

ANALYTICAL CONFIDENCE BOUNDARIES FOR NON-GAUSSIAN UNCERTAINTY IN PERTURBED SPACECRAFT DYNAMICS

Niccolò Michelotti*, Ethan R. Burnett[†] and Francesco Toppo[‡]

This work investigates nonlinear uncertainty propagation in perturbed astrodynamics, focusing on the rapid characterization of non-Gaussian distributions and the construction of three-dimensional “banana-shaped” confidence boundaries. To bridge the gap between computationally intensive high-fidelity methods and inaccurate linear approximations, this paper introduces a fully analytical, sample-free framework for higher-order moments extraction. Leveraging Differential Algebra to bypass repeated numerical integration, statistical moments are extracted analytically via Isserlis’ theorem and a monomial-to-Hermite basis transformation. A pair-product projection strategy is exploited to overcome the severe computational bottleneck of full fourth-order tensor contractions and compute only relevant terms via efficient polynomial algebra. The extracted skewness and kurtosis components directly parameterize non-elliptical confidence geometries that capture spatial bending and out-of-plane coupling of typical non-Gaussian distributions in astrodynamics. The approach is validated in high-fidelity environments including a cislunar Near-Rectilinear Halo Orbit and close-proximity trajectories around Apophis during Earth’s flyby, where the analytical approach achieves geometric accuracy comparable to expensive Monte Carlo simulations while reducing computational runtime by orders of magnitude.

1 INTRODUCTION

Reliable and efficient uncertainty quantification (UQ) is a persistent challenge in modern astrodynamics. As mission architectures increasingly target highly nonlinear and weakly stable regimes, capturing the true evolution of state uncertainty becomes critical for the robustness and efficiency of the mission. In cislunar environments, such as Near-Rectilinear Halo Orbits (NRHOs), spacecraft are subjected to sensitive multi-body gravitational fields and repeated low-altitude perilune passages.¹ Similarly, close-proximity operations around small bodies, such as the asteroid Apophis, are dominated by irregular gravity fields and rapid dynamical timescales. In these environments, initial Gaussian state uncertainties quickly deform into non-Gaussian distributions driven by strong localized gravitational gradients and resonances.² Consequently, standard geometric descriptions based uniquely on the mean and covariance fail to capture relevant features of the distribution. Accurately bounding the spacecraft’s state is a fundamental requirement for robust trajectory optimization and space domain awareness. In the presence of non-Gaussian distributions, this requires the

*PhD Student, Department of Aerospace Science and Technology, Politecnico di Milano, via La Masa, 34, Milan, 20156, Italy. niccolo.michelotti@polimi.it

[†] Assistant Professor, Department of Aerospace Engineering and Engineering Mechanics, University of Texas at Austin, 2617 Wichita St., Austin TX, 78712, USA. ethan.burnett@utexas.edu.

[‡] Full Professor, Department of Aerospace Science and Technology, Politecnico di Milano, via La Masa, 34, Milan, 20156, Italy. francesco.toppo@polimi.it. AIAA Senior Member.

evaluation of higher-order statistical moments, such as skewness and kurtosis, to properly represent the asymmetry and tail-stretching of the probability density function.

Historically, Monte Carlo (MC) methods have served as the highest-fidelity benchmark for capturing these non-Gaussian features. However, their reliance on thousands of numerical integrations makes them computationally expensive for rapid and preliminary analyses, large-scale trajectory design campaigns, or onboard applications for autonomous spacecraft. Conversely, linear covariance (LinCov) propagation is computationally efficient but fundamentally misrepresents uncertainty distribution in highly nonlinear flows, often underestimating the actual confidence boundaries.³ To balance accuracy and computational cost, higher-fidelity nonlinear propagation techniques have been investigated thoroughly.⁴ Advanced techniques such as Conjugate Unscented Transform (CUT),⁵ Polynomial Chaos Expansion (PCE)⁶ and Gaussian Mixture Models (GMMs)⁷ have been implemented in several applications for astrodynamics, orbital determination and conjunction assessment. Nevertheless, all these methods still require repeated numerical evaluations of the full dynamics to compute local expansions. For repetitive UQ queries, this repeated sampling remains a relevant computational limitation.

To circumvent the need for repeated numerical integration, Differential Algebra (DA)⁸ and State Transition Tensors (STTs)⁹ have emerged as effective tools for orbital uncertainty propagation. By computing high-order Taylor expansions of the dynamical flow map, the complex numerical integration is replaced by the evaluation of a polynomial map. When combined with mathematical identities such as Isserlis’ theorem (or Wick’s probability theorem),¹⁰ DA enables the sample-free extraction of statistical moments directly from the expansion coefficients.¹¹ However, the extraction of high-order moments inherently demands the construction and contraction of dense high-dimensional tensors. This combinatorial growth becomes particularly relevant starting from the fourth-order moment (kurtosis), introducing an $\mathcal{O}(N^4)$ computational and memory cost, limiting the scalability of the method for higher-order spatial reconstructions in full 6-degree-of-freedom state spaces.

This work proposes a fully analytical sample-free optimized framework for non-Gaussian uncertainty approximation and three-dimensional confidence boundary reconstruction. First, the paper implements and compares sample-free methods for extracting higher-order moments from DA maps, highlighting the specific sparsity advantages of transforming the standard monomial basis into an orthogonal probabilist’s Hermite basis. Second, to address the $\mathcal{O}(N^4)$ tensor limitation during kurtosis extraction, a “project-then-product” strategy is introduced. By projecting the maps onto principal bending axes and exploiting the orthogonality of the Polynomial Chaos basis, the computational complexity is reduced to one-dimensional polynomial algebra. Finally, the analytically derived skewness and kurtosis are utilized to reconstruct non-elliptical, three-dimensional confidence geometries, extending a previous derivation only provided in two-dimensional slices of the problem variables.¹² This “banana” approximation captures the complex spatial bending and out-of-plane coupling induced by highly perturbed astrodynamical environments. This or analogous analytical methods are best used for non-Gaussian distributions where linear covariance methods are no longer valid, but the distribution remains unimodal and easily parameterizable. Whereas this paper focuses on the efficient real-time computation of higher moments for a non-Gaussian confidence contour, Ref. 13, presented at the same conference, focuses on stochastic guidance applications of the analytic non-Gaussian confidence contour concept. However, that work does not employ differential algebra nor the sample-free methods developed here.

The remainder of this paper is organized as follows. Section 2 outlines the formulation of the 3D

banana-shaped confidence boundary reconstruction. Section 3 details the mathematical framework for analytical higher-order moment extraction, including the DA-Isserlis formulation, the Hermite basis transformation, and the pair-product strategy. Section 4 presents numerical results validating the proposed framework in two representative environments: a cislunar NRHO and close-proximity trajectories around the asteroid Apophis during its upcoming Earth flyby. Results highlight the trade-offs between propagation accuracy and computational runtime of the proposed approach with the performance of state-of-the-art methods. Finally, Section 5 concludes the paper.

2 CONFIDENCE BOUNDARY RECONSTRUCTION

To motivate the development of the analytical moment-extraction framework presented in Section 3, we first outline the target geometric application: the analytical boundary reconstruction. This formulation leverages the higher-order statistical moments (skewness and kurtosis) to deform the principal axes of the local covariance ellipsoid, explicitly capturing the non-Gaussian stretching and folding of the propagated uncertainty. The foundational formulation of this theory was presented by Ref. 12 to approximate two-dimensional “banana-shaped” distributions commonly found in astrodynamics. This section details the reconstruction process and formalizes its extension to three-dimensional spatial volumes to capture out-of-plane nonlinear couplings.

Once the covariance and higher-order moments are computed, the 3D boundary reconstruction proceeds through the following steps:

1. Eigendecomposition and Axis Sorting:

To apply the geometric deformations correctly, a local orthogonal frame must be established along the principal axes of the uncertainty distribution. This is achieved via the eigenvalue decomposition of the 3×3 spatial covariance matrix, $\mathbf{P} = \mathbf{V}\mathbf{\Sigma}\mathbf{V}^\top$. The eigenvectors are sorted based on their associated variance and higher-order couplings. The eigenvector corresponding to the largest eigenvalue (maximum variance) is defined as the primary “length” direction. The secondary “width” axis is selected as the eigenvector exhibiting the strongest skewness coupling (cross-skewness) with the length direction, while the remaining orthogonal eigenvector corresponds to the “thickness” direction. Let $\mathbf{u}$, $\mathbf{v}$, and $\mathbf{w}$ denote the local whitened basis vectors corresponding to these length, width, and thickness directions, respectively.

2. Moment Projection:

The geometric deformation of the boundary is governed by the projection of the skewness and kurtosis tensors onto this local frame. The relevant coupling quantities are mathematically defined as the tensor projections:

$$E_{uuu} := E[\mathbf{u}^3] = \sum_{i,j,k} u_i u_j u_k \mathcal{M}_{ijk}^{(3)} \quad (1a)$$

$$E_{uuv} := E[\mathbf{u}^2 \mathbf{v}] = \sum_{i,j,k} u_i u_j v_k \mathcal{M}_{ijk}^{(3)} \quad (1b)$$

$$E_{uww} := E[\mathbf{u}^2 \mathbf{w}] = \sum_{i,j,k} u_i u_j w_k \mathcal{M}_{ijk}^{(3)} \quad (1c)$$

$$E_{uuuu} := E[\mathbf{u}^4] = \sum_{i,j,k,l} u_i u_j u_k u_l \mathcal{M}_{ijkl}^{(4)} \quad (1d)$$

where E_{uuu} captures the lengthwise asymmetry, E_{uuv} and E_{uuw} capture the in-plane and out-of-plane bending, and E_{uuuu} quantifies the lengthwise elongation. In standard practice, evaluating these specific projections requires the costly construction of the full dense tensors $\mathcal{M}^{(3)}$ and $\mathcal{M}^{(4)}$. Bypassing this computational bottleneck is the primary focus of the mathematical framework developed in Section 3, which computes these exact quantities directly via a highly efficient 1D "project-then-product" strategy.

3. 3D Shape Corrections:

In the local whitened frame, the undeformed confidence boundary is initialized as a sphere parameterized by standard spherical coordinates:

$$u = k \cos \theta \sin \phi, \quad v = k \sin \theta \sin \phi, \quad w = k \cos \phi \quad (2)$$

where k is the desired confidence level parameter (e.g., $k = 3$ for a 3σ equivalent boundary).

The first correction applies an in-plane bending deformation, driven by the cross-skewness E_{uuv} (coupling length and width) and moderated by the kurtosis E_{uuuu} (lengthwise elongation). This is modeled as a quadratic deformation along the width direction:

$$\Delta v_b = \alpha(u^2 - 1), \quad \alpha = \frac{E_{uuv}}{E_{uuuu} - 1} \quad (3)$$

The second correction introduces an asymmetry along the primary length direction to account for the heavy-tailed nature of the distribution. This shift is derived from the Cornish-Fisher expansion:¹⁴

$$\Delta u_{CF} = c_k \frac{u^2}{k^2}, \quad c_k = \frac{1}{6} E_{uuu}(k^2 - 1) \quad (4)$$

The third correction constitutes the extension to the three-dimensional case. Highly perturbed astrodynamical environments often induce spatial torsion, requiring an out-of-plane bending correction driven by E_{uuw} , which captures the coupling between the length and thickness directions. It is computed analogously to the in-plane deformation:

$$\Delta w_b = \beta(u^2 - 1), \quad \beta = \frac{E_{uuw}}{E_{uuuu} - 1} \quad (5)$$

By superimposing these corrections onto the base sphere (Eq. (2)), the deformed local coordinates (u', v', w') are obtained:

$$u' = u + \Delta u_{CF}, \quad v' = v + \Delta v_b, \quad w' = w + \Delta w_b \quad (6)$$

4. Global Frame Transformation:

Finally, the boundary is transformed from the local standard normal space back into the physical state space. The deformed coordinates are scaled by the principal standard deviations ($\Sigma^{1/2}$), rotated back to the global frame using the sorted eigenvector matrix $\mathbf{V}$, and translated by the propagated mean vector $\boldsymbol{\mu}$:

$$\mathbf{x}_{bound} = \boldsymbol{\mu} + \mathbf{V} \Sigma^{1/2} \begin{bmatrix} u' \\ v' \\ w' \end{bmatrix} \quad (7)$$

The outlined approach yields a continuous, analytical 3D confidence boundary that reliably encompasses the true uncertainty geometry in nonlinear regimes, circumventing the limitations of standard ellipsoidal approximations. Furthermore, this geometric framework is inherently modular. While the present implementation relies on quadratic bending and first-order Cornish-Fisher expansions, the formulation can theoretically be extended to incorporate alternative analytical deformation models. By superimposing additional shape corrections directly into the local coordinate frame, more complex topological features could be captured without altering the underlying moment-extraction architecture.

3 ANALYTICAL HIGHER-ORDER MOMENT EXTRACTION

This section details the analytical framework developed to estimate the higher-order moments of propagated uncertainty distributions. The proposed approaches bypass any sampling step and extract statistical moments directly from the algebraic structure of the nonlinear flow maps.

Transitioning from sample-based evaluations to purely analytical formulations offers advantages for both uncertainty quantification and trajectory design. Traditional sampling techniques are affected by statistical noise, slow $\mathcal{O}(1/\sqrt{N})$ convergence rates of MC simulations, and the introduction of non-smooth stochastic gradients. These artifacts are highly problematic for gradient-based robust optimization solvers, which rely on smooth, continuous design spaces to satisfy exact optimality conditions. Furthermore, deterministic moment extraction guarantees exact reproducibility and facilitates the formulation of rigorous chance constraints for safety-critical spaceflight operations.¹³

To achieve these benefits without incurring the severe memory overhead typically associated with dense multidimensional tensors, three analytical formulations of increasing computational efficiency are presented below. First, the fundamental combination of Differential Algebra (DA) techniques and Isserlis’ theorem is presented. Next, a Monomial-to-Hermite basis transformation is introduced to increase the sparsity of the expected value tensors. Finally, a “project-then-product” strategy is derived, enabling the direct computation of only the strictly relevant terms of the higher-order moment tensors.

3.1 Differential Algebra Maps and Isserlis’ Theorem

Consider a general nonlinear dynamical system governed by the ordinary differential equation:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, t) \quad (8)$$

where $\mathbf{x}$ is the state vector. Given a nominal initial state $\mathbf{x}_0$ and an initial deviation $\delta\mathbf{x}_0$, the state at a future time t is defined by the flow map:

$$\mathbf{x}(t) = \phi(t; \mathbf{x}_0 + \delta\mathbf{x}_0, t_0) \quad (9)$$

In standard Monte Carlo methods, this equation requires repeated numerical integration for each sampled initial condition. Differential Algebra provides a mathematical framework to bypass this limitation, substituting integrations of thousands of samples with inexpensive Taylor polynomial evaluations. Initializing the state as a DA variable yields a high-order Taylor expansion of the flow map with respect to the initial deviations:

$$\delta\mathbf{x}(t) = \sum_{|\mathbf{p}| \leq k} \mathbf{M}_{\mathbf{p}}(t) \delta\mathbf{x}_0^{\mathbf{p}} \quad (10)$$

where $M_{\mathbf{p}}(t)$ are the deterministic Taylor coefficients up to order k , and $\mathbf{p}$ is a multi-index representing the variable exponents. To evaluate the statistical moments of the propagated state, let the initial uncertainty be described by a Gaussian random vector $\delta \mathbf{x}_0 \sim \mathcal{N}(\mathbf{0}, \mathbf{P}_0)$. If this correlated deviation is utilized directly, computing higher-order moments requires the general form of Isserlis' theorem (or Wick's probability theorem).¹⁰ The general theorem states that the expected value of the product of zero-mean correlated Gaussian variables is the sum over all possible distinct ways of partitioning the indices into unordered pairs of their covariances:¹¹

$$\mathbb{E}[\delta x_{d_1} \dots \delta x_{d_k}] = \begin{cases} 0 & \text{if } k \text{ is odd} \\ \sum_{\mathbf{p} \in \mathcal{P}} \prod_{\{u,v\} \in \mathbf{p}} (\mathbf{P}_0)_{d_u d_v} & \text{if } k \text{ is even} \end{cases} \quad (11)$$

To simplify the combinatorial evaluation of Eq. (11), a whitening process is performed. The initial deviations are mathematically decorrelated prior to the DA expansion using the Cholesky decomposition $\mathbf{P}_0 = \mathbf{L}\mathbf{L}^T$. This defines a set of independent standard normal variables $\boldsymbol{\xi} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, such that $\delta \mathbf{x} = \mathbf{L}\boldsymbol{\xi}$. The DA engine then maps this linear transformation through the dynamics to yield a high-order Taylor expansion of the final state component x_i entirely within the whitened standard normal space:

$$x_i(\boldsymbol{\xi}) = \sum_{|\mathbf{p}| \leq k} c_{i,\mathbf{p}} \boldsymbol{\xi}^{\mathbf{p}} \quad (12)$$

where $c_{i,\mathbf{p}}$ are the deterministic map coefficients evaluated via DA, $\mathbf{p} \in \mathbb{N}^m$ is a multi-index vector of exponents, and $\boldsymbol{\xi}^{\mathbf{p}} = \prod_{d=1}^m \xi_d^{p_d}$ forms the standard monomial basis.

Because the variables in $\boldsymbol{\xi}$ are independent standard normals, their covariance matrix is the identity matrix $\mathbf{I}$, meaning $\mathbb{E}[\xi_i \xi_j] = \delta_{ij}$. When Isserlis' theorem is applied in this whitened space, any combinatorial pairing in the summation that links two distinct dimensions ($i \neq j$) evaluates to zero and vanishes. The only surviving terms are those that pair a variable strictly with itself ($\mathbb{E}[\xi_i^2] = 1$). For a univariate term ξ_d^n , the number of unique pairings of n identical items is given by the double factorial, resulting in a straightforward 1D expectation:

$$\mathbb{E}[\xi_d^n] = \begin{cases} 0 & \text{if } n \text{ is odd} \\ (n-1)!! & \text{if } n \text{ is even} \end{cases} \quad (13)$$

Furthermore, due to the independence of the components, the expectation of any multidimensional monomial basis term reduces to the product of its univariate expectations:

$$\mathbb{E}[\boldsymbol{\xi}^{\mathbf{p}}] = \mathbb{E} \left[\prod_{d=1}^m \xi_d^{p_d} \right] = \prod_{d=1}^m \mathbb{E}[\xi_d^{p_d}] \quad (14)$$

By applying the linearity of the expectation operator to the DA expansion in Eq. (12), the raw statistical moments of the propagated state can be evaluated analytically. The first two raw moments are obtained as:

$$\mu_i = \mathbb{E}[x_i(\boldsymbol{\xi})] = \sum_{|\mathbf{p}| \leq k} c_{i,\mathbf{p}} \mathbb{E}[\boldsymbol{\xi}^{\mathbf{p}}] \quad (15a)$$

$$\mathcal{R}_{ij}^{(2)} = \mathbb{E}[x_i(\boldsymbol{\xi}) x_j(\boldsymbol{\xi})] = \sum_{|\mathbf{p}| \leq k} \sum_{|\mathbf{q}| \leq k} c_{i,\mathbf{p}} c_{j,\mathbf{q}} \mathbb{E}[\boldsymbol{\xi}^{\mathbf{p}+\mathbf{q}}] \quad (15b)$$

where the combined expectation term $\mathbb{E}[\xi^{p+q}]$ is evaluated using the double factorial property defined in Eqs. (13) and (14). The central covariance matrix is then recovered as $P_{ij} = \mathcal{R}_{ij}^{(2)} - \mu_i \mu_j$.

From Eq. (15), it is possible to generalize the moment extraction process. The evaluation of any n^{th} order raw moment requires an n -fold summation over the product of n map coefficients, weighted by the expectation of the combined monomial exponents. To formalize this, a generic n^{th} order *weight tensor* $\mathcal{W}^{(n)}$ is defined to encode the statistical expectations of the combined monomial basis terms:

$$\mathcal{W}_{p_1 \dots p_n}^{(n)} = \mathbb{E}[\xi^{p_1 + \dots + p_n}] \quad (16)$$

Because these weight tensors depend exclusively on the dimensionality m and the expansion order k , they are independent of the orbital dynamics and can be computed offline. Consequently, the evaluation of the third and fourth raw moments ($\mathcal{R}^{(3)}$ and $\mathcal{R}^{(4)}$) becomes a set of tensor contractions between the DA coefficients c and the precomputed weight tensors:

$$\mathcal{R}_{abc}^{(3)} = \sum_{p,q,r} c_{a,p} c_{b,q} c_{c,r} \mathcal{W}_{pqr}^{(3)}, \quad \mathcal{R}_{abcd}^{(4)} = \sum_{p,q,r,s} c_{a,p} c_{b,q} c_{c,r} c_{d,s} \mathcal{W}_{pqrs}^{(4)} \quad (17)$$

Finally, the required central moments of skewness ($\mathcal{M}^{(3)}$) and kurtosis ($\mathcal{M}^{(4)}$) are recovered from these raw multidimensional tensors via standard binomial expansions centered around the propagated mean.¹¹

3.2 Monomial-To-Hermite Basis Transformation

While the Isserlis evaluation of the monomial basis is exact, computational efficiency can be improved by transforming the DA map into an orthogonal probabilist's Hermite polynomial basis. This derives an analytical Polynomial Chaos Expansion (PCE) directly from the DA map, bypassing the need for sampling or regression techniques.¹⁵

Because the map in Eq. (12) is already expressed in terms of independent standard normal variables ξ , any 1D monomial ξ_d^n can be mapped exactly to a linear combination of Hermite polynomials $He_k(\xi_d)$ of the same parity:¹⁶

$$\xi_d^n = \sum_{k \in \{n, n-2, \dots\}} w_{n,k} He_k(\xi_d) \quad (18)$$

where the projection weights $w_{n,k}$ are given by:

$$w_{n,k} = \frac{n!}{k! 2^{\frac{n-k}{2}} (\frac{n-k}{2})!} \quad (19)$$

Applying this mapping to the multivariate case, the output map is rewritten as an analytical Hermite PCE:

$$x_i(\xi) = \sum_{\alpha} c_{i,\alpha}^{\text{PCE}} \Psi_{\alpha}(\xi) \quad (20)$$

where $\Psi_{\alpha}(\xi) = \prod_{d=1}^m He_{\alpha_d}(\xi_d)$ are the multivariate Hermite polynomials, and $c_{i,\alpha}^{\text{PCE}}$ are the analytically transformed coefficients.

A key property of this basis is its orthogonality with respect to the standard normal Gaussian measure. The squared norm of a multivariate basis polynomial Ψ_{α} evaluates directly to the factorial

of its multi-index:

$$\langle \Psi_\alpha, \Psi_\alpha \rangle = \mathbb{E}[\Psi_\alpha^2] = \alpha! = \prod_{d=1}^m \alpha_d! \quad (21)$$

This PCE framework is effective for moment extraction because it leverages this basis orthogonality. For instance, the expectation of the triple product of 1D Hermite polynomials evaluates to exactly zero unless strict degree parity and triangle-inequality conditions are met:¹⁷

$$\mathbb{E}[He_u(\xi)He_v(\xi)He_w(\xi)] = \begin{cases} 0 & \text{if } u+v+w \text{ is odd, or } s < \max(u, v, w) \\ \frac{u!v!w!}{(s-u)!(s-v)!(s-w)!} & \text{otherwise} \end{cases} \quad (22)$$

where $s = \frac{1}{2}(u+v+w)$. Evaluating higher-order moments involves substituting the PCE coefficients c^{PCE} and the corresponding Hermite weight tensors $\mathcal{W}_{\text{PCE}}$ into the generalized contraction equations (Eq. (17)).

While the triple product efficiently populates the third-order tensor $\mathcal{W}_{\text{PCE}}^{(3)}$ for skewness, evaluating kurtosis requires the expectation of quadruple products to populate $\mathcal{W}_{\text{PCE}}^{(4)}$. Although analytical formulas exist for quadruple products, the combinatorial size of the resulting full tensor remains computationally demanding in high-dimensional spaces. Nevertheless, the Hermite weight tensors are significantly sparser than their monomial counterparts. For a 6D state and a 3rd-order DA expansion, the fill factor of the weight tensors (i.e., the percentage of non-zero elements) is approximately 0.6% in the Hermite basis, compared to 2.5% in the monomial basis. This sparsity reduces the computational cost of the tensor contractions, but the explicit construction of the full fourth-order tensor motivates further optimization, as addressed in the following section.

3.3 Pair-product Strategy for Projected Moments

Generalized tensor contractions provide a complete statistical characterization of the propagated state; however, constructing the full fourth-order moment tensor $\mathcal{W}^{(4)}$ for a 6D state space incurs a massive $\mathcal{O}(N^4)$ computational and memory cost. When the objective is strictly to characterize geometric features of the uncertainty distribution, such as extracting moments along the principal bending directions for boundary reconstruction, evaluating the entire dense tensor is mathematically redundant. To circumvent this limitation, a novel “project-then-product” strategy is introduced. Previous works, such as the efficient polynomial expansion strategies by Lefebvre (2021),¹⁸ have demonstrated how to algebraically reduce the cost of computing high-order moments for scalar outputs by exploiting the orthogonality of the expansion basis. However, extending these scalar methods to evaluate the high-order cross-moments of a multi-dimensional state vector still triggers a combinatorial explosion of the cross-terms. Instead of computing these multi-dimensional cross-moments directly, this approach leverages geometric dimensionality reduction. To isolate specific non-Gaussian characteristics, first the multi-dimensional PCE polynomials are projected onto these specific 1D principal directions. Then, the recursive polynomial product identity are applied strictly to this 1D projection.¹⁸ By prioritizing the geometric projection, this strategy completely circumvents the full 6D tensor contraction, reducing a highly complex multi-dimensional moment problem to a lightweight 1D polynomial algebra operation.

Let u , v , and w be the selected principal directions, defined by the eigenvectors corresponding to the length, width, and thickness of the spatial covariance distribution. Due to the linearity of

the projection operator, the full multidimensional expansion can be mapped directly into 1D, mean-centered scalar polynomial maps:

$$z_u(\boldsymbol{\xi}) = \mathbf{u}^\top (\mathbf{x}(\boldsymbol{\xi}) - \boldsymbol{\mu}) = \sum_{|\mathbf{i}| \leq k} \gamma_{u,i} \Psi_i(\boldsymbol{\xi}) \quad (23)$$

where $\gamma_{u,i}$ are the projected coefficients obtained by taking the dot product of the analytical Hermite PCE coefficients c_i^{PCE} with the principal direction $\mathbf{u}$. Because the map is mean-centered, the zeroth-degree coefficient $\gamma_{u,0}$ evaluates strictly to zero. Analogous definitions to Eq. (23) with respect to the other principal directions $\mathbf{v}$ and $\mathbf{w}$, yield the projected coefficient vectors γ_v and γ_w and the scalar maps z_v and z_w .

To calculate the specific higher-order projected moments required for the confidence boundary deformation (e.g., $E_{uuu}, E_{uuv}, E_{uww}, E_{uuuu}$), a standard evaluation would still require the contraction of the full third- and fourth-order weight tensors. Rather than performing these computationally expensive high-order contractions, complex statistical moments can be algebraically rewritten as the expected value of lower-order polynomials. Following the recursive decomposition strategy,¹⁸ this is achieved by defining an intermediate *pair-product polynomial*. For instance, the squared deviation along the primary bending axis $\mathbf{u}$ is defined as:

$$p_{uu}(\boldsymbol{\xi}) = z_u(\boldsymbol{\xi})z_u(\boldsymbol{\xi}) \quad (24)$$

Because the Hermite polynomials form a complete orthogonal basis, this squared polynomial can be expressed as its own expansion, $p_{uu}(\boldsymbol{\xi}) = \sum c_{p_{uu},\mathbf{k}} \Psi_{\mathbf{k}}(\boldsymbol{\xi})$. The unknown coefficients $c_{p_{uu}}$ are found by multiplying both sides of Eq. (24) by the $\mathbf{k}$ -th basis function and taking the expected value. Substituting the definition of z_u and applying the linearity of the expectation operator yields:

$$\mathbb{E}[p_{uu} \Psi_{\mathbf{k}}] = \mathbb{E}[z_u z_u \Psi_{\mathbf{k}}] = \sum_{i,j} \gamma_{u,i} \gamma_{u,j} \mathbb{E}[\Psi_i \Psi_j \Psi_{\mathbf{k}}] \quad (25)$$

Recognizing that $\mathbb{E}[\Psi_i \Psi_j \Psi_{\mathbf{k}}]$ matches the analytical definition of the third-order weight tensor $\mathcal{W}_{ijk}^{(3)}$, the coefficients of the pair-product polynomial are recovered analytically:

$$c_{p_{uu},\mathbf{k}} = \frac{1}{\mathbb{E}[\Psi_{\mathbf{k}}^2]} \sum_{i,j} \gamma_{u,i} \gamma_{u,j} \mathcal{W}_{ijk}^{(3)} \quad (26)$$

where $\mathbb{E}[\Psi_{\mathbf{k}}^2] = \mathbf{k}!$ is the squared norm of the multi-dimensional Hermite basis term. This formulation demonstrates that the squared polynomial p_{uu} can be assembled explicitly using only the precomputed third-order tensor.

The primary computational advantage of this projection strategy lies in the extensive reuse of the 1D p_{uu} polynomial. Once the coefficients $c_{p_{uu}}$ are evaluated, the extraction of all primary and cross-coupling moments required for the boundary reconstruction reduces to computing the expected value of their product. Because $\Psi_{\mathbf{k}}$ is an orthogonal basis, $\mathbb{E}[\Psi_i \Psi_{\mathbf{k}}] = 0$ for $i \neq \mathbf{k}$, which

naturally zeroes out all cross-terms and yields:

$$E_{uuu} = \mathbb{E}[p_{uu}z_u] = \sum_{\mathbf{k}} c_{p_{uu},\mathbf{k}} \gamma_{u,\mathbf{k}} \mathbf{k}! \quad (27a)$$

$$E_{uuv} = \mathbb{E}[p_{uu}z_v] = \sum_{\mathbf{k}} c_{p_{uu},\mathbf{k}} \gamma_{v,\mathbf{k}} \mathbf{k}! \quad (27b)$$

$$E_{uuw} = \mathbb{E}[p_{uu}z_w] = \sum_{\mathbf{k}} c_{p_{uu},\mathbf{k}} \gamma_{w,\mathbf{k}} \mathbf{k}! \quad (27c)$$

$$E_{uuuu} = \mathbb{E}[p_{uu}p_{uu}] = \sum_{\mathbf{k}} c_{p_{uu},\mathbf{k}}^2 \mathbf{k}! \quad (27d)$$

Because p_{uu} naturally encodes the nonlinear variance distribution along the primary bending axis u , the cross-coupling terms (E_{uuv} and E_{uuw}) are retrieved directly by multiplying the p_{uu} coefficients with the 1D coefficients of the secondary and tertiary axes (z_v and z_w). By projecting the maps before calculating the expectations, this methodology bypasses the evaluation of the dense fourth-order tensor $\mathcal{W}^{(4)}$ entirely.

It is necessary to address a theoretical discrepancy introduced by the pair-product approach. Mathematically, squaring a projected polynomial map of degree k yields a polynomial of degree $2k$. However, because the pair-product coefficients $c_{p_{uu}}$ are evaluated using the precomputed k -degree tensor $\mathcal{W}^{(3)}$, the squared polynomial is inherently truncated back to degree k . The higher-degree terms (degrees $k+1$ to $2k$) are omitted from the inner product, introducing a minor truncation error where the evaluated kurtosis slightly underestimates the exact full-tensor value.

For typical problems in astrodynamics, the magnitude of higher-order polynomial coefficients decays exponentially. Because the omitted terms contribute almost nothing to the overall statistical moment, the resulting truncation error is mathematically negligible (empirically observed to be on the order of 10^{-5} for the evaluated kurtosis). This minimal loss in accuracy is a highly favorable trade-off for the computational speedup and memory reduction achieved by bypassing the full fourth-order tensor contraction.

The flowchart in Figure 1 delineates the computational architecture of the three analytical approaches presented in Section 3. This three-phase structure isolates the heaviest computational burdens, offering specific advantages for onboard GNC and trajectory optimization applications. First, the generation of the high-dimensional weight tensors (Phase 1) is strictly a function of the state dimension m and the expansion order k . Because these tensors are entirely independent of the underlying dynamical system or the specific trajectory, they can be computed offline and stored in onboard memory. Second, the numerical integration of the dynamics via the DA flow map (Phase 2) is decoupled from the statistical parameters of the uncertainty. From a mission operations perspective, this means the map only needs to be recomputed when the spacecraft alters its nominal reference trajectory or when navigation errors exceed the polynomial radius of convergence. Consequently, engineers can rapidly evaluate thousands of different initial uncertainty distributions (e.g., during navigation filter tuning or covariance analysis) without ever re-integrating the equations of motion. Lastly, the online UQ evaluation (Phase 3) is restricted to lightweight polynomial algebra and 1D projections. This characteristic makes the "project-then-product" strategy highly suitable for robust trajectory optimization and stochastic guidance algorithms. Specifically, it provides the capability to evaluate accurate, non-Gaussian chance constraints, such as those required for strict

collision avoidance or precision proximity operations, in real-time, without compromising onboard computational resources.




Figure 1. Flowchart of the three main phases for the proposed analytical frameworks, highlighting the offline and online contributions.

4 NUMERICAL RESULTS

In this section, the proposed analytical framework is applied to representative scenarios in cislunar and small-body proximity regimes, where non-Gaussian uncertainty distributions naturally emerge. The performance of each method is evaluated in terms of computational cost and geometric accuracy against a Monte Carlo (MC) reference. A specific focus is placed on the framework's ability to efficiently capture non-elliptical confidence geometries via the analytical three-dimensional banana boundary reconstruction.

The comparative analysis includes standard and DA-based implementations of LinCov, UT, GMM, and PCE, alongside the fully analytical higher-order moment extraction leveraging Isserlis' theorem, the Hermite transformation, and the pair-product projection strategy.

4.1 Dynamical Environment

The dynamical environment is modeled to provide a medium-to-high-fidelity representation of the perturbations experienced in cislunar space and asteroid proximity operations. The perturbed N -body problem is formulated as:³

$$\dot{\mathbf{r}} = \mathbf{v}, \quad \dot{\mathbf{v}} = \mathbf{a}_{\text{grav}} + \mathbf{a}_{\text{SRP}} + \mathbf{a}_{3B,\odot} + \mathbf{a}_{3B,\oplus} + \mathbf{a}_{3B,\text{Moon}} \quad (28)$$

The first term, $\mathbf{a}_{\text{grav}} = \nabla U_{\text{grav}}$, encompasses the central-body gravitational acceleration. To account for irregular mass distributions, the gravitational potential is expanded via spherical harmonics:¹⁹

$$U_{\text{grav}} = \frac{\mu}{r} \sum_{l=0}^{\infty} \sum_{m=0}^l \left(\frac{R_e}{r} \right)^l \bar{P}_{l,m}(\sin \phi) [\bar{C}_{l,m} \cos m\lambda + \bar{S}_{l,m} \sin m\lambda] \quad (29)$$

where μ is the central body gravitational parameter, R_e is the reference equatorial radius, and $\bar{P}_{l,m}$, $\bar{C}_{l,m}$, $\bar{S}_{l,m}$ are the associated Legendre polynomials and normalized harmonic coefficients, respec-

tively. The resulting acceleration is subsequently rotated from the body-fixed frame to the inertial frame.

The $\mathbf{a}_{3B}$ terms represent the point-mass third-body perturbations from the Sun, Earth, and Moon. These are modeled as the differential Keplerian gravity acting between the central body and the spacecraft:²⁰

$$\mathbf{a}_{3B,k} = \mu_k \left(\frac{\mathbf{r}_{k-sc}}{\|\mathbf{r}_{k-sc}\|^3} - \frac{\mathbf{r}_{k-cb}}{\|\mathbf{r}_{k-cb}\|^3} \right) \quad (30)$$

where μ_k is the gravitational parameter of the k -th perturbing body, while $\mathbf{r}_{k-sc}$ and $\mathbf{r}_{k-cb}$ denote its relative position vector with respect to the spacecraft and the central body, respectively.

Finally, $\mathbf{a}_{SRP}$ is the solar radiation pressure acceleration computed via a standard cannonball model:²

$$\mathbf{a}_{SRP} = -\frac{(1 + \rho)P_0 A_{sc} AU}{c M_{sc}} \frac{\mathbf{r}_{sc-Sun}}{\|\mathbf{r}_{sc-Sun}\|^3} \quad (31)$$

where P_0 is the solar flux at 1 AU, c is the speed of light, and the spacecraft is characterized by its mass M_{sc} , cross-sectional area A_{sc} , and reflectance coefficient ρ . The values considered in this work are related to a representative spacecraft which resembles a 6U Deep Space CubeSat inspired by Milani CubeSat ($M_{sc} = 12$ kg, $A_{sc} = 0.5\text{m}^2$, $\rho = 0.3$)^{*}.

The planetary and asteroidal ephemerides required to evaluate these perturbations are retrieved using NAIF SPICE kernels and JPL Horizons system.

4.2 Cislunar NRHO

Near-Rectilinear Halo Orbits (NRHOs) around the Earth-Moon L_2 point are of critical interest for current and future cislunar architectures. In this scenario, uncertainty propagation is analyzed along a 1.5 days orbital arc passing through the perilune. The initial state, extracted from the Lunar Gateway SPICE kernel[†], and the initial Gaussian distribution are detailed in Table 1. The initial uncertainty is deliberately scaled to exceed nominal cislunar navigation baselines, of comparable order to stress-case dispersions used to force nonlinear covariance growth in CR3BP-based cislunar trajectory studies.²¹ This serves as a computational stress test, designed to force the rapid manifestation of severe nonlinear spatial deformations over a short propagation window. The Moon's gravitational spherical harmonics up to fourth-order are included for this analysis.

Table 1. Initial state and uncertainty for Cislunar NRHO scenario.

Scenario	Cislunar NRHO
Init. Position $\mathbf{r}_0$ (km)	$[-4.731 \times 10^2 \quad 1.897 \times 10^4 \quad -3.546 \times 10^4]^T$
Init. Velocity $\mathbf{v}_0$ (km/s)	$[1.26 \times 10^{-2} \quad -3.64 \times 10^{-2} \quad 3.60 \times 10^{-1}]^T$
Pos. Std $\sigma_{\mathbf{r}_0}$ (km)	$diag(10 \quad 1 \quad 1)$
Vel. Std $\sigma_{\mathbf{v}_0}$ (km/s)	$diag(15 \times 10^{-3} \quad 1 \times 10^{-3} \quad 1 \times 10^{-3})$

Figure 2 illustrates the nominal trajectory and the evolution of the uncertainty distribution. As the spacecraft approaches perilune, the strong local gravitational gradient induces shear across the probability density function. The initial small Gaussian distribution is rapidly stretched into a non-Gaussian one, showing the typical banana shape. Note that the cross-track orientation of the banana

^{*}<https://www.hera mission.space/hera-mission-milani-cubesat>. Last accessed July 1, 2026.

[†]https://naif.jpl.nasa.gov/pub/naif/misc/MORE_PROJECTS/DSG/. Last accessed July 1, 2026.

is induced by the asymmetrical initial covariance, which is characterized by larger uncertainties along the x direction. The right panel of Figure 2 focuses on the final time instant, overlaying the analytical 3D banana surface reconstruction against the MC samples. Different viewing angles are shown to highlight the capability of the analytical approximation to accurately capture the bending of the distribution in the entire three-dimensional space, both in-plane and out-of-plane.




Figure 2. Cislunar NRHO scenario. Left: Nominal trajectory with uncertainty propagation distribution at different instants. Right: Different viewing angles of the analytical boundary reconstruction informed by higher moments at final time.

Figure 3 compares the final propagated distributions obtained using standard astrodynamics UQ methods: LinCov, UT, GMM, and PCE. The traditional covariance ellipsoids produced by LinCov and UT do not capture the extended tails of the actual final distribution. The GMM is implemented using 10 Gaussian components distributed along the principal axis of the initial uncertainty with a bell-shaped weight distribution. In the plot, this model is visualized by the individual component ellipsoids, with opacity proportional to their respective weights. This fixed-axis splitting strategy yields a low-fidelity approximation of the highly non-Gaussian shape, in particular regarding the approximation of higher-order moments; achieving higher accuracy would require a more computationally expensive recursive splitting algorithm. In contrast, the 4th-order non-intrusive PCE accurately reconstructs the full non-Gaussian shape of the distribution with high accuracy.

Table 2 summarizes the computational runtimes and statistical moment accuracies for the evaluated methods. All simulations were performed on a 2021 MacBook Pro (Apple M1 chip, 16 GB RAM). Relative errors are computed against a reference Monte Carlo (MC) simulation of $N = 25,000$ samples, which required 503.15 s for numerical integration. Expanding the DA flow map to the 4th order required an offline computational time of 10.27 s. Note that the computational cost of the DA integration is highly dependent on the degree and order of the spherical harmonics gravity model, as higher degrees significantly increase the number of polynomial operations required per integration step. Consequently, the computational overhead of DA integration relative to standard floating-point integration widens as the gravity model fidelity increases.

Covariance errors are computed using the Frobenius norm of the covariance matrix. Skewness




Figure 3. Comparison of different UQ methods at the final time instant, with respect to the MC solution (gray dots) and the DA analytical approximation (black solid line).

and kurtosis errors are instead computed as the relative norm of the marginal moments,

$$\varepsilon_\gamma = \frac{\|\gamma_{\text{method}} - \gamma_{\text{MC}}\|}{\|\gamma_{\text{MC}}\|}, \quad \varepsilon_\kappa = \frac{\|\kappa_{\text{method}} - \kappa_{\text{MC}}\|}{\|\kappa_{\text{MC}}\|}, \quad (32)$$

where γ and κ denote the vectors of marginal skewness and kurtosis, extracted from the diagonal of the respective third- and fourth-order moment tensors, and the subscript MC denotes the Monte Carlo reference. For the DA “project-then-product” approach, the reported skewness and kurtosis errors instead correspond to the mean of the errors of each of the extracted components (e.g., E_{uuu} , E_{uuv} , E_{uuw} , E_{uuuu}).

Table 2. Comparison of computational runtime and relative accuracy against standard Monte Carlo (25,000 samples) at final time t_f .

Method	Runtime [s]		Relative Error vs. Monte Carlo			
	Online Eval.	Total	Mean	Covariance	Skewness	Kurtosis
<i>Standard Frameworks</i>						
LinCov	0.2496	0.2496	7.34×10^{-3}	1.36×10^{-2}	—	—
UT	0.2726	0.2726	6.26×10^{-4}	3.30×10^{-3}	—	—
PCE (p=4)	9.3443	9.4864	6.16×10^{-4}	1.41×10^{-2}	2.20×10^{-2}	3.49×10^{-2}
GMM (K=10)	2.6434	2.6434	6.17×10^{-4}	1.13×10^{-2}	1.51×10^{-1}	1.34×10^{-1}
Monte Carlo	503.1478	503.1478	Ref.	Ref.	Ref.	Ref.
<i>Differential Algebra Frameworks</i>						
DA Full Tensors (4th order)	5.5289	17.1013	6.17×10^{-4}	1.22×10^{-2}	1.83×10^{-2}	5.49×10^{-2}
DA Project-then-product (4th order)	0.4868	12.0592	6.17×10^{-4}	1.22×10^{-2}	2.31×10^{-2}	2.53×10^{-2}

Note: Total runtime for DA frameworks includes the expansion of the whitened DA flow map. Skewness and kurtosis moments are inherently not extracted for standard LinCov and UT frameworks.

The fully analytical higher-order moments are extracted by exploiting the 4th-order DA flow map. The computational performance of the different approaches described in Section 3 is detailed in Figure 4, which illustrates the runtime breakdown for higher-order moment recovery without sampling. Because the Isserlis and Hermite weight tensors are entirely independent of the orbital dynamics,

they can be computed offline and stored. Consequently, the onboard computational requirement is strictly limited to the DA map expansion, when necessary for updating the reference trajectory, and the subsequent tensor contractions for the uncertainty evaluation. By leveraging the inherent sparsity of the Hermite basis transformation (e.g., a 0.6% fill factor), the online computational burden is significantly reduced compared to applying Isserlis’ theorem directly to the standard monomial basis. Finally, the project-then-product strategy enables the targeted computation of only the specific skewness and kurtosis terms required to parameterize the confidence boundary, entirely bypassing the need for full tensor contractions or discrete sample evaluations.

Note that the reduction in offline computation time achieved by the project-then-product approach stems directly from the pair-product strategy detailed in Section 3.3. By enabling the derivation of fourth-order kurtosis terms through the contraction of lower-order tensors, this strategy eliminates the necessity of computing the more expensive four-dimensional kurtosis tensor. Conversely, the marginal increase in expansion time for the Hermite and pair-product methods is due to the necessary basis conversion from the monomial representation to the Hermite basis.

The computational advantage of the project-then-product method over full-tensor approaches becomes increasingly evident as the DA expansion order increases. In DA, the computational effort scales combinatorially: for instance, in a system with $n = 6$ dimensions, a second-order expansion requires the integration of 258 equations, whereas a third-order expansion necessitates 1554. This scaling is intensified when extracting higher-order statistical moments; because moments are computed via the expectation of products of state variables, the underlying monomial degree grows rapidly. Specifically, resolving m -th order moments from a k -th order DA map requires tracking monomials up to degree $m \times k$. For example, resolving 4th-order kurtosis components from a 3rd-order DA map requires order 12, whereas from a 4th-order map it necessitates order 16. By projecting the dynamics onto the principal axes, the project-then-product method bypasses the need to construct the entire high-degree tensors, mitigating the dimensionality bottleneck.




Figure 4. Runtime breakdown for analytical higher-order moments recovery approaches using DA framework.

Results of the performance comparison with state-of-the-art methods are also shown in Figures 5 and 6. Figure 5 details the computational time for a single UQ query. For the DA-analytical approaches, this online evaluation time excludes the offline DA map expansion, as a new expansion is only required when the estimated truncation error of the polynomial map exceeds a prescribed

threshold. Figure 6 presents the relative errors of the statistical moments compared to the Monte Carlo (MC) ground truth. Note that the GMM method implemented in this work does not actively optimize the splitting of the Gaussian components and their weights, and therefore it shows low computational time but relatively high error with respect to other high-fidelity methods. Together, these plots demonstrate that the DA-analytical approaches estimate the mean, covariance, and higher-order moments with an accuracy comparable to high-fidelity methods like PCE. Furthermore, the project-then-product strategy significantly speeds up the online UQ queries, achieving runtimes similar to lower-fidelity methods such as LinCov and UT. Therefore, this reduction in computational overhead is achieved while maintaining the geometric accuracy of non-Gaussian high-fidelity methods. The analytically recovered marginal skewness and kurtosis closely match the fully integrated MC reference, validating the formulation of the Hermite basis transformation and the pair-product strategy in highly nonlinear regimes.




Figure 5. Computational times of different UQ methods for cislunar scenario. Note that DA methods times do not include the DA expansion time, but are related to the single UQ query. DA Full Tensors refers to the faster Hermite-DA approach.

4.3 Apophis Proximity Trajectories

To further validate the proposed framework, similar analyses are conducted for close-proximity operations around the asteroid Apophis. This environment presents unique UQ challenges driven by the asteroid’s highly irregular gravity field, relevant Solar Radiation Pressure, and rapid dynamical timescales. Two specific scenarios are evaluated: a revolution during deep-space proximity operations, and a highly perturbed trajectory during the asteroid’s predicted 2029 close approach to Earth.³ Apophis’s gravity spherical harmonics to the 4th order are used for this analysis.

The initial state vectors are detailed in Table 3. For both scenarios, the initial Gaussian distribution is parameterized by a spherical position variance of $\sigma_r = 10$ m and spherical velocity variance of $\sigma_v = 0.3 \times 10^{-3}$ m/s. The propagation time horizon is set to 2 days for the deep-space scenario and 0.75 days for the accelerated Earth flyby scenario. As in the previous case study, the refer-




Figure 6. Statistical moments relative errors with respect to the MC reference for different UQ methods for cislunar scenario. Note that the error of skewness and kurtosis of DA project-then-product approach refers only to the extracted terms.

ence Monte Carlo simulations utilizes $N = 25,000$ samples, and the DA map is expanded to the 4th-order.

Table 3. Initial state for Apophis proximity scenarios.

Scenario	Apophis deep-space
Init. Position $\mathbf{r}_0$ (km)	$[0.85 \ 0 \ 0]^T$
Init. Velocity $\mathbf{v}_0$ (km/s)	$[2.7229 \times 10^{-5} \ -2.1036 \times 10^{-5} \ 4.1277 \times 10^{-5}]^T$
Scenario	Apophis Earth's flyby
Init. Position $\mathbf{r}_0$ (km)	$[-0.3255 \ -1.4633 \ -0.0520]^T$
Init. Velocity $\mathbf{v}_0$ (km/s)	$[-2.8502 \times 10^{-5} \ 1.9168 \times 10^{-5} \ -0.1889 \times 10^{-5}]^T$

The resulting uncertainty propagation is illustrated in Figure 7. In both cases, the complex and perturbed dynamics cause the initial Gaussian sphere to shear into a non-elliptical volume. The plots overlay the nominal trajectory, the sample distributions, and the corresponding 3D banana boundary reconstruction at the final time instant. The analytical approximation, parameterized by the analytically recovered skewness and kurtosis, successfully envelopes the non-Gaussian features of the distribution, providing a strict geometric bound that standard covariance ellipsoids fail to capture. In three-dimensional space, the theoretical containment fraction for a 3σ Gaussian confidence boundary is 97.07%. The LinCov ellipsoid includes only 69.43% of the MC samples, significantly underestimating the true dispersion. The UT contains 94.54% of the samples, achieving a closer containment fraction as expected from a nonlinear method. However, because UT still restricts the confidence boundary to a Gaussian distribution, it completely misses its spatial curvature and includes large regions of empty space while excluding the extended tails. In contrast, the pro-

posed analytical approximation accurately reconstructs the non-Gaussian banana shape, containing 96.28% of the samples.




Figure 7. Apophis proximity scenarios. Left: During deep-space. Right: During Earth's flyby.

The computational times of the standard UQ methods and DA-based analytical implementations for the Earth's flyby scenario are summarized in Figure 8. The relative errors, with respect to the reference MC simulation, of the propagated statistical moments at the final time instants are presented in Figure 9. The results confirm that the DA-based analytical implementations maintain geometric and statistical accuracy equivalent to their standard numerical counterparts, validating that the profound reduction in computational time does not compromise the fidelity of the uncertainty quantification.




Figure 8. Computational times of different UQ methods for Apophis (Earth's flyby scenario). Note that DA methods times do not include the DA expansion time, but are related to the single UQ query. DA Full Tensors refers to the faster Hermite-DA approach.




Figure 9. Statistical moments relative errors with respect to the MC reference for different UQ methods for Apophis (Earth’s flyby scenario). Note that the error of skewness and kurtosis of DA project-then-product approach refers only to the extracted terms.

5 CONCLUSION

This work introduces an analytical framework for non-Gaussian uncertainty propagation and three-dimensional confidence boundary reconstruction in highly perturbed dynamical environments. To mitigate the high computational cost and statistical noise of traditional sample-based methods, the proposed methodology leverages a fully analytical approach to extract statistical moments. A differential algebra (DA) formulation is used to replace the numerical integration of the dynamics with high-order polynomial flow maps. Then, higher-order moments are extracted from the DA coefficients (i.e., without sampling) by exploiting Isserlis’ theorem and the mapping of standard monomials into an orthogonal probabilist’s Hermite basis, which significantly improves the tensor sparsity. Additionally, a “project-then-product” approach is used to bypass the computational and memory scaling associated with high-order moments tensor contractions, and to retrieve only the relevant terms. The analytically extracted skewness and kurtosis are then utilized to parameterize a three-dimensional analytical banana-shaped confidence boundary. This geometric reconstruction accurately bounds the spatial, bending, and out-of-plane torsion caused by propagation in nonlinear dynamics, which traditional linear methods fail to capture. It also has the benefit of being analytic, explicitly parameterized by angles.

The proposed framework is tested and validated in two challenging nonlinear regimes: a dynamically realistic cislunar near-rectilinear halo orbit (NRHO), and close-proximity trajectories around the asteroid Apophis in deep space and during Earth flyby. The numerical results demonstrate that the fully analytical approach, exploiting tensor sparsity and projection techniques, achieves geometric and statistical accuracy equivalent to a large-batch Monte Carlo study, or other high-fidelity uncertainty propagation methods, while simultaneously reducing computational times. This sample-free methodology provides an efficient and mathematically rigorous tool for onboard uncer-

tainty quantification (UQ), robust trajectory optimization, and space domain awareness for modern astrodynamics applications. The methodology can be augmented with trajectory optimization techniques making use of polynomial maps like DA or the related state transition tensors (STTs),^{22–24} whereby the same maps that are used for nominal trajectory correction can be used for nonlinear UQ operations.

REFERENCES

- [1] T. Wolf, E. M. Zucchelli, and B. A. Jones, “Multi-Fidelity Uncertainty Propagation for Objects in Cislunar Space,” *AIAA SCITECH 2022 Forum*, San Diego, CA & Virtual, American Institute of Aeronautics and Astronautics, Jan. 2022, 10.2514/6.2022-1774.
- [2] D. J. Scheeres, *Orbital motion in strongly perturbed environments: applications to asteroid, comet and planetary satellite orbiters*. Springer, 2012.
- [3] N. Michelotti, A. Rizza, C. Giordano, and F. Topputo, “Uncertainty Propagation Performance in Proximity Operations Around Small Bodies,” *The Journal of the Astronautical Sciences*, Vol. 71, Dec. 2024, p. 55, 10.1007/s40295-024-00472-5.
- [4] Y.-z. Luo and Z. Yang, “A review of uncertainty propagation in orbital mechanics,” *Progress in Aerospace Sciences*, Vol. 89, 2017, pp. 23–39.
- [5] N. Adurthi, P. Singla, and T. Singh, “The conjugate unscented transform—an approach to evaluate multi-dimensional expectation integrals,” *2012 American Control Conference (ACC)*, IEEE, 2012, pp. 5556–5561.
- [6] B. A. Jones, A. Doostan, and G. H. Born, “Nonlinear propagation of orbit uncertainty using non-intrusive polynomial chaos,” *Journal of Guidance, Control, and Dynamics*, Vol. 36, No. 2, 2013, pp. 430–444.
- [7] V. Vittaldev, R. P. Russell, and R. Linares, “Spacecraft uncertainty propagation using gaussian mixture models and polynomial chaos expansions,” *Journal of Guidance, Control, and Dynamics*, Vol. 39, No. 12, 2016, pp. 2615–2626.
- [8] M. Valli, R. Armellini, P. Di Lizia, and M. R. Lavagna, “Nonlinear Mapping of Uncertainties in Celestial Mechanics,” *Journal of Guidance, Control, and Dynamics*, Vol. 36, Jan. 2013, pp. 48–63, 10.2514/1.58068.
- [9] S. Boone and J. McMahon, “Directional state transition tensors for capturing dominant nonlinear effects in orbital dynamics,” *Journal of Guidance, Control, and Dynamics*, Vol. 46, No. 3, 2023, pp. 431–442.
- [10] L. Isserlis, “On a formula for the product-moment coefficient of any order of a normal frequency distribution in any number of variables,” *Biometrika*, Vol. 12, No. 1/2, 1918, pp. 134–139.
- [11] G. Acciarini, N. Baresi, D. J. B. Lloyd, and D. Izzo, “Nonlinear Propagation of Non-Gaussian Uncertainties,” *Journal of Guidance, Control, and Dynamics*, Vol. 48, Apr. 2025, pp. 903–913, 10.2514/1.G008717.
- [12] E. Burnett and S. Boone, “Analytic Non-Gaussian Chance Constraint Method for Stochastic Trajectory Control,” *IEEE Conference on Decision and Control*, Preprint pending appearance on arXiv in 4/2026.
- [13] E. R. Burnett, S. Boone, and N. Michelotti, “An Efficient Non-Gaussian Chance Constraint Method for Stochastic Nonlinear Problems in Spaceflight,” *AAS/AIAA Astrodynamics Specialist Conference*, No. AAS 26-899, American Astronautical Society, July 2026.
- [14] E. A. Cornish and R. A. Fisher, “Moments and cumulants in the specification of distributions,” *Revue de l’Institut international de Statistique*, 1938, pp. 307–320.
- [15] B. Sudret, “Global sensitivity analysis using polynomial chaos expansions,” *Reliability engineering & system safety*, Vol. 93, No. 7, 2008, pp. 964–979.
- [16] F. W. Olver, *NIST handbook of mathematical functions hardback and CD-ROM*. Cambridge university press, 2010.
- [17] R. Azor, J. Gillis, and J. D. Victor, “Combinatorial applications of Hermite polynomials,” *SIAM Journal on Mathematical Analysis*, Vol. 13, No. 5, 1982, pp. 879–890.
- [18] T. Lefebvre, “On moment estimation from polynomial chaos expansion models,” *IEEE Control Systems Letters*, Vol. 5, No. 5, 2021, pp. 1519–1524.
- [19] D. A. Vallado, *Fundamentals of Astrodynamics and Applications*. Hawthorne, CA, USA: Microcosm Press, Springer, 2013.
- [20] F. Ferrari, V. Franzese, M. Pugliatti, C. Giordano, and F. Topputo, “Trajectory options for Hera’s Milani CubeSat Around (65803) Didymos,” *The Journal of the Astronautical Sciences*, Vol. 68, No. 4, 2021, pp. 973–994, 10.1007/s40295-021-00282-z.

- [21] S. Kelly and D. Geller, “Robust Cislunar Trajectory Optimization in the Presence of Stochastic Errors,” *The Journal of the Astronautical Sciences*, Vol. 71, No. 4, 2024, p. 30.
- [22] E. R. Burnett and F. Topputo, “Rapid Nonlinear Convex Guidance Using a Monomial Method,” *Journal of Guidance, Control, and Dynamics*, Vol. 48, No. 4, 2025, pp. 736–756, 10.2514/1.G008512.
- [23] O. Regantini, E. R. Burnett, A. Rizza, A. Morselli, and F. Topputo, “Convex Trajectory Optimization via Monomial Coordinates Transcription for Cislunar Rendezvous,” *Journal of Guidance, Control, and Dynamics*, 2026, pp. 1–15, 10.2514/1.G009332.
- [24] T. Caleb, R. Armellin, S. Boone, and S. Lizy-Destrez, “Taylor polynomial-based constrained solver for fuel-optimal low-thrust trajectory optimisation,” 2026, 10.48550/arXiv.2502.00398.