Bifidelity Uncertainty Propagation with Directional Splitting for Space Domain Awareness

Kyle J. DeMars,[1] Maaninee Gupta,[1] Renato Zanetti,[2] and Kristen Michaelson[2]

***Abstract*— Accurate uncertainty propagation is key for effective space domain awareness, especially for sparse data problems where longer periods of uncertainty propagation are required. When the dynamical system is nonlinear, guarantees of Gaussianity cannot be made, leading to non-Gaussian uncertainties. This work develops a bifidelity approach to uncertainty propagation in which Gaussian mixture representations of uncertainty are adapted online using directional splitting. Full-fidelity models are used for state propagation, and reduced-fidelity models are used for error propagation. It is found that the proposed combination of directional splitting and bifidelity dynamics is successful in propagating the uncertainty in a cislunar test case.**

I. Introduction

Uncertainty propagation is the process of determining the temporal evolution of the probability density function (pdf)–or some approximation thereof–of the state of a dynamical system. For problems in space domain awareness, the state of the system is usually taken to be Cartesian coordinates (i.e., position and velocity vectors) or some set of orbital elements (e.g., Keplerian elements). The space domain carries several complexities that make uncertainty propagation challenging. First, and foremost, the dynamics that represent the governing equations of motion for the states are generally nonlinear, implying that assumptions of Gaussianity cannot be guaranteed. Secondly, the faithful representation of orbital motion is characterized by complex models that represent the gravitational attraction of multiple celestial bodies, solar radiation pressure, atmospheric drag, etc., such that computational complexity is not always trivial. Finally, data with which to perform inference for space objects can be sparse, meaning that there can be long periods of time during which it is necessary to propagate uncertainty.

The earliest methods for quantifying uncertainty when dealing with orbital motion began with applying techniques like the Kalman filter [1] to the problem of estimating the trajectories of satellites [2] and to the circumlunar navigation problem [3], [4]. A Bayesian connection to the Kalman filter was established shortly thereafter [5], paving the way for key advancements in tractable, non-Gaussian uncertainty representations through the introduction and development of Gaussian mixture (GM) pdfs [6], [7]. Several recent methods that leverage GM representations have been developed and applied to the space object uncertainty propagation problem [8]–[10]. Mixture-based approaches facilitate approximations that address both the lack of Gaussian guarantees and the need for long propagation arcs for space object tracking.

Multifidelity methods, such as [11]–[13], provide methods by which the computational burden of high-fidelity dynamical system modeling can be addressed. These methods tend to use a set of selected high-fidelity state propagations to then identify corrections that can be made (e.g., using stochastic collocation, differential algebra, etc.) to correct low-fidelity state propagations. In doing so, multifidelity methods reduce the computational burden by making judicious use of a reduced-fidelity model of the dynamics. A different approach that is used in onboard spacecraft navigation is a bifidelity method in which the state estimate is propagated at a high rate, while the covariance is propagated at a lower rate [14]. This method naturally lends itself to treating the propagation equations separately and using different levels of fidelity. As an example, high-rate inertial measurement unit data can be used to propagate the state estimate, and a simplified accumulation of the data can be used to propagate the covariance.

In this paper, adaptive GM uncertainty propagation is combined with a bifidelity dynamics model. Adaptation of the GM representation of uncertainty is performed by splitting a Gaussian into multiple Gaussians after the effects of nonlinearity have been detected. The means of the mixture are propagated using a full-fidelity dynamical model, and the covariances of the mixture are propagated using a reduced-fidelity dynamical model.

In this work, scalar quantities are represented by lower- and uppercase Latin and Greek characters, vector quantities are represented by lowercase, bold Latin and Greek characters, and matrix quantities are represented by uppercase, bold Latin and Greek characters. The trace of a matrix is denoted by $\mathrm{trace}\left\{\cdot\right\}$, the determinant of a matrix is denoted by $\left|\cdot\right|$, the matrix inverse is denoted by $(\cdot)^{-1}$, and the vector/matrix transpose is denoted by $(\cdot)^T$.

II. Uncertainty Propagation

Consider a continuous-time dynamical system of the form

$$\dot{\boldsymbol{x}}(t) = \boldsymbol{f}(\boldsymbol{x}(t), t)\,, \tag{1}$$

where $\boldsymbol{f} : \mathbb{R}^n \times \mathbb{R} \to \mathbb{R}^n$ is a deterministic nonlinear function describing the dynamics of the problem. Uncertainty is introduced via the initial conditions, $\boldsymbol{x}_0 = \boldsymbol{x}(t_0) \sim p(\boldsymbol{x}_0)$, where $p(\boldsymbol{x}_0)$ is taken to be a known pdf. The objective of uncertainty propagation is to find $p(\boldsymbol{x}(t))$, given some $t > t_0$. This problem has a well-known solution that is given by Liouville's equation (which is equivalently the Fokker–Plank

[1]Department of Aerospace Engineering, Texas A&M University, College Station, TX 77843, USA

[2]Department of Aerospace Engineering and Engineering Mechanics, The University of Texas at Austin, Austin, TX 78712, USA

equation or the forward Kolmogorov equation in the absence of process noise) [15]. These solutions, however, do not, in general, admit a finite parameterization of the pdf. Instead, approximations are often required.

A. Fixed-Component Approaches

In this work, GM models are leveraged to represent the state pdf; as such, the pdf is given by an L_x-component mixture of the form

$$p(\boldsymbol{x}) = \sum_{\ell=1}^{L_x} w_x^{(\ell)} p_g(\boldsymbol{x}; \boldsymbol{m}_x^{(\ell)}, \boldsymbol{P}_{xx}^{(\ell)}) \,, \tag{2}$$

where $w_x^{(\ell)}$, $\boldsymbol{m}_x^{(\ell)}$, and $\boldsymbol{P}_{xx}^{(\ell)}$ represent the weights, means, and covariances of each component within the mixture. The weights are constrained to be non-negative, and the covariance matrices are required to be symmetric and positive definite.

If the dynamical system is linear and the state distribution is a GM (of which the Gaussian pdf is a special case), then it is well-established that the propagation of uncertainty can be carried out exactly, where the weights are constant and the means and covariances evolve in a similar manner to the Kalman filter [6]. The dynamical system of (1) can be expressed as

$$\dot{\boldsymbol{x}}(t) = \boldsymbol{F}(t)\boldsymbol{x}(t) \,,$$

and the governing equations of motion for the GM parameters are

$$\begin{aligned} \dot{w}_x^{(\ell)}(t) &= 0 \\ \dot{\boldsymbol{m}}_x^{(\ell)}(t) &= \boldsymbol{F}(t)\boldsymbol{m}_x^{(\ell)}(t) \\ \dot{\boldsymbol{P}}_{xx}^{(\ell)}(t) &= \boldsymbol{F}(t)\boldsymbol{P}_{xx}^{(\ell)}(t) + \boldsymbol{P}_{xx}^{(\ell)}(t)\boldsymbol{F}^T(t) \,, \end{aligned}$$

which are applied for each component. Given linear dynamics (and no process noise or exactly Gaussian process noise), the number of components in the GM remains constant. While the approach developed in [6] is formulated for discrete dynamical systems, it is straightforward to extend the approach to the continuous dynamical system representation that is used in this work via an analogous approach used for the continuous-discrete Kalman filter [16]. Additionally, it is worth noting that process noise is not considered in the dynamical system model in this work.

In the more general case where the dynamical system is nonlinear, as in (1), it is common to employ some type of approximation, such as [7]

$$\begin{aligned} \dot{w}_x^{(\ell)}(t) &= 0 \\ \dot{\boldsymbol{m}}_x^{(\ell)}(t) &= \boldsymbol{f}(\boldsymbol{m}_x^{(\ell)}(t), t) \\ \dot{\boldsymbol{P}}_{xx}^{(\ell)}(t) &= \boldsymbol{F}(\boldsymbol{m}_x^{(\ell)}(t), t)\boldsymbol{P}_{xx}^{(\ell)}(t) + \boldsymbol{P}_{xx}^{(\ell)}(t)\boldsymbol{F}^T(\boldsymbol{m}_x^{(\ell)}(t), t) \,, \end{aligned}$$

which is a first-order Taylor series expansion about the mean of each component of the GM. Other approximations, such as a per-component application of the unscented transform (UT) [17]–[19] or Gauss–Hermite (GH) quadrature [20], [21], can also be leveraged. Importantly, all of the aforementioned methods are approximate; the evolution of the GM parameters is no longer guaranteed to exactly represent the forward evolution of the pdf. As part of the approximation, it is assumed that the number of components in the GM remains fixed, but this also cannot be guaranteed. When using the UT or GH methods, the transformation of the sigma/quadrature points is performed through numerical integration of the continuous-time nonlinear dynamical system, similar to [22].

B. Adaptive Approaches

One method to reduce approximation error is to employ an adaptive GM approach, wherein the number of mixands and/or the parameters of the mixands are adapted online [8], [10]. Adaptive approaches are still approximations, but they seek to lessen the restrictions of the fixed-component methods presented previously. In this work, an adaptive approach is presented where adaptation is triggered by monitoring agreement between different approximations of the uncertainty propagation. The adaptation is then to "split" a component of the GM into several smaller components in order to reduce the approximation error.

C. Splitting Criterion

The Kullback–Leibler (KL) divergence, which is defined as [23], [24]

$$D_{\mathrm{KL}}[p_1||p_2] = \int p_1(\boldsymbol{x}) \log\Big(p_1(\boldsymbol{x})/p_2(\boldsymbol{x})\Big) \mathrm{d}\boldsymbol{x} \,, \tag{3}$$

quantifies the information lost in representing one pdf, $p_1(\boldsymbol{x})$, by another pdf, $p_2(\boldsymbol{x})$. The KL divergence satisfies two key properties: 1) it is self-identifying ($D_{\mathrm{KL}}[p||p] = 0$), and 2) it is non-negative ($D_{\mathrm{KL}}[p_1||p_2] \geq 0$). Unfortunately, the KL divergence is not symmetric (in general, $D_{\mathrm{KL}}[p_1||p_2] \neq D_{\mathrm{KL}}[p_2||p_1]$), and it does not satisfy the triangle inequality. As such, the KL divergence does not satisfy the properties to be a metric; nevertheless, self-identification and non-negativity are powerful properties of the KL divergence.

In this work, the KL divergence is used to detect departures of one method of uncertainty propagation from another method of uncertainty propagation. That is, letting p_1 denote a reference density that is governed by a higher-order propagation a GM pdf, p_2 represents a lower-order propagation of the same pdf, and the KL divergence is used to quantify the discrepancy between each component of p_1 and p_2.

In particular, the components of p_1 are propagated forward in time using the unscented transform (UT) [17]–[19], and the components of p_2 are propagated forward in time using a first-order Taylor series expansion approximation, yielding

$$p_1^{(\ell)}(\boldsymbol{x}) = p_g(\boldsymbol{x}; \boldsymbol{m}_{x,\mathrm{uns}}^{(\ell)}, \boldsymbol{P}_{xx,\mathrm{uns}}^{(\ell)}) \tag{4}$$

and

$$p_2^{(\ell)}(\boldsymbol{x}) = p_g(\boldsymbol{x}; \boldsymbol{m}_{x,\mathrm{lin}}^{(\ell)}, \boldsymbol{P}_{xx,\mathrm{lin}}^{(\ell)}) \,. \tag{5}$$

The difference between the different methods of propagating the mean and covariance is quantified using the KL divergence. Substituting the Gaussian pdfs of Eqs. (4) and (5) into

the KL divergence of (3), it follows that

$$D_{\mathrm{KL}}[p_1^{(\ell)}||p_2^{(\ell)}] = \frac{1}{2}\Bigg[\log\{|\boldsymbol{P}_{xx,\mathrm{lin}}^{(\ell)}|/|\boldsymbol{P}_{xx,\mathrm{uns}}^{(\ell)}|\} + \mathrm{trace}\left\{\boldsymbol{P}_{xx,\mathrm{uns}}^{(\ell)}(\boldsymbol{P}_{xx,\mathrm{lin}}^{(\ell)})^{-1}\right\} + (\boldsymbol{m}_{x,\mathrm{lin}}^{(\ell)} - \boldsymbol{m}_{x,\mathrm{uns}}^{(\ell)})^T(\boldsymbol{P}_{xx,\mathrm{lin}}^{(\ell)})^{-1}(\boldsymbol{m}_{x,\mathrm{lin}}^{(\ell)} - \boldsymbol{m}_{x,\mathrm{uns}}^{(\ell)}) - n\Bigg]. \tag{6}$$

When the KL divergence between the two methods exceeds a threshold, i.e., when [25], [26]

$$D_{\mathrm{KL}}[p_1^{(\ell)}||p_2^{(\ell)}] > \tau\,, \tag{7}$$

action should be taken, as the effects of nonlinearity are causing departures in the propagation of the mean and covariance. The action taken is that the ℓ^{th} component of the GM is split into a set of smaller Gaussians. This criterion is applied to all of the components of the GM on a per-component basis.

The threshold τ is chosen to limit the maximum allowable divergence between the two distributions. Consider two Gaussian pdfs of the form

$$q_1(\boldsymbol{x}) = p_g(\boldsymbol{x};\boldsymbol{\mu},\boldsymbol{\Sigma})$$
$$q_2(\boldsymbol{x}) = p_g(\boldsymbol{x};\boldsymbol{\mu} + c\boldsymbol{\Sigma}^{1/2}\boldsymbol{v},\boldsymbol{\Sigma}/k)\,,$$

where $\boldsymbol{v}$ is some unit vector, $c > 0$, and $k > 1$. The first distribution, q_1, represents the higher-order approximation of the true distribution, and the second distribution, q_2, represents the lower-order approximation of the true distribution. Given q_1 and q_2, the KL divergence for this comparison reduces to

$$\tau = D_{\mathrm{KL}}[q_1||q_2] = \frac{1}{2}(n(k - \log k - 1) + c^2 k)\,. \tag{8}$$

The value of k, which is constant, is chosen as the maximum divergence that is permitted between the two covariances, while the constant c is set as the number of standard deviations by which the means are allowed to diverge. Once these parameters are set, (7), in conjunction with (6), is used to detect when splitting is required. As the threshold does not depend on the specific solutions at each step, it can be precomputed and stored.

D. Splitting Direction

After determining that linearized uncertainty propagation is no longer acceptable and that a component split is needed, the next step is to determine the direction of the split, i.e., the direction of maximum nonlinearity. The approach used in this work follows the approach of [27], which expands upon the method of [25] to split nonlinear vector functions of the state.

Let $\boldsymbol{z} \in \mathbb{R}^m$ be a nonlinear function of the state $\boldsymbol{x} \in \mathbb{R}^n$, i.e.,

$$\boldsymbol{z} = \boldsymbol{g}(\boldsymbol{x})\,.$$

The $m \times n$ Jacobian is denoted by $\boldsymbol{G}$, which is comprised of row vectors, $\boldsymbol{G}_i$, i.e., $\boldsymbol{G}_i$ is the Jacobian of the i^{th} scalar component of the nonlinear function. A linear system possesses constant first-order derivatives; the rate of change of the Jacobian at $\bar{\boldsymbol{x}}$ in direction $\boldsymbol{u}$ is therefore a reasonable quantification of nonlinearity. The directional derivative is given by

$$\nabla_{\boldsymbol{u}}\boldsymbol{G}_i(\boldsymbol{x})\big|_{\boldsymbol{x}=\bar{\boldsymbol{x}}} = \lim_{\alpha\to 0}\frac{\boldsymbol{G}_i(\bar{\boldsymbol{x}} + \alpha\boldsymbol{u}) - \boldsymbol{G}_i(\bar{\boldsymbol{x}})}{\alpha}\,. \tag{9}$$

Taking a first-order Taylor series of the Jacobian, it can be shown that [27]

$$\nabla_{\boldsymbol{u}}\boldsymbol{G}_i(\boldsymbol{x})\big|_{\boldsymbol{x}=\bar{\boldsymbol{x}}} = \boldsymbol{u}^T\boldsymbol{H}_i(\bar{\boldsymbol{x}})^T\,,$$

where $\boldsymbol{H}_i(\boldsymbol{x}) \in \mathbb{R}^{n\times n}$ is the Hessian of the i^{th} component of $\boldsymbol{g}(\boldsymbol{x})$. Assembling the components of the directional derivative, it can be shown that the squared Frobenius norm is [27]

$$\left\|\nabla_{\boldsymbol{u}}\boldsymbol{G}(\boldsymbol{x})\big|_{\boldsymbol{x}=\bar{\boldsymbol{x}}}\right\|_F^2 = \boldsymbol{u}^T\boldsymbol{E}(\bar{\boldsymbol{x}})\boldsymbol{u}\,, \tag{10}$$

where

$$\boldsymbol{E}(\bar{\boldsymbol{x}}) = \sum_{i=1}^{m}\boldsymbol{H}_i^T(\bar{\boldsymbol{x}})\boldsymbol{H}_i(\bar{\boldsymbol{x}})\,. \tag{11}$$

Equation (10) provides a measure of the nonlinearity in a particular direction, $\boldsymbol{u}$. The objective of splitting a Gaussian into smaller Gaussians is to mitigate the effects of nonlinearity. It is also important to take into account the amount of uncertainty in a given direction. As such, an appropriate cost function is

$$J(\boldsymbol{u}) = \left\|\nabla_{\boldsymbol{u}}\boldsymbol{G}(\boldsymbol{x})\big|_{\boldsymbol{x}=\bar{\boldsymbol{x}}}\right\|_F^2\,\sigma_u^2 = \frac{\boldsymbol{u}^T\boldsymbol{E}(\bar{\boldsymbol{x}})\boldsymbol{u}}{\boldsymbol{u}^T\boldsymbol{P}_{xx}^{-1}\boldsymbol{u}}\,, \tag{12}$$

which weights (10) by the uncertainty represented by the covariance matrix $\boldsymbol{P}_{xx}$ in the direction $\boldsymbol{u}$.

The direction $\boldsymbol{u}^*$ that maximizes (12) can be easily computed after a change of variables. Let $\boldsymbol{S}_{xx}$ be a square-root factor of $\boldsymbol{P}_{xx}$, such that $\boldsymbol{P}_{xx} = \boldsymbol{S}_{xx}\boldsymbol{S}_{xx}^T$. Let $\boldsymbol{v} = \boldsymbol{S}_{xx}^{-1}\boldsymbol{u}$, such that (12) can be equivalently written as

$$\tilde{J}(\boldsymbol{v}) = \frac{\boldsymbol{v}^T\boldsymbol{S}_{xx}^T\boldsymbol{E}(\bar{\boldsymbol{x}})\boldsymbol{S}_{xx}\boldsymbol{v}}{\boldsymbol{v}^T\boldsymbol{v}}\,. \tag{13}$$

The Rayleigh–Ritz inequality provides an upper bound for (13) that is the maximum eigenvalue of matrix $\boldsymbol{S}_{xx}^T\boldsymbol{E}(\bar{\boldsymbol{x}})\boldsymbol{S}_{xx}$, with the associated eigenvector, $\boldsymbol{v}^*$, being the maximizing argument. Changing variables back to the original problem, the chosen direction for splitting is $\boldsymbol{u}^* = \boldsymbol{S}_{xx}\boldsymbol{v}^*$.

E. Splitting Library

Once it is determined that the uncertainty requires refinement through splitting, and once the direction in which the splitting is to be performed is found, a splitting library is used to replace a single Gaussian component with several smaller Gaussian components. This work leverages the variance-preserving splitting library of [27], which is found by minimizing the KL divergence between a GM and a standard Gaussian, subject to the constraints that 1) the weights are all positive, 2) the weights sum to one, 3) the overall mean of the GM is zero, and 4) the overall variance of the GM is one. The result of this optimization problem for a three-component mixture is summarized in Table I, which is taken from [27].

TABLE I: Variance-preserving splitting library

Comp. #	Weight	Mean	Std. Dev.
-1	0.1616701997	-1.0908000117	0.78439476713
0	0.6766596007	0	0.78439476713
1	0.1616701997	1.0908000117	0.78439476713

F. Model Fidelity

Many dynamical systems admit lower-fidelity approximations of the true dynamics. For real-world systems, even the modeled "true" dynamics are approximations of reality. In situations where high-fidelity dynamics are computationally expensive, leveraging lower-fidelity models can ameliorate computational requirements without deleterious effects to algorithm performance.

Consider the dynamical system given by (1). Let $\tilde{\boldsymbol{f}}(\cdot,\cdot)$ be a reduced-fidelity approximation of $\boldsymbol{f}(\cdot,\cdot)$. In this work, the full-fidelity model is used for the propagation of the means of the GM in both linearized and unscented implementations. The reduced-fidelity model is used in the computation of the Jacobian, Hessian, combined Hessian of (11), and the solution to the optimization problem of (13).

It is worth noting that the model fidelity is different than the specific method of uncertainty propagation. In Section II-C, the terms "higher-order" and "lower-order" are used to refer to the method of propagating means and covariances. The use of "higher-fidelity" and "lower-fidelity" refers to the specific dynamics function that governs the system equations of motion. That is, a higher-order propagation method can be used with a low-fidelity model, and vice versa.

III. Adaptive Splitting Demonstration

To demonstrate the proposed approach for adaptive splitting during uncertainty propagation, consider a bivariate system in which the state is comprised of the semi-major axis (SMA), a, and the mean longitude, ℓ, representing the orbital parameters of an object in orbit about Earth. For unperturbed Keplerian motion, the corresponding dynamics of this system are given by

$$\dot{\boldsymbol{x}} = \boldsymbol{f}(\boldsymbol{x}) = \begin{bmatrix} 0 \\ \sqrt{\mu/a^3} \end{bmatrix}, \quad \text{where} \quad \boldsymbol{x} = \begin{bmatrix} a \\ \ell \end{bmatrix}, \tag{14}$$

and μ is the gravitational parameter of Earth. Note, in particular, that the semi-major axis is constant. In this example, the full model fidelity is used throughout the uncertainty propagation process. The Jacobian associated with the dynamical system of (14) is

$$\boldsymbol{G}(\boldsymbol{x}) = \frac{\partial \boldsymbol{f}(\boldsymbol{x})}{\partial \boldsymbol{x}} = \begin{bmatrix} 0 & 0 \\ -\frac{3}{2}\sqrt{\mu/a^5} & 0 \end{bmatrix}.$$

The Hessians corresponding to the first and second rows of the Jacobian, i.e., $\boldsymbol{H}_1(\boldsymbol{x})$ and $\boldsymbol{H}_2(\boldsymbol{x})$, respectively, can be found in a rather straightforward manner such that, with $\boldsymbol{E}(\boldsymbol{x})$ as defined in (11),

$$\boldsymbol{E}(\boldsymbol{x}) = \begin{bmatrix} \frac{225}{16}\frac{\mu}{a^7} & 0 \\ 0 & 0 \end{bmatrix}.$$

Five propagation methods are implemented and compared: the extended Kalman filter (EKF), polynomial chaos (PC) [28], AEGIS [10], the adaptive Gaussian sum filter (AGSF) [8], and the proposed method of this work. The EKF and PC are used to approximate the first and second central moments of the distribution, and AEGIS, the AGSF, and the proposed method are used to approximate distribution itself via Gaussian mixture representations. Additionally, a Monte Carlo (MC) simulation with 1×10^5 samples is used to represent the true distribution. Each method is applied to propagate either the moments or the distribution for a period of one day, starting from an initial distribution that is Gaussian with mean and covariance given, respectively, by

$$\boldsymbol{m}_0 = \begin{bmatrix} 42164.172 \text{ km} \\ 0 \text{ deg} \end{bmatrix} \quad \text{and}$$
$$\boldsymbol{P}_0 = \begin{bmatrix} (5000 \text{ km})^2 & 0 \\ 0 & (5 \text{ deg})^2 \end{bmatrix}.$$

Hermite polynomials are utilized for both random variables in PC, with a total expansion order of three. AEGIS is implemented using the three-component splitting library detailed in [10], whereas the proposed method leverages the splitting library of Table I. When using the UT within AEGIS, the "$2n$" cubature rule is used [18]. The AGSF approach uses a fixed number of Gaussian components in its mixture. There are a variety of ways in which these components can be chosen, but to maintain similar computational cost to AEGIS and the proposed approach, 25 components are used for the AGSF method. The initial mixture is determined by applying the 5-component splitting library of [10] to each marginal univariate distribution of the initial Gaussian distribution. At the final time (one day after the specified initial condition), AEGIS uses 27 components in its Gaussian mixture representation, while the proposed method uses only 11 components.

To analyze the performance, estimates for each of the states (the semi-major axis and the mean longitude) are determined by computing the conditional mean and covariance from each method. The mean estimation errors are determined by taking the absolute value of the difference of each method from the MC result. These results are illustrated in Fig. 1, from which, it is clear that the AGSF approach performs the worst in the SMA error. All of the other approaches, on the other hand, are clustered around a value of 10^{-10}. This performance of the AGSF occurs because, unlike the other GM-based approaches, the weights in the AGSF representation are adjusted via an online solution to a quadratic programming problem, which provides no guarantee with respect to preserving the overall mean. AEGIS and the proposed approach, however, use mean-preserving libraries. Since the SMA is constant, AEGIS and the proposed approach are able to retain the correct solution. The mean longitude error, on the other hand, shows more dynamic trends, with PC consistently performing the best and the EKF consistently performing the worst.

Standard deviation estimation errors are determined in a similar way to the mean estimation errors, albeit from the

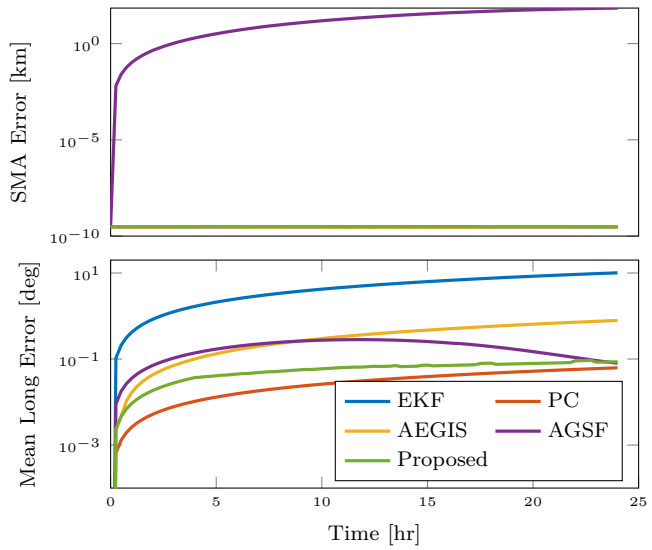

Fig. 1: Mean estimation error

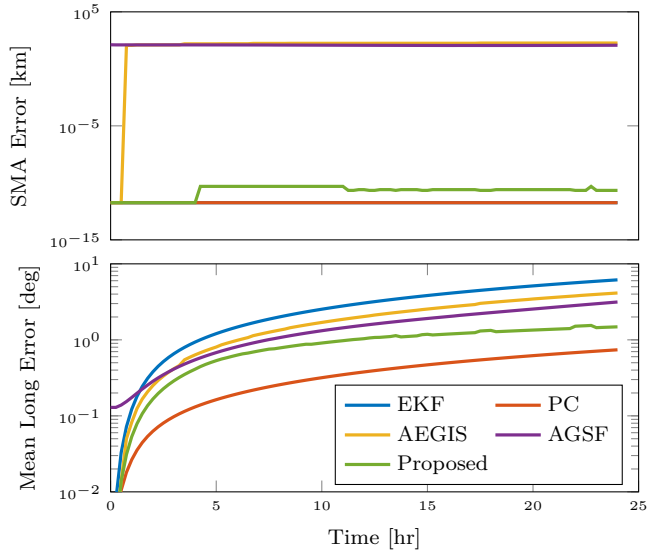

Fig. 2: Standard deviation estimation error

estimated covariances of each method. These results are provided in Fig. 2. Noting again that the SMA is constant, Fig. 2 illustrates that both AEGIS and the AGSF have comparatively poor performance in propagating the standard deviation of the semi-major axis. This occurs because neither approach guarantees covariance continuity in their weight determination (AGSF) or splitting application (AEGIS). For the mean longitude standard deviation error, PC and the EKF bookend the performance, with the proposed approach consistently outperforming both AEGIS and the AGSF.

IV. Results and Discussion

To highlight the effectiveness of using reduced fidelity modeling, consider the scenario in which an object is nominally in a 9:2 Near Rectilinear Halo Orbit (NRHO), which is the orbit that serves as the current baseline for NASA's Gateway lunar outpost.

A. Full Fidelity Model

The "full fidelity" model of the dynamical system is written in a form compatible with (1) as

$$\begin{aligned}\dot{\boldsymbol{r}} &= \boldsymbol{v}\\ \dot{\boldsymbol{v}} &= \boldsymbol{a}_{\text{moon}} + \boldsymbol{a}_{\text{earth}} + \boldsymbol{a}_{\text{sun}} + \boldsymbol{a}_{\text{srp}}\,,\end{aligned}$$

where $\boldsymbol{r}$ is the inertial position of the object, $\boldsymbol{v}$ is the inertial velocity of the object, $\boldsymbol{a}_{\text{body}}$ is the gravitational acceleration acting on the object due to the Moon, Earth, or Sun, and $\boldsymbol{a}_{\text{srp}}$ is the acceleration due to solar radiation pressure acting on the object. Solar radiation pressure is determined based on the physical characteristics of NASA's Gateway lunar outpost, with an area of approximately $350\ \text{m}^2$ and mass $39,000$ kg [29]. The coefficient of reflectivity is 1.3 [30].

Fig. 3 visualizes the motion of the object using the full-fidelity model, with Fig. 3a illustrating the object's trajectory in an Earth-Moon rotating frame and Fig. 3b illustrating the object's trajectory in a Moon-centered inertial frame. Propagation is initiated at apolune and terminates near perilune, approximately 3.27 days later. All quantities are nondimensionalized prior to propagation using the values provided in Table II.

TABLE II: Nondimensionalizing quantities.

Characteristic Quantity	Value
Length	384747.991979046
Time	375699.859037759
Mass	$6.04582557449506 \times 10^{24}$

B. Reduced Fidelity Model

To assess the influence of the each of the four terms in the full-fidelity model, Fig. 3c illustrates the magnitude of acceleration along the NRHO over the propagated time. Due to the object's proximity to the Moon, lunar gravity dominates, followed by the point mass gravity of the Earth and the Sun, respectively. Solar radiation pressure along the orbit remains constant. Given the relatively minor effects of solar gravity and solar radiation pressure, the reduced fidelity model of the dynamical system is taken to include only the effects of the gravity of the Moon and Earth.

Based on the reduced-fidelity model, the propagation of uncertainty (including all Jacobian and Hessian calculations) is based on the dynamical system described by

$$\begin{aligned}\dot{\boldsymbol{r}} &= \boldsymbol{v}\\ \dot{\boldsymbol{v}} &= \boldsymbol{a}_{\text{moon}} + \boldsymbol{a}_{\text{earth}}\,.\end{aligned}$$

For brevity, the Jacobians and Hessians required to formulate (11) and to solve the optimization problem of (13) are not derived here.

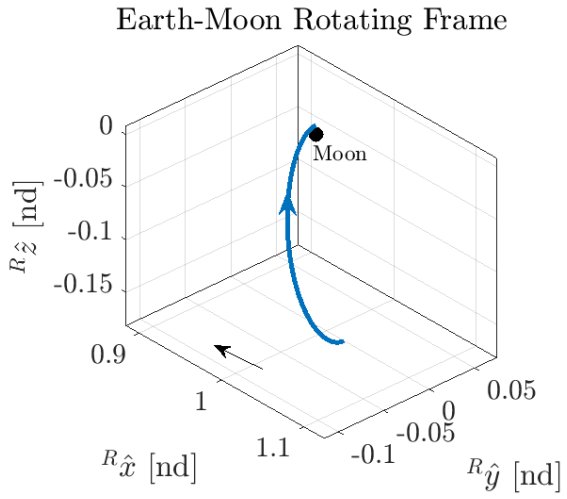

(a) Earth-Moon rotating frame

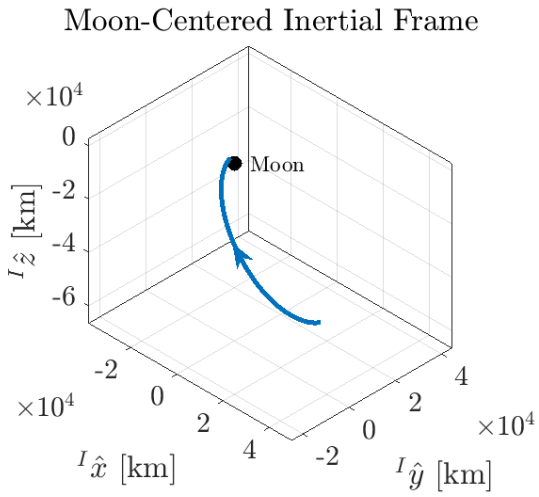

(b) Moon-centered inertial frame

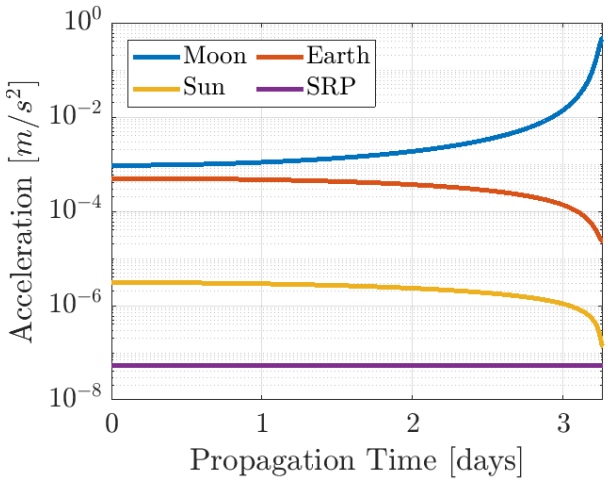

(c) Acceleration magnitude

Fig. 3: The 9:2 NRHO propagated for 3.27 days.

C. Uncertainty Propagation Results

Three configurations for uncertainty propagation using directional splitting, detailed in Table III, are investigated. Case 1 represents the baseline where the full-fidelity model is used throughout. Case 2 is the proposed bifidelity approach that balances computational expense against accuracy of the uncertainty propagation. Case 3 is included to demonstrate the need for retaining high-fidelity dynamics to some degree.

The directional splitting approach uses the the "$2n$" cubature rule UT as the higher-order method and analytical linearization (i.e., the propagation step of the EKF) as the lower-order method. When the full-fidelity model is employed according to the "mean dynamics" column of Table III, all of the sigma points of the UT and the mean of the EKF are propagated with the full-fidelity model. When the reduced-fidelity model is employed, the sigma points and EKF mean use the reduced-fidelity model. The fidelity of the "covariance dynamics" indicated in Table III dictates the model used for computing the Jacobian and Hessian related to the EKF and the directional splitting process.

TABLE III: Simulation details.

	Mean Dynamics	**Covariance Dynamics**
Case 1	Full	Full
Case 2	Full	Reduced
Case 3	Reduced	Reduced

The initial distribution is taken to be Gaussian with 1σ values of uncertainty of 10 km and 0.1 m/s in the inertial frame. In each case, the results are compared against a Kernel Density Estimate (KDE) computed from $100,000$ MC samples propagated using the full-fidelity dynamical model.

The results of each case are illustrated as pair plots in Fig. 4, where a pair plot visualizes the marginal distribution of each univariate element of the state and all bivariate combinations of the state elements. Each pair plot compares the result of a KDE with the results obtained using the proposed GM-based directional splitting approach, where the proposed approach is illustrated using contours of uncertainty on top of the KDE result.

The results of Case 1 and Case 2, which are shown in Figs. 4a and 4b, respectively, demonstrate good agreement between the GM representation of uncertainty and the KDE. These results also demonstrate good agreement between one another, validating the bifidelity approach to uncertainty propagation. The results of Case 3 (shown in Fig. 4c), where the reduced-fidelity dynamics are used in both the propagation of the means and the covariances (including the directional splitting process) shows that reduced-fidelity dynamics cannot fully replace the full-fidelity dynamics.

V. Conclusions

Accurate uncertainty propagation is critical for space domain awareness. Having reliable knowledge of the state distribution between sparse measurements facilitates conjunction screening and data processing, among other things.

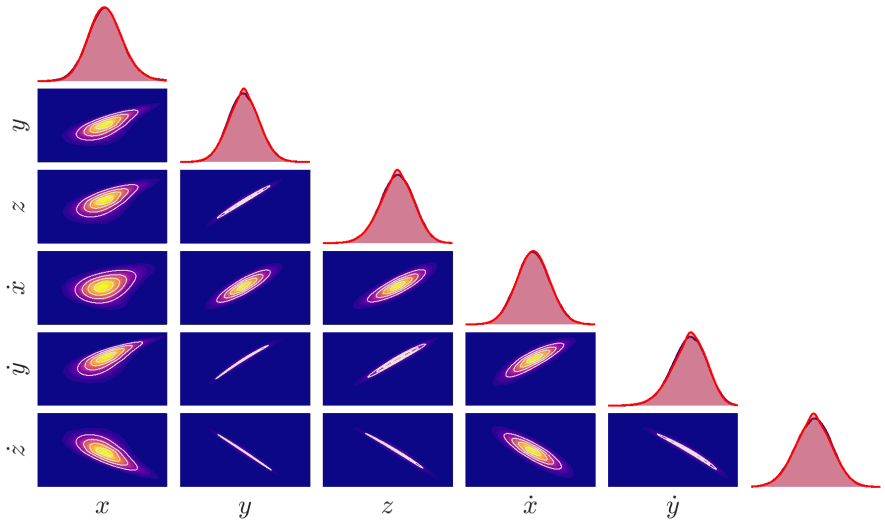

(a) Case 1

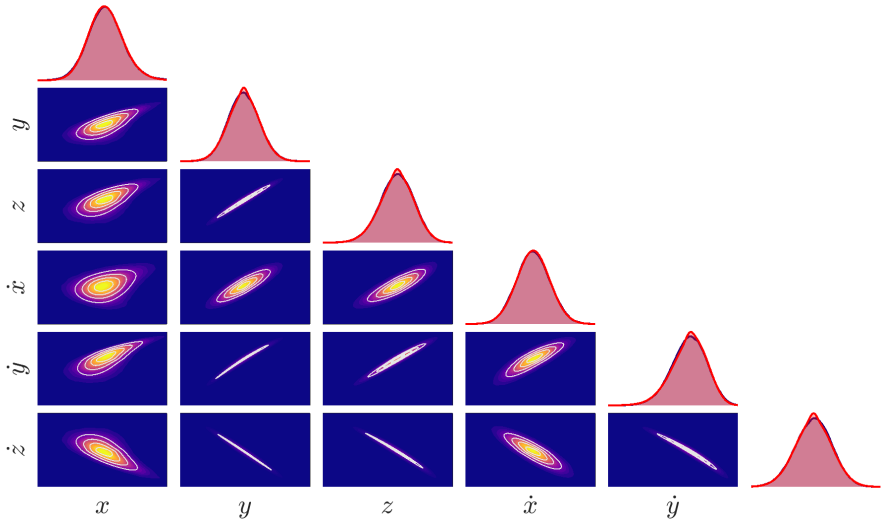

(b) Case 2

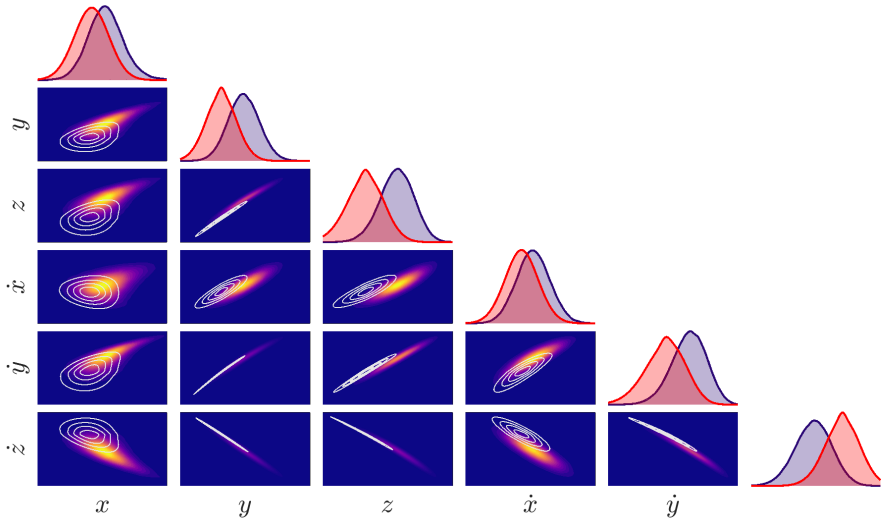

(c) Case 3

Fig. 4: Pair plot representations of propagated uncertainty for the 9:2 NRHO

This work investigates the application of bifidelty dynamics modeling for uncertainty propagation of space objects. To facilitate non-Gaussian uncertainties, an adaptive Gaussian mixture (GM) approach is employed that uses the Kullback–Leibler divergence to both detect the need for splitting and to formulating the splitting of Gaussian components into smaller components. The splitting direction is determined by solving a weighted optimization problem to find the uncertainty-weighted direction of maximum nonlinearity.

The directional-splitting approach for adaptive uncertainty propagation is applied to a representative, low-dimensional orbital dynamics problem. It is found that the performance of the proposed approach exceeds that of other adaptive GM methods. The proposed method produces more accurate predictions of the moments of the distribution and it does so while using fewer mixture components than similar methods. A more challenging problem of propagating the uncertainty along a nominal Near Rectilinear Halo Orbit (NRHO) is then considered. The bifidelity approach is applied, with the full fidelity used to propagate the means of the GM, and the reduced fidelity used to propagate the uncertainty. It is found that full fidelity is needed for the mean propagation, but that the reduced fidelity models are capable of representing the errors, i.e., the covariances. Combining multi-fidelity modeling with directional splitting in an adaptive GM uncertainty propagation framework is a viable method for long-term space object uncertainty propagation.

A quantitative analysis of the computing reduction requirements is forthcoming after algorithm optimization is carried out to ensure runtime predictions are meaningful. Future work will also investigate error bounds between the multiple fidelity levels employed in the method. Error bounds of this type can provide theoretical guarantees on the achievable precision, thereby informing the required level of fidelity used in the lower fidelity models.

VI. Acknowledgment

This work was partially supported by the Air Force Office of Scientific Research (AFOSR) under agreement number FA9550-23-1-0646, *Create the Future Independent Research Effort (CFIRE)*. The authors would like to thank the anonymous reviewers for their helpful feedback.

References

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

[2] P. Swerling, "First order error propagation in a stagewise smoothing procedure for satellite observations," *Journal of the Astronautical Sciences*, vol. 6, pp. 46–52, 1959.

[3] G. L. Smith, S. F. Schmidt, and L. A. McGee, "Application of statistical filter theory to the optimal estimation of position and velocity on board a circumlunar vehicle," NASA, Technical Report R-135, 1962.

[4] S. F. Schmidt, "Applications of state space methods to navigation problems," *in C. T. Leondes, Editor, Advanced Control Systems*, vol. 3, pp. 293–340, 1966.

[5] Y. C. Ho and R. Lee, "A Bayesian approach to problems in stochastic estimation and control," *IEEE Transactions on Automatic Control*, vol. 9, no. 4, pp. 333–339, 1964.

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

[7] D. L. Alspach and H. W. Sorenson, "Nonlinear Bayesian estimation using Gaussian sum approximations," *IEEE Transactions on Automatic Control*, vol. 17, no. 4, pp. 439–448, 1972.

[8] G. Terejanu, P. Singla, T. Singh, and P. D. Scott, "Adaptive Gaussian sum filter for nonlinear Bayesian estimation," *IEEE Transactions on Automatic Control*, vol. 56, no. 9, pp. 2151–2156, 2011.

[9] J. T. Horwood, N. D. Aragon, and A. B. Poore, "Gaussian sum filters for space surveillance: Theory and simulations," *Journal of Guidance, Control, and Dynamics*, vol. 34, no. 6, pp. 1839–1851, 2011.

[10] K. J. DeMars, R. H. Bishop, and M. K. Jah, "Entropy-based approach for uncertainty propagation of nonlinear dynamical systems," *Journal of Guidance, Control, and Dynamics*, vol. 36, no. 4, pp. 1047–1057, 2013.

[11] B. A. Jones and R. Weisman, "Multi-fidelity orbit uncertainty propagation," *Acta Astronautica*, vol. 155, pp. 406–417, 2019.

[12] B. Peherstorfer, K. Willcox, and M. Gunzburger, "Survey of multi-fidelity methods in uncertainty propagation, inference, and optimization," *SIAM Review*, vol. 60, no. 3, pp. 550–591, 2018.

[13] A. Fossà, R. Armellin, E. Delande, and F. Sanfedino, "Efficient multifidelity uncertainty propagation in the presence of process noise," *Journal of Guidance, Control, and Dynamics*, 2025, (to appear).

[14] T. Crain *et al.*, "Approach-phase precision landing with hazard relative navigation: Terrestrial test campaign results of the Morpheus/ALHAT project," in *Proceedings of the AIAA Guidance, Navigation, and Control Conference*, San Diego, CA, 2016, AIAA 2016-0099.

[15] A. H. Jazwinski, *Stochastic processes and filtering theory*. Courier Corporation, 1970.

[16] A. Gelb, *Applied Optimal Estimation*. Cambridge, Massachusetts: The M.I.T. Press, 1974.

[17] S. J. Julier and J. K. Uhlmann, "Unscented filtering and nonlinear estimation," *Proceedings of the IEEE*, vol. 92, no. 3, pp. 401–422, 2004.

[18] R. van der Merwe, "Sigma-point kalman filters for probabilistic inference in dynamic state-space models," Ph.D. dissertation, Oregon Health and Science University, 2004.

[19] S. Särkkä, *Bayesian Filtering and Smoothing*. Cambridge, 2013.

[20] K. Ito and K. Xiong, "Gaussian filters for nonlinear filtering problems," *IEEE Transactions on Automatic Control*, vol. 45, no. 5, pp. 910–927, 2000.

[21] I. Arasaratnam *et al.*, "Discrete-time nonlinear filtering algorithms using Gauss-Hermite quadrature," *Proceedings of the IEEE*, vol. 95, no. 5, pp. 953–977, 2007.

[22] S. Särkkä, "On unscented Kalman filtering for state estimation of continuous-time nonlinear systems," *IEEE Transactions on Automatic Control*, vol. 52, no. 9, pp. 1631–1641, 2007.

[23] S. Kullback and R. A. Leibler, "On information and sufficiency," *The Annals of Mathematical Statistics*, vol. 22, no. 1, pp. 79–86, 1951.

[24] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed. Wiley-Interscience, 2006.

[25] K. Tuggle, R. Zanetti, and C. D'Souza, "A Splitting Gaussian Mixture Formulation for a Nonlinear Measurement Update," in *Proceedings of the AAS/AIAA Space Flight Mechanics Meeting*, San Antonio, TX, Feb 5-9, 2017 2017, AAS 17-430.

[26] K. Tuggle and R. Zanetti, "Automated splitting gaussian mixture nonlinear measurement update," *Journal of Guidance, Control, and Dynamics*, vol. 41, no. 3, pp. 725–734, March 2018, doi: 10.2514/1.G003109.

[27] R. Zanetti, K. J. DeMars, D. Tuggle, K. Michaelson, and M. Gupta, "Uncertainty quantification using directional splitting and Gaussian mixture models with applications to orbital dynamics," in *Proceedings of the AAS/AIAA Astrodynamics Specialist Conference*, Boston, MA, Aug 10-14, 2025 2025, AAS 25-872.

[28] C. L. Schmid, "Generalization of polynomial chaos for estimation of angular random variables," Ph.D. dissertation, Missouri University of Science and Technology, 2020.

[29] C. P. Newman, J. R. Hollister, D. C. Davis, and E. M. Zimovan-Spreen, "Investigating solar radiation pressure modeling for operations in near rectilinear halo orbit," in *Astrodynamics Specialist Conference*, Charlotte, North Carolina, 2022.

[30] S. T. Scheuerle, "Low-energy lunar transfers in the bicircular restricted four-body problem," Ph.D. dissertation, Purdue University, West Lafayette, Indiana, 2024.