

NON-GAUSSIAN RECURSIVE BAYESIAN FILTERING FOR OUTER PLANETARY ORBILANDER NAVIGATION

Benjamin L. Hanson*, Aaron J. Rosengren†, Thomas R. Bewley‡, and Todd A. Ely§

Gaussian estimation filters have successfully aided spacecraft navigation for decades. However, future deep-space missions plan to operate orbilanders in unstable quasi-periodic orbits and perform low-altitude flybys of outer planetary moons. These complex trajectories may necessitate non-Gaussian filtering for accurate estimation over realistic measurement cadences. To mitigate the inherent risk associated with testing novel navigation software, non-Gaussian filters must be accurate, efficient, and robust. A novel Eulerian approach, Grid-based, Bayesian Estimation Exploiting Sparsity, along with the contemporary landscape of filters, are evaluated on these criteria through a Bayesian investigation, wherein the state uncertainty of a Saturn-Enceladus Distant Prograde Orbit is propagated.

I. INTRODUCTION

Navigation has long adhered to an established assumption that is often underappreciated: spacecraft (S/C) uncertainty remains Gaussian on the time scale of the measurement updates.¹ This assumption, while nontrivial, is operationally validated via frequent telecommunication with the S/C, dividing the nominal (nonlinear) trajectory into rectilinear segments. Adequate measurement cadences are assured for near-Earth S/C, but deep space missions pose advanced challenges. To thoroughly explore outer planetary systems, mission designers often leverage unstable, periodic orbits, the dynamics of which are often highly chaotic.^{2,3} Such chaoticity, dominated by third- and fourth-body perturbations, make linearizations accurate on much shorter time scales, thus requiring more frequent measurement updates and station-keeping maneuvers. However, in deep space regimes, ground-based communication has significant round-trip light-time delays, making communication with the Deep Space Network (DSN) potentially insufficient.

A complement to the DSN, Autonomous Navigation (AutoNav), an onboard optical-based navigation system,⁴ has been successfully utilized for multiple small body flyby missions, but application to missions beyond the asteroid belt prove more challenging due to the limited number of optical targets. Moreover, the S/C that utilized AutoNav followed near-linear, cruising trajectories, in contrast to the highly nonlinear tours conducted by outer planetary explorers like Cassini,⁵ Juno⁶ and Juice.⁷ Another such explorer, a proposed (now cancelled) Europa orbilander, planned to operate in a 5:6 resonance trajectory in the Jupiter-Europa system prior to approaching Europa for

*Graduate Student Researcher, University of California San Diego, La Jolla, CA, United States, blhanson@ucsd.edu

†Assistant Professor, University of California San Diego, La Jolla, CA, United States, arosengren@ucsd.edu

‡Professor, University of California San Diego, La Jolla, CA, United States, tbewley@ucsd.edu

§Principal Navigation Engineer, Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, United States, todd.a.ely@jpl.nasa.gov




Figure 1. The proposed Europa lander trajectory, where the leveraging maneuver results in a low-energy capture. Image sourced from [Hernandez-Doran et al., 2021](#).

either direct landing or capture.⁸ Figure 1 demonstrates the single leveraging maneuver required to insert into the L_2 Lyapunov orbit along the stable manifold. From here, an exit maneuver along the unstable manifold is performed to directly land on, or be captured by, Europa. However, from the leveraging maneuver to L_2 insertion, state uncertainty becomes highly non-Gaussian. For this study, orbit determination was performed via UD filter.⁹ To avoid divergence, the S/C planned to receive frequent measurement updates and perform small, statistical maneuvers at the most sensitive points of the orbit. To summarize, the mission design for the Europa orbilander was constrained by an insufficiently robust estimation strategy, in turn limiting the potential for scientific return.

The decision to operate S/C in unstable trajectories that would require new estimation techniques relies on the delicate balance of the potential for scientific discovery with the inherent added risk of instability. Historically, deep space S/C have flown on ballistic trajectories to reach their destination, transferring to stable, two-body orbits upon arrival. However, future missions to outer planetary systems, such as the Enceladus mission, conclude that the potential outweighs the risk. The Enceladus S/C is expected to operate in an unstable, near-rectilinear halo orbit (NRHO) and perform a flyby within 10 km of the surface to perform *in situ* analysis of the plumes of gas ejecting from the moon's surface.¹⁰ As such, it is likely that the instability of these conditions will cause the state uncertainty to become non-Gaussian faster than can be corrected by measurement updates.

Another orbit family of interest that leverages the chaotic dynamics of the circular restricted three-body problem (CR3BP) to efficiently explore the three-body system is the Distant Prograde Orbits (DPOs).¹¹ DPOs are planar, M_2 -centered, stable/unstable periodic orbits whose invariant manifolds provide heteroclinic connections between L_1 and L_2 Lyapunov orbits, meaning a substantial volume of the three-body system may be explored by forming chains of unstable periodic orbits connected by low-energy transfers.^{12,13} Figure 2 provides two trajectories in the Earth-Moon planar CR3BP (or PCR3BP), one where the unstable manifold of an L_1 orbit is connected directly to the stable manifold of an L_2 orbit and another where a DPO is chained between the two. Depending on the energy of the third-body, the DPO may be unstable, but operation there may be advantageous, as the L_1 –DPO– L_2 chain efficiently traverses different volumes of phases space with the option to idle for multiple orbits between L_1 and L_2 . Due to the instability, uncertainty propagation for this orbit family is best handled by non-Gaussian filters.

This investigation is arranged into six sections. Section I introduces the limitations of the current




Figure 2. (left) An $L_1 - L_2$ heteroclinic connection and (right) an $L_1 - \text{DPO} - L_2$ heteroclinic chain in the Earth-Moon PCR3BP. Image sourced from Gupta et al., 2021.

S/C navigation methods exposed by future deep-space missions. Section II presents a comprehensive analysis of the contemporary landscape of filtering methods, classifying the filters based on their respective methodologies. Section III provides rationale for the filters compared in this study, while also presenting a novel Bayesian approach to comparing the results of the approximations propagated by filters, an alternative to the frequentist approach that is often followed. Section IV defines the nonlinear dynamics and measurement models used in the Saturn-Enceladus DPO example, the numerical results of which are presented in Section V. Section VI discusses in-depth the results of the investigation and comments on the efficacy of the selected filters.

II. OVERVIEW OF THE LANDSCAPE OF RECURSIVE BAYESIAN FILTERS

Recursive Bayesian Filters (RBFs) are typically categorized as either *linear* or *nonlinear* based on their prediction step. However, the ‘nonlinear’ label can be misleading, as it inherently implies that the filter can adequately propagate uncertainty for nonlinear systems. The Unscented Kalman Filter (UKF),¹⁴ for instance, is often classified as a nonlinear filter due to its prediction step integrating the true, nonlinear dynamics of the system. While more effective than a linear KF which uses an *analytical* linearization for prediction, such as the Extended Kalman Filter (EKF) does, the UKF performs a *statistical* linearization at the correction step which limits its efficacy when uncertainty becomes non-Gaussian.^{15,16} Therefore, we deem it appropriate to more broadly label RBFs as either *Gaussian* or *non-Gaussian*, indicating whether they can represent uncertainty as non-Gaussian. This distinction will be utilized throughout the remainder of the paper.

The evolution of the probability density function (PDF) $p(\mathbf{x}, t)$ of a stochastic process $\mathbf{X}(t) \in \mathbb{R}^d$ governed by a combination of deterministic and random forces can be described by the Fokker-Planck equation:

$$\frac{\partial p(\mathbf{x}, t)}{\partial t} = - \sum_{i=1}^d \frac{\partial f_i(\mathbf{x}, t) p(\mathbf{x}, t)}{\partial x_i} + \frac{1}{2} \sum_{i=1}^d \sum_{j=1}^d \frac{\partial^2 Q_{ij}(\mathbf{x}, t) p(\mathbf{x}, t)}{\partial x_i \partial x_j} \quad (1)$$

where $\mathbf{x} = (x_1, \dots, x_d)$, $f_i(\mathbf{x}, t)$ is the i^{th} component of the deterministic system dynamics $\mathbf{f}$ at realization $\mathbf{x}$ and time t , and $Q_{ij}(\mathbf{x}, t)$ is the $(i, j)^{\text{th}}$ element of the stochastic state disturbances Q .

Recursive Bayesian Filters




Figure 3. A thorough categorization of the contemporary landscape of RBFs. Underlined labels are the methodologies, *italicized labels* further subdivide the filters within their groupings, and large labels represent the filters compared in this paper.

In general, Equation (1) is non-integrable, and $p(\mathbf{x}, t)$ cannot be described by a finite number of parameters. The optimal solution assimilates data via Bayes' theorem:

$$p(\mathbf{x}, t^{(k+)}) = \frac{p(\mathbf{y}^{(k)}|\mathbf{x}) p(\mathbf{x}, t^{(k-)})}{C}, \quad (2)$$

where $p(\mathbf{x}, t^{(k+)})$ is the *a posteriori*, $p(\mathbf{y}^{(k)}|\mathbf{x})$ is the measurement likelihood, $p(\mathbf{x}, t^{(k-)})$ is the *a priori*, and C is a normalization constant. Again, in general, the number of parameters necessary to represent $p(\mathbf{x}, t^{(k+)})$ is not finite. Thus, the goal of the RBF is to approximate and propagate the full PDF with as few parameters as possible while incorporating information from measurements updates. We now outline the different methodologies that shape the contemporary landscape of RBFs, depicted schematically in Figure 3.

Kalman Approach

In the Kalman approach, uncertainty is approximated as Gaussian, or as a mixture of Gaussians.¹ For linear systems with linear measurements, Gaussian uncertainty remains Gaussian globally. In this case, the optimal solution may be represented as the first and second central moments of a

Gaussian, $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$, where the d -dimensional PDF of a Gaussian is defined as

$$p(\mathbf{x} | \boldsymbol{\mu}; \boldsymbol{\Sigma}) = \mathcal{N}(\mathbf{x} | \boldsymbol{\mu}; \boldsymbol{\Sigma}) = \frac{1}{\sqrt{(2\pi)^d |\boldsymbol{\Sigma}|}} \exp \left(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1}(\mathbf{x} - \boldsymbol{\mu}) \right),$$

$$\text{i.e. } \boldsymbol{\mu} \triangleq \mathcal{E}\{\mathbf{x}\} = \int_{\mathbb{R}^d} \mathbf{x} \mathcal{N}(\mathbf{x} | \boldsymbol{\mu}; \boldsymbol{\Sigma}) d\mathbf{x}$$

$$\text{and } \boldsymbol{\Sigma} \triangleq \mathcal{E}\{(\mathbf{x} - \boldsymbol{\mu})(\mathbf{x} - \boldsymbol{\mu})^T\} = \int_{\mathbb{R}^d} (\mathbf{x} - \boldsymbol{\mu})(\mathbf{x} - \boldsymbol{\mu})^T \mathcal{N}(\mathbf{x} | \boldsymbol{\mu}; \boldsymbol{\Sigma}) d\mathbf{x}.$$

In the presence of nonlinearities, Gaussian filters are nonoptimal, but uncertainty may remain nearly Gaussian should the measurement update frequency be high relative to the speed with which $\mathbf{x}$ changes. Analytical linearizations of the dynamics and/or measurement models, *a la* the EKF or the UD filter, or statistical linearizations of the point distribution representation, *a la* the UKF or the Ensemble Kalman Filter (EnKF),¹⁷ are performed to approximate the full PDF as Gaussian. However, when measurement updates are sparse, uncertainty can become highly non-Gaussian, and Gaussian filters are no longer adequate. Mixture filters propagate sets of first and second moments, $\boldsymbol{\mu}_i$ and $\boldsymbol{\Sigma}_i$, representing non-Gaussian uncertainty as the weighted superposition of the Gaussian components:

$$p(\mathbf{x}) = \sum_{i=1}^M \alpha_i \mathcal{N}(\mathbf{x} | \boldsymbol{\mu}_i; \boldsymbol{\Sigma}_i) \quad \text{where} \quad \sum_{i=1}^M \alpha_i = 1 \quad \text{and} \quad \alpha_i \geq 0.$$

The standard of these mixture variants, the Gaussian Sum Filter (GSF),¹⁸ predicts and corrects each Gaussian component separately via the Kalman formulation. When the number of Gaussian components M is sufficiently high, the standard GSF provides a near-optimal approximation to the full PDF. A limitation of this filter is that the weights of the Gaussian components α_i are invariant. Adaptive entropy-based Gaussian mixture information synthesis (AEGIS)¹⁹ addresses this limitation via a splitting algorithm triggered by an entropy flag that monitors the nonlinearity of the dynamics. In this manner, the effect of nonlinearity is treated by systematically increasing the number of Gaussian components.

Lagrangian Approach

In the Lagrangian approach, uncertainty is modeled as an ensemble of particles and propagated via the evolution of the particles subject to the underlying dynamics.²⁰ In particular, the Monte Carlo (MC) method approximates the full PDF as a point-mass representation, i.e., a weighted superposition of Dirac delta functions:

$$p(\mathbf{x}) = \sum_{i=1}^M \alpha_i \delta(\mathbf{x} - \mathbf{x}_i) \quad \text{where} \quad \sum_{i=1}^M \alpha_i = 1 \quad \text{and} \quad \alpha_i \geq 0.$$

Standard MC assigns an equal weight ($\alpha_i = 1/M$) to each sample $\mathbf{x}_i$, but requires knowledge of the true distribution to sample from. To address this, Sequential Importance Sampling (SIS) algorithms assign weight based on an importance sampling distribution. Measurement updates are incorporated in the form of weight adjustments, but this often leads to sample degeneracy, where the majority of the probability is concentrated to a small number of samples. Particle filters use a resampling step to disperse probability across a new, focused ensemble; variants differ mainly in either their choice of importance sampling distribution or resampling method. Bootstrap Particle Filtering (BPF),²¹




Figure 4. 2D schematic depicting the notation of GBEES.

for instance, chooses the transition prior as the importance distribution. Another variation, Auxiliary Particle Filtering (APF),²² exploits the latest measurement in the prior prediction step to favor particles with higher likelihoods when sampling the *a priori*. Ensemble Gaussian Mixture Filtering (EnGMF),²³ a combination of a mixture filter and an SIS filter (thus the overlap in Figure 3), associates a Gaussian kernel with each particle. The Gaussian kernels are updated in the same manner as the GSF, and resampling is performed on the resultant Gaussian mixture.

Eulerian Approach

The Eulerian approach considers fixed points in space and evolves probability at said points.²⁰ Assuming $d\mathbf{X}(t)$ is near-deterministic, i.e. Q is relatively small (as is the case for most astrodynamic systems), Equation (1) is hyperbolic and satisfies the conservative form of the d -dimensional

advection equation:

$$\frac{\partial p(\mathbf{x}, t)}{\partial t} + \sum_{i=1}^d \frac{\partial f_i(\mathbf{x}, t) p(\mathbf{x}, t)}{\partial x_i} \approx 0, \quad (3)$$

The fluid mechanics community has devoted significant effort to developing sufficient methods for numerically discretizing and time-marching hyperbolic PDEs, but rarely have such methods been applied to uncertainty propagation. One such filter, Grid-based Bayesian Estimation Exploiting Sparsity (GBEES),²⁴ implements a Godunov-type finite volume method, a fully discrete, flux-differencing technique that uses a Taylor series approximation of Equation (3). For clarity, a two-dimensional example depicting how time derivatives are fully converted into spatial ones is provided:

$$\frac{p_{(i,j)}^{(k+1)} - p_{(i,j)}^{(k)}}{\Delta t} = - \frac{F_{(i+1/2,j)}^{(k)} - F_{(i-1/2,j)}^{(k)}}{\Delta x} - \frac{G_{(i,j+1/2)}^{(k)} - G_{(i,j-1/2)}^{(k)}}{\Delta y}, \quad (4)$$

where k represents the time step, Δx and Δy are the grid widths in the x - and y -directions respectively, $p_{(i,j)}^{(k)}$ is the probability at grid cell (i, j) at $t = k$, and fluxes F and G are defined at the interfaces of the grid, thus the half-step indexing (Figure 4 displays an example two-dimensional grid schematically). GBEES treats probability as a fluid, flowing the discretized PDF through phase space subject to the dynamics of the system. Time-marching high-dimensional, discretized PDFs is generally computationally infeasible, but GBEES exploits the fact that the PDF is near zero over most of phase space, thoroughly decreasing the computation cost of the algorithm. Through the inclusion of higher-order corrections and flux limiters, the filter is 2nd-order accurate and total variation diminishing. The filter avoids degeneracy by matching the uncertainty growth with the grid growth, thus maintaining resolution through measurement updates.

III. BAYESIAN FRAMEWORK FOR FILTER EVALUATION

Having extensively covered the landscape of RBFs, we now list those compared in this analysis, as well as the rationale for their inclusion. Lagrangian methods converge to the optimal solution as the number of particles approach infinity. Propagating a nearly-optimal number of particles may be infeasible (as well as ambiguous) for S/C estimation. Therefore, a large smooth bootstrap is used as a benchmark for both accuracy and efficiency (exact implementation of this model is discussed later in this section). The UKF is compared as a Gaussian control, as its nonlinear prediction step makes it more accurate than other linear, Gaussian filters. The non-Gaussian RBFs compared are the BPF, the EnGMF, and GBEES. The BPF is the conventional approach to particle filtering, while the other two methods are more state-of-the-art. Recent efforts have focused on reducing the computational cost of determining the optimal bandwidth parameter for the EnGMF samples.^{25,26} Similarly, GBEES has been significantly computationally optimized such that it may be feasibly applied to astrodynamics uncertainty propagation.²⁷ In spite of these improvements, none of these non-Gaussian filters have been used operationally as the primary method for S/C navigation. To achieve their practical application, non-Gaussian RBFs must be accurate enough to outperform their Gaussian counterparts in fast-changing nonlinear systems, robust enough to handle diverse measurement cadences while avoiding degeneracy, and efficient enough for real-time operation. In this investigation, the performances of the selected RBFs are evaluated based on these criteria. Specifically, accuracy is assessed on probability distribution similarity, degeneracy is assessed on uncertainty resolution maintenance, and efficiency is assessed on computation time.




Figure 5. Measurement update in truth model vs. UKF, BPF, EnGMF, and GBEES. Equal information is provided to the BPF, EnGMF, and GBEES (i.e., $M_{\text{particles}} = M_{\text{grid cells}}$).

Measurement update

The assimilation of a nonlinear measurement with a Gaussian *a priori* distribution may not result in a Gaussian *a posteriori*. A critical drawback of Gaussian RBFs is their limited ability to represent nonlinear measurement functions, especially when the resulting distribution deviates significantly from Gaussianity. In Figure 5, a nonlinear measurement update comparison is performed on a two-dimensional test problem. A highly-refined discretized implementation of Equation (2), used as a truth model, is compared with the results of the selected RBFs (UKF, BPF, EnGMF, and GBEES). The *a priori* distribution $p(\mathbf{x}, t^{(0-)})$ is a Gaussian with statistics

$$\boldsymbol{\mu} = \begin{bmatrix} -3.5 \\ 0 \end{bmatrix} \quad \text{and} \quad \boldsymbol{\Sigma} = \begin{bmatrix} 1 & 0.5 \\ 0.5 & 1 \end{bmatrix},$$

and the nonlinear measurement function is

$$y = h(\mathbf{x}, v) = |\mathbf{x}| + v, \quad \text{where} \quad v \sim \mathcal{N}(\mathbf{x}' | 0; R),$$

and $|\cdot|$ represents the l^2 -norm. In general, R is the measurement noise covariance matrix (in this case it is scalar), i.e., v is uncorrelated, zero-mean Gaussian noise. The measurement statistics are

$$y^{(0)} = 1 \quad \text{and} \quad R^{(0)} = 0.05.$$

The update step is performed for each model to obtain the *a posteriori* $p(\mathbf{x}, t^{(0+)})$. To ensure equivalent information is provided to the two non-Gaussian RBFs, we set the size of the ensemble for the BPF and EnGMF equal to the number of grid cells for GBEES. Qualitatively, it is obvious that the non-Gaussian filters return more accurate *a posteriori* than the UKF compared with the truth model. In fact, the truth model update step is identical to that of GBEES, just at a lower level of refinement, further emphasizing the fundamental nature of GBEES. Exact quantification of distribution similarity is now covered.

Distribution similarity metric

The frequentist approach often pursued in these types of analyses assumes that many Monte Carlo runs of the filter may be performed to calculate an average accuracy and consistency. Practically,

Algorithm 1 Calculate Bhattacharyya Coefficient of (P, Q)

Inputs: $P, Q : \mathbb{R}^d \rightarrow [0, 1], n > 1$

```
bounds  $\leftarrow [\infty \cdot \mathbf{1}_{d \times 1}, -\infty \cdot \mathbf{1}_{d \times 1}]$ 
grid  $\leftarrow \mathbf{0}_{d \times n}$  ▷ finding bounding region
for  $i = 1$  to  $d$  do
    bounds $[i, 1] \leftarrow \min \{b[i, 1], \min\{\mathbf{x}[i] \in \text{support}(P \cup Q)\}\}$ 
    bounds $[i, 2] \leftarrow \max \{b[i, 2], \max\{\mathbf{x}[i] \in \text{support}(P \cup Q)\}\}$ 
    grid $[i] \leftarrow \{\text{linear set from bounds}[i, 1] \text{ to bounds}[i, 2] \text{ of size } n\}$ 
end for

 $\chi \leftarrow \mathbf{0}_{n^d \times d}$ 
 $P^*, Q^* \leftarrow \mathbf{0}_{n^d \times 1}$ 
count  $\leftarrow 1$  ▷ discretizing  $P, Q$  over  $\chi$ 
for  $x_1 = 1$  to  $n$  do
     $\vdots$ 
    for  $x_d = 1$  to  $n$  do
         $\chi[\text{count}] \leftarrow \{\text{grid}[1, x_1], \dots, \text{grid}[d, x_d]\}$ 
         $P^*[\text{count}] \leftarrow P(\chi[\text{count}])$ 
         $Q^*[\text{count}] \leftarrow Q(\chi[\text{count}])$ 
        count  $\leftarrow \text{count} + 1$ 
    end for
     $\vdots$ 
end for
 $P^* \leftarrow P^* / \text{sum}(P^*)$ 
 $Q^* \leftarrow Q^* / \text{sum}(Q^*)$ 

BC  $\leftarrow 0$  ▷ calculating  $BC(P^*, Q^*)$ 
for  $i = 1$  to  $n^d$  do
    BC  $\leftarrow BC + \sqrt{P^*[i]Q^*[i]}$ 
end for
```

this may not be feasible due to time constraints. Instead, we perform a deterministic, Bayesian comparison of the filter-estimated distributions with the representative truth model. The Bhattacharyya coefficient (BC),²⁸ a metric of similarity of two probability distributions P and Q , is defined as

$$BC(P, Q) = \sum_{\mathbf{x} \in \chi} \sqrt{P(\mathbf{x})Q(\mathbf{x})} \quad \text{where } \chi \subseteq \mathbb{R}^d, \quad |\chi| < \infty, \quad \text{and } 0 \leq BC \leq 1; \quad (5)$$

$BC(P, Q) = 0$ represents no overlap in P and Q while $BC(P, Q) = 1$ represents perfect overlap in P and Q . This metric is applicable to the comparison of any discrete distributions, either Gaussian or non-Gaussian, making it well-suited for this application. The process of calculating BC for two probability distributions is summarized in Algorithm 1.

Algorithm 1 ensures that the $BC(P, Q)$ may be calculated in the cases where P and Q are either both continuous, both discrete, or one of each. A critical step is the creation of P^* and Q^* , as it ensures that the distributions are both unitary over χ . Utilizing this framework, a quantitative




Figure 6. Algorithm 1 applied to a two-dimensional test problem ($n = 10$ and $BC(P, Q) = 0.1595$).

measure is calculated for the two-dimensional test problem shown in Figure 6. P is taken to be a discrete point mass distribution, while Q is continuous Gaussian distribution. As P is discrete, nearest interpolation is utilized for calculating $P(x)$ when $x \notin \text{support}(P)$. As Q is continuous, a finite bounding region may be constrained to consider the 3σ -ellipse defined by the mean and covariance. A Cartesian grid χ is formed over the bounding region, and $BC(P, Q)$ is calculated.

This procedure is used to calculate a quantitative metric for the test problem from Figure 5. Figure 7 gives the BC for the UKF, BPF, EnGMF, and GBEES compared with truth distribution, both pre- and post-measurement update. As expected, the non-Gaussian filters outperform the UKF, and the EnGMF and GBEES outperform the BPF. This metric is used as a measure of accuracy for the remainder of the paper.

Uncertainty resolution metric

Particle degeneracy is a key drawback to Lagrangian methods, and a metric for measuring degeneracy is uncertainty resolution. For this investigation, uncertainty resolution is defined as

$$\text{Uncertainty resolution} = \frac{\text{Number of particles/grid cells}}{\text{Volume of uncertainty}}; \quad (6)$$

high uncertainty resolution represents a non-degenerate ensemble, while low uncertainty resolution represents a near-degenerate ensemble. Standard Lagrangian methods are generally either (a) degenerate over long propagation windows or (b) inefficient. In case (a), too few ensemble members are initialized at the start of the propagation period, and as uncertainty grows, resolution shrinks. At measurement update epochs, probability is concentrated to a few particles, and without substantial process noise, the repetition of this process results in degeneracy. In case (b), degeneracy is planned for by initializing too many ensemble members, with the expectation that the resolution will shrink to an acceptable level at the measurement update. This results in inefficiencies at the initial stages of the propagation period. This metric is used as a measure of degeneracy for the the non-Gaussian RBFs analyzed.

Truth implementation

Filters are often qualitatively compared to large MC simulations meant to represent point-mass approximations of the truth distribution. There are a couple problems with this. First, the standard MC implementation does not allow for measurement updates, and resampling directly from the




Figure 7. *BC* for the *a priori* and *a posteriori* distributions of the UKF, BPF, EnGMF, and GBEES compared with the truth model from Figure 5.

measurement likelihood results in loss of the *a priori*. Second, equally-weighted MC ensembles provide less information for comparison than weighted ensembles based on likelihood. To address both issues, we represent the truth distribution as a large smooth bootstrap.^{29,30} The steps for implementing the truth model are now explained in detail.

To begin, ensemble members $\mathbf{x}_i^{(k-1)}$ are sampled from the initial Gaussian *a priori* uncertainty $\mathcal{N}(\mathbf{x} | \boldsymbol{\mu}^{(k-1)}; \boldsymbol{\Sigma}^{(k-1)})$ and probability $P_i^{(k-1)}$ is calculated:

$$\mathbf{x}_i^{(k-1)} \sim \mathcal{N}(\mathbf{x} | \boldsymbol{\mu}^{(k-1)}, \boldsymbol{\Sigma}^{(k-1)}) \quad \text{and} \quad P_i^{(k-1)} = \mathcal{N}(\mathbf{x}_i^{(k-1)} | \boldsymbol{\mu}^{(k-1)}; \boldsymbol{\Sigma}^{(k-1)}),$$

where $\sum_{i=1}^M P_i^{(k-1)} = 1$. By attributing probability to the initial ensemble members, more accurate comparisons between filter-estimated and truth distributions may be made. Once initialized, the ensemble members are then propagated via the system dynamics up to a measurement update epoch $t^{(k)}$. At this point, the members are resampled from the estimated *a posteriori* $p(\mathbf{x}, t^{(k+)})$. Each *a posteriori* datum is treated as a Gaussian kernel with covariance $\boldsymbol{\Sigma}_*^{(k)}$ determined by Silverman's rule of thumb,³¹

$$\boldsymbol{\Sigma}_*^{(k)} = \left(\frac{4}{d+2} \right)^{\frac{2}{d+4}} M_*^{(k)-\frac{2}{d+4}} \boldsymbol{\Sigma}^{(k)}$$

where $M_*^{(k)}$ is the number of datum in $p(\mathbf{x}, t^{(k+)})$ and $\boldsymbol{\Sigma}^{(k)}$ is the covariance of $p(\mathbf{x}, t^{(k+)})$. New members are sampled from $p(\mathbf{x}, t^{(k+)})$ and zero-mean Gaussian noise $\mathcal{N}(\mathbf{x} | \mathbf{0}; \boldsymbol{\Sigma}_*^{(k)})$ is superimposed. From here, the new samples $\mathbf{x}_i^{(k)}$ are propagated until the next measurement update epoch.

Resampling from the estimated *a posteriori* ensures that information from the *a priori* is combined with the information from the measurement, while also reintroducing random Gaussian noise into the simulation. As GBEES implements Equation (2) most fundamentally, its estimated *a posteriori* is used as the model for resampling from at each measurement step. We highlight that, while sampling from the estimated distribution may introduce some bias towards GBEES, the measurements in this investigation have small uncertainties relative to the *a priori*, making the GBEES *a posteriori* distributions nearly identical to the likelihood distributions.




Figure 8. The discretized zero-velocity curves in the three-body problem.

IV. SYSTEM MODELS

Dynamics model

The state and orbital dynamics of a S/C in the PCR3BP in the synodic, non-dimensional frame are

$$\mathbf{x} = \begin{bmatrix} x \\ y \\ \dot{x} \\ \dot{y} \end{bmatrix} \quad \text{and} \quad \dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) = \begin{bmatrix} \dot{x} \\ \dot{y} \\ 2\dot{y} + x - \frac{(1-\mu)(x+\mu)}{r_1} - \frac{\mu(x-1+\mu)}{r_2} \\ -2\dot{x} + y - \frac{(1-\mu)y}{r_1} - \frac{\mu y}{r_2} \end{bmatrix}, \quad (7)$$

where $\mu = \mu_2/(\mu_1 + \mu_2)$ is the mass ratio, μ_i represents the gravitational parameter of body m_i , and r_i is the distance to body m_i .

Jacobi bounding. One integral of motion exists for the PCR3BP, the Jacobi constant, and is defined as

$$C = x^2 + y^2 + \frac{2(1-\mu)}{r_1} + \frac{2\mu}{r_2} + \mu(1-\mu) - \dot{x}^2 - \dot{y}^2; \quad (8)$$

because GBEES is a 2nd-order accurate numerical scheme, Equation (8) is not necessarily conserved. To compensate for this, the requirement can be hardcoded into the grid generation. The initial discretized PDF has a minimum and maximum C . As the grid grows in phase space, forbidden cells are those that fall outside of this bound and are not created. Admissible cells fall within the initial Jacobi bounds and are inserted into the grid as needed. Figure 8 presents an example of this cell differentiation for the zero-velocity curves, i.e. where $\dot{x} = \dot{y} = 0$. Using Jacobi bounding, the conservation of C is artificially ensured in GBEES.

Measurement model

The measurement model is

$$\mathbf{y} = \begin{bmatrix} \rho \\ \theta \\ \dot{\rho} \end{bmatrix} = \mathbf{h}(\mathbf{x}) = \begin{bmatrix} \sqrt{(x-1+\mu)^2 + y^2} \\ \tan^{-1}\left(\frac{y}{x-1+\mu}\right) \\ \frac{(x-1+\mu)\dot{x} + y\dot{y}}{\rho} \end{bmatrix}.$$




Figure 9. A Saturn-Enceladus DPO in the synodic frame; *(left)* position and *(right)* velocity.

where ρ is the range, θ is the azimuth angle, and $\dot{\rho}$ is the range-rate, all relative to m_2 . Measurements are taken to be the true state at epoch, transformed by $\mathbf{h}$ with zero-mean Gaussian noise.

V. NUMERICAL RESULTS

Trajectory properties

The selected RBFs are evaluated on a Saturn-Enceladus DPO with initial state

$$\mathbf{x}^{(0)} = \begin{bmatrix} 1.001471 & (\text{LU}) \\ -1.751810\text{e-}5 & (\text{LU}) \\ 7.198783\text{e-}5 & (\text{LU/TU}) \\ 1.363392\text{e-}2 & (\text{LU/TU}) \end{bmatrix} = \begin{bmatrix} 238879.876159 & (\text{km}) \\ -4.178575 & (\text{km}) \\ 9.079038\text{e-}4 & (\text{km/s}) \\ 1.719497\text{e-}2 & (\text{km/s}) \end{bmatrix};$$

propagating with the dynamics from Equation (7) results in the synodic frame and Saturn-barycentric inertial frame trajectories in Figure 9; JPL’s Mission analysis, Operations, and Navigation Toolkit Environment (MONTE)³² is used to convert between the two frames. Other properties of the trajectory are provided in Table 1.

Table 1. Saturn-Enceladus DPO properties

μ	LU (km)	TU (s)	C	SI	T (hr)
1.901110e-7	238529	18913	3.0 + 7.809821e-5	3.018700e+2	19.58109

Filter parameters

The filter parameters are provided in Table 2. For all filters, process noise is negligible, i.e., $Q = 0$. The range, azimuth, and range-rate uncertainty magnitudes are consistent with those derived

from optical limb measurements beyond the asteroid belt.³³ Note that the initial size $M^{(0)}$ of the BPF, EnGMF, GBEES are equal, while the size of the truth ensemble is an order of magnitude larger. The measurement cadence Δt_y evenly splits the propagation period into 4 segments, ending the simulation at the initial condition. As a rule of thumb, p_{thresh} is set such that $> 75\%$ of the initial grid cells exceed p_{thresh} , and the grid width vector Δ is set to half of the initial Cartesian standard deviation in all directions, or 10 km in position and 1e-3 km/s in velocity. The truth, UKF, BPF, and EnGMF are numerically integrated with Runge-Kutta 8(7).³⁴

Table 2. Saturn-Enceladus DPO filter parameters

Model	Measurement Uncertainty and Cadence	$M^{(0)}$	Misc.
Truth	$\sigma_\rho = 20 \text{ km}^*$	5e+5	Propagation scheme: RK8(7)
UKF	$\sigma_\theta = 1.74533\text{e-}2 \text{ rad}^*$	N/A	Propagation scheme: RK8(7), $\alpha = 1\text{e-}3$, $\beta = 2$, $\kappa = 0$
BPF	$\sigma_{\dot{\rho}} = 2.0\text{e-}3 \text{ km/s}^*$	28561	Propagation scheme: RK8(7)
EnGMF	$\Delta t_y = T/4 = 4.895 \text{ hrs}^*$	28561	Propagation scheme: RK8(7)
GBEES	*parameter applies to all models	28561	$p_{\text{thresh}} = 1\text{e-}7$, $\Delta = [10 \text{ km}, 10 \text{ km}, 1\text{e-}3 \text{ km/s}, 1\text{e-}3 \text{ km/s}]$




Figure 10. Saturn-Enceladus DPO true synodic state uncertainty.

Uncertainty propagation

With the framework set, the state uncertainty of the Saturn-Enceladus DPO is propagated using the selected RBFs with conditions and properties from Tables 1 and 2. Figure 10 displays the true uncertainty distribution propagated by the large, smooth bootstrap. While the MC members are 4D, the states are separated and plotted by position and velocity. Both position and velocity curves become highly non-Gaussian at the particular measurement cadence, as expected.




Figure 11. Saturn-Enceladus DPO synodic state uncertainty propagated by the UKF, BPF, EnGMF, and GBES.

The color changes throughout indicate when a measurement update has occurred, where the *a posteriori* $p(\mathbf{x}, t^{(k+)})$ is the resultant assimilation of the *a priori* $p(\mathbf{x}, t^{(k-)})$ and the likelihood distribution $p(\mathbf{y}^{(k)}|\mathbf{x})$. The distributions shown are not separated by equal time intervals; instead, they are spaced to optimize visualization.

For brevity, the results of the selected RBFs are presented in Figure 11. Qualitatively, these results may be compared with the truth model in Figure 10. The UKF implementation diverges in the prediction step after the first nonlinear measurement update, thus the absent PDFs beyond this epoch. The BPF implementation becomes highly degenerate after the three measurement updates, resulting in only 6 unique particles by the end of the propagation period. The EnGMF ensemble means, covariances, and weights are utilized to create isocontours by sampling points across the domain space and superimposing the probability from each member. The isocontour levels are set to $[0.68, 0.95, 0.997]$, *a la* the empirical rule. Although sparse, these contours at first glance align more consistently with the truth distributions.

The GBEES-propagated distribution is a discretized, 4D PDF, but integrating over the velocity- and position-spaces returns the 2D position and velocity PDFs, respectively. This is done via a numerical implementation of the following formulae:

$$p(x, y) = \int_{\Omega(\dot{x}, \dot{y})} p(\mathbf{x}) d\dot{x} d\dot{y} \quad \text{and} \quad p(\dot{x}, \dot{y}) = \int_{\Omega(x, y)} p(\mathbf{x}) dx dy.$$

The isocontours for GBEES are also set to the empirical rule values. The GBEES contours are generally more complete than the EnGMF, and align more closely with the true ensembles around epoch $t = t^{(3)}$. We now aim to quantitatively compare these results using the defined metrics.

Figure 12 displays the time history of the quantitative metrics of comparison for the UKF, BPF, EnGMF, and GBEES. The position/velocity error is taken to be the l^2 -norm of the difference between the weighed position/velocity mean of the truth distribution and the filter distribution. Table 3 lists the final position error, velocity error, and BC for each filter.

Table 3. Final position error, velocity error, and BC for selected RBFs

Model	Position Error (km)	Velocity Error (km/s)	BC
UKF	1.835640e+6	1.141540e+2	0.000
BPF	1.056313e+2	3.016343e-2	0.708
EnGMF	1.801216e+1	3.789827e-3	0.845
GBEES	1.414935e+1	3.960708e-3	0.937

The uncertainty resolution time histories of the non-Gaussian RBFs are compared in Figure 13. The volume of uncertainty from Equation 6 is calculated using a d -dimensional implementation of the convex hull³⁵ applied to the 4D truth ensembles from Figure 10. The volumes are then normalized by the initial volume at $t^{(0)}$, thus the initial uncertainty resolutions for each of the non-Gaussian RBFs at $t^{(0)}$ is equal to $M^{(0)}$. Due to the lack of process noise, for the BPF $M^{(k+1)} \leq M^{(k)}$, meaning each measurement update pushes the BPF towards degeneracy. Because the EnGMF resamples at each measurement epoch using the Gaussian sum filter update step, $M^{(k+1)} = M^{(k)}$. GBEES adaptively grows the discretized grid with the uncertainty, thus it maintains uncertainty resolution better throughout the propagation period.

Finally, the computational efficiency metric is normalized computation time, provided in Figure 14. All models are run in C or C++ on a single core of a 3.49 GHz Apple M2 Max chip. The




Figure 12. (*top*) Position error, (*middle*) velocity error, and (*bottom*) BC time histories for UKF, BPF, EnGMF, and GBEES implementations. Measurement update epochs are marked by vertical dotted lines.




Figure 13. Uncertainty resolution of the BPF, EnGMF, and GBEES implementations. Measurement update epochs are marked by vertical dotted lines.

filter computation times are normalized by the truth model computation time. For reference, the truth model takes ~ 4.00 hours to perform the ~ 19.58 -hour propagation. As expected, the UKF implementation is the fastest. The BPF is about $5\times$ faster than GBEES, and the EnGMF is slightly slower than the BPF due to the intricate update step, but is much more accurate with the same size of ensemble. To match the accuracy of GBEES using Lagrangian methods would require more members, though the exact number of members necessary is ambiguous. This highlights an issue with the Lagrangian approach; achieving specific accuracy thresholds requires an inefficient trial-and-error process. Although GBEES is the slowest among the filters analyzed, it is still nearly $4\times$ faster than the truth model, achieving a $BC = 0.937$ by the end of the propagation period. Depending on whether accuracy or efficiency is prioritized, either the EnGMF or GBEES may be utilized to effectively propagate non-Gaussian uncertainty.




Figure 14. Normalized computation time of the UKF, BPF, EnGMF, and GBEES implementations relative to the truth model.

VI. CONCLUSION

State-of-the-art S/C navigation algorithms approximate uncertainty as Gaussian. This approximation has sufficed for all missions hitherto due to frequent measurement correction via telecommunication. However, future missions plan to operate outer planetary orbilanders in unstable regimes where the state-of-the-art may be insufficient due to fast-changing state uncertainty and round-trip light-time delays; in these cases, S/C navigation may be enabled through the use of non-Gaussian filters. The contemporary non-Gaussian recursive Bayesian filters, while more accurate than their Gaussian counterparts, can be inefficient when applied to high-dimensional systems, require ambiguous resampling/splitting procedures to avoid degeneracies, or both. For practical onboard application, more robust solutions are required.

In this investigation, the capabilities of a fundamentally different non-Gaussian filter, Grid-based Bayesian Estimation Exploiting Sparsity, are compared to the contemporary landscape via a Bayesian inference framework. First, the measurement update step is demonstrated on a nonlinear measurement function for the UKF, BPF, EnGMF, and GBEES implementations. In the 2D test example, the UKF suffers from an inability to represent non-Gaussian uncertainty and the BPF faces major particle degeneracy, while the EnGMF and GBEES return accurate approximations of the truth *a posteriori*, with the results of GBEES being slightly more accurate when provided the same amount of information. The measurement update step is crucial to the success of the non-Gaussian filter, and this test example highlights the strengths and weaknesses of the selected filters.

In addition to showcasing the capabilities of GBEES, this paper presents an alternative method to quantifying the accuracy of filter-estimated distributions. Historically, a frequentist approach has been taken to determining accuracy, where filters are run numerous time and success is deemed a measure of the average position error. In these studies, a lack of emphasis is placed on matching full probability distributions, opting for first and second moment accuracy instead. For highly non-Gaussian distributions, these low-order moments do not convey complete information. The Bayesian approach followed here evaluates accuracy on full probability distribution similarity quantified by the Bhattacharyya coefficient. A smooth bootstrap represents the truth distribution, with resampling from estimated *a posteriori* performed after measurement updates for the purpose of incorporating prior knowledge. This model also serves as the baseline for efficiency, a key consideration when evaluating operational practicality.

Finally, the Bayesian framework is applied to the state uncertainty propagation of an unstable Saturn-Enceladus Distant Prograde Orbit. The selected recursive Bayesian filters propagate the initially Gaussian uncertainty for a full period, or about 19.58 hours, receiving nonlinear measurement updates every fourth of the period, or about 4.895 hours. The measurement function includes range, azimuth, and range-rate, and the magnitude of the measurement uncertainties reflect those expected from optical navigation measurements, as round-trip light-time delays from the Deep Space Network may render measurement corrections through this channel unreliable. The UKF implementation diverges shortly after performing the first measurement update. The BPF implementation estimates the uncertainty more accurately for longer than the UKF, but slowly succumbs to particle degeneracy with each measurement correction. The EnGMF utilizes a Gaussian sum update step with kernel density estimation to resample the mixture model from members of higher likelihood. This procedure disperses the probability more evenly across the distribution, making it an effective technique when process noise is low. Lastly, GBEES performs most consistently, returning a BC value near 1 for the entirety of the propagation, finishing the propagation period with position and velocity error nearly within 1σ of the measurement uncertainty. GBEES also demonstrates the capability of more aptly maintaining uncertainty resolution throughout the prediction step, while most Lagrangian methods tend towards particle degeneracy. While GBEES is more computationally intensive than the other filters compared in this paper, it is much more efficient than the large truth model used for comparison.

Future work will focus on enhancing the efficiency of GBEES, making it computationally competitive with other non-Gaussian filters when applied to six-dimensional systems. This may be achieved through parallelization, as many subprocesses within GBEES are embarrassingly parallelizable. To extend beyond low-fidelity models, effort is also being focused on embedding JPL's MONTE within the algorithm, such that acceleration may be sourced from ephemeris-quality models. These two advancements may enable onboard autonomous estimation and navigation for outer planetary orbilanders. GBEES may also serve as an accurate truth model for assessing when Gaussian filters are likely to diverge.

APPENDIX

GBEES is open-source and can be accessed at <https://github.com/bhanson10/GBEES>. The Saturn-Enceladus DPO example from this investigation is provided in the repository.

ACKNOWLEDGEMENTS

This investigation was supported by the NASA Space Technology Graduate Research Opportunities Fellowship (80NSSC23K1219). The research was carried out in part at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (80NM0018D0004).

REFERENCES

- [1] R. E. Kalman, "A New Approach to Linear Filtering and Prediction Problems," *Journal of Basic Engineering*, Vol. 82, No. 1, 1960, pp. 35–45.
- [2] S. Ross *et al.*, *Dynamical Systems, the Three-Body Problem, and Space Mission Design*. World Scientific, 2022.
- [3] S. Campagnola *et al.*, "Mission design for the exploration of Neptune and Triton," *IEEE Aerospace and Electronic Systems Magazine*, Vol. 30, No. 7, 2015, pp. 6–17.
- [4] S. Bhaskaran, "Autonomous navigation for deep space missions," *SpaceOps 2012*, p. 1267135, 2012.

- [5] F. Peralta *et al.*, “Cassini interplanetary trajectory design,” *Control Engineering Practice*, Vol. 3, No. 11, 1995, pp. 1603–1610.
- [6] S. J. Bolton *et al.*, “The juno mission,” *Space Science Reviews*, Vol. 213, 2017, pp. 5–37.
- [7] O. Grasset *et al.*, “JUPITER ICy moons Explorer (JUICE): An ESA mission to orbit Ganymede and to characterise the Jupiter system,” *Planetary and Space Science*, Vol. 78, 2013, pp. 1–21.
- [8] S. Hernandez-Doran *et al.*, “Navigating low-energy trajectories to land on the surface of Europa,” tech. rep., Jet Propulsion Laboratory, National Aeronautics and Space Administration, 2021.
- [9] G. J. Bierman, “Measurement updating using the UD factorization,” *Automatica*, Vol. 12, No. 4, 1976, pp. 375–382.
- [10] Z. Martins *et al.*, “Report of the Expert Committee for the Large-class mission in ESA’s Voyage 2050 plan covering the science theme “Moons of the Giant Planets”,” 2024.
- [11] M. Hénon, “Numerical exploration of the restricted problem, V,” *Astronomy and Astrophysics*, vol. 1, p. 223-238 (1969)., Vol. 1, 1969, pp. 223–238.
- [12] G. Mingotti *et al.*, “Mixed low-thrust invariant-manifold transfers to distant prograde orbits around Mars,” *Journal of guidance, control, and dynamics*, Vol. 33, No. 6, 2010, pp. 1753–1764.
- [13] M. Gupta *et al.*, “Earth-Moon multi-body orbits to facilitate cislunar surveillance activities,” *AIAA/AAS Astrodynamics Specialist Conference*, 2021, pp. 17–18.
- [14] S. J. Julier *et al.*, “Unscented filtering and nonlinear estimation,” *Proceedings of the IEEE*, Vol. 92, No. 3, 2004, pp. 401–422.
- [15] A. Gelb, *Applied optimal estimation*. MIT press, 1974.
- [16] B. Ristic *et al.*, *Beyond the Kalman filter : particle filters for tracking applications*. Artech House, 2004.
- [17] G. Evensen, “The ensemble Kalman filter: Theoretical formulation and practical implementation,” *Ocean dynamics*, Vol. 53, 2003, pp. 343–367.
- [18] D. Alspach *et al.*, “Nonlinear Bayesian estimation using Gaussian sum approximations,” *IEEE transactions on automatic control*, Vol. 17, No. 4, 1972, pp. 439–448.
- [19] K. J. DeMars *et al.*, “Entropy-based approach for uncertainty propagation of nonlinear dynamical systems,” *Journal of Guidance, Control, and Dynamics*, Vol. 36, No. 4, 2013, pp. 1047–1057.
- [20] G. K. Batchelor, *An introduction to fluid dynamics*. Cambridge university press, 2000.
- [21] J. V. Candy, “Bootstrap particle filtering,” *IEEE Signal Processing Magazine*, Vol. 24, No. 4, 2007, pp. 73–85.
- [22] M. K. Pitt *et al.*, “Filtering via simulation: Auxiliary particle filters,” *Journal of the American statistical association*, Vol. 94, No. 446, 1999, pp. 590–599.
- [23] J. L. Anderson *et al.*, “A Monte Carlo implementation of the nonlinear filtering problem to produce ensemble assimilations and forecasts,” *Monthly weather review*, Vol. 127, No. 12, 1999, pp. 2741–2758.
- [24] T. R. Bewley *et al.*, “Efficient grid-based Bayesian estimation of nonlinear low-dimensional systems with sparse. non-Gaussian PDFs,” *Automatica*, Vol. 48, No. 7, 2012, pp. 1286–1290.
- [25] S. Yun *et al.*, “Kernel-based ensemble gaussian mixture filtering for orbit determination with sparse data,” *Advances in Space Research*, Vol. 69, No. 12, 2022, pp. 4179–4197.
- [26] A. A. Popov *et al.*, “An adaptive covariance parameterization technique for the ensemble Gaussian mixture filter,” *SIAM Journal on Scientific Computing*, Vol. 46, No. 3, 2024, pp. A1949–A1971.
- [27] B. L. Hanson *et al.*, “State Estimation of Chaotic Trajectories: A Higher-Dimensional, Grid-Based, Bayesian Approach to Uncertainty Propagation,” *AIAA SCITECH 2024 Forum*.
- [28] K. Fukunaga, *Introduction to statistical pattern recognition*. Elsevier, 2013.
- [29] B. Efron, “Bootstrap methods: another look at the jackknife,” *Breakthroughs in statistics: Methodology and distribution*, pp. 569–593, Springer, 1992.
- [30] S. Wang, “Optimizing the smoothed bootstrap,” *Annals of the Institute of Statistical Mathematics*, Vol. 47, 1995, pp. 65–80.
- [31] B. W. Silverman, *Density estimation for statistics and data analysis*. Routledge, 2018.
- [32] S. Evans *et al.*, “MONTE: the next generation of mission design and navigation software,” *CEAS Space Journal*, Vol. 10, 2018, pp. 79–86.
- [33] S. B. Broschart *et al.*, “Kinematic approximation of position accuracy achieved using optical observations of distant asteroids,” *Journal of Spacecraft and Rockets*, Vol. 56, No. 5, 2019, pp. 1383–1392.
- [34] J. R. Dormand *et al.*, “A family of embedded Runge-Kutta formulae,” *Journal of computational and applied mathematics*, Vol. 6, No. 1, 1980, pp. 19–26.
- [35] C. B. Barber *et al.*, “The quickhull algorithm for convex hulls,” *ACM Transactions on Mathematical Software (TOMS)*, Vol. 22, No. 4, 1996, pp. 469–483.