \hat{x} in the space \mathbb{X} does not require the numerical solution to the differential equations of the force models, this makes the problem formulation tractable. Furthermore, accurate propagation will yield convergence in a mean square sense, which is stronger than convergence in density.

Additionally, we will simplify the optimization problem by considering only one low-fidelity model over the space \mathbb{X} (over the entire cislunar domain). This not only eliminates a decision variable, but also allows us to use the trajectory from the low-fidelity propagation to inform the model choice for the high-fidelity propagation. This is possible within the multi-fidelity framework since the low-fidelity propagation is always done before the high-fidelity propagation, meaning the high-fidelity model does not yet need to be determined when the low-fidelity propagation is performed. Moreover, a truncated force model that is constant across cislunar space may

still account for the dominant forces, and as demonstrated by the numerical simulations in Section 4, yields sufficient multi-fidelity accuracy and runtime, making it justified from a dynamical and computational standpoint.

The reformulated optimization problem is then

$$\begin{aligned} \min_{f^H \in F} \quad & t_R(f^L, f^H) \\ \text{subject to} \quad & a_{\text{error}} < \varepsilon_a, \end{aligned} \quad (26)$$

where a_{error} is the expected error in acceleration magnitude and ε_a is a user-defined value. The high-fidelity force model that solves this optimization problem is the one that includes only the minimal set of perturbations necessary to satisfy the accuracy requirement for a given region of cislunar space, as incorporating additional perturbations generally increases runtime. Since the high-fidelity model is selected as a function of position in cislunar space, only perturbations that vary spatially need to be adjusted – specifically, the non-spherical gravitational perturbations of the Earth and Moon. As a space object approaches one of these bodies, its non-spherical gravity terms become more significant and should be included in the model. Conversely, as the object moves farther away, those terms diminish in influence and can be truncated.

Among orbital perturbations, non-spherical gravity is the primary contributor to computational cost, as evaluating high-degree spherical harmonic expansions at each integration step is computationally expensive. In contrast, perturbations such as cannonball solar radiation pressure (SRP) and the Sun’s point-mass gravity remain relatively constant across cislunar space and impose negligible computational overhead compared to spherical harmonic evaluations. Therefore, these perturbations are kept constant in both the low- and high-fidelity models. Additionally, Earth atmospheric drag is not considered, as trajectories that encounter significant atmospheric effects are better suited to an Earth-centric multi-fidelity framework rather than a cislunar one.

Since non-spherical gravity is the only perturbation that varies in the high-fidelity model, we focus on determining the appropriate spherical harmonic expansion degree and order as a function of position in space and acceptable acceleration error. The solution is obtained separately for each celestial body’s gravity using a lookup table from [12], which provides the appropriate expansion degree and order for a given altitude to a celestial body and user-defined acceleration noise, thereby defining the mapping in Eq. 25.

For a spherical harmonic expansion of degree and order L , [12] defines the total acceleration error as

$$\mathbb{E}[\|\nabla U_{T,L}\|^2] = \mathbb{E}[\|\nabla U_{O,L}\|^2] + \mathbb{E}[\|\nabla U_{C,L}\|^2], \quad (27)$$

where $\mathbb{E}[\|\nabla U_{O,L}\|^2]$ the omission error, $\mathbb{E}[\|\nabla U_{C,L}\|^2]$ is the commission error, and U is the gravitational potential. The omission error represents the expected acceleration error due to truncation of spherical harmonic terms beyond degree and order L , while the commission error accounts for the expected acceleration error due to published uncertainties in the included spherical harmonic terms up to and including degree and order L . The mean omission and commission errors are computed by averaging the squared acceleration errors in Cartesian coordinates for the celestial body in question.

The x -, y -, and z -components of error are assumed to be Gaussian distributed, meaning the squared error in Eq. 27 follows a chi-squared distribution. Reference [12] then computes the upper bound of a 99.7% confidence interval for all expansion degrees to produce the lookup table. Our adaptive high-fidelity model selection method uses the low-fidelity trajectory to determine the space object's altitude relative to the Earth or Moon. It then queries the appropriate lookup table to select the expansion degree and order to use in the high-fidelity model for the current position in cislunar space.

A limitation of the present implementation is that it considers non-spherical gravity from only one celestial body at a time – either the Earth or Moon. If the space object is within a predefined radius of the Moon, only the Moon's non-spherical gravity is included; otherwise, the Earth's is used. Incorporating non-spherical gravity from both bodies simultaneously is an area for future work and could be implemented by querying both lookup tables, provided that the propagator used can accommodate non-spherical gravity from multiple bodies.

In regions of cislunar space far from either body, the recommended expansion degree and order for both the Earth and Moon may approach zero. When this occurs, the low- and high-fidelity models become sufficiently similar that no multi-fidelity correction is needed. Therefore, at a specified threshold for the high-fidelity model's expansion degree, our implementation performs only a low-fidelity propagation and skips the multi-fidelity procedure.

To summarize the sequence of steps for the adaptive multi-fidelity method:

1. A low-fidelity propagation of all particles is performed from the initial time to the final time.
2. The altitude and corresponding gravity expansion degree and order are computed at discrete intervals along the low-fidelity trajectories.
3. The high-fidelity propagation is then performed, with the gravity expansion changing at each interval used in the previous step.

4. The multi-fidelity solution is obtained at the final time.

To obtain the time steps and associated state vectors used in the multi-fidelity algorithm's snapshot matrix (Eq. 1), we propagate the particles backward in time N time steps rather than store their histories to reduce memory footprint, as was done in [13]. This is beneficial in the context of MTT since storing the time histories for multiple different tracks can become overly memory intensive. For applications where speed is desired over memory, the time histories could instead be stored from the initial forward propagation, rather than doing the backward propagation.

Lastly, the number of time steps serves as a user-defined tuning parameter. Since the rank of the snapshot matrix limits the number of important samples that can be identified, increasing the number of time steps generally raises the matrix rank, enabling the identification of more important samples and enhancing the accuracy of the multi-fidelity solution. However, this comes at the cost of increased computation, as additional important samples must be propagated, and the improvement in accuracy exhibits diminishing returns. Eventually, the state vectors introduced from additional time steps lack sufficient linear independence to increase the rank of the snapshot matrix. In general, longer propagation times warrant the use of more time steps.

4. NUMERIC SIMULATIONS

4.1. Test Case Descriptions

This section provides the test case scenarios and parameters that are used for assessing the adaptive multi-fidelity model selection method in the next two subsections. Four scenarios from [13] are used: a distant retrograde orbit (DRO), a near-rectilinear halo orbit (NRHO), a lunar transfer orbit (LTO), and a low Lunar flyby. Orbits such as these are likely to be used in upcoming cislunar missions, and are therefore relevant to cislunar SSA. Furthermore, these orbits span large swaths of cislunar space, making them challenging test cases for the adaptive method.

In the tracking results presented in Section 4.3, each of the four scenarios includes a cluster of five ASOs. The ASOs in each cluster all have the same initial state PDF, with the means provided in Tab. 1 and diagonal covariances of $10^{-8} \mathbf{I}_6$ in normalized units (LU^2 , LU^2/TU^2) in the rotating barycentric frame. Fig. 1 depicts the nominal trajectory for each of the four scenarios, computed from a deterministic propagation of the mean for ten days – the length of the scenario. The nominal trajectory of the Lunar flyby scenario has a close approach with the Moon of 133 km, assuming a spherical Moon. However, due to random realizations of the initial conditions, the closest approaches in the test cases will vary.

Table 1. Initial means for each scenario in normalized units (LU, LU/TU), expressed in the rotating barycentric frame.

Scenario	Mean					
	x	y	z	\dot{x}	\dot{y}	\dot{z}
DRO	0.806	0.000	0.000	0.000	0.519	0.000
NRHO	1.022	0.000	-0.182	0.000	-0.103	0.000
LTO	-0.112	0.000	0.000	2.194	-3.440	0.000
Flyby	0.949	-0.019	0.304	-0.006	0.064	0.003

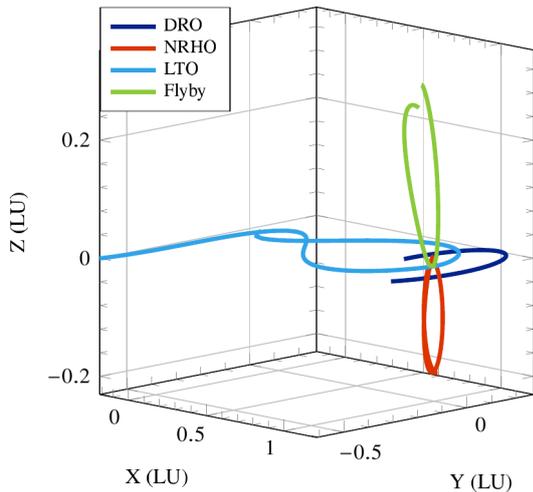


Figure 1. Nominal trajectories for each cluster in the rotating barycentric frame.

As mentioned in Section 3, the current implementation considers non-spherical gravity for only one body at a time. The radius in which to use non-spherical Lunar gravity is defined as the distance along the Earth-Moon vector at which the total (omission plus commission) error from each body is equal. This distance was found to be 91,264 km. When the space object is within this radius, the Moon’s non-spherical gravity is used; otherwise, the Earth’s is used.

The acceleration noise limit from Eq. 26 is set to $\varepsilon_a = 10^{-15}$ km/s² for all simulations. This value was found to provide a suitable balance between increased computation time from additional spherical harmonic gravity terms and accuracy in propagated states.

To evaluate the adaptive multi-fidelity method, three non-adaptive multi-fidelity baselines are considered, each using a constant gravity expansion for the high-fidelity model throughout the entire propagation: 30×30 , 60×60 , and 90×90 . In a non-adaptive approach, an operator tasked with propagating uncertainty or tracking space objects may not have prior knowledge of the specific regions of cislunar space through which the objects will pass, and therefore would have to select the gravity expansion heuristically. These three non-adaptive methods are chosen to represent the range of possible gravity expansions that might be used in the absence of such prior

knowledge and fall within the set of gravity expansions used by the adaptive method in Sec 4.2 and 4.3. All other aspects of the force models in the non-adaptive method are identical to the force models in the adaptive method.

The low-fidelity, high-fidelity, and “truth” propagators all use a 4th order accurate Runge-Kutta integrator with a 5th order error estimate and adaptive step size [20]. The low-fidelity model uses Earth and Moon point-mass gravity and the “truth” model uses 120×120 gravity for both bodies. The degree and order for the adaptive method never exceeds that of the truth. The Earth and Moon use the EGM2008 [21] and LP165P [22] gravity fields, respectively, as well as the IAU2006 coordinate system reduction [23]. Models of all fidelity have cannonball SRP and Sun point-mass gravity. Additional force model parameters may be found in Tab. 2. The positions of celestial bodies are determined by the JPL DE430 ephemerides [24].

Table 2. Force model and satellite parameters.

Parameter	Value
Satellite mass	500 kg
Satellite SRP area	1 m ²
Reflectivity coefficient	1.5
Epoch time	2455200.5 UTC

Runtime tests are generated on a Dell Precision Tower 3430 desktop computer running Red Hat Enterprise Linux 8.10 and the Linux 4.18.0 kernel with a 3.2 GHz Intel Core i7-8700 processor and 16 GB of random-access memory. The propagation software is written in C and all other software (tracker and adaptive model selection) is written in C++ and compiled with the GCC 8.5.0 compiler.

4.2. Results: Multi-Fidelity Propagation

This section evaluates the performance of the adaptive multi-fidelity method solely in the context of propagation, without incorporating it into a tracking framework. Analyzing the propagation in isolation is informative due to its broad applicability to uncertainty propagation. Here, a particle ensemble is propagated using a Monte Carlo approach, though the multi-fidelity method can also be applied to Gaussian mixture propagation via the unscented transform [11]. Additionally, this analysis can be viewed as the prediction step of a filtering process. The integration of the method into a full filtering and tracking framework is presented in the subsequent section.

Process noise is excluded from the propagation in this section to isolate the accuracy of the adaptive multi-fidelity method from the confounding effects that stochastic processes may introduce. The degree and order of the spherical harmonic gravity expansion is updated in 1 h time increments and the multi-fidelity snapshot matrix

is constructed using seven time steps in 1 h increments at the end of the propagation. All test cases use 1000 low-fidelity samples with the number of high-fidelity important samples being the determined value of r (see Eq. 8).

Runtime and accuracy statistics may be found in Tab. 3 and 4, respectively. Each set of results reflects the average of 50 Monte Carlo simulations, with each simulation having a different set of random initial states sampled from the distribution described for each test case in Section 4.1. For each Monte Carlo simulation, the same set of initial states is used for all propagation methods (low-fidelity, high-fidelity, non-adaptive multi-fidelity, and adaptive multi-fidelity). The root-mean-square error (RMSE) is computed for each propagation method using the “truth” defined in Section 4.1. In Tab. 3 and 4, the low-fidelity results reflect using only the low-fidelity model to propagate the entire particle ensemble. Likewise, the high-fidelity results reflect using only the high-fidelity model to propagate all particles, with the results presented for all three (30×30 , 60×60 , and 90×90) non-adaptive high-fidelity models described in Section 4.1. The non-adaptive multi-fidelity results reflect using those three non-adaptive high-fidelity models in the multi-fidelity framework. The non-adaptive and adaptive multi-fidelity runtime and accuracy results are plotted in Fig. 2 and 3, respectively. Additionally, the highest gravity degree and order used by the adaptive method for each scenario may be found in Tab 5.

Table 3. Runtime (sec) for all test cases. LF, HF, and MF are shorthand for low-fidelity, high-fidelity, and multi-fidelity, respectively.

Method	Scenario			
	DRO	NRHO	LTO	Flyby
LF	1.280	2.170	2.168	1.814
30×30 HF	19.09	27.49	62.90	27.64
60×60 HF	42.92	86.66	128.2	74.48
90×90 HF	82.19	185.3	239.1	150.6
30×30 MF	2.264	3.792	6.052	3.220
60×60 MF	3.290	6.260	8.496	5.218
90×90 MF	4.920	10.44	12.55	8.385
Adaptive MF	1.804	2.834	4.711	3.016

Table 4. RMSE (km) for all test cases.

Method	Scenario			
	DRO	NRHO	LTO	Flyby
LF	0.4824	202.5	299.3	42.72
30×30 HF	0	$4.144 \cdot 10^{-3}$	$2.000 \cdot 10^{-10}$	0.8186
60×60 HF	0	$6.640 \cdot 10^{-5}$	0	0.1833
90×90 HF	0	$1.490 \cdot 10^{-5}$	0	$5.517 \cdot 10^{-3}$
30×30 MF	$4.156 \cdot 10^{-4}$	0.1127	1.497	0.8212
60×60 MF	$4.156 \cdot 10^{-4}$	0.1129	1.497	0.1838
90×90 MF	$4.156 \cdot 10^{-4}$	0.1129	1.497	$2.704 \cdot 10^{-2}$
Adaptive MF	$4.126 \cdot 10^{-4}$	0.1104	1.473	$2.850 \cdot 10^{-2}$

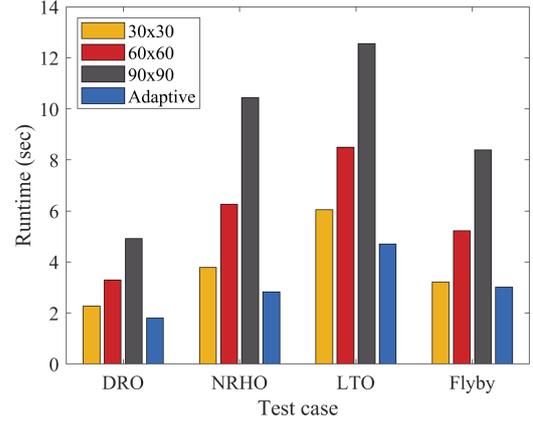


Figure 2. Runtime of non-adaptive and adaptive multi-fidelity methods for all test cases.

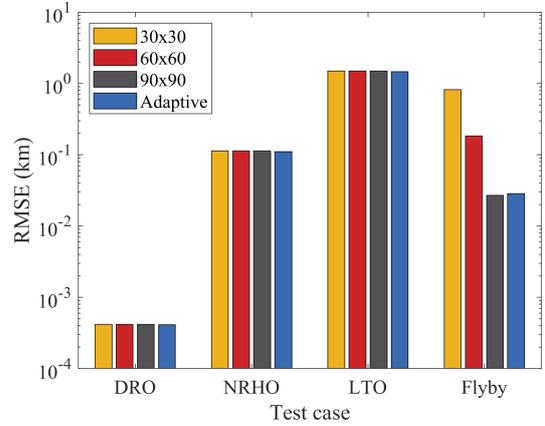


Figure 3. RMSE of non-adaptive and adaptive multi-fidelity methods for all test cases. Note the logarithmic scaling of the vertical axis.

The multi-fidelity algorithm achieves a multiple-order-of-magnitude improvement in the RMSE compared to the low-fidelity solution across all test cases, yielding kilometer to sub-kilometer accuracy, depending on the scenario. Additionally, the use of multi-fidelity propagation significantly reduces runtime relative to high-fidelity propagation, demonstrating an approximately tenfold speedup across test cases. The difference between high- and multi-fidelity runtime is amplified as the degree and order of the gravity expansion used in the high-fidelity model increase. The RMSE also generally decreases with more terms in the gravity expansion, though Tab. 4 indicates diminishing returns. Notably, for the DRO and LTO test cases, the high-fidelity propagation reaches zero RMSE relative to the truth model. This arises because these test cases remain at high altitudes, where, beyond a certain degree and order expansion, non-spherical gravitational accelerations fall below machine precision. This highlights a benefit of the adaptive method: by dynamically adjusting the gravity expansion based on an acceleration error threshold, it prevents the inclusion of unrec-

Table 5. Maximum spherical harmonic gravity degree and order used by the adaptive propagation method for each test case.

Scenario	DRO	NRHO	LTO	Flyby
Max deg/ord	3	39	20	91

essary terms, thereby reducing computational cost.

As shown in Tab. 3 and 4, and Fig. 2 and 3, the adaptive multi-fidelity method consistently achieves equal or superior accuracy compared to the non-adaptive multi-fidelity approach while maintaining a lower runtime. This is most evident in the flyby test case, where the non-adaptive method continues to improve in RMSE with increasing gravity expansion order but at the cost of greater runtime. In contrast, the adaptive method achieves the accuracy of the non-adaptive approach (from the 90×90 case) while maintaining a significantly faster runtime. This may be attributed to the flyby case having the largest altitude variation of all test cases. At perilune, a large gravity expansion is necessary to maintain accuracy, but at higher altitudes during the remainder of the trajectory, these additional terms become extraneous. This indicates that the adaptive method is most beneficial for cases where the distance to the Earth or Moon varies significantly over the trajectory. Similar behavior may also be observed in the NRHO test case.

Fig. 4 depicts the nominal NRHO trajectory colored by the degree and order of the spherical harmonic gravity expansion. As expected, the degree and order increase closer to the Moon, reaching a maximum at perilune as the higher order terms gain significance. This effect is further demonstrated in Fig. 5, which plots the degree and order as a function of time and exhibits a sharp peak at perilune.

4.3. Results: Multi-Target Tracking

This section presents Monte Carlo analyses of the multi-target filter’s performance in five simulated cislunar tracking scenarios. Each scenario involves tracking one or more clusters of space using the GLMBF with each of the following propagation methods: low-fidelity propagation with no correction, non-adaptive multi-fidelity propagation, and adaptive multi-fidelity propagation. In four of these scenarios, a single cluster of five objects have their initial states sampled from one of the Gaussian PDFs described in the previous subsection, corresponding to the DRO, NRHO, LTO, and flyby trajectories. In the other scenario, all four clusters are tracked simultaneously, for a total population of 20 objects on a diverse set of cislunar trajectories. This is repeated 20 times and the average results are presented in this subsection.

The low-fidelity, truth, and adaptive multi-fidelity propagators are the same as in the previous subsection, but when using the adaptive multi-fidelity method, the

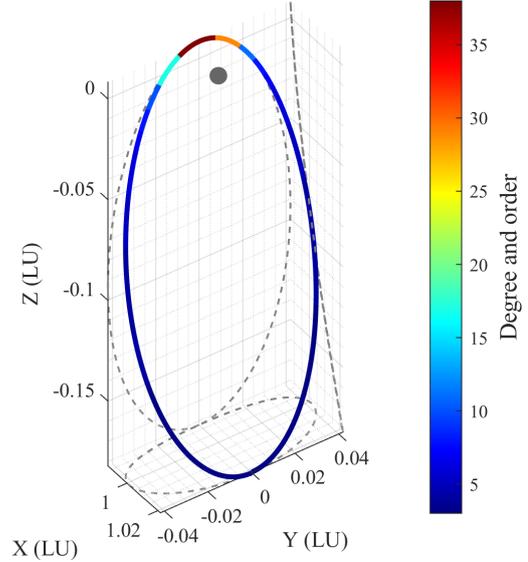


Figure 4. NRHO trajectory colored by spherical harmonic gravity degree and order. Note that the Moon is not shown to scale.

EnGMF skips the multi-fidelity correction when the maximum recommended gravity model resolution has stayed less than 5×5 since the last correction. The non-adaptive multi-fidelity propagator in this subsection uses a 10×10 gravity field model for both bodies. The EnGMF uses 200 particles to parameterize each track.

These simulations use a single optical sensor located at the L_1 Earth–Moon Lagrange point to obtain measurements of right ascension and declination and their respective rates. Sensor parameters from [13] are used, with the angle and angle-rate measurements possessing Gaussian distributed errors with a mean of zero and standard deviation of 0.1 arcsec for the angles and 0.001 arcsec/s for the angle-rates. The measurement errors are assumed to be uncorrelated. The sensor generates measurements every 12 h after the start of the simulation, with probability of detection $P_D = 0.95$. The simulated sensor does not generate clutter returns and the multi-target filter assumes a constant clutter intensity of $\kappa = 10^{-9}$.

The multi-target filter assumes a survival probability near one for all objects, discards hypotheses with weights below 10^{-3} , and employs Murty’s algorithm to cap the number of hypotheses at 50. This limit is distributed among prior GLMB components in proportion to the square root of their respective weights. The filter timestep is 1 h. The snapshot matrix is again constructed using the last seven timesteps at the time of multi-fidelity correction. Additionally, the EnGMF assumes a process noise covariance $\mathbf{Q} = \sigma_a^2 \mathbf{I}_3$, where the acceleration standard deviation is $\sigma_a = 10^{-12}$ km/s².

Fig. 6 shows the GLMBF’s tracking error using each method in the scenario with all 20 objects. The error is defined here as the optimal subpattern assignment (OSPA) multi-object distance between the GLMBF’s es-

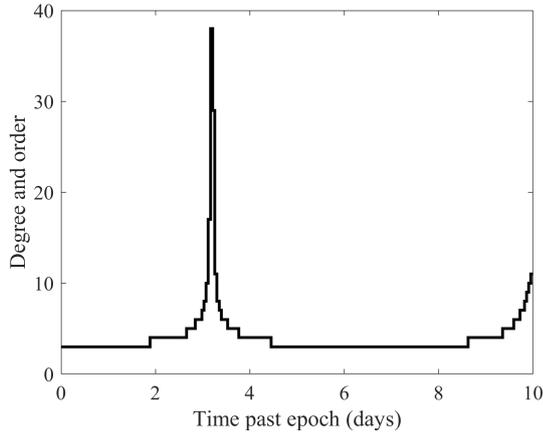


Figure 5. Spherical harmonic gravity degree and order of the NRHO trajectory as a function of time.

estimate of the states of the objects and their true states [25]. The OSPA metric used here is based on the Euclidean distance in Cartesian position space and has order 2 and cutoff 100 km. As the figure shows, all three methods initially perform well, but the low-fidelity-only method begins to diverge around four days into the simulation. This is primarily due to difficulty in maintaining custody of the objects in the NRHO and flyby clusters, as shown in Fig. 7 and 8, due to the strong effect of non-spherical gravity at the time of closest approach to the Moon. On the other hand, both the non-adaptive and adaptive methods have almost the same tracking error over the course of the simulation.

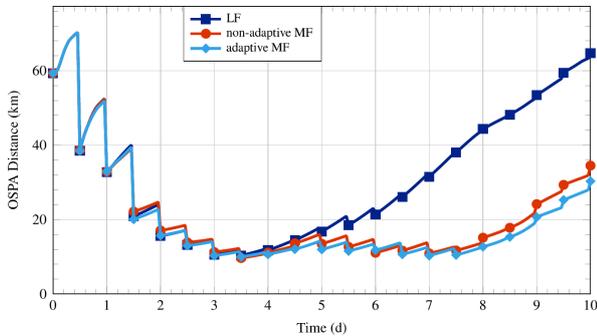


Figure 6. Plot of tracking error over time for the scenario with all twenty objects. Markers coincide with measurement updates.

The increase in the two adaptive methods' errors from $t = 8$ d to the end of the simulation is driven by difficulty tracking the DRO cluster due to poor observability. This can be mitigated by changing the sensor's location or adding another sensor. This is demonstrated in Figs. 9 and 10, which show tracking error for the DRO case with the sensor located at L_1 and at the Moon, respectively.

Tab. 6 shows the average tracking error for each scenario and propagation method. These results show that adaptive multi-fidelity propagation yields a tracking ac-

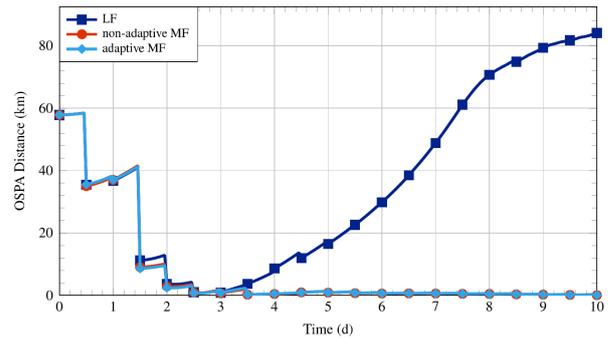


Figure 7. Plot of tracking error over time for the scenario with five objects on NRHOs. Markers coincide with measurement updates.

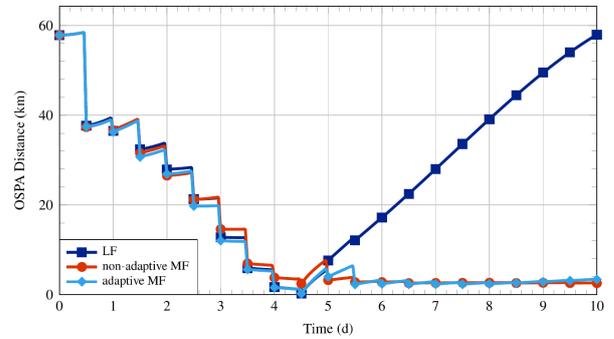


Figure 8. Plot of tracking error over time for the scenario with five objects on close flyby trajectories. Markers coincide with measurement updates.

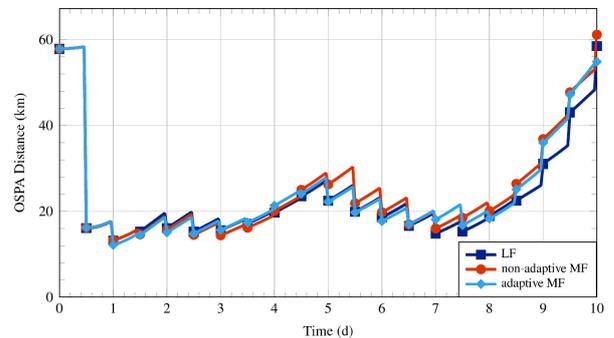


Figure 9. Plot of tracking error over time for the scenario with five objects on DROs and the sensor at L_1 . Markers coincide with measurement updates.

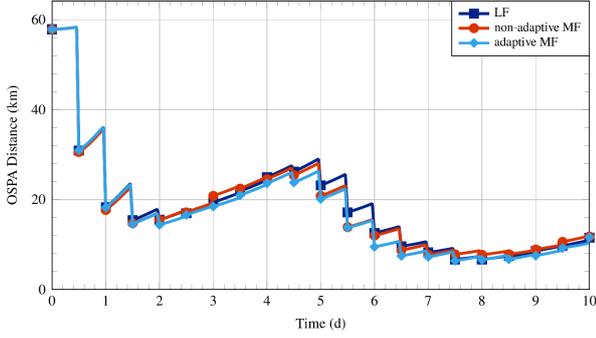


Figure 10. Plot of tracking error over time for the alternate scenario with five objects on DROs and the sensor located at the Moon. Markers coincide with measurement updates.

curacy similar to the non-adaptive multi-fidelity propagation in all five scenarios. As shown in Tab. 7, the adaptive method is able to achieve this level of accuracy while reducing total computation time by 34–62% compared to the non-adaptive method.

Table 6. Average OSPA tracking error (km) for each propagation method and scenario.

Method	All Clusters	DRO	NRHO	LTO	Flyby
LF	31.82	23.32	37.14	15.23	28.71
10×10 MF	21.15	24.70	7.78	10.85	13.61
Adaptive MF	19.88	24.15	7.76	10.82	13.19

Table 7. Average total runtime cost (sec) of multi-target filter for each propagation method and scenario.

Method	All Clusters	DRO	NRHO	LTO	Flyby
LF	28.51	7.70	10.40	15.98	8.94
10×10 MF	72.83	22.24	30.82	49.15	25.04
Adaptive MF	47.03	8.56	23.14	31.44	16.61

To explain these differences in filtering time, Fig. 11 shows the time cost of each filter step for each method in the scenario with all 20 objects. As expected, the cost spikes on measurement updates due to the GLMBF joint prediction and update and the multi-fidelity correction (when using one of the multi-fidelity methods). These spikes are always relatively high for the non-adaptive method because it is always performing a correction for every track. On the other hand, the adaptive method sometimes produces runtimes similar to those of the low-fidelity method because it is frequently able to use a lower degree and order expansion or even skip the correction entirely for many of its tracks. In this case, the adaptive method is only more expensive than the non-adaptive method at $t = 4.5$ d, where it takes a fraction of a second longer. This is around the time of the flyby cluster’s closest approach to the Moon and demonstrates how the adaptive method is able to recognize that a higher degree and order expansion is required for those objects at this time. The other significant spikes in the adaptive method’s run-

time correspond to other clusters’ closest approaches to either the Earth (for the LTO) or the Moon.

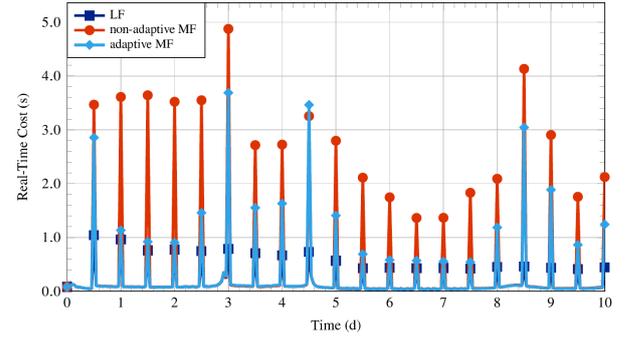


Figure 11. Plot of computational cost of each filter step for the scenario with all twenty objects. Markers coincide with measurement updates.

5. CONCLUSION

This paper developed an adaptive method for multi-fidelity uncertainty propagation in cislunar space by posing the issue as an optimization problem to minimize runtime subject to an upper bound on the expected acceleration error. The solution was obtained by varying the perturbations included in the high-fidelity model as a function of position via a precomputed library of acceleration errors for differing gravity expansion degrees and orders. This permits rapid and accurate uncertainty propagation of space objects in the cislunar domain. The effectiveness of the method for uncertainty propagation was demonstrated via Monte Carlo analysis. For ASO tracking, the propagator is integrated in a combined EnGMF and GLMBF multiple-hypothesis, multi-target filter. In simulated test cases, the adaptive approach yields a significant reduction in runtime when compared to a non-adaptive approach while maintaining equivalent or superior accuracy. Furthermore, the test scenarios demonstrated the method’s applicability to orbits relevant to upcoming cislunar missions and SSA, including a DRO, NRHO, LTO, and low Lunar flyby. Future work may entail extending the adaptive method to incorporate nonspherical gravity from both the Earth and Moon simultaneously or dynamically adjusting the time intervals at which the high-fidelity model is updated, allowing for more rapid adaptation in regions with faster-varying dynamics.

ACKNOWLEDGMENTS

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

necessarily representing the official policies or endorsements, either expressed or implied, of AFRL and/or the U.S. Government. AFRL Public Release Case Number: AFRL-2025-1167.

REFERENCES

1. M. J. Holzinger, C. C. Chow, and P. Garretson. *A Primer on Cislunar Space*. Air Force Research Laboratory, 2021.
2. Sam Wishnek, Marcus J. Holzinger, and Patrick Handley. Robust cislunar initial orbit determination. In *Advanced Maui Optical and Space Surveillance Technologies Conference*, September 2021.
3. Mark Bolden et al. Probabilistic initial orbit determination and object tracking in cislunar space using optical sensors. In *Advanced Maui Optical and Space Surveillance Technologies Conference*, September 2022.
4. Surabhi Bhadauria and Carolin Frueh. Optical observation regions in cislunar space using the bi-circular restricted four body problem geometry. In *Advanced Maui Optical and Space Surveillance Technologies Conference*, September 2022.
5. Jeffrey L. Anderson and Stephen L. Anderson. A Monte Carlo implementation of the nonlinear filtering problem to produce ensemble assimilations and forecasts. *Monthly Weather Review*, 127(12):2741–2758, December 1999.
6. Sehyun Yun, Renato Zanetti, and Brandon A. Jones. Kernel-based ensemble Gaussian mixture filtering for orbit determination with sparse data. *Advances in Space Research*, 69(12):4179–4197, April 2022.
7. Benjamin L. Reifler, Andrey A. Popov, Brandon A. Jones, and Renato Zanetti. Large-scale space object tracking in a proliferated LEO scenario. In *Proceedings of the 26th International Conference on Information Fusion*, June 2023.
8. Ba-Tuong Vo and Ba-Ngu Vo. Labeled random finite sets and multi-object conjugate priors. *IEEE Transactions on Signal Processing*, 61(13), July 2013.
9. Ba-Ngu Vo, Ba-Tuong Vo, and Dinh Phung. Labeled random finite sets and the Bayes multi-target tracking filter. *IEEE Transactions on Signal Processing*, 62(24), December 2014.
10. Ba-Ngu Vo, Ba-Tuong Vo, and Hung Gia Hoang. An efficient implementation of the generalized labeled multi-Bernoulli filter. *IEEE Transactions on Signal Processing*, 65(8), April 2017.
11. Brandon A. Jones and Ryan Weisman. Multi-fidelity orbit uncertainty propagation. *Acta Astronautica*, 155:406–417, 2019.
12. Sean McArdle. *Gravity Modeling for Lunar Orbits*. PhD thesis, The University of Texas at Austin, 2022.
13. Benjamin L. Reifler and Brandon A. Jones. Efficient cislunar multi-target tracking with adaptive multi-fidelity propagation. In *Advanced Maui Optical and Space Surveillance Technologies Conference*, September 2024.
14. Trevor N. Wolf, Enrico M. Zucchelli, and Brandon A. Jones. Multi-fidelity uncertainty propagation for objects in cislunar space. In *Proceedings of the 2022 AIAA SciTech Forum*, January 2022.
15. Akil Narayan, Claude Gittelsohn, and Dongbin Xiu. A stochastic collocation algorithm with multifidelity models. *SIAM Journal on Scientific Computing*, 36(2):A495–A521, 2014.
16. Bernard W. Silverman. *Density Estimation for Statistics and Data Analysis*. Routledge, 1998.
17. Rudolph van der Merwe and Eric A. Wan. The square-root unscented Kalman filter for state and parameter-estimation. *Proceedings of the 2001 IEEE International Conference on Acoustics, Speech, and Signal Processing*, pages 1–4, May 2001.
18. R. A. Broucke and P. J. Cefola. On the equinoctial orbital elements. *Celestial Mechanics*, 5:303–310, 1972.
19. Katta G. Murty. An algorithm for ranking all the assignments in order of increasing cost. *Space Operations Research*, 16(3):682–687, 1968.
20. J.R. Dormand and P.J. Prince. A family of embedded Runge–Kutta formulae. *Journal of Computational and Applied Mathematics*, 6(1):19–26, 1980.
21. Nikolaos K. Pavlis, Simon A. Holmes, Steve C. Kenyon, and John K. Factor. The development and evaluation of the Earth Gravitational Model 2008 (EGM2008). *Journal of Geophysical Research: Solid Earth*, 117(B4), 2012.
22. A.S. Konopliv, S.W. Asmar, E. Carranza, W.L. Sjogren, and D.N. Yuan. Recent gravity models as a result of the Lunar Prospector mission. *Icarus*, 150(1):1–18, 2001.
23. G. Petit and B. Luzum. IERS conventions (2010). IERS Technical Note 36, International Earth Rotation and Reference Systems Service (IERS), Frankfurt am Main, Germany, 2010.
24. W.M. Folkner, J.G. Williams, D.H. Boggs, R.S. Park, and P. Kuchynka. The planetary and lunar ephemerides DE430 DE431. *IPN Progress Report*, 42-196, February 2009.
25. Dominic Schuhmacher, Ba-Tuong Vo, and Ba-Ngu Vo. A consistent metric for performance evaluation of multi-object filters. *IEEE Transactions on Signal Processing*, 56(8):3447–3457, August 2008.