ESA GNC & ICATT Conference 2026
28 Sep. – 2 Oct. 2026, Sevilla, Spain
https://atpi.eventsair.com/gncicatt/

Confidence Boundary Methods for Non-Gaussian Mission Analysis and Stochastic Guidance

Ethan R. Burnett[1,*], **Niccolò Michelotti**[2], **and Francesco Topputo**[2]

[1]*Department of Aerospace Engineering and Engineering Mechanics, University of Texas at Austin, 2617 Wichita St, Austin TX, 78712, USA*
[2]*Dept. of Aerospace Science and Technology, Politecnico di Milano, Via La Masa 34, 20156 Milano, Italy*
[*]*Corresponding author: ethan.burnett@utexas.edu*

ABSTRACT

Traditional stochastic guidance and control in spaceflight often assumes Gaussian statistics for vehicle state uncertainties. However, a Gaussian prior is not guaranteed to remain Gaussian under nonlinear dynamics. This paper outlines state-of-the-art methods for computing and using approximations of confidence boundaries in state space, or projections thereof, at a prescribed cumulative probability level, while relaxing the Gaussian assumption, allowing for its corruption under the action of weak nonlinearity. New arguments are developed to better motivate these methods, including a second-order flow-map interpretation of a recently developed non-Gaussian boundary correction. Practical challenges are also examined, including enforcement of the resulting non-convex constraints in chance-constrained control and, for sufficiently non-Gaussian distributions, projection of confidence boundaries from full state space onto an operational subspace. This paper also reviews other recent efforts from uncertainty quantification and reachability literature, and previews ideas from our own recent investigations. Results from applications to uncertainty quantification and stochastic trajectory optimization problems are then summarized.

1 Introduction

Spacecraft guidance problems are naturally stochastic: maneuvers must be planned not only to steer a nominal trajectory, but also to manage the distribution of possible outcomes resulting from state uncertainty, maneuver execution error, and other disturbances. Linear covariance ("LinCov") methods are attractive in this setting because they propagate uncertainty cheaply and lead naturally to chance-constrained guidance formulations [1, 2]. Their accuracy, however, depends on the propagated distribution remaining reasonably close to Gaussian. Over long propagation intervals or in strongly nonlinear dynamical regimes, this assumption may fail as an initially compact distribution is stretched, skewed, and bent by the dynamics [3, 4]. Fig. 1 illustrates a motivating problem considered in this context: a stochastic two-maneuver planning problem. At the first maneuver, the vehicle state may be well estimated and its uncertainty approximately Gaussian, but the next planned maneuver may lie far downstream. The intervening dynamics therefore map the present uncertainty into a wide and potentially non-Gaussian distribution of possible future states, which must be considered at the initial maneuver time, despite possibly strong navigation performance between maneuvers. This issue is particularly relevant to long-horizon cislunar and deep-space maneuver planning.

Accurately recovering non-Gaussian distributions is often done by Monte Carlo methods, which are straightforward but expensive. Many works have thus explored computationally middle-of-the-road approaches between Monte Carlo and LinCov, including the conjugate unscented transform (CUT) [5], polynomial chaos expansions [6], and Gaussian mixture approaches [7]. Polynomial flow maps obtained through differential algebra or state transition tensors provide another useful route: the expensive nonlinear integration is performed once, after which states or statistical information can be evaluated inexpensively from the resulting local map [8, 9]. Furthermore,

© *Ethan R. Burnett, Niccolò Michelotti, and Francesco Topputo*

higher-order moments can be efficiently recovered directly from such maps without sampling [10]. Complementary work [11] develops a Taylor-map approximation of the Fokker–Planck solution in the presence of process noise, providing a grid-free description capable of retaining non-Gaussian deformation and asymmetric tails.

For many guidance applications, however, consideration of the complete probability density is unnecessary: the principal object of interest is the confidence boundary that separates statistically acceptable from unacceptable violation of mission constraints. One may seek a region $\mathcal{A}_t(P)$ containing a prescribed cumulative probability $P \in (0,1)$ and require its boundary, or a projection thereof, to respect keep-out zones and other constraints. Recent work has therefore used higher-order moments to construct analytic corrections to the Gaussian covariance ellipse for weakly non-Gaussian, "banana-shaped" distributions [12, 13, 10]. The present paper summarizes, unifies, and contextualizes these developments. We review efficient map-based uncertainty propagation and moment recovery, develop a flow map-based justification of the banana shape, and discuss direct transport of confidence boundaries without moment reconstruction. We then consider two practical issues that can become important once the Gaussian assumption is relaxed: projection of a full state-space boundary into lower-dimensional spaces, and enforcement of the resulting potentially non-convex chance constraints. The resulting methods are illustrated for non-Gaussian uncertainty quantification and stochastic maneuver planning.

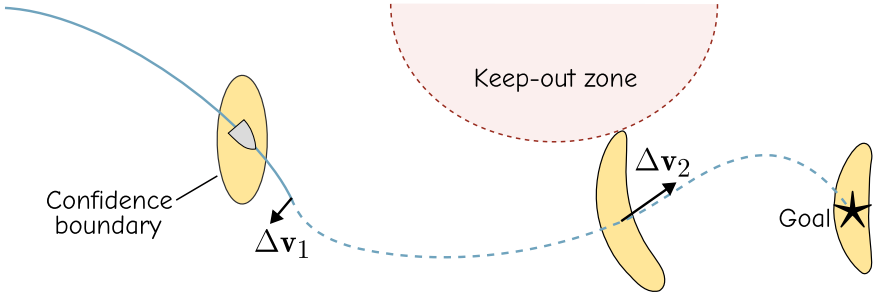

Figure 1 *Stochastic Maneuver Planning with Non-Gaussian Statistics*

2 Polynomial flow maps for uncertainty propagation

Consider a nonlinear dynamical system

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

where $\boldsymbol{x} \in \mathbb{R}^n$ is the state vector. Given an initial state $\boldsymbol{x}_0$ at time t_0, the state at a later time t is obtained via the flow of the dynamics

$$\boldsymbol{x}(t) = \boldsymbol{\varphi}(t; \boldsymbol{x}_0, t_0). \tag{2}$$

In the context of uncertainty propagation, we are interested in the evolution of a statistical distribution of states from time t_0 to t, but propagating the requisite probability density function $p(\boldsymbol{x}, t)$ is generally a computationally burdensome task for nonlinear dynamical systems. Instead, a standard approach is to evaluate Eq. (2) for many realizations of $\boldsymbol{x}_0$ sampled from $p(\boldsymbol{x}, t_0)$, using the result to estimate the statistical moments at time t. An alternative is to approximate the dependence of the propagated state on the initial perturbation with a local polynomial flow map. Then, applying such a map, one may evaluate the final statistics from either a large ensemble of points, as in a Monte Carlo method, or from a small set of *sigma points*, which are carefully selected to provide maximal statistical information per point. One such sigma point-based method is the conjugate unscented transform (CUT) [5]. The polynomial approximation of the flow of the dynamics, recovered by employing techniques such as state transition tensors (STT) [14, 15] or differential algebra (DA) [16, 8], can be used as a surrogate of full nonlinear propagation of Eq. (1) each time the flow of Eq. (2) must be computed. A one-time cost is paid to generate the polynomial flow map from t_0 to t, whereas evaluating this flow map repeatedly is comparatively numerically inexpensive. This can offer significant computational time savings in propagating states and even statistical moments onboard, particularly if the same propagation is to be repeated many times, as in optimization or filtering contexts.

Let the initial condition be written as $\boldsymbol{x}_0 = \overline{\boldsymbol{x}}_0 + \delta\boldsymbol{x}_0$, where $\overline{\boldsymbol{x}}_0$ denotes a nominal state and $\delta\boldsymbol{x}_0$ is a random deviation. Differential algebra (DA) carries a truncated Taylor expansion through the numerical integration of the dynamics starting from the nominal state $\overline{\boldsymbol{x}}_0$. The propagated deviation from the nominal trajectory may then be

expressed as

$$\delta\boldsymbol{x}(t) = \sum_{1\le|\boldsymbol{p}|\le k} \boldsymbol{M}_{\boldsymbol{p}}(t)\,\delta\boldsymbol{x}_0^{\boldsymbol{p}}, \tag{3}$$

where $\boldsymbol{p} \in \mathbb{N}^n$ is a multi-index, $|\boldsymbol{p}| = \sum_i p_i$, and $\delta\boldsymbol{x}_0^{\boldsymbol{p}} = \prod_i \delta x_{0,i}^{p_i}$. The coefficients $\boldsymbol{M}_{\boldsymbol{p}}(t)$ contain the local derivatives of the nonlinear flow through order k. Often times the shorthand notation $\delta\boldsymbol{x} = \mathcal{M}(\delta\boldsymbol{x}_0)$ may also be used in lieu of Eq. (3) [8].

For a sufficiently compact distribution, the vast majority of the probability mass will be located for some finite time within the region of validity of a single flow map. In such circumstances, all Monte Carlo or sigma point samples can reuse this same flow map, then be used to estimate the desired statistical moments. It is then natural to wonder whether the information contained in the flow map can simply be used directly for propagating statistical moments, avoiding any random (or for that matter, deterministic) sampling at all. We explore this in Reference [10], following earlier work by [8] and [9], which showed that this can be done directly from the polynomial coefficients.

Reference [8] developed their approach as the "DA higher-order" (DAHO) method. Their coefficient-expectation formulation is formally applicable whenever the requisite moments of the initial distribution are known; for Gaussian initial uncertainty, these moments are evaluated using Isserlis' theorem. The difficulty is that the required algebra grows rapidly: an rth-order propagated moment obtained from a kth-order flow map involves initial monomial moments through degree rk, together with the corresponding products of map coefficients. This quickly becomes burdensome or intractable in practice with even fairly modest values of r or k. Reference [9] subsequently provided a systematic extension beyond the Gaussian case by obtaining the required initial moments from the moment-generating function $\mathcal{G}_{\delta\boldsymbol{x}_0}$,

$$\mathbb{E}\big[\delta\boldsymbol{x}_0^{\boldsymbol{p}}\big] = \left.\frac{\partial^{|\boldsymbol{p}|}\mathcal{G}_{\delta\boldsymbol{x}_0}(\boldsymbol{t})}{\partial t_1^{p_1}\cdots\partial t_n^{p_n}}\right|_{\boldsymbol{t}=\boldsymbol{0}}, \tag{4}$$

so that these distribution-dependent quantities can be generated separately from the dynamics and reused with the polynomial flow coefficients.

In Reference [10], we focus on making this direct moment extraction efficient for initially Gaussian uncertainty, with particular focus on the third- and fourth-order moments. We now briefly review the key findings. The initial deviations are mathematically decorrelated through a whitening process: by writing $P_0 = LL^T$ and $\delta\boldsymbol{x}_0 = L\boldsymbol{\xi}$, with $\boldsymbol{\xi} \sim \mathcal{N}(\boldsymbol{0}, I_n)$, the polynomial map can be formed directly in independent standard normal variables.

$$x_i(\boldsymbol{\xi}) = \sum_{|\boldsymbol{p}|\le k} c_{i,\boldsymbol{p}}\boldsymbol{\xi}^{\boldsymbol{p}}, \tag{5a}$$

$$\mathbb{E}[\boldsymbol{\xi}^{\boldsymbol{p}}] = \prod_{\ell=1}^{n}\begin{cases}0, & p_\ell \text{ odd},\\ (p_\ell-1)!!, & p_\ell \text{ even}.\end{cases} \tag{5b}$$

Thus the general r^{th} raw moment is simply a contraction of the flow-map coefficients with precomputable Gaussian weights obtained through Isserlis' theorem,

$$\mathcal{R}^{(r)}_{i_1\cdots i_r} = \sum_{\boldsymbol{p}_1,\ldots,\boldsymbol{p}_r}\left(\prod_{j=1}^{r} c_{i_j,\boldsymbol{p}_j}\right)\mathcal{W}^{(r)}_{\boldsymbol{p}_1\cdots\boldsymbol{p}_r}, \qquad \mathcal{W}^{(r)}_{\boldsymbol{p}_1\cdots\boldsymbol{p}_r} = \mathbb{E}\big[\boldsymbol{\xi}^{\boldsymbol{p}_1+\cdots+\boldsymbol{p}_r}\big]. \tag{6}$$

The remaining computational burden lies in these contractions, particularly the fourth-order tensor $\mathcal{W}^{(4)}$. To improve computational efficiency, Reference [10] transforms the whitened DA map to the probabilists' Hermite basis, by mapping any monomial to a linear combination of Hermite polynomials, and resulting in a polynomial chaos expansion (PCE):

$$x_i(\boldsymbol{\xi}) = \sum_{\boldsymbol{\alpha}} c^{\mathrm{H}}_{i,\boldsymbol{\alpha}}\Psi_{\boldsymbol{\alpha}}(\boldsymbol{\xi}), \qquad \Psi_{\boldsymbol{\alpha}}(\boldsymbol{\xi}) = \prod_{d=1}^{n}\mathrm{He}_{\alpha_d}(\xi_d), \qquad \mathbb{E}\big[\Psi_{\boldsymbol{\alpha}}\Psi_{\boldsymbol{\beta}}\big] = \boldsymbol{\alpha}!\,\delta_{\boldsymbol{\alpha}\boldsymbol{\beta}}. \tag{7}$$

This orthogonality property yields weight tensors $\mathcal{W}^{(r)}$ that are significantly sparser than those derived directly from the monomial DA expansion, reducing both the storage and computational cost of the moment contractions.

Moreover, when only specific projected third- and fourth-order moments are required (as in the methods to be discussed in Sections 3.1 and 3.2), the multivariate Hermite expansion can first be projected onto the directions of interest, after which intermediate pair-product polynomials are formed from the resulting scalar expansions. This reduces the required higher-order moment evaluation from a full multidimensional tensor contraction to low-dimensional polynomial algebra and inner products. For a projection direction $\boldsymbol{a}$, define the scalar map

$$z_a(\boldsymbol{\xi}) = \boldsymbol{a}^T\big(\boldsymbol{x}(\boldsymbol{\xi}) - \boldsymbol{\mu}\big) = \sum_{|\boldsymbol{\alpha}|\le k} \gamma_{a,\boldsymbol{\alpha}} \Psi_{\boldsymbol{\alpha}}(\boldsymbol{\xi}), \tag{8}$$

where $\gamma_{a,\boldsymbol{\alpha}}$ are the projected coefficients obtained by taking the dot product of the analytical PCE coefficients $\boldsymbol{c}_{\boldsymbol{i}}^{\mathrm{H}}$ with the principal direction $\boldsymbol{a}$. Let the degree-k projection of its squared polynomial be

$$\Pi_k[z_a^2] = \sum_{|\boldsymbol{\alpha}|\le k} d_{a,\boldsymbol{\alpha}} \Psi_{\boldsymbol{\alpha}}(\boldsymbol{\xi}). \tag{9}$$

The unknown coefficients $d_{a,\boldsymbol{\alpha}}$ are analytically recovered [10]:

$$d_{a,\boldsymbol{\alpha}} = \frac{1}{\boldsymbol{\alpha}!} \sum_{|\boldsymbol{i}|\le k} \sum_{|\boldsymbol{j}|\le k} \gamma_{a,\boldsymbol{i}} \gamma_{a,\boldsymbol{j}} \mathcal{W}^{(3)}_{\boldsymbol{ij\alpha}}, \tag{10}$$

Consider some z_b defined analogously for a second projection direction $\boldsymbol{b}$. Since z_b also lies in the retained degree-k polynomial space, Hermite orthogonality gives the projected third-order moment exactly as

$$\mathbb{E}[z_a^2 z_b] = \sum_{|\boldsymbol{\alpha}|\le k} d_{a,\boldsymbol{\alpha}} \gamma_{b,\boldsymbol{\alpha}} \boldsymbol{\alpha}!, \tag{11}$$

while the corresponding fourth-order moment is approximated by

$$\mathbb{E}[z_a^4] \approx \sum_{|\boldsymbol{\alpha}|\le k} d_{a,\boldsymbol{\alpha}}^2 \boldsymbol{\alpha}!. \tag{12}$$

The latter approximation arises because z_a^2 generally has degree $2k$, whereas only its projection back onto the retained degree-k basis is used. This "project-then-product" construction therefore bypasses the full fourth-order tensor and leaves the online uncertainty calculation primarily as polynomial projection and inner-product operations, at the price of a small truncation error in the fourth-order moment. The result is a sample-free means of recovering the higher-order moments, joining our basket of techniques for calculating such information.

Fig. 2 illustrates a number of different techniques for uncertainty propagation with and without DA-computed Taylor maps. A confidence boundary for a particular statistical distribution, at a particular confidence level, is illustrated with the solid black closed curves. The boundary of the region of validity of the Taylor map, for some error metric, is illustrated by the red dashed lines. The confidence boundary must entirely be contained within the region of validity of the map. Many of these methods were explored in Ref. [17]. Ref. [10] proposes a fast new implementation of sample-free, map-based moment recovery (bottom-right panel). All methods in the bottom row require some type of operator overloading to enable DA for computing the Taylor map. Sometimes this is not possible or desirable. One may consider alternate techniques that are still fast, such as the top-center panel (employed in e.g. [13]).

3 Approximations of non-Gaussian confidence boundaries

3.1 Functional Form Method and banana-shaped confidence regions

References [12] and [13] introduced a simple geometric approach for approximating confidence boundaries when a propagated distribution is appreciably non-Gaussian but remains sufficiently regular that its dominant geometric features can be identified. The basic idea, summarized in Fig. 3, is to posit low-order functional forms corresponding to these features, then determine their coefficients from available statistical moments, informing analytic departures from a Gaussian baseline. In lieu of attempting to reconstruct the full probability density, the method seeks a direct approximation of the confidence boundary itself, envisioned mainly for stochastic guidance

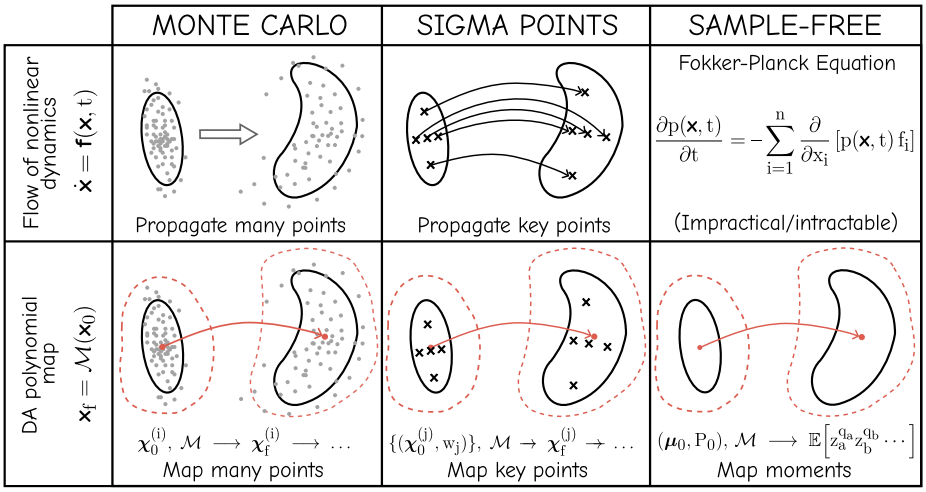

Figure 2 *Uncertainty Propagation Techniques (See Refs. [10], [17])*

OBSERVED GEOMETRIC FEATURE	FUNCTIONAL FORM	COEFFICIENTS FROM MOMENTS	BOUNDARY PARAMETERIZATION
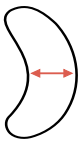 *Quadratic bending along short axis*	$\hat{v} \approx \beta + \alpha \hat{u}^2$	$\alpha = \dfrac{\mathbb{E}[\hat{v}\hat{u}^2]}{\mathbb{E}[\hat{u}^4] - 1},$ $\beta = -\alpha$	$\hat{u}_G = k \cos t$ $\Downarrow$ $\Delta\hat{v}(t) = \alpha\left(k^2 \cos^2 t - 1\right)$
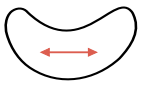 *Quartile shift along long axis*	$\Delta\hat{u}(\hat{u}_G) = a_0 + a_2 \hat{u}_G^2$	$c(k) = \dfrac{k^2 - 1}{6}\mathbb{E}[\hat{u}^3],$ $a_0 = 0,\ a_2 = \dfrac{c(k)}{k^2}$	$\hat{u}_G = k \cos t$ $\Downarrow$ $\Delta\hat{u}(t) = c(k) \cos^2 t$

Figure 3 *Functional Form Method, Applied to Banana-Shaped Distributions*

applications. It has been observed that in many different dynamical contexts in spaceflight, an initially Gaussian distribution can yield similar banana-shaped distributions and confidence boundaries. These include cislunar orbits [10], aerocapture [17], and in the classical Keplerian case [18]. While the underlying physics differ in each case, the same functional forms have been used with success for all cases.

The fundamental approach is now summarized. Consider a two-state slice $\boldsymbol{r}$ with mean $\boldsymbol{\mu}$ and covariance

$$\Sigma = R\Lambda R^T, \qquad \Lambda = \mathrm{diag}(\lambda_1, \lambda_2), \tag{13}$$

where $\lambda_1 \geq \lambda_2$. In normalized principal coordinates,

$$\begin{pmatrix} \hat{u} \\ \hat{v} \end{pmatrix} = W(\boldsymbol{r} - \boldsymbol{\mu}), \qquad W = \Lambda^{-1/2} R^T, \tag{14}$$

the Gaussian confidence contour at level k is simply

$$\hat{u}_G = k \cos t, \qquad \hat{v}_G = k \sin t. \tag{15}$$

For the banana-shaped distributions of interest here, two departures from the Gaussian baseline are proposed. First, bending of the distribution is represented by the lowest-order transverse coupling

$$\hat{v} \approx \beta + \alpha \hat{u}^2. \tag{16}$$

The coefficients are chosen to minimize the mean-square discrepancy,

$$(\alpha, \beta) = \arg\min_{\alpha,\beta} \mathbb{E}\left[\left(\hat{v} - \beta - \alpha\hat{u}^2\right)^2\right]. \tag{17}$$

Since the coordinates are centered and normalized, $\mathbb{E}[\hat{v}] = 0$ and $\mathbb{E}[\hat{u}^2] = 1$, giving

$$\beta = -\alpha, \qquad \alpha = \frac{\mathbb{E}[\hat{v}\hat{u}^2]}{\mathbb{E}[\hat{u}^4] - 1}. \tag{18}$$

Defining necessary partitions of the whitening matrix W,

$$\boldsymbol{a} = W^T \boldsymbol{e}_1, \qquad \boldsymbol{b} = W^T \boldsymbol{e}_2, \tag{19}$$

this becomes

$$\alpha = \frac{\sum_{i,j,k} b_i a_j a_k M^{(3)}_{ijk}}{\sum_{i,j,k,l} a_i a_j a_k a_l M^{(4)}_{ijkl} - 1}. \tag{20}$$

Second, asymmetry along the long axis is modeled from the first-order Cornish–Fisher quantile correction (e.g., [19]). In the normalized long-axis coordinate,

$$c(k) = \frac{k^2 - 1}{6}\,\mathbb{E}[\hat{u}^3] = \frac{k^2 - 1}{6} \sum_{i,j,k} a_i a_j a_k M^{(3)}_{ijk}. \tag{21}$$

Combining the two corrections and composing with the parameterization of a Gaussian circle in whitened coordinates, $u = k\cos t$, $v = k\sin t$, and then undoing the normalization gives the banana contour

$$u(t) = k\sqrt{\lambda_1}\cos t + c(k)\sqrt{\lambda_1}\cos^2 t, \tag{22a}$$

$$v(t) = k\sqrt{\lambda_2}\sin t + \alpha\sqrt{\lambda_2}\left(k^2\cos^2 t - 1\right), \tag{22b}$$

with

$$\boldsymbol{r}(t) = \boldsymbol{\mu} + R\begin{pmatrix} u(t) \\ v(t) \end{pmatrix}. \tag{23}$$

When the projected skewness vanishes and the Gaussian fourth-moment identity is recovered, $c(k) = 0$ and $\alpha = 0$, reducing the result to the standard covariance ellipse. The two corrections to the Gaussian ellipse to account for the banana shape induce a non-convex constraint when trying to enforce the chance constraint, if the concave part of the boundary is facing towards the half-plane. The strategy for robustly enforcing this is briefly discussed in Section 4.2. See Section 5 for simulation results.

Key commentary about this approach is as follows. First, it can be paired with any computationally efficient method of computing higher-order moments, such as DA-based methods of pushing forward the statistical moments [10], or a simple 4th-order CUT approach [5]. Second, this approach is not specific to banana distributions: other recognizable geometric departures may be represented by alternative low-order functions and related to the corresponding statistical moments. Reference [13] develops a particular construction from intuition about observed geometry in astrodynamics problems. This motivates a third observation: the functional form method presupposes that a lot is known about the shape of a distribution, relying on particular informed choices of the engineer. However, expected geometric features can also be directly discerned from the flow of the dynamics – to be discussed in Section 3.2.

Finally, while this derivation is for a 2D slice, Reference [10] introduced a strategy for constructing 3D confidence boundaries by composing multiple 2D parameterizations. The justification is perturbative: because the functional-form corrections are retained only to first order, separate u–v and u–w deformations sharing the same dominant u direction may be superimposed, while products of their correction amplitudes are higher-order terms and are neglected. This allows the projected moments $E[u^2v]$ and $E[u^2w]$ to independently parameterize bending in the two orthogonal transverse directions. This precludes application to extremely non-Gaussian cases, which in general are not our goal.

3.2 A second-order flow map interpretation

We can explain the initial appearance of the banana shape as a consequence of quadratic terms in the flow of the general nonlinear dynamics. In a LinCov setting, the linearized flow shears and stretches the covariance ellipse while preserving its elliptical geometry: lines remain lines, and the propagated covariance remains characterized by straight principal axes. From a differential geometric perspective, the restriction of a second-order flow map to an arbitrary line, such as the principal axis of the initially Gaussian distribution, is generically a quadratic curve under the action of the Hessian of the flow. This perspective already provides some useful intuition. The propagated probability distribution, however, is not determined by this geometric deformation alone: the nonlinear flow also redistributes probability, while projection combines contributions from different directions in the initial uncertainty. The derivation that follows accounts for these effects and shows that the same quadratic structure emerges in the leading-order bending of the distribution regardless of the particular nature of the flow. Thus, the recurring banana geometry can be understood not as a feature specific to any one physical model, but as a common leading-order signature of weak nonlinearity in systems that exhibit a dominant stretching direction.

Preliminaries

Consider one two-state slice of the propagated uncertainty. Let

$$\delta \boldsymbol{x}_0 = L\boldsymbol{\xi}, \qquad \boldsymbol{\xi} \sim \mathcal{N}(\mathbf{0}, I_n), \qquad P_0 = LL^T. \tag{24}$$

Let $\boldsymbol{\varphi}$ denote the nonlinear flow from the initial to the final state, so that $\boldsymbol{x}_f = \boldsymbol{\varphi}(t; \boldsymbol{x}_0, t_0)$, and let $S \in \mathbb{R}^{2\times n}$ select two components of the final state. Define

$$\Phi = \frac{\partial \boldsymbol{\varphi}}{\partial \boldsymbol{x}_0}, \qquad H_i = \frac{\partial^2 [S\boldsymbol{\varphi}]_i}{\partial \boldsymbol{x}_0^2}, \qquad C = S\Phi L, \qquad Q_i = L^T H_i L, \qquad i = 1, 2. \tag{25}$$

The selected final deviation from its mean is $\delta \boldsymbol{y} = S\boldsymbol{x}_f - \mathbb{E}[S\boldsymbol{x}_f]$. Introducing the scalar ε to track the quadratic contribution, truncating the flow at second order:

$$\delta y_i = (C\boldsymbol{\xi})_i + \frac{\varepsilon}{2}\left(\boldsymbol{\xi}^T Q_i \boldsymbol{\xi} - \operatorname{tr} Q_i\right). \tag{26}$$

The mixed linear–quadratic covariance terms can be shown to vanish:

$$\mathbb{E}\left[\xi_r(\boldsymbol{\xi}^T Q_i \boldsymbol{\xi} - \operatorname{tr} Q_i)\right] = 0, \tag{27}$$

thus $\operatorname{Cov}(\delta \boldsymbol{y}) = CC^T + \mathcal{O}(\varepsilon^2)$. Let

$$CC^T = R\Lambda R^T, \qquad \Lambda = \operatorname{diag}(\lambda_1, \lambda_2), \qquad W = \Lambda^{-1/2} R^T, \tag{28}$$

where the columns of R are the principal directions and W is a whitening transformation. With $\boldsymbol{e}_1 = (1, 0)^T$ and $\boldsymbol{e}_2 = (0, 1)^T$, set

$$A = W_{11}Q_1 + W_{12}Q_2, \qquad B = W_{21}Q_1 + W_{22}Q_2, \tag{29a}$$

$$F = \frac{1}{2}(\boldsymbol{\xi}^T A \boldsymbol{\xi} - \operatorname{tr} A), \qquad G = \frac{1}{2}(\boldsymbol{\xi}^T B \boldsymbol{\xi} - \operatorname{tr} B). \tag{29b}$$

Applying W to the linear and second-order deviations yields

$$\begin{pmatrix} X \\ Y \end{pmatrix} = WC\boldsymbol{\xi}, \qquad \mathbb{E}\left[\begin{pmatrix} X \\ Y \end{pmatrix}\begin{pmatrix} X & Y \end{pmatrix}\right] = WCC^T W^T = I_2, \tag{30a}$$

$$\begin{pmatrix} \hat{u} \\ \hat{v} \end{pmatrix} = W\delta \boldsymbol{y} = \begin{pmatrix} X \\ Y \end{pmatrix} + \varepsilon \begin{pmatrix} F \\ G \end{pmatrix}. \tag{30b}$$

Thus X and Y are independent $\mathcal{N}(0, 1)$ variables. We now define some needed projections:

$$\boldsymbol{p}^T = \boldsymbol{e}_1^T WC, \qquad \boldsymbol{q}^T = \boldsymbol{e}_2^T WC, \tag{31a}$$

$$A_{uu} = \boldsymbol{p}^T A \boldsymbol{p}, \qquad A_{uv} = \boldsymbol{q}^T A \boldsymbol{p}, \qquad B_{uu} = \boldsymbol{p}^T B \boldsymbol{p}. \tag{31b}$$

Then one can show $X = \boldsymbol{p}^T\boldsymbol{\xi}$, $Y = \boldsymbol{q}^T\boldsymbol{\xi}$, $\boldsymbol{p}^T\boldsymbol{p} = \boldsymbol{q}^T\boldsymbol{q} = 1$, and $\boldsymbol{p}^T\boldsymbol{q} = 0$.

For fixed $X = u$, write

$$\boldsymbol{\xi} = \boldsymbol{p}u + \boldsymbol{\zeta}, \qquad \mathbb{E}[\boldsymbol{\zeta}] = \mathbf{0}, \qquad \mathbb{E}[\boldsymbol{\zeta}\boldsymbol{\zeta}^T] = I_n - \boldsymbol{p}\boldsymbol{p}^T, \qquad Y = \boldsymbol{q}^T\boldsymbol{\zeta}. \tag{32}$$

Defining $\mathrm{He}_2(s) = s^2 - 1$, the second probabilists' Hermite polynomial, and $F_X \equiv \frac{\partial F}{\partial X} = \boldsymbol{p}^T A\boldsymbol{\xi}$, we compute multiple expectations needed for subsequent calculations:

$$\mathbb{E}[F \mid X = u] = \frac{1}{2}(u^2 A_{uu} + \mathrm{tr}(A(I_n - \boldsymbol{p}\boldsymbol{p}^T)) - \mathrm{tr}\,A) = \frac{A_{uu}}{2}\mathrm{He}_2(u), \tag{33a}$$

$$\mathbb{E}[G \mid X = u] = \frac{B_{uu}}{2}\mathrm{He}_2(u), \tag{33b}$$

$$\mathbb{E}[YF \mid X = u] = \frac{1}{2}\mathbb{E}\left[(\boldsymbol{q}^T\boldsymbol{\zeta})(2u\boldsymbol{p}^T A\boldsymbol{\zeta} + \boldsymbol{\zeta}^T A\boldsymbol{\zeta})\right] = A_{uv}u, \tag{33c}$$

$$\begin{aligned}\mathbb{E}[F_X \mid X = u] &= A_{uu}u,\\ \mathbb{E}[YF_X \mid X = u] &= A_{uv}.\end{aligned} \tag{33d}$$

Transverse bend

The "centerline" of the distribution for a two-state slice may be defined by $m_v(u) = \mathbb{E}[\hat{v} \mid \hat{u} = u]$. At a fixed slice $\hat{u} = u$, $m_v(u)$ is the mean value of $\hat{v}$ with respect to the probability distributed along that slice. Its calculation therefore requires accounting for how the second-order map redistributes initially Gaussian probability, rather than only how individual points are mapped.

Recall that $\hat{u} = X + \varepsilon F$. Fixing the variables orthogonal to $\boldsymbol{p}$, consider a small interval $[u, u + \mathrm{d}u]$ in the final coordinate $\hat{u}$. Then we may write $\mathrm{d}X = \left|\frac{\partial X}{\partial u}\right| \mathrm{d}u$. Since $X \sim \mathcal{N}(0,1)$, the probability contained in $\mathrm{d}X$ is

$$\mathrm{d}P = \rho(X)\,\mathrm{d}X = \rho(X)\left|\frac{\partial X}{\partial u}\right| \mathrm{d}u. \tag{34}$$

Setting $\hat{u} = u$ and perturbatively inverting $u = X + \varepsilon F$ gives

$$X = u - \varepsilon F + \mathcal{O}(\varepsilon^2), \tag{35a}$$

$$\rho(X)\left|\frac{\partial X}{\partial u}\right| = \rho(u)\,(1 + \varepsilon\,(uF - F_X)) + \mathcal{O}(\varepsilon^2), \tag{35b}$$

with the variables orthogonal to $\boldsymbol{p}$ held fixed. The conditional mean is then obtained by averaging over these remaining variables. Applying Eq. (35) together with $\hat{v} = Y + \varepsilon G$ gives

$$\begin{aligned} m_v(u) &= \frac{\mathbb{E}[(Y + \varepsilon G)(1 + \varepsilon(uF - F_X)) \mid X = u]}{\mathbb{E}[1 + \varepsilon(uF - F_X) \mid X = u]} + \mathcal{O}(\varepsilon^2)\\ &= \varepsilon\,(\mathbb{E}[G \mid X = u] + u\mathbb{E}[YF \mid X = u] - \mathbb{E}[YF_X \mid X = u]) + \mathcal{O}(\varepsilon^2)\\ &= \varepsilon\left(A_{uv} + \frac{1}{2}B_{uu}\right)\mathrm{He}_2(u) + \mathcal{O}(\varepsilon^2). \end{aligned} \tag{36}$$

Now we may write $m_v(u) = \alpha\,\mathrm{He}_2(u) + \mathcal{O}(\varepsilon^2)$ where $\alpha = \varepsilon\left(A_{uv} + \frac{1}{2}B_{uu}\right)$. Since $\mathrm{He}_2(u) = u^2 - 1$, writing the same centerline as $\beta + \alpha u^2$ gives $\beta = -\alpha$, recovering the familiar functional form of Eq. (16). Thus, the quadratic form assumed in the prior subsection for the bending mode was *not* an ad hoc choice. It is the leading-order form that appears in the case of weak nonlinearity, when third-order and higher contributions are still subdominant.

To obtain a simple contour model, retain the bending mode derived above while otherwise preserving the Gaussian (X, Y) geometry:

$$\hat{u} = X, \qquad \hat{v} = Y + \alpha\,\mathrm{He}_2(X). \tag{37}$$

Let k be the radius of the reference Gaussian confidence circle and t its angular parameter,

$$X = k\cos t, \qquad Y = k\sin t.$$

Substitution into Eq. (37) recovers the contour form for just the quadratic bending:

$$\hat{u}(t) = k\cos t, \qquad \hat{v}(t) = k\sin t + \alpha(k^2\cos^2 t - 1). \tag{38}$$

Long-axis quantile shift

Let $p_{\hat{u}}(u)$ be the marginal density of $\hat{u}$. Integrating Eq. (35) over the remaining components of $\boldsymbol{\xi}$ gives

$$\begin{aligned} p_{\hat{u}}(u) &= \rho(u)(1+\varepsilon[u\mathbb{E}[F \mid X=u] - \mathbb{E}[F_X \mid X=u]]) + O(\varepsilon^2) \\ &= \rho(u)\left(1+\varepsilon\frac{A_{uu}}{2}(u^3-3u)\right) + O(\varepsilon^2). \end{aligned} \tag{39}$$

Since $\frac{d}{du}(\rho(u)(u^2-1)) = -\rho(u)(u^3-3u)$,

$$\Pr(\hat{u} \le u) = \int_{-\infty}^{u} \rho(s)\,ds - \varepsilon\frac{A_{uu}}{2}\rho(u)(u^2-1) + O(\varepsilon^2). \tag{40}$$

Let $P = \int_{-\infty}^{z} \rho(s)\,ds$ and let the corresponding quantile of $\hat{u}$ be $u_P = z + \Delta(z)$. Substitution into Eq. (40) gives

$$0 = \rho(z)\left(\Delta(z) - \varepsilon\frac{A_{uu}}{2}(z^2-1)\right) + O(\varepsilon^2), \qquad \Delta(z) = \varepsilon\frac{A_{uu}}{2}(z^2-1) + O(\varepsilon^2). \tag{41}$$

Also, using Eq. (33),

$$\mathbb{E}[\hat{u}^2] = 1 + 2\varepsilon\frac{A_{uu}}{2}\mathbb{E}[X(X^2-1)] + O(\varepsilon^2) = 1 + O(\varepsilon^2), \tag{42a}$$

$$\mathbb{E}[\hat{u}^3] = 3\varepsilon\frac{A_{uu}}{2}\mathbb{E}[X^2(X^2-1)] + O(\varepsilon^2) = 3\varepsilon A_{uu} + O(\varepsilon^2). \tag{42b}$$

The standardized skewness of $\hat{u}$ is

$$\gamma_1 \equiv \frac{\mathbb{E}[\hat{u}^3]}{\mathbb{E}[\hat{u}^2]^{3/2}} = 3\varepsilon A_{uu} + O(\varepsilon^2). \tag{43}$$

Hence the map-derived quantile correction can be written as

$$\Delta(z) = \frac{\gamma_1}{6}(z^2-1) + O(\varepsilon^2), \tag{44}$$

which is the first-order Cornish–Fisher correction [19]. At $z = \pm k$, let $c(k)$ denote the longitudinal displacement of the k-level endpoints. Then we recover the expected coefficient

$$c(k) = \frac{\gamma_1}{6}(k^2-1) = \varepsilon\frac{A_{uu}}{2}(k^2-1) + O(\varepsilon^2). \tag{45}$$

Again a choice from Section 3.1 is recovered. We thus observe that consideration of the flow map itself can be used to rigorously select functional forms for describing the distribution and confidence boundaries.

3.3 Direct parameterization of the confidence boundary

While not the primary focus of this work, we note that another possibility is to propagate a representation of the confidence boundary itself, rather than reconstructing it from propagated statistical moments. Assume first that Eq. (1) represents a Hamiltonian system. Let $\rho(\boldsymbol{x}, t)$ denote the probability density associated with the Hamiltonian flow $\boldsymbol{\varphi}_t$. For a fixed density level $\rho_q > 0$, define the superlevel region

$$\mathcal{A}_t(\rho_q) = \{\boldsymbol{x} : \rho(\boldsymbol{x}, t) \ge \rho_q\}, \tag{46}$$

whose boundary is the corresponding iso-density set, $\partial\mathcal{A}_t(\rho_q) = \{\boldsymbol{x} : \rho(\boldsymbol{x}, t) = \rho_q\}$. Provided ρ is sufficiently smooth and $\nabla_{\boldsymbol{x}}\rho \neq \boldsymbol{0}$ on this set, the boundary is a smooth hypersurface. For Hamiltonian dynamics, Liouville's theorem gives $\nabla_{\boldsymbol{x}} \cdot \boldsymbol{f} = 0$. Together with the continuity equation,

$$\frac{\partial\rho}{\partial t} + \nabla_{\boldsymbol{x}} \cdot (\rho \boldsymbol{f}) = 0, \tag{47}$$

this implies that the density is constant along trajectories, $\rho(\boldsymbol{\varphi}_t(\boldsymbol{x}_0), t) = \rho(\boldsymbol{x}_0, 0)$. Consequently, the superlevel region is transported by the flow,

$$\mathcal{A}_t(\rho_q) = \boldsymbol{\varphi}_t\big(\mathcal{A}_0(\rho_q)\big)\,. \tag{48}$$

The probability content enclosed by this material region is likewise preserved. Equivalently, applying the Reynolds transport theorem together with the continuity equation gives

$$\int_{\mathcal{A}_t(\rho_q)} \rho(\boldsymbol{x}, t)\, d\boldsymbol{x} = \int_{\mathcal{A}_0(\rho_q)} \rho(\boldsymbol{x}, 0)\, d\boldsymbol{x} = P. \tag{49}$$

Thus, the density level ρ_q may equivalently be associated with the enclosed confidence level P. Finally, if $\boldsymbol{\varphi}_t$ is locally diffeomorphic over the region of interest, boundaries are mapped to boundaries,

$$\partial\mathcal{A}_t(P) = \boldsymbol{\varphi}_t(\partial\mathcal{A}_0(P))\,. \tag{50}$$

This provides the basis for propagating a confidence boundary directly, rather than reconstructing it from propagated statistical moments. In the context of this work, we assume that the boundary is to be (approximately) transported via Taylor maps of the flow of the dynamics (computed via DA). Figure 4a illustrates the transport.

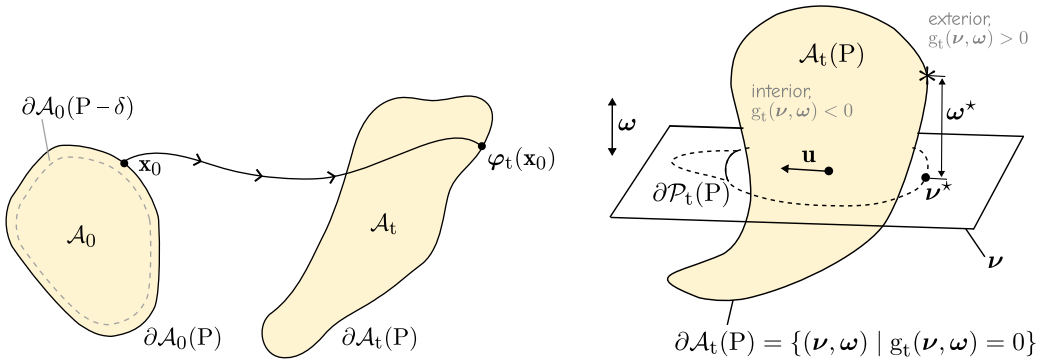

(a) *Direct transport of a confidence boundary under the nonlinear flow*

(b) *Projection of a full-state confidence boundary into operational coordinates (see Section 4.1)*

Figure 4 *Confidence-boundary transport and projection under nonlinear flow.*

We now discuss implicit and explicit parameterizations of the confidence boundary. Suppose first that the initial boundary is locally represented by the implicit equation

$$g_0(\boldsymbol{x}) = 0. \tag{51}$$

If the local inverse of the flow exists, an implicit representation at time t follows immediately by change of variables,

$$g_t(\boldsymbol{x}) = g_0\big(\boldsymbol{\varphi}_t^{-1}(\boldsymbol{x})\big) = 0. \tag{52}$$

For an initially Gaussian distribution, for example, g_0 is simply the familiar quadratic equation defining the covariance ellipsoid. Under nonlinear flow, however, composition with $\boldsymbol{\varphi}_t^{-1}$ generally produces a nonlinear boundary of correspondingly higher complexity.

Alternatively, let $\boldsymbol{x}_0(\boldsymbol{s})$ directly parameterize the initial boundary. The propagated boundary is then obtained using only the forward flow,

$$\boldsymbol{x}_t(\boldsymbol{s}) = \varphi_t(\boldsymbol{x}_0(\boldsymbol{s}))\,. \tag{53}$$

For example, if the initial uncertainty is Gaussian with $P_0 = LL^T$, its k-level boundary may be written

$$\boldsymbol{x}_0(\boldsymbol{s}) = \boldsymbol{\mu}_0 + kL\boldsymbol{s}, \qquad \boldsymbol{s} \in \mathbb{S}^{n-1}. \tag{54}$$

The same parameter $\boldsymbol{s}$ therefore continues to label corresponding points of the boundary after propagation. A polynomial flow map such as those discussed in Section 2 may be substituted directly for $\boldsymbol{\varphi}_t$, making this construction especially simple when an accurate local map is already available.

The implicit and parametric descriptions are alternative representations of the same transported boundary, and are naturally related. The parameterization approach avoids inversion of the flow map and provides the boundary directly through forward evaluation. We have noted that the implicit approach may face numerical challenges related to the region of validity of the Taylor map for the full composition, Eq. (52). For a Gaussian boundary under linear flow, however, Eq. (52) reduces to the usual linear covariance construction. Although we do not pursue this approach further here, it provides another perspective on confidence boundaries independent from the moment-based constructions of Sections 3.1–3.2, while similarly making use of DA. However, having relaxed the assumption of weak non-Gaussianity, the marginalization process for a distribution in full state space down into lower-dimensional operational coordinates must now be mentioned, along with the related problem of projecting the boundary.

4 Projection and constraint geometry

This section addresses common challenges for boundary parameterizations and approximations for non-Gaussian probability distributions. First, we discuss challenges with boundary projection into lower-dimensional operational coordinates (which affects only significantly non-Gaussian distributions), and nonconvexities encountered in the use of the confidence boundaries in enforcement of inequalities (which affect all methods).

4.1 Projection into operational coordinates

The confidence boundaries discussed thus far are naturally defined in the state space in which the uncertainty is propagated. In many applications, however, the operational constraint depends on only a subset of the state. For example, a collision or keep-out constraint may depend on position but not velocity. Consider therefore a partition

$$\boldsymbol{x} = \begin{pmatrix} \boldsymbol{\nu} \\ \boldsymbol{\omega} \end{pmatrix}, \tag{55}$$

where $\boldsymbol{\nu}$ denotes the coordinates of operational interest and $\boldsymbol{\omega}$ the remaining coordinates. Suppose that a full-state confidence region is represented implicitly as

$$\mathcal{A}_t(P) = \{(\boldsymbol{\nu}, \boldsymbol{\omega}) \,:\, g_t(\boldsymbol{\nu}, \boldsymbol{\omega}) \leq 0\}. \tag{56}$$

Its projection into the operational coordinates is

$$\mathcal{P}_t(P) = \{\boldsymbol{\nu} \,:\, \exists\, \boldsymbol{\omega} \text{ where } g_t(\boldsymbol{\nu}, \boldsymbol{\omega}) \leq 0\}. \tag{57}$$

Determining the boundary of Eq. (57) is not equivalent to simply projecting arbitrary points from $\partial\mathcal{A}_t(P)$. Many such points may project into the interior of $\mathcal{P}_t(P)$. More generally, boundaries of lower-dimensional images of nonlinear maps may arise at critical points of the corresponding map, and candidate boundary branches need not all belong to the true outer boundary. Ref. [20], a reachability study, explores this and other relevant projective phenomena. In the present case, the projection $(\boldsymbol{\nu}, \boldsymbol{\omega}) \longmapsto \boldsymbol{\nu}$ is itself full rank, so the projected boundary originates from $\partial\mathcal{A}_t(P)$. At a smooth candidate point $(\boldsymbol{\nu}_*, \boldsymbol{\omega}_*)$ on its projected envelope, however, the restriction of this projection to the full-state boundary must become critical. For the implicit representation above, this gives

$$g_t(\boldsymbol{\nu}_*, \boldsymbol{\omega}_*) = 0, \qquad \nabla_{\omega} g_t(\boldsymbol{\nu}_*, \boldsymbol{\omega}_*) = \mathbf{0}. \tag{58}$$

An equivalent interpretation follows from considering an extremum of the projected set in some unit direction $\boldsymbol{u}$. The constrained problem

$$\min_{\boldsymbol{\nu}, \boldsymbol{\omega}} \boldsymbol{u}^T \boldsymbol{\nu} \qquad \text{s.t.} \qquad g_t(\boldsymbol{\nu}, \boldsymbol{\omega}) \leq 0 \tag{59}$$

has, at a regular active solution, the first-order conditions

$$g_t(\boldsymbol{\nu}_*, \boldsymbol{\omega}_*) = 0, \tag{60a}$$

$$\nabla_{\omega} g_t(\boldsymbol{\nu}_*, \boldsymbol{\omega}_*) = \mathbf{0}, \tag{60b}$$

$$\boldsymbol{u} + \lambda_{\boldsymbol{u}} \nabla_{\boldsymbol{\nu}} g_t(\boldsymbol{\nu}_*, \boldsymbol{\omega}_*) = \mathbf{0}, \qquad \lambda_{\boldsymbol{u}} > 0. \tag{60c}$$

Thus, $\nabla_\omega g_t = \mathbf{0}$ identifies the points at which the full-state boundary becomes tangent to the eliminated coordinates, while the final condition aligns the projected boundary normal with $\boldsymbol{u}$. Multiple stationary branches may satisfy these equations, and not every branch need form part of the desired outer boundary. This issue, together with the additional complications introduced by non-convexity, is addressed in Section 4.2. Figure 4b, in the previous section, provides a conceptual illustration of these concepts, whereby the coordinates to be marginalized, $\boldsymbol{\omega}$, are assumed one-dimensional to facilitate a three-dimensional cartoon.

The familiar linear covariance result follows as a special case. For a Gaussian state let

$$g_t(\boldsymbol{x}) = \begin{pmatrix} \delta\boldsymbol{\nu} \\ \delta\boldsymbol{\omega} \end{pmatrix}^T \Sigma^{-1} \begin{pmatrix} \delta\boldsymbol{\nu} \\ \delta\boldsymbol{\omega} \end{pmatrix} - \kappa \tag{61}$$

and partition $\Lambda \equiv \Sigma^{-1}$ according to the $(\boldsymbol{\nu}, \boldsymbol{\omega})$ state partition. Applying $\nabla_\omega g_t = \mathbf{0}$ gives

$$\delta\boldsymbol{\omega}^* = -\Lambda_{\omega\omega}^{-1}\Lambda_{\omega\nu}\delta\boldsymbol{\nu}. \tag{62}$$

Substitution into the quadratic form yields $\delta\boldsymbol{\nu}^T \left(\Lambda_{\nu\nu} - \Lambda_{\nu\omega}\Lambda_{\omega\omega}^{-1}\Lambda_{\omega\nu}\right)\delta\boldsymbol{\nu} = \kappa$. By the Schur-complement identity,

$$\delta\boldsymbol{\nu}^T\Sigma_{\nu\nu}^{-1}\delta\boldsymbol{\nu} = \kappa. \tag{63}$$

Eq. (63) recovers the usual geometric projection of the covariance ellipsoid.

4.2 Nonconvexity and support

Reference [13] discusses extensively the numerical details of constraint reinforcement for a parameterized non-convex confidence boundary, and the interested reader is directed there. For constraint enforcement, it is useful to view a confidence region through its support function. Consider a local half-space constraint

$$\boldsymbol{n}^T\boldsymbol{\nu} - b_0 \leq 0. \tag{64}$$

For a projected confidence region $\mathcal{P}_t(P;\boldsymbol{\chi})$, where $\boldsymbol{\chi}$ denotes the relevant guidance or optimization variables, exact geometric enforcement is equivalent to

$$g(\boldsymbol{\chi}) \equiv \max_{\boldsymbol{\nu}\in\mathcal{P}_t(P;\boldsymbol{\chi})} \boldsymbol{n}^T\boldsymbol{\nu} - b_0 \leq 0. \tag{65}$$

Thus, only the point or points attaining the support in direction $\boldsymbol{n}$ are relevant to a given half-space constraint. Moreover, the support of a set is identical to that of its convex hull. Concave portions of a non-convex confidence region therefore do not directly affect Eq. (65); their practical effect instead appears through changes in which boundary point is active.

For an explicitly parameterized boundary $\boldsymbol{\nu} = \boldsymbol{\nu}(\boldsymbol{\chi}, s)$, Eq. (65) reduces to a low-dimensional maximization over the boundary parameter s. Let $s^\star$ denote the point producing the maximum constraint violation. If this maximizer is unique and smooth, then

$$g(\boldsymbol{\chi}) = \boldsymbol{n}^T\boldsymbol{\nu}(\boldsymbol{\chi}, s^\star) - b_0. \tag{66}$$

As the techniques in this work are developed for use with nonlinear dynamics, in general their implementation in maneuver and trajectory optimization is expected to be sequential in nature, whereby derivatives of Eq. (66) with respect to the current iterate $\boldsymbol{\chi}[k]$ are used to inform an update on the free variables $\Delta\boldsymbol{\chi}[k+1]$:

$$\nabla_\chi g = \nabla_\chi \left[\boldsymbol{n}^T\boldsymbol{\nu}(\boldsymbol{\chi}, s^\star) - b_0\right]. \tag{67}$$

The complication arises when two or more boundary points give the same support value. The active point may then switch from one branch of the boundary to another, producing a non-smooth change in the constraint gradient. Reference [13] examines this behavior for the banana-shaped boundaries of Section 3.1: the relevant switching occurs between the two ends of the contour in a narrow range of constraint orientations. Both local gradient smoothing and conservative smooth surrogates were proposed to enforce the resulting chance constraint.

5 Numerical examples

5.1 Sample-free moment recovery via DA, & 3D boundaries

We revisit a scenario first introduced in Reference [10]: an Apophis encounter with propagation of initially Gaussian spacecraft state uncertainty until it becomes fairly non-Gaussian, using the methods of Sections 2 and 3.1. Details on the spacecraft characteristics, initial states, and covariances for the analyzed scenarios are publicly available in Reference [10]. The perturbed dynamics incorporate spherical harmonics, third-body perturbations, and solar radiation pressure. Figure 5a shows the trajectory and confidence boundary computed via a LinCov approach, unscented transform-based covariance ellipse, and the banana-shaped contour developed via the functional forms method. Figure 5b details the computational time breakdown for the proposed methodologies compared to other approaches. Notably, the online evaluation of the sample-free moment pushforward is rendered highly efficient by the increased sparsity of the Hermite weight tensors (second-to-last-column), and the targeted computation of only the relevant terms using the "project-then-product" strategy (final column). While the offline weight tensor setup (computed once) and map computation (computed with each reference trajectory change) take several seconds to complete, real-time UQ operation is well within an order of magnitude of LinCov and unscented transform-based approaches, despite being much more accurate.

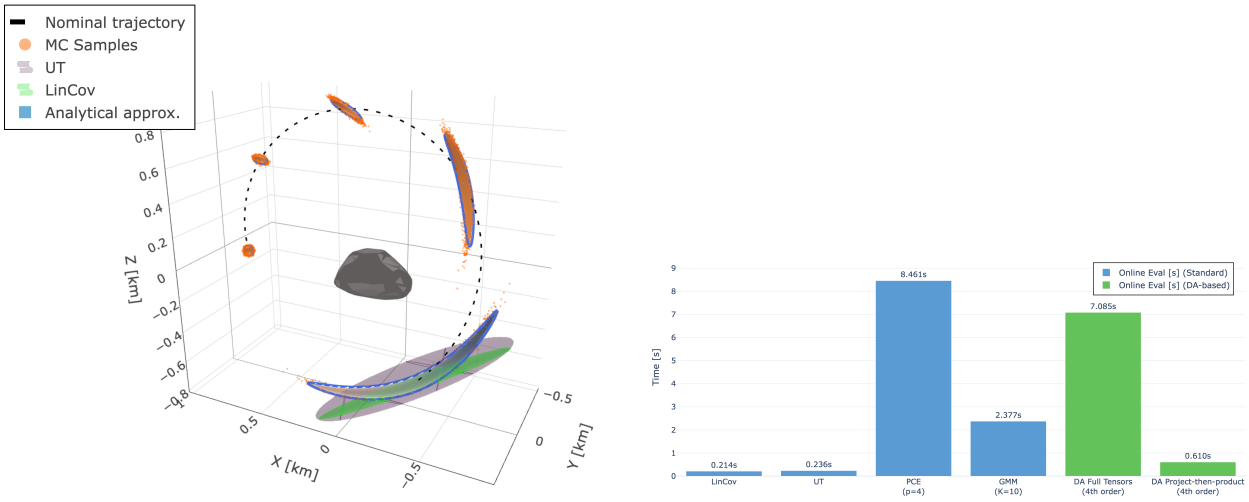

(a) *Apophis trajectory and uncertainty evolution* **(b)** *Computational performance*

Figure 5 *Apophis Example Numerical Results*

5.2 Lunar free-return stochastic maneuver design

We next revisit a stochastic maneuver planning example based on a planar CR3BP analog of an Artemis II-like lunar free-return. All parameters are publicly available in Reference [13]. Two correction maneuvers are considered on the post-lunar-flyby return leg, with navigation and maneuver execution errors treated statistically. The first maneuver occurs 0.25 days after lunar flyby, with 3.62 days remaining until Earth entry. The goal is to correct the initial dispersion as much as possible without violating *passive safety* at Earth encounter, should the second retargeting maneuver be missed. Four one-sided 3σ chance constraints enforce the flyby/arrival passive safety corridor (which is much wider than an entry corridor):

$$150 < h < 400 \text{ km}, \qquad -6^\circ < \gamma < 8^\circ, \tag{68}$$

where h and γ denote entry altitude and flight-path angle, respectively. Maneuver execution uncertainty is incorporated through a Gates-type model [21], such that the commanded maneuver affects both the nominal trajectory and post-maneuver covariance.

A naive retargeting maneuver requires 6.485 m/s and violates the arrival constraints in 62.86% of a 5000-sample Monte Carlo validation. LinCov reduces the maneuver to 3.352 m/s and the observed violation fraction to 0.56%,

but the propagated distribution is sufficiently non-Gaussian that the covariance ellipse does not accurately describe its outer geometry. The banana policy instead requires 2.867 m/s with a 0.20% violation fraction. Figure 6 shows the resulting trajectory and arrival dispersions. Note a fair number of shallow encounter failure cases from the LinCov policy in subfigure 6c. By contrast, the moment-informed confidence boundary closely follows the Monte Carlo distribution in both position and velocity, while the corresponding covariance ellipses visibly misrepresent its shape. Relative to LinCov, the resulting maneuver is 14.4% smaller and the observed violation fraction is reduced by approximately 64%. The banana confidence boundary-based optimization required 4.29 s, compared with 0.25 s for LinCov and 27.6 s for the 5000-sample Monte Carlo validation. A second correction performed only five hours before entry produces a nearly Gaussian final dispersion; in this case the LinCov and banana policies essentially coincide, as expected. Those results are omitted for the sake of brevity.

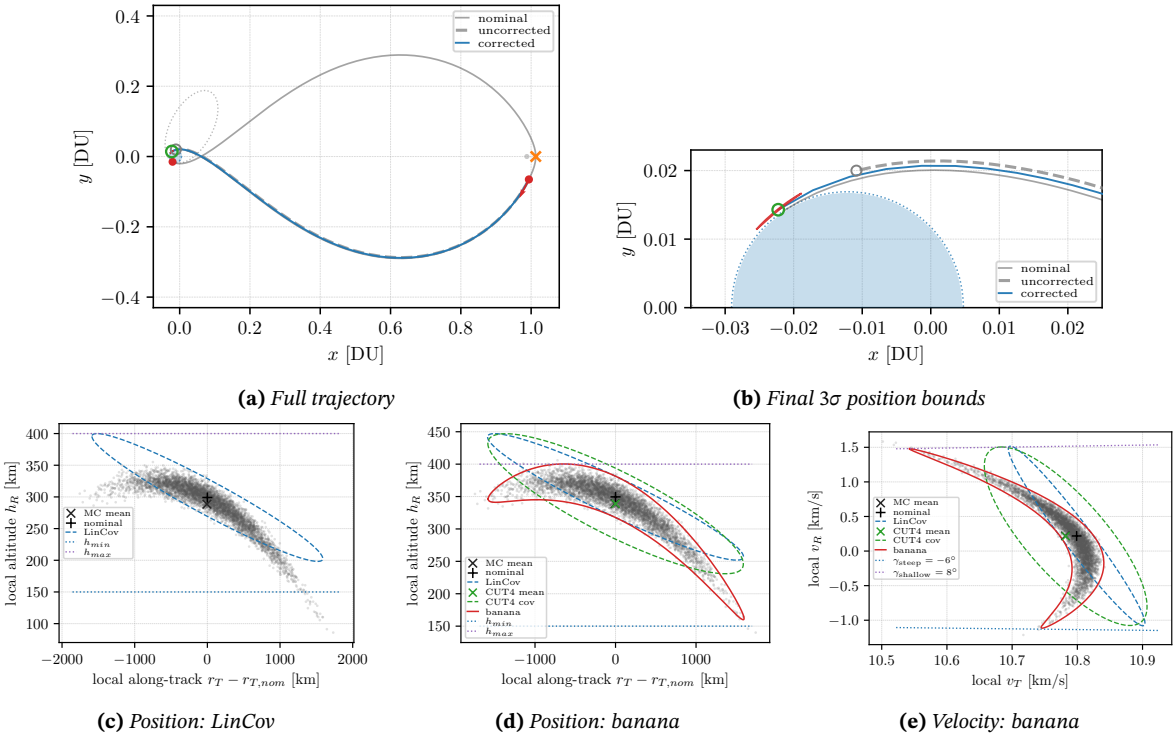

(a) *Full trajectory* **(b)** *Final 3σ position bounds*

(c) *Position: LinCov* **(d)** *Position: banana* **(e)** *Velocity: banana*

Figure 6 *Lunar free-return stochastic maneuver design and arrival constraint satisfaction.*

6 Discussion

The methods considered in this paper occupy a useful computational middle ground. Higher-order moments can efficiently inform low-order corrections to the Gaussian confidence boundary; the analysis of Section 3.2 provides a dynamical basis for this construction under weak nonlinearity. For stronger deformation, direct propagation of the boundary provides a more general alternative, but introduces additional challenges of projection and non-convex constraint enforcement. In either case, polynomial flow maps provide a common computational foundation, though it is worth reiterating that they are not explicitly required by the functional form method.

References

[1] A. W. Berning Jr., E. R. Burnett, and S. Bieniawski, "Chance-Constrained, Drift-Safe Guidance for Spacecraft Rendezvous," *AAS Rocky Mountain Guidance, Navigation, and Control Conference*, American Astronautical Society, 2023, 10.48550/arXiv.2401.11077.

[2] K. Oguri, "Chance-Constrained Control for Safe Spacecraft Autonomy: Convex Programming Approach," *2024 American Control Conference (ACC)*, 2024, pp. 2318–2324, 10.48550/arXiv.2403.04062.

[3] D. J. Scheeres, *Orbital motion in strongly perturbed environments: applications to asteroid, comet and planetary satellite orbiters*. Springer, 2016.

[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)*, 2012, pp. 5556–5561, 10.1109/ACC.2012.6314970.

[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, 10.2514/1.57599.

[7] S. Boone and J. W. McMahon, "Non-Gaussian Chance-Constrained Trajectory Control Using Gaussian Mixtures and Risk Allocation," *2022 IEEE 61st Conference on Decision and Control (CDC)*, 2022, 10.1109/CDC51059.2022.9993274.

[8] M. Valli, R. Armellin, P. Di Lizia, and M. R. Lavagna, "Nonlinear Mapping of Uncertainties in Celestial Mechanics," *Journal of Guidance, Control, and Dynamics*, Vol. 36, No. 1, 2013, pp. 48–63, 10.2514/1.58068.

[9] 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.

[10] N. Michelotti, E. R. Burnett, and F. Topputo, "Analytical Confidence Boundaries for Non-Gaussian Uncertainty in Perturbed Spacecraft Dynamics," *AAS/AIAA Astrodynamics Specialist Conference*, No. AAS 26-909, American Astronautical Society, July 2026.

[11] J. A. Rebollo, R. Vázquez, and C. Bombardelli, "Closed-Form Solutions to the Fokker-Planck Equation for Orbital Uncertainty Propagation," *AAS/AIAA Astrodynamics Specialist Conference*, No. AAS 26-861, American Astronautical Society, July 2026.

[12] E. Burnett and S. Boone, "Analytic Non-Gaussian Confidence Boundary Method for Chance-Constrained Trajectory Control," arXiv, 2026, 10.48550/arXiv.2604.04304.

[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] R. S. Park and D. J. Scheeres, "Nonlinear Mapping of Gaussian Statistics: Theory and Applications to Spacecraft Trajectory Design," *Journal of Guidance, Control, and Dynamics*, Vol. 29, No. 6, 2006, pp. 1367–1375, 10.2514/1.20177.

[15] S. Boone and J. McMahon, "Orbital Guidance Using Higher-Order State Transition Tensors," *Journal of Guidance, Control, and Dynamics*, Vol. 44, No. 3, 2021, pp. 493–504, 10.2514/1.G005493.

[16] M. Berz, "Chapter 2 - Differential Algebraic Techniques," *Modern Map Methods in Particle Beam Physics*, Vol. 108 of *Advances in Imaging and Electron Physics*, pp. 81–117, San Diego, CA: Academic Press, 1999, 10.1016/S1076-5670(08)70228-3.

[17] E. R. Burnett and S. Boone, "Efficient Nonlinear Uncertainty Quantification for Spaceflight Leveraging Nonlinear Expansions," *International Symposium on Space Flight Dynamics*, 2026, 10.48550/arXiv.2605.24147.

[18] J. L. Junkins, "Adventures on the Interface of Dynamics and Control," *AIAA Journal of Guidance, Control, and Dynamics*, Vol. 20, Nov.–Dec. 1997, pp. 1058–1071, 10.2514/2.4176.

[19] M. Abramowitz and I. A. Stegun, *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. New York: Dover Publications, Inc., 1965.

[20] I. Lyss-Loren and C. Frueh, "State-Transition-Tensor-Based Reachability Methods Applied to the Cislunar Region," *AAS/AIAA Astrodynamics Specialist Conference*, No. AAS 26-686, American Astronautical Society, July 2026.

[21] C. R. Gates, "A Simplified Model of Midcourse Maneuver Execution Errors," Tech. Rep. JPL-TR-32-504, NASA-CR-53032, Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, Oct. 1963.

© *Ethan R. Burnett, Niccolò Michelotti, and Francesco Topputo*