

MONTE CARLO CONJUNCTION ANALYSIS USING IMPORTANCE SAMPLING AND TAYLOR POLYNOMIAL EXPANSION

Shripad Sant*, Ryan Ketzner†, Tarek A. Elgohary‡

Accurate and efficient collision probability estimation is critical for conjunction analysis. While traditional Monte Carlo methods provide unbiased estimates, they are computationally expensive, especially in rare-event scenarios where collision probabilities are very low. This article proposes a fast surrogate-assisted Monte Carlo framework for rare-event estimation that combines high-order Taylor expansions with importance sampling guided by the cross-entropy method. The Taylor expansions, constructed using generalized dual numbers, provide analytical surrogates that approximate the nonlinear mapping from uncertain initial conditions to the distance and time of closest approach, enabling rapid evaluation of propagated samples without repeatedly solving the equations of motion numerically. Importance sampling is used to bias the sample distribution toward the most relevant regions of the initial probability density functions, reducing the number of samples needed and the variance of the resulting estimate. Numerical experiments are conducted on representative short-duration conjunction scenarios drawn from the CelesTrak SOCRATES catalog and long-duration scenarios from Alfano's test set. The results demonstrate that surrogate-assisted methods can be combined with cross-entropy importance sampling without any reduction in accuracy of the probability of collision estimate, while substantially reducing computational cost.

INTRODUCTION

The proliferation of artificial satellites, particularly due to collisions and the recent deployment of several large mega-constellations, has resulted in an increasingly congested orbital environment, leading to a substantial rise in the frequency of conjunction events between active satellites and space debris. To prevent further damage to the space environment, it is crucial to monitor such events closely. NASA's Conjunction Assessment and Risk Analysis (CARA) program highlights the importance of timely assessment and mitigation of such threats to ensure the safety and sustainability of space operations.¹ Conjunction analysis relies on propagating the orbits of space objects and evaluating the collision risk over a given time window. As the number of conjunction events continue to rise, it is more critical than ever to improve both the speed and accuracy of conjunction assessments to prevent catastrophic collisions and minimize unnecessary collision avoidance maneuvers.

Conjunctions are generally classified according to the duration of the encounter (i.e., the time interval of substantial probability accumulation), and specialized analytical and numerical methods

*Doctoral Student, Mechanical and Aerospace Engineering, University of Central Florida, 12760 Pegasus Drive, Orlando, FL 32816.

†Doctoral Student, Mechanical and Aerospace Engineering, University of Central Florida, 12760 Pegasus Drive, Orlando, FL 32816.

‡Associate Professor, Mechanical and Aerospace Engineering, University of Central Florida, 12760 Pegasus Drive, Orlando, FL 32816 and AIAA Member.

have been developed for both short and long-encounter conjunctions. Short encounters occur on times scales of no more than a few seconds, and were the focus of much of the early work on probability of collision (PC) computation, since they are simpler to model and occur more often in practice than long encounters. The most commonly used methods for computing the PC for short encounters rely on an integral in the relative position at the time of closest approach (TCA).

$$P_c = \frac{1}{2\pi\sqrt{|\Sigma_m|}} \int \int_A \exp(-\frac{1}{2}\Delta r_m^T \Sigma_m^{-1} \Delta r_m) dA \quad (1)$$

Where, r_m is the relative position vector in the encounter plane (normal to the relative velocity direction), and Σ_m is the relative position covariance matrix marginalized in the relative velocity direction. The solution to Equation 1 has been the focus of most research on short-encounter PC estimation, as developed by Foster *et al.*,² Akella and Alfriend,³ Patera,⁴ Alfano,⁵ Chan,⁶ Xu *et al.*,⁷ Serra *et al.*,⁸ and Garcia-Pelayo *et al.*,⁹ among others. These works are discussed and compared in detail in the recent review article by Xu *et al.*⁷

Another class of analytical and semi-analytical methods have been developed for conjunctions where the relative motion of the satellites near the TCA cannot be credibly modeled as linear, which is generally the case for long-duration conjunctions. These methods compute an integral of the collision rate over time by considering the flux of trajectories over a sphere fixed in the relative position space of the two objects. This approach allows for arbitrary relative motion between the two satellites, and models the effect of uncertainty in the velocity over the course of the encounter. This technique was first introduced by Coppola,¹⁰ and later extended to model more complex scenarios (such as non-Gaussian distributions) by DeMars *et al.*,¹¹ Shelton and Junkins,¹² and Hall.¹³ Hall's method is used operationally by NASA to estimate PC for long-encounter conjunctions¹⁴.

Although the rate-based methods are generally more capable than the earlier 2D-PC methods, they still require several simplifying assumptions. First, the probability density function must be represented in Cartesian coordinates to integrate in the relative position space, which requires a transformation from the more commonly used equinoctial representation.¹³ This is typically accomplished via a linear transformation of the covariance, which can potentially bias the resulting PC value severely.¹² Additionally, the rate-based methods are derived using the explicit assumption that trajectories may only collide a single time, so they struggle to effectively model multi-encounter conjunctions.^{1,10}

When any of these simplifying assumptions break down, Monte Carlo simulation becomes the tool of choice, since it in principal requires virtually no simplifying assumptions. NASA's operational conjunction assessment tool employs a multi-step approach which applies PC estimation algorithms in order of increasing fidelity, starting with 2D-PC, then moving onto the rate-based methods, and finally, if necessary, employing brute-force Monte Carlo simulation when it is determined that no other method can accurately model the conjunction.¹⁴ A typical brute-force Monte Carlo simulation takes several hours to estimate the PC, since it involves propagating millions of satellite trajectories individually using a high-fidelity propagator.¹

Strategies for reducing the runtime of these simulations generally fall into one of two categories: either to reduce the propagation time for each sample by employing a surrogate model for the propagation step, or to reduce the number of samples using rare-event estimation techniques. Surrogate models previously employed in the literature have included neural network propagators,¹⁵ polynomial chaos expansion,¹⁶⁻¹⁸ and Taylor polynomial expansions obtained via automatic differ-

entiation.^{19–21} Of these, Taylor polynomial expansions have proved especially promising for Monte Carlo PC estimation. Morselli et al. demonstrate that, for typical initial uncertainties and a maximum propagation window of one week, the average error in the Taylor expansion of the distance of closest approach between two satellites was less than one meter using a third-order expansion.¹⁹ They later demonstrated that Taylor polynomial expansions could be used to reduce the runtime for short-encounter PC estimation by several orders of magnitude without compromising accuracy.²⁰

On the other hand, rare event estimation techniques applied to Monte Carlo conjunction analysis have included importance sampling,²² line sampling,²⁰ and subset sampling.^{20,23,24} Importance sampling is especially appealing since it requires only a small deviation from the conventional Monte Carlo method: rather than drawing samples directly from the original probability density function (called the *target*), samples are instead drawn from a proxy distribution (called the *proposal*) which exaggerates the frequency of the rare event. Then, an unbiased estimate of the PC can be recovered by weighting these samples by the likelihood ratio of the target and proposal. Ketzner and Elgohary describe how the cross-entropy method can be used to find the optimal importance sampling distribution for PC estimation, thereby drastically reducing the number of samples required for convergence.²⁵

In this work, we propose using a Taylor polynomial surrogate model together with importance sampling and the cross-entropy method to rapidly compute the PC. The multivariate Taylor maps are generated using generalized dual numbers, which encode the system’s sensitivity to small perturbations in the initial states and final time. When embedded directly into the Monte Carlo framework, these expansions allow efficient computation of the time and distance of closest approach for each sample using polynomial evaluation only. This provides a scalable and accurate foundation for high-fidelity conjunction analysis with significantly reduced computational effort. Since the accuracy of a Taylor polynomial expansion quickly decreases away from the expansion point, the choice of the expansion point is of interest. In the absence of additional information, it is reasonable to center the expansions about the mean of the target PDF. However, the optimal proposal distribution produced by the cross-entropy method is centered closer to the event set, which is the set of all points in the initial random variable space leading to collision. Thus, centering the Taylor polynomial expansion on the mean of the importance sampling distribution could potential improve performance by insuring that the Taylor polynomial is most accurate in the regions where collisions are likely to occur. The effect of the choice of expansion point on accuracy and performance is evaluated in our numerical experiments.

The remainder of the article is organized as follows. In the first section, we describe the use of multivariate Taylor expansions in orbit propagation and their automatic computation using the truncated Taylor algebra. In the second section, we give an overview of importance sampling and the cross entropy method. Next, we demonstrate the performance and accuracy of our method combining Taylor polynomials with importance sampling on both short-duration and long-duration conjunctions; the output of the importance sampling estimator and Taylor-accelerated importance sampling estimator are compared to assess the performance and efficiency of each method.

MULTIVARIATE TAYLOR EXPANSIONS AND THE TRUNCATED TAYLOR ALGEBRA

The Taylor series expansion is a mathematical tool that approximates a smooth function about the expansion point using the derivatives of the function. While the one-dimensional Taylor expansion uses successive derivatives of the function with respect to a single variable, the multivariate Taylor expansion generalizes the concept to a function of several variables. This is specifically useful when

analyzing systems where the outputs depend on multiple inputs. The second-order multivariate Taylor series approximation of a scalar function near a point $\mathbf{x}_0$ can be expressed compactly using the gradient and Hessian as follows²⁶

$$f(\mathbf{x}) \approx f(\mathbf{x}_0) + \nabla f(\mathbf{x}_0)^T (\mathbf{x} - \mathbf{x}_0) + \frac{1}{2} (\mathbf{x} - \mathbf{x}_0)^T H_f(\mathbf{x}_0) (\mathbf{x} - \mathbf{x}_0) \quad (2)$$

where ∇f is the gradient vector and H_f is the Hessian matrix of second partial derivatives.

This general formulation is now applied to the propagation of a satellite's state using a nonlinear mapping based on the two-body problem dynamics. The motion of a satellite can be described by a state $\mathbf{y} = [x, y, z, v_x, v_y, v_z]^T \in \mathbb{R}^6$. Assuming the two-body problem dynamics, this state evolves over time under the influence of gravity. Let $\mathbf{y}_0$ denote the state of the satellite at time $t = 0$, and $\mathbf{y}(T)$ the final state after a fixed time interval T . The evolution of the satellite state over time T can be described by a nonlinear function

$$\mathbf{y}(T) = \Phi(\mathbf{y}_0) \quad (3)$$

where $\Phi : \mathbb{R}^6 \rightarrow \mathbb{R}^6$ represents the solution of the integrated equations of motion from $t = 0$ to $t = T$. Estimates of the satellite's state are subject to errors that are most often modeled as a Gaussian distribution about a nominal trajectory. Propagating this uncertainty by drawing many samples and numerically integrating each initial state is computationally expensive. Instead, employing a high-order Taylor series approximation to map initial deviation $\Delta \mathbf{y}_0$ directly is far more efficient.

$$\Delta \mathbf{y}_0 = \mathbf{y}_0 - \mathbf{y}_{\text{ref}} \quad (4)$$

Substituting into the function Φ , the final state becomes

$$\mathbf{y}(T) = \Phi(\mathbf{y}_0) = \Phi(\mathbf{y}_{\text{ref}} + \Delta \mathbf{y}_0) \quad (5)$$

Assuming the function Φ is smooth, it is expanded using a multivariate Taylor series about the nominal reference state $\mathbf{y}_{\text{ref}}$, giving

$$\mathbf{y}(T) = \Phi(\mathbf{y}_{\text{ref}} + \Delta \mathbf{y}_0) \approx \Phi(\mathbf{y}_{\text{ref}}) + D\Phi(\mathbf{y}_{\text{ref}})\Delta \mathbf{y}_0 + \frac{1}{2} \Delta \mathbf{y}_0^T D^2\Phi(\mathbf{y}_{\text{ref}})\Delta \mathbf{y}_0 + \dots \quad (6)$$

where

- $D\Phi(\mathbf{y}_{\text{ref}}) \in \mathbb{R}^{6 \times 6}$ is the Jacobian matrix

$$D\Phi(\mathbf{y}_{\text{ref}}) = \frac{\partial \Phi_i}{\partial \mathbf{y}_j} = \begin{bmatrix} \nabla \phi_1(\mathbf{y}_{\text{ref}}) \\ \nabla \phi_2(\mathbf{y}_{\text{ref}}) \\ \nabla \phi_3(\mathbf{y}_{\text{ref}}) \\ \nabla \phi_4(\mathbf{y}_{\text{ref}}) \\ \nabla \phi_5(\mathbf{y}_{\text{ref}}) \\ \nabla \phi_6(\mathbf{y}_{\text{ref}}) \end{bmatrix} = \begin{bmatrix} \frac{\partial \phi_1}{\partial x_{\text{ref}}} & \frac{\partial \phi_1}{\partial y_{\text{ref}}} & \dots & \frac{\partial \phi_1}{\partial v_{z_{\text{ref}}}} \\ \frac{\partial \phi_2}{\partial x_{\text{ref}}} & \frac{\partial \phi_2}{\partial y_{\text{ref}}} & \dots & \frac{\partial \phi_2}{\partial v_{z_{\text{ref}}}} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial \phi_6}{\partial x_{\text{ref}}} & \frac{\partial \phi_6}{\partial y_{\text{ref}}} & \dots & \frac{\partial \phi_6}{\partial v_{z_{\text{ref}}}} \end{bmatrix} \quad (7)$$

- $D^2\Phi(\mathbf{y}_{\text{ref}}) \in \mathbb{R}^{6 \times 6 \times 6}$ is a third-order tensor

$$[D^2\phi(\mathbf{y}_{\text{ref}})]_{i,j,k} = \frac{\partial^2 \Phi_i}{\partial \mathbf{y}_{\text{ref},k} \partial \mathbf{y}_{\text{ref},j}} = \{H_1 \quad H_2 \quad H_3 \quad H_4 \quad H_5 \quad H_6\}^T \quad (8)$$

$$\begin{aligned}
H_1 &= \begin{bmatrix} \frac{\partial}{\partial x_{\text{ref}}} \\ \frac{\partial}{\partial y_{\text{ref}}} \\ \frac{\partial}{\partial z_{\text{ref}}} \\ \frac{\partial v_{x_{\text{ref}}}}{\partial} \\ \frac{\partial v_{y_{\text{ref}}}}{\partial} \\ \frac{\partial v_{z_{\text{ref}}}}{\partial} \end{bmatrix} \otimes \begin{bmatrix} \frac{\partial \phi_1}{\partial x_{\text{ref}}} & \frac{\partial \phi_1}{\partial y_{\text{ref}}} & \frac{\partial \phi_1}{\partial z_{\text{ref}}} & \frac{\partial \phi_1}{\partial v_{x_{\text{ref}}}} & \frac{\partial \phi_1}{\partial v_{y_{\text{ref}}}} & \frac{\partial \phi_1}{\partial v_{z_{\text{ref}}}} \end{bmatrix} \\
&= \begin{bmatrix} \frac{\partial^2 \phi_1}{\partial x_{\text{ref}}^2} & \frac{\partial^2 \phi_1}{\partial x_{\text{ref}} \partial y_{\text{ref}}} & \cdots & \frac{\partial^2 \phi_1}{\partial x_{\text{ref}} \partial v_{z_{\text{ref}}}} \\ \frac{\partial^2 \phi_1}{\partial y_{\text{ref}} \partial x_{\text{ref}}} & \frac{\partial^2 \phi_1}{\partial y_{\text{ref}}^2} & \cdots & \frac{\partial^2 \phi_1}{\partial y_{\text{ref}} \partial v_{z_{\text{ref}}}} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial^2 \phi_1}{\partial v_{z_{\text{ref}}} \partial x_{\text{ref}}} & \frac{\partial^2 \phi_1}{\partial v_{z_{\text{ref}}} \partial y_{\text{ref}}} & \cdots & \frac{\partial^2 \phi_1}{\partial v_{z_{\text{ref}}}^2} \end{bmatrix} \quad (9)
\end{aligned}$$

Generalized Dual Numbers for Higher-Order Derivatives

A generalized dual number of order m is the truncated Taylor expansion of a smooth function f about a point a . They are the multivariate generalization of dual numbers, which are used in forward-mode automatic differentiation of univariate functions.²⁷ Generalized dual numbers provide a convenient way to automatically compute high-order Taylor approximations of multivariate functions in software. The basic idea is to overload all arithmetic operations (addition, subtraction, multiplication, division) and any necessary elementary functions (trigonometric functions, logarithms, etc.) to operate on truncated Taylor polynomials instead of ordinary floating point numbers. These arithmetic operations are straightforward to define, since they are identical to ordinary polynomial arithmetic, except that all terms beyond the truncation order m are discarded. Each arithmetic operation on truncated polynomials corresponds exactly to the truncation of that operation on the original functions: that is, given two functions f and g , and their order- m truncated Taylor expansions T_f^m and T_g^m , $T_f^m \pm T_g^m = T_{f \pm g}^m$, $T_f^m * T_g^m = T_{f * g}^m$, and $T_f^m / T_g^m = T_{f/g}^m$.²⁷ These properties allow high-order partial derivatives to be exactly propagated algebraically, since arithmetic on generalized dual numbers automatically carries forward all derivative information about the expansion point up to order m .

To illustrate the concept of generalized dual numbers, consider the function $f(x, y) = x^2y + y^3$ expanded about the point $(x, y) = (1, 2)$ and truncated to second order. First, the variables x and y are represented as truncated Taylor polynomials

$$x = 1 + dx \cdot \epsilon_x \quad y = 2 + dy \cdot \epsilon_y \quad (10)$$

where ϵ_x, ϵ_y are infinitesimal elements. Then, each arithmetic operation is applied step-by-step using the truncated Taylor polynomial algebra

$$x^2 = (1 + dx \cdot \epsilon_x)^2 = 1 + 2dx \cdot \epsilon_x + dx^2 \cdot \epsilon^2 \quad (11)$$

$$y^3 = 8 + 12dy \cdot \epsilon_y + 6dy^2 \cdot \epsilon_y^2 \quad (12)$$

Substituting back into the original function yields

$$f(x, y) = x^2y + y^3 = 10 + 4dx \cdot \epsilon_x + 13dy \cdot \epsilon_y + 2dxdy \cdot \epsilon_x\epsilon_y + 2dx^2 \cdot \epsilon_x^2 + 6dy^2 \cdot \epsilon_y^2 \quad (13)$$

which is the second-order multivariate Taylor expansion of the original function. Note that the truncation to order 2 is imposed for equations (12) and (13) by discarding ϵ -terms whose total order is higher than 2.

Generalized Dual Number Compatibility with Numerical Solvers

Many standard numerical integrators for solving ordinary differential equations (ODEs) are constructed using only elementary arithmetic operations. For instance, fixed-step explicit Runge-Kutta methods, unlike their adaptive-stepping counterparts, do not require any control flow statements. These methods compute the state evolution over time by evaluating a set of function calls that include additions, multiplications, and scalar-vector operations between states and their derivatives. If the algorithm operates on generalized dual numbers instead of ordinary floating point numbers, the ODE solver will automatically propagate the partial derivatives of the state together with the nominal solution. After a single integration pass, the final state vector is returned as a high-order multivariate polynomial that maps perturbations in the initial conditions to perturbations in the final states.²⁷ Hence, numerical integrators are transformed into high-order Taylor map generators with minimal effort.

High-Order State Mapping via Generalized Dual Numbers

To construct a high-order Taylor map for a satellite's final state, generalized dual numbers are utilized to encode perturbations in the initial conditions. The satellite dynamics are governed by a system of ODEs, and $\mathbf{y}(T) \in \mathbb{R}^6$, is computed via numerical integration from a known initial state $\mathbf{y}_0$. Due to observational uncertainties, the initial state is subject to small deviations $\Delta\mathbf{y}_0$ about a reference condition $\mathbf{y}_{\text{ref}}$ and the final state can be expressed as:

$$\mathbf{y}(T) = \Phi(\mathbf{y}_{\text{ref}} + \Delta\mathbf{y}_0) \quad (14)$$

where Φ is a smooth, nonlinear function. Each component of the initial state is represented using a generalized dual number

$$\mathbf{y}_0 = \begin{bmatrix} x_{\text{ref}} + dx\epsilon_1 \\ y_{\text{ref}} + dy\epsilon_2 \\ z_{\text{ref}} + dz\epsilon_3 \\ v_{x\text{ref}} + dv_x\epsilon_4 \\ v_{y\text{ref}} + dv_y\epsilon_5 \\ v_{z\text{ref}} + dv_z\epsilon_6 \end{bmatrix} \quad (15)$$

where $\epsilon_1, \dots, \epsilon_6$ are infinitesimal elements satisfying the truncation property $\epsilon_i^{N+1} = 0$ and all cross-term products $\epsilon_i^n \epsilon_j^m$ whose total order exceeds $n + m \geq N$ are discarded. N is the desired order of expansion of the Taylor polynomial.

In addition to encoding the initial condition, the time variable is also represented using a generalized dual number

$$t = T + \Delta t \epsilon_t \quad (16)$$

where T is the nominal TCA, Δt is the size of the time perturbation. This representation enables the construction of a time-expanded Taylor map that describes how the state evolves locally around

time T .

The generalized dual number-encoded initial condition is passed into a standard numerical integration routine such as RK6, which is composed entirely of elementary arithmetic operators. The integration routine propagates the nominal state and simultaneously tracks all mixed and higher-order interactions between perturbations through the algebraic rules of the generalized dual number system.

As a result, each component of $\mathbf{y}_i(T)$ of the final states is expressed as a multivariate Taylor polynomial. The polynomials encode the full sensitivity of the system with respect to both the initial condition deviations and time. Collectively, they form a high-order multivariate Taylor map¹⁹

$$\mathbf{y}(T + \Delta t) = \Phi(\Delta \mathbf{y}_0, \Delta t) \quad (17)$$

Time of Closest Approach

The high-order Taylor maps constructed through generalized dual numbers are utilized to estimate the time and distance of closest approach between two satellites. Assuming that the nominal TCA has been obtained, it is then treated as the expansion point about which the satellite's state vector is integrated using generalized dual number arithmetic and returns a Taylor map for each satellite's orbit evolution.

From the Taylor-expanded position vectors, the relative displacement between both satellites is computed

$$\Delta \mathbf{r}(t) = \mathbf{r}_1(t) - \mathbf{r}_2(t) \quad (18)$$

and the squared relative distance is defined

$$D^2(t) = \|\Delta \mathbf{r}(t)\|^2 \quad (19)$$

Since each component of $\Delta \mathbf{r}(t)$ is a polynomial in Δt , the squared distance is also a scalar polynomial in time. The TCA corresponds to the point where this distance is minimized

$$\frac{d}{dt} D^2(t) = 0 \quad (20)$$

The Newton-Raphson method is applied to this polynomial to solve for the time offset Δt^* that satisfies the stationary condition. The refined TCA is then given by

$$t_{TCA} = T + \Delta t^* \quad (21)$$

IMPORTANCE SAMPLING AND THE CROSS ENTROPY METHOD

The brute-force Monte Carlo estimator for the probability of collision P_c using n random samples is given as,

$$\hat{P}_c = \frac{1}{n} \sum_{i=1}^n f(x_i), \quad x_i \stackrel{\text{i.i.d.}}{\sim} p(x) \quad (22)$$

where $p(x)$ is the initial combined (12D) probability density function for the satellite states and $f(x)$ is an indicator function for the collision event. The importance sampling estimator gives an alternative way to estimate the same quantity more efficiently by drawing samples from a proposal distribution $q(x)$ which exaggerates the event frequency, and reweighting by the likelihood ratio.

$$\hat{P}_c = \frac{1}{n} \sum_{i=1}^n \frac{p(x_i)}{q(x_i)} f(x_i), \quad x_i \stackrel{\text{i.i.d.}}{\sim} q(x) \quad (23)$$

If the importance sampling (proposal) distribution is selected optimally, it will converge much faster than the brute-force Monte Carlo estimator.²⁵ The cross entropy method provides a standard algorithm for finding the optimal proposal distribution within a parametric family (e.g., the family of Gaussian distributions). To be compatible with the cross entropy method, the event indicator function must be represented in the form $S(x) < R$, where $S(x)$ is a scoring function representing the “distance” to the event set, and R is some threshold value. For the P_c estimation problem, $S(x)$ is taken as the distance of closest approach over the time interval of interest and R as the combined hard body radius.

The cross entropy algorithm works using successive rounds of importance sampling, iteratively approaching the optimal proposal distribution by fitting a new proposal distribution to the top-scoring fraction ρ of samples at each iteration. This process continues until the fraction ρ consists entirely of samples less than the target threshold R . Hence, the final proposal distribution output by the cross-entropy method will have been constructed by fitting a Gaussian distribution to samples drawn exclusively from the event set. The algorithm can be summarized as follows:²⁵

1. Define the initial proposal distribution q^0 as the target distribution p . Set the iteration counter $k \leftarrow 0$.
2. Generate N samples from q^k and compute the performance function $S(x)$ for each. Sort the samples by descending order of $S(x)$. Then, locate the $(1 - \rho)$ quantile R^k of the performances. If $R^k < R$, set R^k to R .
3. Use the samples from the previous step to find the weighted-sample maximum likelihood estimate for the parameters of q^{k+1} over the samples which satisfy the indicator function $I_{\{S(x) \leq R^k\}}$. For the Gaussian family, this consists of finding the weighted mean and covariance of these samples, using the likelihood ratios $\frac{p(x)}{q^k(x)}$ as weights.
4. Terminate if either $R^k = R$ or $k = K$; else set $k \leftarrow k + 1$ and return to step 2.

The control parameters for the algorithm are N , the number of samples drawn at each iteration, ρ , the fraction of samples used at each iteration to fit the next iteration’s proposal distribution, and K , the maximum number of iterations.

RESULTS & DISCUSSION

Short-Encounter Conjunctions

This section presents and analyzes the results of four simulations designed to evaluate the accuracy and computational efficiency of Taylor polynomial approximations of the distance of closest approach in conjunction with importance sampling (IS) for satellite collision probability estimation. The differential algebra system for Taylor polynomial computation is implemented in PyAudi, an open-source Python library.²⁷ Simulation 1 validates the Python implementation by replicating a reference numerical experiment documented by Ketzner and Elgohary.²⁵ The subsequent simulations assess Taylor polynomial surrogates as alternatives to direct numerical integration, exploring

the effect of the expansion point location and truncation order on accuracy and performance. The conjunctions analyzed are short-duration encounters from the CelesTrak SOCRATES catalog.²⁸

All simulations used the same encounter window, force model, and initial covariance. The satellite states are provided by the SOCRATES catalogue at the time of closest approach. For each case, initial states were generated by propagating backward from the TCA by either 10,000 seconds or 2 days. Satellite trajectories were propagated using a two-body model with J2 perturbation. Numerical integration was performed using the RK6 method with a fixed 15-second step size, which was confirmed to yield sub-5cm position error over the full propagation interval. The initial probability distribution was defined as a 12-dimensional Gaussian centered at the initial states. The covariance structure was block-diagonal, with position variances set to $1 \times 10^{-2} \text{km}^2$ and velocity variances to $1 \times 10^{-8} \text{km}^2/\text{s}^2$ for both satellites. The CE method selected the top 10% of samples ($\rho = 0.1$) from $N = 5000$ candidates per iteration and was run for a maximum of $K = 10$ iterations. For the IS step used to produce the P_c estimate, the same total number of samples was used across all simulations. For Simulation 1, CE and IS phases were parallelized using 24 logical processors to reduce computation time. Simulations 2-4 were executed without parallelization, as their runtimes were sufficiently low. To allow consistent comparison of wall-clock times across simulations, the reported time for Simulation 1 was scaled by a factor of 24 to estimate the equivalent non-parallel runtime.

Simulation 1 was conducted to validate the implementation of the propagation and sampling framework by replicating results reported by Ketzner and Elgohary.²⁵ For each sample, forward propagation was performed using RK6 from the initial epoch (10,000s or 2 days before the nominal TCA) until 100 seconds prior to the expected TCA. From this point, the relative radial velocity between the satellites was monitored at each integration step to detect a sign change, indicating a local extremum in relative distance. Once detected, Brent’s method was used to refine the time of closest approach to within 10^{-6} seconds, and the corresponding minimum separation distance was computed. This distance served as the score function for both the CE and IS procedures. The number of IS samples was adaptively chosen to ensure a 1% relative error at 95% confidence. To enable consistent comparisons across subsequent simulations, the CE random seed, the resulting proposal distribution, and the IS random seed were saved and reused in Simulation 2-4.

Simulation 2 was designed to evaluate the accuracy and computational performance of the Taylor map surrogate when used in place of direct numerical propagation during IS. The CE procedure was not rerun; instead, the proposal distribution obtained from Simulation 1 was reused. To ensure a fair comparison, the same IS seed and total number of samples were used, resulting in an identical sample set to Simulation 1. The only change from Simulation 1 was the replacement of RK6 propagation with a Taylor map centered at the nominal initial state (the mean of the target distribution). Taylor maps of order 2, 3, and 4 were pre-generated and tested. For each sample, the displacement from the nominal state was substituted into the Taylor-expanded expression for squared relative distance as a function of time offset. The TCA was determined by minimizing this expression using Newton’s method, initialized at the nominal TCA with a solver tolerance of 10^{-10} seconds and a maximum of 100 iterations. The resulting distance of closest approach served as the score function during IS.

Simulation 3 was designed to evaluate whether centering the Taylor map at the mean of the optimal proposal distribution—rather than at the mean of the target distribution—could substantially improve accuracy. The CE procedure was not rerun; the saved optimal proposal distribution from Simulation 1 was reused to define the new expansion point. All other elements remained unchanged:

the same IS seed and total number of samples were used, resulting in an identical sample set to Simulation 1. Since the proposal distribution mean defines a different trajectory than the target distribution mean, the corresponding TCA was first determined using the same RK6-based procedure implemented in Simulation 1. This updated TCA was then used as the expansion point for generating Taylor maps of order 2, 3, and 4.

Simulation 4 aimed to evaluate the effectiveness of using a fixed Taylor map surrogate throughout both the CE and IS phases. The surrogate was centered at the target distribution mean, and Taylor maps of order 2, 3, and 4 were pre-generated. The CE procedure was rerun using the same random seed saved from Simulation 1, but with the surrogate-based score function replacing numerical propagation. This allowed the CE process to select elite samples based on minimum separation distances computed directly from the Taylor-expanded expression, using Newton’s method. The IS phase reused the same seed and total number of samples as in Simulation 1, resulting in an identical sample set for all simulations. This setup enabled a direct evaluation of the computational speedup and accuracy achieved when the Taylor map was used across the entire simulation process.

Tables 1-4 present results for each short-duration encounter. Each table specifies the NORAD ID, propagation duration (10,000 seconds or 2 days prior to the nominal TCA), and the number of samples used for IS. For all simulations and Taylor map orders n , the tables report the estimated collision probability $\hat{P}_c$, variance estimate $\hat{\sigma}^2$, relative error ϵ_r , and the individual computation times (in seconds) for the cross-entropy (CE) phase, Taylor map generation (TM), and the importance sampling (IS) phase.

Table 1. NORAD ID: 51772-39433, 10,000s Propagation, IS sample size 33000.

Simulation	n	$\hat{P}_c$	$\hat{\sigma}^2$	ϵ_r (%)	CE (s)	TM (s)	IS (s)
1	–	3.853×10^{-5}	1.176×10^{-9}	0.960	4927	-	7118
	2	3.853×10^{-5}	1.268×10^{-9}	0.997	-	9.39	1.89
2	3	3.853×10^{-5}	1.176×10^{-9}	0.960	-	26.10	4.02
	4	3.853×10^{-5}	1.176×10^{-9}	0.960	-	75.30	18.99
3	2	3.853×10^{-5}	1.178×10^{-9}	0.961	-	8.53	1.37
	3	3.853×10^{-5}	1.176×10^{-9}	0.960	-	26.96	4.15
	4	3.853×10^{-5}	1.176×10^{-9}	0.960	-	75.90	19.83
4	2	3.836×10^{-5}	1.176×10^{-9}	0.964	0.90	8.70	1.33
	3	3.847×10^{-5}	1.281×10^{-9}	1.004	2.50	26.54	4.16
	4	3.843×10^{-5}	1.279×10^{-9}	1.004	12.20	77.00	19.83

Simulation 1 validated the full Python-based implementation by successfully reproducing collision probability estimates, variance estimates, and relative errors consistent with those reported by Ketzner and Elgohary²⁵ across all test cases and propagation durations. This confirms the correctness of the propagation model, root-finding procedure, and CE/IS framework. As such, Simulation 1 establishes a reliable reference baseline for evaluating the accuracy and efficiency of Taylor polynomial surrogates in Simulation 2-4.

Simulation 2 evaluated whether Taylor maps centered at the nominal initial state could replace numerical propagation during IS without sacrificing accuracy. Results show that for all conjunction cases and propagation durations, third- and fourth-order Taylor maps produced collision probability estimates, variances, and relative errors that closely matched those from Simulation 1. However, third-order accuracy was not consistently sufficient across all cases—only the fourth-order expan-

Table 2. NORAD ID: 51772-39433, 2 days Propagation, IS sample size 49500.

Simulation	n	$\hat{P}_c$	$\hat{\sigma}^2$	ϵ_r (%)	CE (s)	TM (s)	IS (s)
1	–	3.021×10^{-6}	8.510×10^{-12}	0.851	97536	-	189900
	2	4.128×10^{-6}	3.482×10^{-9}	12.593	-	153.83	2.01
2	3	3.470×10^{-6}	4.454×10^{-10}	5.358	-	452.79	6.93
	4	3.021×10^{-6}	8.510×10^{-12}	0.851	-	1311.39	28.64
3	2	3.023×10^{-6}	6.750×10^{-12}	0.757	-	154.11	2.07
	3	3.477×10^{-6}	4.470×10^{-10}	5.357	-	459.55	6.13
	4	3.021×10^{-6}	8.510×10^{-12}	0.851	-	1342.34	30.23
4	2	3.810×10^{-6}	2.390×10^{-11}	1.130	0.90	149.91	1.77
	3	3.392×10^{-6}	1.326×10^{-11}	0.946	3.60	457.04	7.15
	4	3.023×10^{-6}	8.618×10^{-12}	0.856	14.90	1334.45	29.39

Table 3. NORAD ID: 43184-46583, 10,000s Propagation, IS sample size 48000.

Simulation	n	$\hat{P}_c$	$\hat{\sigma}^2$	ϵ_r (%)	CE (s)	TM (s)	IS (s)
1	–	2.558×10^{-5}	6.563×10^{-10}	0.896	4366	-	10270
	2	2.434×10^{-5}	6.483×10^{-10}	0.895	-	8.89	2.03
2	3	2.559×10^{-5}	6.568×10^{-10}	0.896	-	26.44	5.94
	4	2.558×10^{-5}	6.563×10^{-10}	0.896	-	75.00	27.35
3	2	2.557×10^{-5}	6.557×10^{-10}	0.896	-	8.65	1.96
	3	2.559×10^{-5}	6.568×10^{-10}	0.896	-	26.82	5.92
	4	2.558×10^{-5}	6.563×10^{-10}	0.896	-	76.42	28.81
4	2	2.536×10^{-5}	5.900×10^{-10}	0.857	0.80	9.00	1.91
	3	2.540×10^{-5}	6.220×10^{-10}	0.879	2.50	27.36	6.11
	4	2.538×10^{-5}	6.200×10^{-10}	0.878	11.80	77.52	27.97

sion reliably reproduced the Simulation 1 baseline in every scenario. In contrast, second-order Taylor expansions showed noticeable degradation in accuracy, particularly for the 2-day propagation cases, where increased nonlinearity rendered quadratic approximation insufficient.

Replacing RK6 propagation with fourth-order Taylor map surrogates during the IS phase led to a substantial speedup in runtime while preserving accuracy. Across all test cases, this substitution accelerated IS computation by several orders of magnitude. For scenarios involving 10,000s propagation, the Taylor map approach achieved a speedup by a factor of 375, while in 2-day propagation cases, the speedup exceeded a factor of 6600, compared to the full RK6-based IS in Simulation 1. The runtime required for IS increased with the order of the Taylor expansion. Higher-order maps involve the evaluation of more complex multivariate polynomial expressions, which increases computational effort. However, even for fourth-order expansions, the total IS cost remained significantly lower than that of full numerical propagation, demonstrating the efficiency of using Taylor maps as surrogates.

Simulation 3 evaluated whether centering the Taylor map at the mean of the proposal distribution, rather than the nominal initial state, would improve the accuracy of the surrogate model and enable lower-order expansions to match the RK6 baseline. The results show that this centering strategy did not lead to meaningful accuracy improvements at any order. Second-order Taylor expansions continued to show degraded accuracy for the 2-day propagation scenarios due to the increased non-

Table 4. NORAD ID: 43184-46583, 2 days Propagation, IS sample size 42000.

Simulation	n	$\hat{P}_c$	$\hat{\sigma}^2$	ϵ_r (%)	CE (s)	TM (s)	IS (s)
1	–	1.625×10^{-6}	3.197×10^{-12}	1.053	101405	-	167038
	2	3.424×10^{-6}	5.060×10^{-8}	62.825	-	153.74	1.71
2	3	1.857×10^{-6}	7.761×10^{-12}	1.435	-	448.65	5.94
	4	1.625×10^{-6}	3.197×10^{-12}	1.053	-	1315.18	24.28
3	2	1.633×10^{-6}	2.672×10^{-12}	0.957	-	157.11	1.78
	3	1.814×10^{-6}	5.798×10^{-12}	1.270	-	456.69	5.21
	4	1.625×10^{-6}	3.197×10^{-12}	1.053	-	1502.42	25.29
4	2	2.416×10^{-6}	1.178×10^{-11}	1.359	1.00	154.29	1.67
	3	1.886×10^{-6}	4.865×10^{-12}	1.118	3.70	462.00	6.22
	4	1.622×10^{-6}	3.156×10^{-12}	1.048	14.60	1339.65	24.43

linearity, while third- and fourth-order expansions produced results nearly identical to those in Simulation 2. As in Simulation 2, speedups were observed exclusively in the IS phase. For fourth-order expansions, Simulation 3 achieved a speedup by a factor of 356 for the 10,000s cases and over 6280 for the 2-day scenarios, compared to the full RK6-based IS in Simulation 1. Overall, Simulation 3 demonstrated that shifting the expansion point to the proposal mean offers no measurable benefit in accuracy over using the nominal initial state.

Simulation 4 tested whether numerical propagation could be completely replaced by a single fixed Taylor map applied across both the CE and IS phases. The results show that third- and fourth-order Taylor expansions maintained accuracy comparable to Simulation 1 across all cases, while second-order expansions exhibited degraded accuracy, particularly for the 2-day propagation scenarios due to stronger nonlinearities. Simulation 4 achieved the greatest reduction in total computation time. CE runtimes dropped from thousands of seconds in Simulation 1 to just a few seconds with the fourth-order map. IS runtimes saw similar reductions to Simulations 2-3.

The results confirm that Taylor maps can enable substantial computational speedups when applied across both CE and IS phases. However, maintaining accuracy especially for long-duration propagation requires using a sufficiently high expansion order.

Long-Encounter Conjunctions

This section analyzes two long-duration conjunction scenarios drawn from Alfano’s test set,²⁹ which were also examined by Ketzner and Elgohary.²⁵ These cases were selected to evaluate the performance of the Taylor polynomials under increased orbital nonlinearity and extended propagation intervals.

All simulation settings remain the same as those used for the short-duration encounters, except for the following differences. First, the dynamics model includes only two-body motion without J2 perturbation, consistent with the original Alfano test cases.²⁹ Second, the propagation duration was fixed at two days prior to the nominal TCA. During the CE phase, each iteration used 10,000 samples as documented by Ketzner and Elgohary.²⁵ For Simulation 1, the minimum separation was computed using the same root-finding procedure as before, but with a ± 1419 second window centered around the nominal TCA, consistent with the setup in Ketzner and Elgohary²⁵ and Alfano.²⁹

Tables 5-6 present results for each long-encounter scenario. Each table specifies the test case, the

propagation duration (2 days), and the number of samples used for IS. For all simulations and Taylor map orders n , the tables report the estimated collision probability $\hat{P}_c$, variance estimate $\hat{\sigma}^2$, relative error ϵ_r , and the individual computation times (in seconds) for the cross-entropy (CE) phase, Taylor map generation (TM), and the importance sampling (IS) phase.

Table 5. Alfano Case 6, 2 days Propagation, IS sample size 13500.

Simulation	n	$\hat{P}_c$	$\hat{\sigma}^2$	ϵ_r (%)	CE (s)	TM (s)	IS (s)
1	–	4.314×10^{-3}	6.533×10^{-6}	0.999	133330	-	58754
	2	4.312×10^{-3}	6.515×10^{-6}	0.999	-	82.41	0.52
2	3	4.315×10^{-3}	6.544×10^{-6}	1.000	-	232.87	2.15
	4	4.314×10^{-3}	6.533×10^{-6}	0.999	-	636.40	8.54
3	2	4.313×10^{-3}	6.520×10^{-6}	0.999	-	84.21	0.57
	3	4.315×10^{-3}	6.544×10^{-6}	1.000	-	238.55	2.18
	4	4.314×10^{-3}	6.533×10^{-6}	0.999	-	658.49	8.76
4	2	4.328×10^{-3}	6.526×10^{-6}	0.996	1.30	85.85	0.59
	3	4.330×10^{-3}	6.721×10^{-6}	1.010	6.00	234.10	2.17
	4	4.341×10^{-3}	6.646×10^{-6}	1.002	19.90	633.73	8.81

Table 6. Alfano Case 7, 2 days Propagation, IS sample size 28500.

Simulation	n	$\hat{P}_c$	$\hat{\sigma}^2$	ϵ_r (%)	CE (s)	TM (s)	IS (s)
1	–	1.635×10^{-4}	4.425×10^{-8}	1.494	195809	-	130901
	2	1.595×10^{-4}	1.600×10^{-8}	0.921	-	82.65	1.12
2	3	1.664×10^{-4}	4.873×10^{-8}	1.540	-	233.22	4.63
	4	1.637×10^{-4}	4.457×10^{-8}	1.498	-	639.72	18.86
3	2	1.624×10^{-4}	1.641×10^{-8}	0.916	-	86.82	1.18
	3	1.659×10^{-4}	4.542×10^{-8}	1.492	-	245.25	4.67
	4	1.636×10^{-4}	4.456×10^{-8}	1.498	-	663.99	19.07
4	2	1.602×10^{-4}	1.533×10^{-8}	0.898	1.80	84.21	1.16
	3	1.658×10^{-4}	1.819×10^{-8}	0.945	9.70	242.51	4.53
	4	1.628×10^{-4}	1.729×10^{-8}	0.938	28.90	638.41	18.95

Simulation 1 validated the Python-based implementation for the long-duration Alfano scenarios by reproducing collision probability estimates, variance estimates, and relative errors consistent with those reported by Ketzner and Elgohary.²⁵ In contrast to short-duration encounter results, Simulations 2, 3, and 4 showed that even the second-order Taylor map surrogates produced collision probability estimates in close agreement with the RK6 baseline.

Centering the Taylor map at the proposal distribution mean (Simulation 3) did not yield meaningful accuracy gains for any expansion order. Simulation 4 delivered the greatest computational savings by fully replacing numerical propagation with a surrogate model across both CE and IS phases, accelerating total runtimes for fourth-order scenarios by a factor of 290 for case 6 and 470 for case 7, relative to Simulation 1 while preserving accuracy.

CONCLUSION

This work evaluated the use of Taylor polynomial maps together with cross-entropy importance sampling to improve the efficiency of satellite collision probability estimation. The approach aimed

to reduce computational cost while maintaining accuracy comparable to high-fidelity numerical propagation. Results from numerical experiments on representative short-duration conjunctions from SOCRATES and long-duration nonlinear conjunctions from Alfano’s test set demonstrated that third- and fourth-order Taylor expansions accurately replicated baseline estimates while enabling significant computational savings. For short-duration encounters, second-order expansions exhibited degraded performance due to increased nonlinearity, which rendered quadratic approximations insufficient. In contrast, for long-duration conjunctions, even second-order expansions produced collision probability estimates in close agreement with the baseline.

These findings confirm that Taylor maps can substantially accelerate conjunction analysis in a rare-event estimation framework without any loss in accuracy, achieving total computational speedups ranging from a factor of 100 to nearly 300 across a variety of short- and long-duration conjunction scenarios when used throughout the cross-entropy and importance sampling phases.

REFERENCES

- [1] F. J. Krage, “Nasa spacecraft conjunction assessment and collision avoidance best practices handbook,” tech. rep., National Aeronautics and Space Administration, 2023.
- [2] J. L. Foster and H. S. Estes, *A parametric analysis of orbital debris collision probability and maneuver rate for space vehicles*. NASA, National Aeronautics and Space Administration, Lyndon B. Johnson Space ..., 1992.
- [3] M. R. Akella and K. T. Alfriend, “Probability of collision between space objects,” *Journal of Guidance, Control, and Dynamics*, Vol. 23, No. 5, 2000, pp. 769–772.
- [4] R. P. Patera, “General method for calculating satellite collision probability,” *Journal of Guidance, Control, and Dynamics*, Vol. 24, No. 4, 2001, pp. 716–722.
- [5] S. Alfano, “A numerical implementation of spherical object collision probability,” *The Journal of the Astronautical Sciences*, Vol. 53, 2005, pp. 103–109.
- [6] F. K. Chan, *Spacecraft collision probability*. American Institute of Aeronautics and Astronautics, Inc., 2008.
- [7] X.-L. Xu and Y.-Q. Xiong, “A method for calculating probability of collision between space objects,” *Research in Astronomy and Astrophysics*, Vol. 14, No. 5, 2014, p. 601.
- [8] R. Serra, D. Arzelier, M. Joldes, J.-B. Lasserre, A. Rondepierre, and B. Salvy, “Fast and accurate computation of orbital collision probability for short-term encounters,” *Journal of Guidance, Control, and Dynamics*, Vol. 39, No. 5, 2016, pp. 1009–1021.
- [9] R. García-Pelayo and J. Hernando-Ayuso, “Series for collision probability in short-encounter model,” *Journal of Guidance, Control, and Dynamics*, Vol. 39, No. 8, 2016, pp. 1904–1912.
- [10] V. T. Coppola *et al.*, “Including velocity uncertainty in the probability of collision between space objects,” *AAS/AIAA Spaceflight Mechanics Meeting*, American Astronautical Soc, 2012, pp. 2012–2247.
- [11] K. J. DeMars, Y. Cheng, and M. K. Jah, “Collision probability with Gaussian mixture orbit uncertainty,” *Journal of Guidance, Control, and Dynamics*, Vol. 37, No. 3, 2014, pp. 979–985.
- [12] C. T. Shelton and J. L. Junkins, “Probability of collision between space objects including model uncertainty,” *Acta Astronautica*, Vol. 155, 2019, pp. 462–471.
- [13] D. T. Hall, “Expected collision rates for tracked satellites,” *Journal of Spacecraft and Rockets*, Vol. 58, No. 3, 2021, pp. 715–728.
- [14] D. T. Hall, L. G. Baars, and S. J. Casali, “A Multistep Probability of Collision Computation Algorithm,” *AAS/AIAA Astrodynamics Specialist Conference*, 2023.
- [15] H. Quebedeaux and T. A. Elgohary, “PHYSICS-INFORMED NEURAL NETWORKS FOR ORBIT PROPAGATION IN ASTRODYNAMICS PROBLEMS,”
- [16] B. A. Jones, A. Doostan, and G. Born, “Conjunction assessment using polynomial chaos expansions,” *Proceedings of the 23rd International Symposium and Space Flight Dynamics*, Jet Propulsion Laboratory Pasadena, 2012, pp. 1–18.
- [17] B. A. Jones and A. Doostan, “Satellite collision probability estimation using polynomial chaos expansions,” *Advances in Space Research*, Vol. 52, No. 11, 2013, pp. 1860–1875.
- [18] 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.

- [19] A. Morselli, R. Armellin, P. Di Lizia, and F. B. Zazzera, “A high order method for orbital conjunctions analysis: sensitivity to initial uncertainties,” *Advances in Space Research*, Vol. 53, No. 3, 2014, pp. 490–508.
- [20] A. Morselli, R. Armellin, P. Di Lizia, and F. B. Zazzera, “A high order method for orbital conjunctions analysis: Monte Carlo collision probability computation,” *Advances in Space Research*, Vol. 55, No. 1, 2015, pp. 311–333.
- [21] A. Zollo, C. Parigini, R. Armellin, J. F. San Juan Díaz, S. Aida, and R. Kahle, “Long-term collision probability computation through high-order Taylor polynomials evaluation,” 2024.
- [22] J. Dolado, P. Legendre, R. Garmier, B. Reveling, and X. Pena, “Satellite collision probability computation for long term encounters,” *Adv. Astronaut. Sci.*, Vol. 142, 2011.
- [23] R. Pastel, “Estimating satellite versus debris collision probabilities via the adaptive splitting technique,” *Proc. 3rd Int. Conf. Comput. Modeling Simulation*, 2011, pp. 1–6.
- [24] Y. Sztamfater Garcia, M. Sanjurjo Rivo, J. Miguez, and G. Escibano, “Novel method for the Computation of In-Orbit Collision Probability by Multilevel Splitting and Surrogate Modelling,” *AIAA SCITECH 2024 Forum*, 2024, p. 1813.
- [25] R. Ketzner and T. A. Elgohary, “FAST MONTE CARLO CONJUNCTION ANALYSIS USING IMPORTANCE SAMPLING AND THE CROSS-ENTROPY METHOD,”
- [26] M. E. Taylor, *Multivariable Calculus*. University of North Carolina at Chapel Hill, 2020. Lecture notes for Math 233H.
- [27] D. Izzo, “Taylor Algebra Theory,” 2020. Accessed: 2025-04-14.
- [28] T. S. Kelso and S. Alfano, “Satellite orbital conjunction reports assessing threatening encounters in space (SOCRATES),” *Modeling, Simulation, and Verification of Space-based Systems III*, Vol. 6221, SPIE, 2006, p. 622101.
- [29] S. Alfano, “Satellite conjunction monte carlo analysis,” *Adv. Astronaut. Sci.*, Vol. 134, 2009, pp. 2007–2024.