

Finite-Time Transport Boundaries and Family-Constrained Robustness in the Earth–Moon CR3BP

Ayush Samanta

Independent Researcher

Bloho Labs

Abstract

Finite-time transport in the circular restricted three-body problem (CR3BP) can change qualitatively under perturbations to initial conditions, making robustness difficult to characterize using smooth trajectory sensitivities alone. This work investigates whether geometric distance to an outcome-change boundary can provide a useful measure of transport robustness for a two-dimensional family of invariant-manifold seeds generated from planar L_1 Lyapunov orbits in the Earth–Moon CR3BP. A 160×33 parameter atlas comprising 5280 initial conditions was classified by first-passage outcome, and adaptive refinement produced 484 numerical samples of the resulting outcome boundaries. To avoid arbitrary Euclidean distance in the family parameters, the seed-state map was used to induce a local pullback metric, defining a family-constrained boundary-distance diagnostic ρ_η . The metric construction was numerically consistent, and ρ_η revealed strongly nonuniform proximity to the resolved transport interfaces. Its ordering, however, depended on the prescribed state weighting, with Spearman correlations ranging from approximately 0.794 to 0.999 across the tested metrics. Direct perturbation tests also rejected a guaranteed-radius interpretation: one outcome change occurred at approximately $0.5\rho_\eta$, while finite directional searches at five representative points recovered no independent minimum flip distance. These results support ρ_η as a family-constrained geometric diagnostic of proximity to resolved finite-time outcome boundaries, but not as a certified outcome-preserving radius or validated transport condition number. A stronger robustness measure requires converged boundary resolution, physically specified uncertainty geometry, and direct computation of minimum outcome-changing perturbations.

1 Introduction

Transport in the circular restricted three-body problem (CR3BP) is strongly shaped by the phase-space structures associated with periodic orbits and their invariant manifolds. In the Earth–Moon system, these structures provide natural pathways between dynamically distinct regions and form the basis of many low-energy trajectory-design strategies. A trajectory initialized near such a pathway, however, can exhibit a qualitatively different finite-time evolution after a comparatively small change in its initial condition. For trajectory design, the existence of a transport pathway is therefore only part of the problem; its sensitivity to admissible perturbations is also important.

This sensitivity is particularly difficult to summarize when the quantity of interest is a discrete transport outcome rather than a smooth trajectory observable. Nearby initial conditions may reach the same destination with different passage times, but they may also terminate through different events altogether. In the latter case, a local derivative of passage time or state deviation does not by itself describe the transition. The relevant structure is instead the boundary separating initial conditions with different

finite-time outcomes.

Outcome and basin maps provide a natural representation of this structure. By classifying a family of initial conditions according to their subsequent behavior, they reveal regions of common outcome and the interfaces separating them. Such boundaries can be geometrically complicated in nonlinear gravitational systems, and their presence suggests a simple robustness question: how far is a nominal trajectory from a perturbation that changes its transport outcome?

Answering that question requires more than measuring coordinate separation in an arbitrarily chosen parameterization. A family of manifold-seeded trajectories may be described by parameters such as orbital phase and periodic-orbit amplitude, but equal displacements in those coordinates do not generally correspond to equal perturbations of the physical spacecraft state. Any distance-based robustness measure therefore requires both a definition of the relevant outcome boundary and a metric connecting parameter changes to physical state changes.

This work examines that problem for a two-dimensional family of invariant-manifold seeds generated from planar L_1 Lyapunov orbits in the Earth–Moon

CR3BP. The central objective is to determine whether distance to a numerically resolved finite-time outcome boundary can serve as a meaningful measure of transport robustness, and to identify the conditions under which that interpretation succeeds or fails.

2 Related Work

The circular restricted three-body problem has long served as a fundamental model for understanding transport near libration points and for constructing trajectories that exploit the natural dynamics of multi-body systems. Classical treatments establish the geometry of the Jacobi integral, Hill regions, equilibrium points, and families of periodic solutions in the rotating frame [16, 15, 9]. Analytical and numerical constructions of periodic motion near the collinear points have subsequently become standard components of libration-point mission design [15, 10, 8].

A central development in modern three-body trajectory design is the use of stable and unstable invariant manifolds associated with periodic orbits. These manifolds organize natural transport between dynamically distinct regions and can be used to construct low-energy transfers and heteroclinic connections. Koon et al. demonstrated heteroclinic connections between periodic orbits surrounding L_1 and L_2 and showed how invariant manifolds provide dynamical channels for transport in the planar CR3BP [11]. Subsequent work has developed these ideas for trajectory design and transfer construction in the Earth–Moon and related three-body systems [12, 8, 7, 13].

The sensitivity of transport is closely related to the structure of escape, collision, and destination basins. Numerical studies of the restricted three-body problem have shown that sets of initial conditions associated with distinct outcomes can possess highly structured and fractal boundaries. In the planar Earth–Moon problem, de Assis and Terra identified escape channels through the neighborhoods of L_1 and L_2 and analyzed the associated escape basins and fractal basin boundaries [4]. Zotos similarly classified bounded, escaping, and collisional motion in the planar restricted three-body problem and reported extensive fractal structure in the corresponding basin boundaries [17]. Three-dimensional extensions likewise display complicated partitions between bounded, escaping, and collisional outcomes [18]. These studies establish that qualitative changes of trajectory outcome can occur across geometrically complex sets of initial conditions.

Outcome-boundary geometry has also been considered more broadly in nonlinear dynamics as a measure of practical stability. In nonlinear mechanical systems, for example, dynamical-integrity measures quantify the extent of safe basins and their erosion as parameters are

varied [14]. Such approaches motivate the use of distance to a basin boundary as a local measure of tolerance to perturbation, while also emphasizing that the resulting quantity depends on the definition of the safe set and on the geometry used to measure perturbations.

A related mathematical viewpoint arises in numerical conditioning. For several classical problems, the condition number is closely related to the reciprocal distance from a well-posed instance to an appropriate set of ill-posed instances [6, 5]. This distance-to-ill-posedness perspective provides a useful conceptual analogy for trajectory transport: a nominal initial condition close to a set across which its qualitative outcome changes may be expected to possess reduced tolerance to perturbation. However, an inverse distance becomes a genuine condition number only when the relevant ill-posed set, metric, and perturbation-response relationship are defined and validated.

Robustness in astrodynamics is also commonly approached through uncertainty propagation, control, and trajectory optimization. Minimum-time and minimum-fuel control problems in the restricted three-body problem have been studied using continuation and optimal control techniques [2, 3]. More recent cislunar trajectory studies explicitly incorporate navigation, dynamical, and actuation uncertainty into robust design [1]. These approaches differ from the present work in that robustness is generally evaluated through controlled trajectory performance, uncertainty propagation, or optimization rather than through geometric proximity to a finite-time outcome-change boundary.

The present study focuses on the intersection of these viewpoints. It constructs a finite-time outcome partition on a two-dimensional family of invariant-manifold seeds, resolves the interfaces between distinct transport outcomes, and measures proximity to those interfaces using a metric induced by the physical seed-state map. The objective is not to claim a new form of invariant-manifold transport or basin analysis. Rather, it is to test whether a family-constrained distance to a resolved outcome boundary can support a condition-like interpretation of finite-time transport robustness, and to determine experimentally where that interpretation breaks down.

3 Mathematical Model

3.1 Planar Circular Restricted Three-Body Problem

We consider spacecraft motion in the planar circular restricted three-body problem (CR3BP), following the standard nondimensional formulation of [12].

The equations are expressed in a synodic coordinate system rotating with the two primaries. Length is nor-




Figure 1: Geometry of the planar circular restricted three-body problem in the synodic frame. The two primaries remain fixed on the x -axis, while the spacecraft moves under their gravitational influence. Distances are expressed in units of the primary separation.

malized by their separation, mass by $m_1 + m_2$, and time such that the angular velocity of the rotating frame is unity. The dynamics are therefore parameterized by the dimensionless mass ratio

$$\mu = \frac{m_2}{m_1 + m_2}. \quad (1)$$

For the Earth–Moon system considered here,

$$\mu = 0.0121505856. \quad (2)$$

With the barycenter at the origin, the primaries remain fixed at $(-\mu, 0)$ and $(1 - \mu, 0)$. For a spacecraft at (x, y) , its distances from the two primaries are

$$r_1 = \sqrt{(x + \mu)^2 + y^2}, \quad (3)$$

$$r_2 = \sqrt{(x - 1 + \mu)^2 + y^2}. \quad (4)$$

The planar spacecraft state is written as

$$\mathbf{x} = [x \ y \ v_x \ v_y]^T, \quad v_x = \dot{x}, \quad v_y = \dot{y}. \quad (5)$$

The restriction to planar motion reduces the spacecraft state to four dimensions. This model does not reproduce the full Earth–Moon environment; in particular, it neglects out-of-plane motion, orbital eccentricity, solar perturbations, and other gravitational bodies. It is used here because it retains the nonlinear transport structures central to the present analysis while providing a controlled setting in which their numerical behavior can be isolated.

3.2 Equations of Motion

In the synodic frame, the gravitational attraction of the two primaries and the centrifugal contribution may be collected into the effective potential

$$\Omega(x, y) = \frac{1}{2}(x^2 + y^2) + \frac{1 - \mu}{r_1} + \frac{\mu}{r_2}. \quad (6)$$

Its spatial derivatives are

$$\Omega_x = x - (1 - \mu) \frac{x + \mu}{r_1^3} - \mu \frac{x - 1 + \mu}{r_2^3}, \quad (7)$$

$$\Omega_y = y - (1 - \mu) \frac{y}{r_1^3} - \mu \frac{y}{r_2^3}. \quad (8)$$

The planar equations of motion are

$$\ddot{x} - 2\dot{y} = \Omega_x, \quad (9)$$

$$\ddot{y} + 2\dot{x} = \Omega_y. \quad (10)$$

The velocity-dependent terms arise from the Coriolis acceleration associated with the rotating coordinate system. In terms of the state defined in Eq. (5), the equations may be written as the autonomous first-order system

$$\dot{\mathbf{x}} = f(\mathbf{x}) = \begin{bmatrix} v_x \\ v_y \\ 2v_y + \Omega_x \\ -2v_x + \Omega_y \end{bmatrix}. \quad (11)$$

Equation (11) defines the dynamical system used for all trajectory propagation in this study.

3.3 Lagrange Equilibria

Equilibrium solutions of the CR3BP correspond to spacecraft states that remain stationary in the synodic frame. At an equilibrium, the velocity and acceleration vanish, so the equations of motion reduce to

$$\Omega_x(x, y) = 0, \quad \Omega_y(x, y) = 0. \quad (12)$$

These equations admit five equilibrium points, conventionally denoted $L_1, \dots, L_5$. Three of them, L_1 , L_2 , and L_3 , lie on the line joining the two primaries, while L_4 and L_5 form equilateral configurations with the primaries.

For a collinear equilibrium, $y = 0$, and the equilibrium condition reduces to the scalar equation

$$x - (1 - \mu) \frac{x + \mu}{|x + \mu|^3} - \mu \frac{x - 1 + \mu}{|x - 1 + \mu|^3} = 0. \quad (13)$$

For the Earth–Moon mass parameter used in this study, the L_1 equilibrium lies between the two primaries. We denote its position by




Figure 2: Locations of the five Lagrange equilibria in the synodic frame. The collinear points L_1 , L_2 , and L_3 lie on the Earth–Moon axis, while L_4 and L_5 form equilateral configurations with the two primaries. The present analysis focuses on the dynamics near L_1 .

$$\mathbf{x}_{L_1} = [x_{L_1} \ 0 \ 0 \ 0]^T. \quad (14)$$

The collinear equilibria are dynamically unstable. Small perturbations near L_1 therefore do not, in general, remain close to the equilibrium indefinitely. The local dynamics nevertheless contain structured families of periodic and asymptotic motions. Of particular interest here is the family of planar Lyapunov periodic orbits surrounding L_1 , which will provide the reference trajectories from which the transport initial conditions are constructed.

3.4 Jacobi Integral and Regions of Possible Motion

The planar CR3BP admits a conserved quantity known as the Jacobi integral. In the nondimensional synodic coordinates introduced above, the Jacobi constant is

$$C_J = 2\Omega(x, y) - (v_x^2 + v_y^2), \quad (15)$$

where $\Omega(x, y)$ is the effective potential defined in Eq. (6). Along an exact CR3BP trajectory, C_J remains constant.

Rearranging Eq. (15) gives

$$v_x^2 + v_y^2 = 2\Omega(x, y) - C_J. \quad (16)$$

Because the squared speed cannot be negative, a spacecraft with fixed Jacobi constant can occupy only positions satisfying

$$2\Omega(x, y) - C_J \geq 0. \quad (17)$$

The equality

$$2\Omega(x, y) = C_J \quad (18)$$

defines the zero-velocity curves of the planar problem. These curves separate energetically accessible regions from regions that cannot be reached by any trajectory possessing the specified Jacobi constant [12].

The opening of such a neck changes the connectivity of the region in which motion is energetically permitted. This provides a useful distinction for the transport analysis developed later in this work: energetic accessibility is a necessary condition for a trajectory to pass between regions, but it does not determine which dynamically admissible trajectories will actually do so.

In particular, two initial conditions with the same Jacobi constant may lie within the same energetically accessible domain yet exhibit different finite-time outcomes. The analysis developed in Sec. 6 therefore concerns a finer structure than the zero-velocity geometry: the partition of a prescribed family of initial conditions according to their realized transport outcomes.

3.5 Planar Lyapunov Periodic Orbits

The dynamics near a collinear equilibrium contain families of periodic solutions. In the planar problem, one such family consists of Lyapunov orbits surrounding L_1 . Rather than treating L_1 itself as the reference object for transport, we construct trajectories from members of this periodic-orbit family.

Let x_{L_1} denote the x -coordinate of the L_1 equilibrium. We parameterize the family by the horizontal displacement

$$A_x = x_{L_1} - x_0, \quad (19)$$

where x_0 is the initial x -coordinate of a periodic orbit at a crossing of the x -axis. The corresponding initial state is written

$$\mathbf{x}_0(A_x) = [x_{L_1} - A_x \ 0 \ 0 \ v_{y,0}]^T. \quad (20)$$

The value of $v_{y,0}$ is selected so that the resulting trajectory satisfies the symmetry and closure conditions of a periodic orbit. For a half-period crossing of the x -axis at time $T/2$, these conditions may be expressed as

$$y\left(\frac{T}{2}\right) = 0, \quad v_x\left(\frac{T}{2}\right) = 0. \quad (21)$$

The reflected second half of the trajectory then completes a periodic orbit of period T .

A linear approximation about L_1 provides an initial estimate for a small-amplitude periodic solution. For finite amplitudes, however, the nonlinear CR3BP equations generally prevent this approximation from closing exactly.

Earth-Moon CR3BP zero-velocity geometry around L_1




Figure 3: Change in the regions of possible motion as the Jacobi constant passes through its value at L_1 . For $C_J > C_{J,L_1}$, the corresponding regions surrounding the two primaries are separated. At $C_J = C_{J,L_1}$, the zero-velocity curve reaches the L_1 equilibrium. For $C_J < C_{J,L_1}$, a neck is open in the vicinity of L_1 , making transit between the adjacent accessible regions energetically possible.

The periodic solutions used in this study are therefore obtained by numerical differential correction, as described in Sec. 4.

The present analysis considers the family interval

$$0.006 \leq A_x \leq 0.014. \quad (22)$$

Each accepted family member is required to satisfy numerical closure and Jacobi-conservation criteria before it is used in the subsequent transport analysis. The periodic orbit itself does not determine a transport destination. Instead, its local stability structure provides directions along which nearby trajectories approach or depart from the orbit.

4 Numerical Methods

4.1 Trajectory Propagation and Numerical Accuracy

All trajectories in this study were propagated by direct numerical integration of the nondimensional CR3BP equations given in Eq. (11). An adaptive high-order integration scheme was used so that the integration step size could vary with the local dynamics rather than being prescribed uniformly in time.

Unless otherwise stated, integrations used relative and absolute tolerances

$$\text{rtol} = 10^{-11}, \quad \text{atol} = 10^{-13}. \quad (23)$$

These tolerances were retained throughout the periodic-orbit, stability, manifold, and first-passage calculations used in the transport analysis. Numerical integrations were not accepted solely because the ODE

solver reported successful termination. Conservation of the Jacobi integral was monitored independently as a trajectory-level diagnostic.

For a trajectory initialized with Jacobi constant $C_J(0)$, we define the instantaneous conservation error

$$\Delta C_J(t) = C_J(t) - C_J(0), \quad (24)$$

and summarize the numerical drift over an integration interval by

$$\varepsilon_C = \max_{0 \leq t \leq t_f} |\Delta C_J(t)|. \quad (25)$$

This quantity provides an independent check on the numerical propagation because the exact CR3BP dynamics conserve C_J . It is used throughout the computational pipeline to distinguish changes in trajectory behavior from obvious integration failure.

The transport calculations are finite-time calculations. A trajectory is therefore propagated only until the first relevant terminal event occurs or until the prescribed horizon is reached. For the final transport atlas, the maximum nondimensional propagation horizon was

$$T_{\max} = 160. \quad (26)$$

This horizon should not be interpreted as an asymptotic statement about the trajectory. A trajectory classified as having no section crossing by $T_{\max}$ is only known not to have produced the specified event during the interval examined.

Event detection was performed during integration rather than by searching a discretely sampled trajectory after propagation. In particular, section crossings and




Figure 4: Representative members of the planar Lyapunov-orbit family about L_1 in the Earth–Moon CR3BP. The family is shown for $A_x = 0.006, 0.008, 0.010, 0.012$, and 0.014 . Each trajectory is a numerically corrected nonlinear periodic solution. Increasing A_x moves the initial x -axis crossing farther from L_1 and changes the orbit geometry.

primary-proximity terminations were represented as continuous event functions whose roots were localized by the integrator. A candidate section crossing was retained only when it satisfied the prescribed crossing direction and transversality criteria. This prevents an initial condition lying on the section, or a near-tangent contact with it, from being counted automatically as a physical first passage.

For the transport experiments, the perturbation scale used to generate nearby manifold trajectories was

$$\epsilon = 10^{-6}, \quad (27)$$

in nondimensional state-space units. The same perturbation convention was retained when comparing trajectories across the periodic-orbit family so that changes in the resulting transport structure were not introduced

by changing the seed magnitude.

The combination of adaptive integration, event-based termination, Jacobi monitoring, and fixed numerical conventions provides the propagation layer used by all subsequent calculations. Periodic-orbit construction introduces an additional numerical requirement: the reference trajectory must close to high accuracy before its stability or associated transport structures are evaluated. The correction and continuation procedure used to obtain those reference trajectories is described next.

4.2 Differential Correction and Lyapunov-Family Continuation

The planar Lyapunov orbits used as reference trajectories were obtained by differential correction of the symmetric initial states introduced in Eq. (20). For a prescribed

amplitude A_x , the initial position is fixed at

$$x_0 = x_{L_1} - A_x, \quad y_0 = 0, \quad v_{x,0} = 0, \quad (28)$$

while $v_{y,0}$ is adjusted to enforce the periodicity condition.

Starting from this state, the trajectory is propagated until the next admissible crossing of the x -axis. Denoting this crossing time by t_h , symmetry requires

$$y(t_h) = 0, \quad v_x(t_h) = 0. \quad (29)$$

The first condition defines the half-period event. The remaining velocity component therefore provides the scalar correction residual,

$$F(v_{y,0}) = v_x(t_h). \quad (30)$$

The corrected initial velocity is obtained iteratively using a Newton update,

$$v_{y,0}^{(k+1)} = v_{y,0}^{(k)} - \frac{F(v_{y,0}^{(k)})}{F'(v_{y,0}^{(k)})}. \quad (31)$$

Because the crossing time itself depends on the initial condition, F' must include the displacement of the terminal event. Let $\Phi(t, 0)$ denote the state-transition matrix and let

$$\Phi_{ij}(t, 0) = \frac{\partial x_i(t)}{\partial x_j(0)}. \quad (32)$$

At the half-period crossing, the event condition $y(t_h) = 0$ gives

$$\frac{\partial t_h}{\partial v_{y,0}} = -\frac{\Phi_{24}(t_h, 0)}{v_y(t_h)}. \quad (33)$$

The derivative of the correction residual is consequently

$$F' = \Phi_{34}(t_h, 0) - \frac{\dot{v}_x(t_h)}{v_y(t_h)} \Phi_{24}(t_h, 0), \quad (34)$$

where the state ordering is $\mathbf{x} = [x, y, v_x, v_y]^T$. The iteration is continued until the half-period residual satisfies the prescribed correction tolerance. The resulting trajectory is then propagated through a complete period and tested independently for closure and Jacobi conservation before being accepted.

Solving every amplitude independently is not sufficient to guarantee that the resulting solutions remain on the same nonlinear family. As the amplitude increases, a Newton iteration initialized from a local linear approximation may converge to a different root or leave the desired branch. We therefore construct the family by predictor-corrector continuation.

The first member of the family is initialized from the linearized center mode about L_1 . After correction and validation, its solution provides the initial estimate for the next amplitude. Once two validated members are available, the initial transverse velocity for the next member is predicted by secant continuation.

For consecutive amplitudes $A_x^{(k-1)}$ and $A_x^{(k)}$, the predictor for a new amplitude $A_x^{(k+1)}$ is

$$\hat{v}_{y,0}^{(k+1)} = v_{y,0}^{(k)} + \frac{A_x^{(k+1)} - A_x^{(k)}}{A_x^{(k)} - A_x^{(k-1)}} (v_{y,0}^{(k)} - v_{y,0}^{(k-1)}). \quad (35)$$

This prediction is used only as the starting point for the same differential-correction procedure; it does not replace the nonlinear periodicity solve. Every family member must therefore satisfy the same closure, conservation, and stability checks regardless of how its initial guess was obtained.

The family used in the final transport calculations contains 33 members uniformly spaced over

$$A_x \in [0.006, 0.014], \quad \Delta A_x = 2.5 \times 10^{-4}. \quad (36)$$

Across this continuation sequence, the corrected initial states and periods vary smoothly. The continuation procedure is therefore used to preserve branch consistency while moving through the finite-amplitude family rather than to impose smoothness on the resulting solutions.

4.3 State-Transition Matrix and Floquet Stability

The linear sensitivity of a trajectory to perturbations in its initial state is described by the state-transition matrix (STM). For the nonlinear system

$$\dot{\mathbf{x}} = f(\mathbf{x}), \quad (37)$$

an infinitesimal perturbation $\delta \mathbf{x}$ evolves according to the variational equations

$$\delta \dot{\mathbf{x}} = A(t) \delta \mathbf{x}, \quad A(t) = \left. \frac{\partial f}{\partial \mathbf{x}} \right|_{\mathbf{x}(t)}. \quad (38)$$

The STM $\Phi(t, 0)$ therefore satisfies

$$\dot{\Phi}(t, 0) = A(t) \Phi(t, 0), \quad \Phi(0, 0) = I_4, \quad (39)$$

where I_4 is the 4×4 identity matrix. The spacecraft state and the 16 STM components are integrated simultaneously, so that the linearized dynamics are evaluated




Figure 5: Differential correction of a symmetric planar L_1 Lyapunov orbit. For fixed A_x , an initial linearized guess is propagated to the next admissible $y = 0$ half-period crossing, where the residual $v_x(t_h)$ is evaluated. The initial transverse velocity $v_{y,0}$ is iteratively corrected using the state-transition matrix until the symmetry condition $v_x(t_h) = 0$ is satisfied. The resulting corrected trajectory forms a periodic orbit under reflection symmetry; full-period closure and Jacobi conservation are verified independently.

along the same nonlinear trajectory used for the reference orbit.

For the planar CR3BP state ordering $\mathbf{x} = [x, y, v_x, v_y]^T$, the Jacobian has the form

$$A(t) = \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ \Omega_{xx} & \Omega_{xy} & 0 & 2 \\ \Omega_{xy} & \Omega_{yy} & -2 & 0 \end{bmatrix}_{\mathbf{x}(t)}, \quad (40)$$

with

$$\Omega_{xx} = 1 - (1 - \mu) \left(\frac{1}{r_1^3} - \frac{3(x + \mu)^2}{r_1^5} \right) - \mu \left(\frac{1}{r_2^3} - \frac{3(x - 1 + \mu)^2}{r_2^5} \right), \quad (41)$$

$$\Omega_{yy} = 1 - (1 - \mu) \left(\frac{1}{r_1^3} - \frac{3y^2}{r_1^5} \right) - \mu \left(\frac{1}{r_2^3} - \frac{3y^2}{r_2^5} \right), \quad (42)$$

$$\Omega_{xy} = 3(1 - \mu) \frac{(x + \mu)y}{r_1^5} + 3\mu \frac{(x - 1 + \mu)y}{r_2^5}. \quad (43)$$

For a periodic orbit of period T , the STM after one complete revolution defines the monodromy matrix,

$$M = \Phi(T, 0). \quad (44)$$

Its eigenvalues,

$$M\mathbf{v}_i = \lambda_i\mathbf{v}_i, \quad (45)$$

are the Floquet multipliers of the periodic orbit. They characterize the linear evolution of infinitesimal perturbations after one period.

For the hyperbolic planar Lyapunov orbits considered here, the relevant nontrivial multipliers form a reciprocal real pair,

$$|\lambda_s| < 1, \quad |\lambda_u| > 1, \quad \lambda_s\lambda_u \simeq 1, \quad (46)$$

where λ_s and λ_u denote the stable and unstable multipliers, respectively. The associated eigenvectors $\mathbf{v}_s$ and $\mathbf{v}_u$ define the corresponding local directions at the reference phase of the orbit.

Because manifold trajectories are seeded at multiple phases rather than only at the initial point, these directions must be transported along the periodic orbit. If $\mathbf{x}_p(t)$ denotes the periodic reference trajectory, the phase-dependent directions are obtained from

$$\tilde{\mathbf{v}}_{s,u}(t) = \Phi(t, 0)\mathbf{v}_{s,u}. \quad (47)$$

They are normalized before perturbations are applied,

$$\hat{\mathbf{v}}_{s,u}(t) = \frac{\tilde{\mathbf{v}}_{s,u}(t)}{\|\tilde{\mathbf{v}}_{s,u}(t)\|}. \quad (48)$$

This construction preserves the phase dependence of the linear stable and unstable directions along the nonlinear periodic orbit.

The magnitude of the unstable multiplier also provides a compact measure of the linear instability of each family member. Across the 33-member continuation sequence used here, the unstable multiplier decreased smoothly from approximately

$$\lambda_u = 2.620 \times 10^3 \quad \text{at} \quad A_x = 0.006 \quad (49)$$

to

$$\lambda_u = 2.337 \times 10^3 \quad \text{at} \quad A_x = 0.014, \quad (50)$$

while the corresponding stable multiplier increased from approximately 3.817×10^{-4} to 4.280×10^{-4} . The smooth variation of these quantities provides an additional diagnostic that continuation has remained on a consistent periodic-orbit branch.

The Floquet eigenvectors are local linear objects and are not themselves finite-time transport trajectories. They are used only to define controlled perturbations away from the periodic reference orbit. The nonlinear evolution of those perturbed states is computed by direct integration of the full CR3BP equations.

4.4 Invariant-Manifold Seed Construction

The stable and unstable Floquet directions provide local directions of contraction and expansion about a hyperbolic periodic orbit. To convert these linear directions into initial conditions for nonlinear transport trajectories, perturbations are applied to the periodic orbit at a set of phases distributed over one complete revolution.

Let the phase along a periodic orbit of period T be parameterized by

$$\theta = 2\pi \frac{t}{T}, \quad 0 \leq \theta < 2\pi. \quad (51)$$

The corresponding state on the periodic orbit is denoted by

$$\mathbf{x}_p(\theta; A_x), \quad (52)$$

where A_x identifies the member of the planar Lyapunov family.

At each phase, the stable and unstable Floquet directions are transported from the reference phase using the state-transition matrix as described in Sec. 4.3. After normalization, the local directions are denoted by $\hat{\mathbf{v}}_s(\theta)$ and $\hat{\mathbf{v}}_u(\theta)$.

Local manifold seeds are then constructed as

$$\mathbf{x}_s^\pm(\theta; A_x) = \mathbf{x}_p(\theta; A_x) \pm \epsilon \hat{\mathbf{v}}_s(\theta), \quad (53)$$

and

$$\mathbf{x}_u^\pm(\theta; A_x) = \mathbf{x}_p(\theta; A_x) \pm \epsilon \hat{\mathbf{v}}_u(\theta), \quad (54)$$

where ϵ is a small perturbation magnitude. Unless otherwise stated, the computations in this study use

$$\epsilon = 10^{-6}. \quad (55)$$

The two signs generate the two local branches associated with each Floquet direction. Thus, at a given phase, the linearized construction produces two stable seeds and two unstable seeds.

The perturbation in Eqs. (53)–(54) is used only to initialize the trajectory. Once a seed has been formed, its subsequent evolution is obtained from the full nonlinear CR3BP equations rather than from the variational dynamics. The resulting curves therefore depart from the linear Floquet directions as they evolve away from the periodic orbit.

For the unstable direction, seeds are propagated forward in time,

$$\mathbf{x}_u^\pm(t; \theta, A_x) = \varphi_t(\mathbf{x}_u^\pm(\theta; A_x)), \quad t > 0, \quad (56)$$

where φ_t denotes the nonlinear CR3BP flow. Stable-manifold trajectories may analogously be obtained by backward-time propagation,


$$\mathbf{x}_s^\pm(t; \theta, A_x) = \varphi_t(\mathbf{x}_s^\pm(\theta; A_x)), \quad t < 0. \quad (57)$$

This time orientation follows directly from the Floquet structure: perturbations along the unstable direction grow under forward propagation, whereas perturbations along the stable direction contract toward the periodic orbit under forward propagation.

It is important to distinguish the local seed construction from the global invariant manifold. The vectors $\hat{\mathbf{v}}_{s,u}$ are obtained from the linearized dynamics and describe only the tangent directions to the corresponding manifolds near the periodic orbit. The finite trajectories obtained after propagation are nonlinear objects and may develop geometries that are not apparent from the local eigendirections alone.

For numerical transport experiments, the orbital phase θ is therefore treated as an independent parameter. Sampling θ over the interval $[0, 2\pi)$ produces a family of nearby initial conditions distributed around the reference periodic orbit. Varying A_x in addition to θ extends this construction across the Lyapunov family and defines the two-parameter initial-condition domain

$$\mathcal{D} = \{(\theta, A_x) : 0 \leq \theta < 2\pi, 0.006 \leq A_x \leq 0.014\}. \quad (58)$$



Floquet multipliers


λ_u	2511.77712
λ_s	0.000398124497

$$A_x = 0.010$$

$$T = 2.717209585$$

Arrows: normalized x-y
projections; length graphical.

Figure 6: Floquet stability of a representative hyperbolic planar L_1 Lyapunov orbit. Integration of the state-transition matrix over one orbital period yields the monodromy matrix $M = \Phi(T, 0)$. Its nontrivial reciprocal Floquet multipliers identify the stable and unstable eigendirections associated with contraction and expansion about the periodic orbit. These eigendirections are transported along the orbit using the state-transition matrix and subsequently used to construct the corresponding invariant-manifold initial conditions.



Unstable seed construction

$$A_x = 0.010$$

$$\theta = \pi/2$$

$$t = T/4$$

$$\epsilon = 10^{-6}$$

Inset magnifies the view only;
the propagated seeds are unscaled.

Figure 7: Construction of local invariant-manifold trajectories from a planar L_1 Lyapunov orbit. At a selected orbital phase θ , the unstable Floquet eigendirection is transported to the corresponding point on the periodic orbit using the state-transition matrix and normalized. Perturbations of magnitude ϵ are applied with both signs to generate local manifold seeds, which are subsequently propagated under the full nonlinear CR3BP dynamics to form the unstable manifold branches.

Each point in $\mathcal{D}$ therefore identifies a periodic-orbit family member and a phase on that orbit, from which a prescribed manifold branch can be initialized.

This parameterization is central to the transport analysis developed below. Rather than studying isolated manifold trajectories, we examine how the finite-time outcome of the nonlinear propagation varies across $\mathcal{D}$. Abrupt changes of outcome between nearby points in this domain define the transport boundaries investigated in the subsequent experiments.

4.5 First-Passage Transport Classification

The manifold construction of Sec. 4.4 generates a nonlinear trajectory for each sampled initial condition (θ, A_x) . To compare these trajectories systematically, a finite-time first-passage problem is defined. Each trajectory is propagated until it satisfies a prescribed terminal event or reaches the maximum integration horizon.

A transport section is introduced at

$$x = x_{\text{sec}}, \quad (59)$$

where x_{sec} is fixed for a given experiment. A section crossing is accepted only when the trajectory reaches this surface with the prescribed crossing orientation and sufficient transversality. If

$$g_{\text{sec}}(\mathbf{x}) = x - x_{\text{sec}}, \quad (60)$$

then a candidate crossing occurs when

$$g_{\text{sec}}(\mathbf{x}(t)) = 0. \quad (61)$$

Because

$$\dot{g}_{\text{sec}} = v_x, \quad (62)$$

the magnitude $|v_x|$ at the event provides a direct measure of transversality. Numerically degenerate crossings are rejected when

$$|v_x| < v_{\text{trans}}, \quad (63)$$

where v_{trans} is the prescribed transversality tolerance.

For an initial condition $\mathbf{x}_0$, the first-passage time is defined by

$$\tau(\mathbf{x}_0) = \inf \{t > 0 : g_{\text{sec}}(\mathbf{x}(t; \mathbf{x}_0)) = 0 \text{ and the crossing is accepted}\} \quad (64)$$

If such an event occurs before any competing terminal condition, the trajectory is assigned the outcome

$$\mathcal{O}(\theta, A_x) = \text{SECTION_CROSSING}. \quad (65)$$

Propagation is also terminated when a trajectory enters a prescribed proximity region surrounding either

primary. Let $r_{1,\text{min}}$ and $r_{2,\text{min}}$ denote the corresponding termination radii. A primary-proximity event occurs when

$$r_1 \leq r_{1,\text{min}} \quad \text{or} \quad r_2 \leq r_{2,\text{min}}. \quad (66)$$

Such a trajectory is assigned

$$\mathcal{O}(\theta, A_x) = \text{PRIMARY_PROXIMITY}. \quad (67)$$

This category is treated as a numerical and dynamical termination class rather than being interpreted as a physical impact model. The CR3BP primaries are idealized point masses, and the proximity threshold is introduced to prevent integrations from proceeding arbitrarily close to the singular gravitational terms.

If neither an accepted section crossing nor a primary-proximity event occurs before the finite integration horizon

$$T_{\text{max}} = 160, \quad (68)$$

the trajectory is classified as

$$\mathcal{O}(\theta, A_x) = \text{NO_CROSSING_BY_160}. \quad (69)$$

The resulting outcome map is therefore a discrete function

$$\mathcal{O} : \mathcal{D} \longrightarrow \left\{ \begin{array}{l} \text{SECTION_CROSSING}, \\ \text{PRIMARY_PROXIMITY}, \\ \text{NO_CROSSING_BY_160} \end{array} \right\}, \quad (70)$$

where $\mathcal{D}$ is the (θ, A_x) initial-condition domain defined in Eq. (58).

The first-passage formulation is intentionally finite-time. A trajectory assigned NO_CROSSING_BY_160 is not claimed to remain non-transiting indefinitely; the classification states only that no accepted section crossing was observed within the specified horizon. Similarly, the classification does not imply that all trajectories belonging to the same outcome have comparable passage times or geometries.

For trajectories assigned SECTION_CROSSING, additional continuous diagnostics are recorded. These include the first-passage time τ , the minimum distance from each primary, the maximum rotating-barycentric excursion, and the time spent in the prescribed L_1 neighborhood. These quantities retain information that is lost when the trajectory is reduced to a discrete outcome label.

Numerical quality is monitored simultaneously with the event classification. In particular, Jacobi-constant drift is evaluated along propagated trajectories, and accepted




Figure 8: Finite-time first-passage classification of manifold-seeded trajectories. Propagation under the nonlinear CR3BP dynamics terminates when a trajectory undergoes an accepted transverse crossing of the prescribed transport section, enters a primary-proximity termination region, or reaches the finite integration horizon $T_{\max} = 160$ without an accepted crossing. Candidate section events that fail the prescribed transversality criterion are rejected. These terminal conditions define the discrete outcome map $\mathcal{O}(\theta, A_x)$.

section events are checked for both section residual and transversality. Event classifications that fail the prescribed numerical criteria are retained as numerical failures rather than silently reassigned to a physical outcome.

This distinction is important near transport boundaries. A change in $\mathcal{O}$ between neighboring initial conditions is interpreted as an observed finite-time outcome change only when both classifications satisfy the numerical acceptance criteria. Numerical integration failure, insufficient event transversality, or unresolved propagation is not itself treated as evidence for a transport boundary.

4.6 Outcome Atlas and Boundary Refinement

The first-passage classification defined in Sec. 4.5 assigns a discrete transport outcome to each manifold initial condition. To examine how these outcomes vary across the Lyapunov family, the two-parameter domain

$$\mathcal{D} = \{(\theta, A_x) : 0 \leq \theta < 2\pi, 0.006 \leq A_x \leq 0.014\} \quad (71)$$

is sampled on a structured grid.

The amplitude coordinate is represented by the con-

tinuation sequence

$$A_x^{(j)} = 0.006 + 0.00025j, \quad j = 0, \dots, 32, \quad (72)$$

giving 33 validated members of the planar Lyapunov family. For each family member, the orbital phase is sampled uniformly over one complete period. The resulting initial conditions form a discrete approximation to the continuous parameter domain $\mathcal{D}$.

For every grid point, the prescribed manifold seed is constructed and propagated independently according to the first-passage procedure of Sec. 4.5. The corresponding discrete atlas may be written as

$$\mathcal{A}_{ij} = \mathcal{O}(\theta_i, A_x^{(j)}), \quad (73)$$

where $\mathcal{O}$ is the finite-time outcome map defined in Eq. (70).

The atlas therefore represents a partition of the sampled (θ, A_x) domain according to realized nonlinear transport outcome. Adjacent samples need not belong to the same outcome class. Changes between neighboring valid samples provide the initial numerical evidence for transport boundaries.

4.6.1 Candidate Boundary Detection

Let two neighboring atlas samples be denoted by parameter vectors

$$\mathbf{p}_a = \begin{bmatrix} \theta_a \\ A_{x,a} \end{bmatrix}, \quad \mathbf{p}_b = \begin{bmatrix} \theta_b \\ A_{x,b} \end{bmatrix}. \quad (74)$$

An edge is marked as a candidate transport boundary when the two endpoint classifications are both numerically valid and satisfy

$$\mathcal{O}(\mathbf{p}_a) \neq \mathcal{O}(\mathbf{p}_b). \quad (75)$$

This condition detects a change in the observed finite-time outcome along the sampled edge. It does not, by itself, establish that the underlying boundary is smooth, unique, or globally connected.

Each candidate edge defines an initial bracket

$$B_0 = [\mathbf{p}_a, \mathbf{p}_b]. \quad (76)$$

For edges admitting a two-outcome bracket, the midpoint

$$\mathbf{p}_m = \frac{1}{2}(\mathbf{p}_a + \mathbf{p}_b) \quad (77)$$

is evaluated using the same manifold construction, nonlinear propagation, and event-classification procedure used for the base atlas.

If the midpoint has the outcome of endpoint a , the endpoint $\mathbf{p}_a$ is replaced by $\mathbf{p}_m$. If it has the outcome of endpoint b , $\mathbf{p}_b$ is replaced instead. Repeated application of this procedure produces a nested sequence of brackets,

$$B_0 \supset B_1 \supset B_2 \supset \dots, \quad (78)$$

that localizes the observed outcome transition.

A refined numerical boundary location may then be represented by the center of the final bracket,

$$\hat{\mathbf{p}}_b = \frac{\mathbf{p}_a^* + \mathbf{p}_b^*}{2}, \quad (79)$$

where the starred endpoints denote the final retained bracket.

The associated numerical boundary-resolution scale is taken from the remaining parameter-space separation,

$$\delta_b = \|\mathbf{p}_a^* - \mathbf{p}_b^*\|, \quad (80)$$

with the parameter normalization used by the computational implementation.

4.6.2 Complex and Unresolved Cells

Not every local outcome structure can be represented reliably by a single two-outcome edge bracket. A sampled cell may contain an additional outcome class, an invalid numerical evaluation, or multiple neighboring outcome transitions whose topology is not resolved by the available sampling.

Such cases are retained explicitly as complex or unresolved structures rather than being forced into a smooth-boundary interpretation.

We therefore distinguish between the existence of an observed outcome change and a stronger claim about the topology of the underlying boundary. The refinement procedure estimates local transitions along sampled edges; it does not reconstruct an exact global separatrix or prove smoothness of the boundary in the continuous parameter domain.

The resulting collection of refined boundary samples provides a numerical representation of the finite-time transport partition in parameter space. These samples are not interpreted as singularities of the equations of motion. Rather, they identify locations at which nearby members of the prescribed initial-condition family produce different finite-time transport outcomes.

This distinction motivates the robustness analysis introduced next. Instead of asking only whether an outcome boundary exists, we seek to quantify how far a nominal initial condition lies from the observed boundary under a physically informed metric on the (θ, A_x) parameter domain.




Figure 9: Construction of the transport-outcome atlas over the two-dimensional Lyapunov-family parameter domain (θ, A_x) . The orbital phase θ and family amplitude A_x determine a manifold-seeded initial state $\mathbf{X}(\theta, A_x)$, which is propagated under the nonlinear CR3BP dynamics and assigned a finite-time first-passage outcome. Repeating this procedure over the sampled parameter grid produces the discrete outcome map $\mathcal{O}(\theta, A_x)$, from which neighboring samples with different outcomes are identified as candidate transport-boundary edges.




Figure 10: Adaptive refinement of finite-time transport-outcome boundaries in the (θ, A_x) parameter domain. Neighboring valid samples assigned different outcomes define an initial boundary bracket. Each midpoint is propagated through the full nonlinear first-passage classifier, after which the subinterval retaining the outcome change is recursively selected. Refinement continues until the prescribed numerical stopping criterion is reached. The final bracket defines a numerical boundary sample together with its finite resolution; locally complex configurations are retained rather than forced into a simple binary boundary.

4.7 Pullback Metric and Robustness Distance

The outcome atlas identifies locations at which nearby members of the prescribed initial-condition family produce different finite-time transport outcomes. A useful measure of robustness, however, requires more than the

coordinate separation from such a boundary. The two parameters θ and A_x have different dynamical meanings and scales, so an ordinary Euclidean distance in the (θ, A_x) plane is not intrinsically meaningful.

We therefore define distance through the initial spacecraft state generated by the parameterization.

Let

$$\mathbf{p} = \begin{bmatrix} \theta \\ A_x \end{bmatrix} \quad (81)$$

denote a point in the parameter domain, and let

$$\mathbf{X}(\mathbf{p}) = \mathbf{X}(\theta, A_x) \in \mathbb{R}^4 \quad (82)$$

denote the corresponding manifold-seed state produced by the construction of Sec. 4.4.

The differential of this map is the 4×2 Jacobian

$$J(\mathbf{p}) = \frac{\partial \mathbf{X}}{\partial \mathbf{p}} = \begin{bmatrix} \frac{\partial \mathbf{X}}{\partial \theta} & \frac{\partial \mathbf{X}}{\partial A_x} \end{bmatrix}. \quad (83)$$

For a small parameter perturbation $\delta \mathbf{p}$,

$$\delta \mathbf{X} = J(\mathbf{p}) \delta \mathbf{p} + \mathcal{O}(\|\delta \mathbf{p}\|^2). \quad (84)$$

To compare perturbations in position and velocity, introduce a positive-definite state-space weighting matrix

$$W_\eta = \begin{bmatrix} I_2 & 0 \\ 0 & \eta I_2 \end{bmatrix}, \quad \eta > 0, \quad (85)$$

where η controls the relative contribution of velocity perturbations to the state-space norm.

The weighted infinitesimal state displacement is

$$\|\delta \mathbf{X}\|_{W_\eta}^2 = \delta \mathbf{X}^T W_\eta \delta \mathbf{X}. \quad (86)$$

Substitution of the local parameterization gives

$$\|\delta \mathbf{X}\|_{W_\eta}^2 \simeq \delta \mathbf{p}^T G_\eta(\mathbf{p}) \delta \mathbf{p}, \quad (87)$$

where

$$G_\eta(\mathbf{p}) = J(\mathbf{p})^T W_\eta J(\mathbf{p}) \quad (88)$$

is the pullback metric induced on the two-dimensional parameter domain.

Thus, G_η converts an infinitesimal displacement in (θ, A_x) into the corresponding first-order weighted displacement of the physical initial state. In general,

$$G_\eta = \begin{bmatrix} G_{\theta\theta} & G_{\theta A} \\ G_{\theta A} & G_{AA} \end{bmatrix}, \quad (89)$$

so the metric need not align with either coordinate axis. Its eigenvectors identify locally preferred parameter-space directions, while its eigenvalues quantify the state-space sensitivity associated with perturbations along those directions.

4.7.1 Distance to an Observed Transport Boundary

Let $\mathcal{B}$ denote the numerically resolved set of transport boundary samples obtained using the procedure of Sec. 4.6. For a nominal parameter point $\mathbf{p}$, we define a local metric distance to a candidate boundary point $\mathbf{b} \in \mathcal{B}$ by

$$d_\eta(\mathbf{p}, \mathbf{b}) = \sqrt{(\mathbf{b} - \mathbf{p})^T G_\eta(\mathbf{p}) (\mathbf{b} - \mathbf{p})}. \quad (90)$$

The estimated robustness radius is then

$$\rho_\eta(\mathbf{p}) = \min_{\mathbf{b} \in \mathcal{B}} d_\eta(\mathbf{p}, \mathbf{b}) \quad (91)$$

subject to the family and numerical-boundary constraints imposed by the computational implementation.

The quantity ρ_η therefore estimates the size of the smallest local weighted perturbation connecting the nominal initial condition to an observed finite-time outcome boundary. Small values indicate that a numerically resolved boundary lies nearby in the induced metric, whereas larger values indicate greater separation from the sampled boundary.

Importantly, ρ_η is not itself a probability of failure, a Lyapunov exponent, or a proof that every perturbation of magnitude greater than ρ_η changes the trajectory outcome. It is a boundary-distance diagnostic defined relative to the selected initial-condition family, finite-time outcome map, numerical boundary resolution, and state-space metric.

4.7.2 Metric Dependence

The velocity weight η is not fixed by the CR3BP equations. Different choices correspond to different definitions of what constitutes an equally important perturbation of the initial state. The robustness calculation is therefore evaluated over the set


$$\eta \in \{0.1, 0.3, 1, 3, 10\}. \quad (92)$$

Rather than selecting a favorable value after observing the transport atlas, the analysis retains the dependence of ρ_η on η . This allows the stability of relative robustness rankings to be examined explicitly.

For two choices η_a and η_b , the corresponding robustness rankings may be compared through their rank correlation,

$$r_s(\eta_a, \eta_b) = \text{corr}_{\text{rank}}(\rho_{\eta_a}, \rho_{\eta_b}). \quad (93)$$

A robustness ordering that remains nearly unchanged over a wide range of η would indicate that the inferred ordering is relatively insensitive to the chosen position-velocity weighting. Conversely, substantial changes in ranking indicate that the robustness measure is metric



Metric is induced only on the prescribed two-dimensional manifold-seeded family.

Figure 11: Geometric interpretation of the local pullback metric on the (θ, A_x) parameter domain. A small parameter perturbation $\delta \mathbf{p}$ is mapped by the Jacobian J of the manifold-seed parameterization to a perturbation $\delta \mathbf{X}$ of the four-dimensional initial state. Pulling the weighted state-space norm back through this map yields $G_\eta = J^T W_\eta J$. Equal-distance contours under this metric are generally elliptical in parameter coordinates, reflecting the unequal physical effect of perturbations in different local directions.

dependent and should not be interpreted as an invariant property of the trajectory family.

4.7.3 Local Metric Validation

Because G_η is obtained from a first-order pullback, its numerical implementation is checked against finite perturbations of the actual manifold-seed map. For a small parameter displacement $\delta \mathbf{p}$, the metric predicts

$$d_{\text{pred}}^2 = \delta \mathbf{p}^T G_\eta \delta \mathbf{p}, \quad (94)$$

while direct evaluation of the perturbed seed gives

$$d_{\text{actual}}^2 = [\mathbf{X}(\mathbf{p} + \delta \mathbf{p}) - \mathbf{X}(\mathbf{p})]^T W_\eta [\mathbf{X}(\mathbf{p} + \delta \mathbf{p}) - \mathbf{X}(\mathbf{p})] \quad (95)$$

Agreement is quantified through the relative error

$$e_G = \frac{|d_{\text{actual}}^2 - d_{\text{pred}}^2|}{d_{\text{actual}}^2}. \quad (96)$$

The perturbation scale is successively reduced to verify the expected local convergence of the linear approximation. This check tests the implementation of the pullback geometry independently of whether ρ_η ultimately proves to be a reliable predictor of an observed outcome change.

4.8 Numerical Validation and Reproducibility

The transport boundaries studied here arise from finite-time classifications of nonlinear trajectories. Numerical error can therefore produce apparent changes of outcome that are difficult to distinguish from genuine dynamical sensitivity. For this reason, the computational pipeline was constructed so that periodic-orbit generation, stability analysis, manifold initialization, trajectory propagation, event detection, and boundary refinement each admit separate numerical checks.

4.8.1 Trajectory Integration and Jacobi Monitoring

All CR3BP trajectories are propagated by adaptive numerical integration of Eq. (11). Unless otherwise stated, the integration tolerances used in the transport calculations are

$$\text{rtol} = 10^{-11}, \quad \text{atol} = 10^{-13}. \quad (97)$$

Because the Jacobi constant is conserved by the exact CR3BP dynamics, its numerical drift provides an independent diagnostic of integration accuracy. For a propagated trajectory $\mathbf{x}(t)$, define

$$\Delta C_J(t) = C_J(\mathbf{x}(t)) - C_J(\mathbf{x}(0)). \quad (98)$$

The trajectory-level diagnostic is

$$E_J = \max_t |\Delta C_J(t)|. \quad (99)$$




Figure 12: Geometric interpretation of the family-constrained boundary-distance diagnostic ρ_η . For a nominal parameter point $\mathbf{p}$, distances to the numerically resolved outcome-boundary samples are measured using the local pullback metric $G_\eta(\mathbf{p})$. The quantity $\rho_\eta(\mathbf{p})$ is defined by the nearest boundary sample under this metric and may be visualized as the smallest metric ellipse centered on $\mathbf{p}$ that reaches the sampled boundary. The construction is local, depends on the weighting parameter η , and is restricted to the prescribed two-dimensional family of manifold-seeded initial conditions.

This quantity is recorded independently of the trajectory's eventual transport classification. A trajectory is therefore not regarded as numerically reliable merely because an event was detected.

4.8.2 Periodic-Orbit Validation

The Lyapunov family used in the transport calculations is generated by differential correction of the symmetric initial condition introduced in Eq. (20). A corrected member must first satisfy the half-period symmetry conditions of Eq. (21).

The resulting state is then propagated through a complete period. Closure is measured by

$$E_{\text{cl}} = \|\mathbf{x}(T) - \mathbf{x}(0)\|_2. \quad (100)$$

Jacobi conservation is checked independently during the same propagation. Only family members satisfying the prescribed closure and conservation criteria are admitted to the subsequent Floquet and manifold calculations.

The family is constructed by predictor-corrector continuation in A_x . The first member is initialized from the local linear dynamics near L_1 . The next corrected member provides an initial continuation direction, after which a secant predictor is formed from the two previously validated family members. Every predicted state is still subjected to the full nonlinear differential correction and validation procedure; continuation therefore supplies an initial guess rather than replacing the periodic-orbit solve.

This continuation step is important because independently initialized finite-amplitude solves can converge to unintended branches or nonphysical correction iterates even when the underlying periodic family remains smooth.

4.8.3 State-Transition and Floquet Validation

For each accepted periodic orbit, the state-transition matrix is integrated together with the reference trajectory according to the variational equations. After one period, the resulting monodromy matrix

$$M = \Phi(T, 0) \quad (101)$$

is used to obtain the Floquet multipliers and the stable and unstable eigendirections employed in the manifold construction.

The computed stability structure is subjected to independent checks before the corresponding eigenvectors are used as transport seeds. These include consistency of the full-period variational propagation and nonlinear propagation of small perturbations along the candidate stable and unstable directions. The latter check verifies that the observed finite-time behavior agrees with the contraction or expansion implied by the corresponding Floquet multiplier.

The accepted continuation family is additionally inspected for smooth variation of the periodic states, periods, and nontrivial Floquet multipliers with A_x .

Abrupt unexplained jumps are treated as evidence of branch switching or failed continuation rather than as physical features of the family.

4.8.4 Event Validation and Outcome Classification

Transport outcomes are assigned using explicit first-passage events rather than trajectory appearance. A section crossing is accepted only when the required crossing orientation and transversality conditions are satisfied.

If $h(\mathbf{x}) = 0$ defines the section, the transversality diagnostic is

$$\gamma = |\nabla h(\mathbf{x}) \cdot f(\mathbf{x})|. \quad (102)$$

Crossings for which γ falls below the prescribed transversality threshold are not silently interpreted as ordinary transport events. Likewise, primary-proximity termination, absence of a crossing before the finite integration horizon, near-tangent encounters, and numerical integration failures are retained as distinct computational outcomes.

This distinction prevents a failed or ambiguous propagation from being converted into an apparent member of a neighboring physical outcome class.

The default finite-time horizon used in the final atlas is

$$T_{\max} = 160. \quad (103)$$

Consequently, the label `NO_CROSSING_BY_160` denotes a finite-time observation rather than a statement that the trajectory can never cross the section.

4.8.5 Boundary Refinement and Resolution

Outcome boundaries are first identified from neighboring atlas samples with different valid classifications. Each accepted edge brackets a one-dimensional transition between its endpoint outcomes. The bracket is then refined by repeated midpoint evaluation until the requested parameter-space resolution is reached or the refinement becomes numerically or topologically ambiguous.

The procedure does not assume that every cell contains a single smooth boundary. Cells involving invalid samples, additional outcome classes, or incompatible edge structure are retained as complex cells rather than being forced into a binary boundary model.

For a refined bracket with endpoints $\mathbf{p}_L$ and $\mathbf{p}_R$, we associate the local numerical resolution

$$\delta_B = \|\mathbf{p}_R - \mathbf{p}_L\|, \quad (104)$$

with the appropriate parameter scaling used by the implementation. Boundary-derived quantities such as

ρ_η must therefore be interpreted relative to this finite numerical resolution.

4.8.6 Robustness Validation

The boundary-distance estimate ρ_η is tested independently of its construction. Representative nominal points are perturbed in multiple directions in the local metric geometry, both inside and outside the estimated radius. The resulting initial conditions are propagated through the same first-passage classifier used to construct the atlas.

In addition, directional searches are performed for an observed change of finite-time outcome. For each selected point, eight parameter-space directions are examined and candidate outcome changes are localized by radial bisection when a suitable bracket is found.

Let ρ_{obs} denote the smallest observed metric distance to a sampled perturbation producing a different outcome. Comparison of

$$\rho_\eta \quad \text{and} \quad \rho_{\text{obs}} \quad (105)$$

provides a direct empirical test of whether distance to the numerically resolved atlas boundary behaves as a useful local robustness estimate.

Failure to observe a flip within the sampled search domain is recorded as an unresolved validation result rather than interpreted as evidence that no such perturbation exists.

4.8.7 Checkpointing and Deterministic Restart

The final atlas requires repeated long-horizon first-passage propagations and adaptive boundary refinements. To prevent interruption from requiring the scientific calculation to be restarted from the beginning, the execution pipeline uses persistent configuration-specific checkpoints.

Checkpoint identities depend on the complete scientific configuration required to reproduce an evaluation, including the CR3BP mass parameter, family coordinate, phase, manifold branch, perturbation magnitude, section definition, propagation horizon, event settings, and numerical tolerances. Previously completed evaluations are reused only when this configuration agrees.

Boundary-refinement states are checkpointed independently. Their saved records include the current bracket, completed midpoint evaluations, refinement depth, complex-cell status, and final boundary estimate when available. Checkpoint records are committed atomically so that an interrupted write cannot be mistaken for a completed scientific result.

Restarting an interrupted calculation therefore resumes from completed evaluations without changing the underlying numerical experiment.

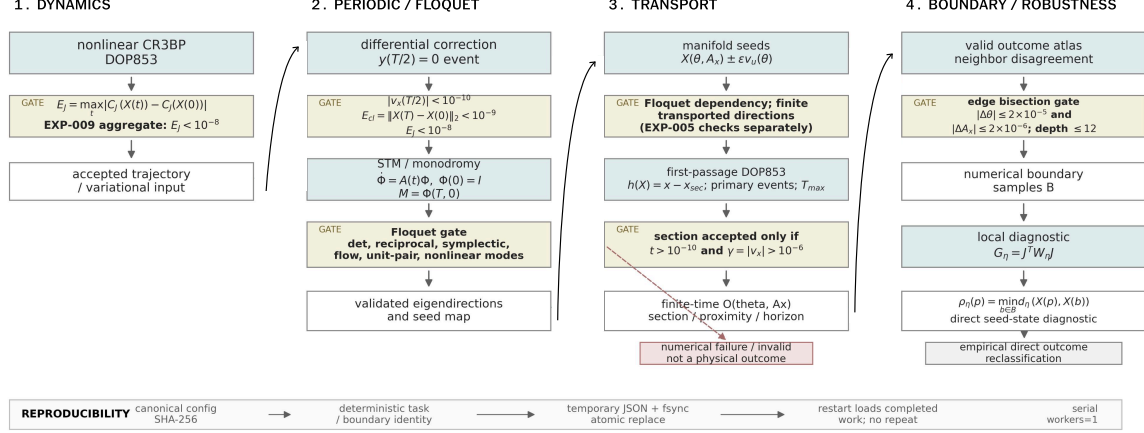


Figure 13: Validation hierarchy of the numerical transport pipeline. CR3BP integrations are checked for Jacobi conservation before validated periodic solutions are subjected to full-period closure and STM/Floquet consistency checks. The resulting eigendirections are used to construct manifold seeds, whose nonlinear trajectories are processed by the finite-time first-passage classifier. Valid outcome samples subsequently enter boundary refinement and pullback-metric analysis. Numerical failures or ambiguous evaluations are retained explicitly rather than being assigned to a physical transport class, preventing invalid computations from propagating through subsequent stages of the analysis.

4.8.8 Reproducibility of Execution

Representative transport evaluations were also compared across serial and process-based execution. For identical scientific inputs, the classification, first-passage time, and Jacobi diagnostics were required to agree with the serial calculation.

Parallel execution did not provide a performance advantage for the present workload because adaptive boundary refinement dominates the total cost. The final calculation was therefore executed with a single worker. This choice changes only the execution strategy and not the scientific definitions, tolerances, propagation horizon, atlas grid, or boundary-refinement criteria.

The complete computational workflow consequently separates three questions that are otherwise easy to conflate:

1. whether a trajectory integration is numerically trustworthy;
2. whether two nearby initial conditions receive different finite-time outcome classifications; and
3. whether a proposed metric quantity reliably characterizes the proximity of such an outcome change.

The first is addressed through integration and dynamical validation, the second through the outcome atlas and boundary-refinement procedure, and the third through the independent robustness tests. This separation is main-

tained in the interpretation of the numerical results that follow.

5 Numerical Validation

The transport calculations rely on several layers of numerical integration, including periodic-orbit correction, state-transition matrix propagation, invariant-manifold construction, and long-horizon first-passage trajectories. Numerical validation was therefore performed at each stage before its output was used in subsequent calculations. The checks reported here address numerical consistency of the computations rather than the physical interpretation of the resulting transport structures.

5.1 Jacobi Conservation

The Jacobi integral provides a direct invariant-based test of the numerical integration of the autonomous CR3BP equations. For a trajectory with initial state $\mathbf{X}_0$, the instantaneous Jacobi error was defined as

$$\Delta C_J(t) = C_J(\mathbf{X}(t)) - C_J(\mathbf{X}_0), \quad (106)$$

and the maximum absolute drift over an integration interval was measured by

$$E_J = \max_{0 \leq t \leq t_f} |\Delta C_J(t)|. \quad (107)$$

This check was applied independently of the geometric or event-based conditions used elsewhere in the pipeline. Consequently, a trajectory could satisfy a section-crossing or periodic-closure condition while still being rejected if its invariant error exceeded the prescribed numerical tolerance.

Jacobi conservation was first verified on dedicated CR3BP integrations and was subsequently monitored during periodic-orbit and transport calculations. Across the 33-member Lyapunov family used in the final parameter sweep, the largest recorded Jacobi error was approximately

$$E_J^{\text{family}} \simeq 4.97 \times 10^{-13}. \quad (108)$$

The longer first-passage integrations naturally accumulated larger numerical errors. In the boundary-scaling calculations preceding the final atlas, the largest observed Jacobi drift remained approximately

$$E_J^{\text{transport}} \simeq 1.02 \times 10^{-9}. \quad (109)$$

The several-order-of-magnitude difference between these values is consistent with the substantially longer and more geometrically varied transport integrations compared with the periodic-orbit calculations. No physical outcome was inferred solely from invariant conservation; rather, the Jacobi error served as an independent numerical consistency check on the propagated trajectories.

5.2 Periodic-Orbit and Continuation Validation

Each planar Lyapunov orbit was accepted only after satisfying the half-period symmetry condition and an independent full-period closure test. If $\mathbf{X}_0$ denotes the corrected initial state and T the computed orbital period, the closure error was defined as

$$E_{\text{cl}} = \|\mathbf{X}(T) - \mathbf{X}_0\|_2. \quad (110)$$

This full-state measure supplements the differential-correction constraint at the half-period crossing. In particular, satisfying $v_x(T/2) \approx 0$ alone was not taken as sufficient evidence of a periodic solution; the corrected state was propagated over the complete period and compared directly with its initial condition.

The family used for the transport calculations contained 33 periodic orbits spanning

$$0.006 \leq A_x \leq 0.014, \quad \Delta A_x = 2.5 \times 10^{-4}. \quad (111)$$

A predictor-corrector continuation procedure was used to remain on the same periodic-orbit branch as A_x increased. The first family member was initialized from

the linear center-mode approximation. The second used the preceding corrected solution as its initial estimate, while subsequent members used a secant prediction based on the two preceding validated solutions. Each predicted state was then subjected to the same nonlinear differential-correction and validation procedure.

All 33 members of the production family passed the periodic-orbit validation checks. The largest full-period closure error over the family was

$$\max E_{\text{cl}} \simeq 1.82 \times 10^{-11}, \quad (112)$$

while the maximum Jacobi error over the same family was approximately 4.97×10^{-13} , consistent with the invariant check reported in Sec. 5.1.

Continuation also produced a smooth variation of the corrected orbit properties across the sampled amplitude interval. The full period increased from approximately

$$T = 2.69998 \quad \text{at} \quad A_x = 0.006 \quad (113)$$

to

$$T = 2.74740 \quad \text{at} \quad A_x = 0.014, \quad (114)$$

without a discontinuous jump between neighboring family members. The corresponding corrected initial transverse velocity likewise varied smoothly, from approximately 0.05263 to 0.13130 over the same interval.

The continuation step was particularly important at the upper end of the amplitude range. Independently initialized corrections could converge toward an unintended branch or encounter the trivial $t = 0$ crossing of the symmetry section. Propagating the validated family by continuation avoided this ambiguity while retaining the original differential-correction equations and independent closure tests. Consequently, the transport atlas was constructed from a single numerically continuous and independently validated Lyapunov family.

5.3 State-Transition Matrix and Floquet Validation

The linearized dynamics associated with each validated periodic orbit were checked independently through the state-transition matrix (STM). The STM was initialized as

$$\Phi(0) = I_4 \quad (115)$$

and propagated together with the nonlinear reference trajectory according to the variational equations. After one complete orbital period, the monodromy matrix was obtained as

$$M = \Phi(T, 0). \quad (116)$$

For the planar CR3BP, the Hamiltonian structure imposes a reciprocal organization on the Floquet spectrum. The nontrivial hyperbolic multipliers were therefore checked for consistency with

$$\lambda_s \lambda_u \approx 1, \quad |\lambda_s| < 1, \quad |\lambda_u| > 1, \quad (117)$$

together with the expected unit-multiplier structure associated with the periodic orbit. These checks were applied only after the corresponding nonlinear orbit had passed the closure and Jacobi-conservation tests.

Every member of the 33-orbit production family retained the expected hyperbolic structure. At the lower end of the family, $A_x = 0.006$, the unstable multiplier was approximately

$$\lambda_u \simeq 2.620 \times 10^3, \quad (118)$$

with the corresponding stable multiplier

$$\lambda_s \simeq 3.82 \times 10^{-4}. \quad (119)$$

At $A_x = 0.014$, these values varied smoothly to approximately

$$\lambda_u \simeq 2.337 \times 10^3, \quad \lambda_s \simeq 4.28 \times 10^{-4}. \quad (120)$$

The evolution of the multipliers was continuous across the sampled family rather than exhibiting isolated jumps between neighboring amplitudes. For example, the representative $A_x = 0.010$ orbit had an unstable multiplier of approximately

$$\lambda_u \simeq 2.512 \times 10^3, \quad (121)$$

consistent with the surrounding family members.

The strong separation between the stable and unstable multipliers also provides a numerical check on the eigendirections used for manifold construction. Stable and unstable eigenvectors extracted from the validated monodromy matrix were transported to the required orbital phases using the STM before local manifold perturbations were applied. Thus, the manifold seeds used in the subsequent transport calculations were derived only from periodic orbits for which both the nonlinear closure conditions and the expected linear stability structure had been verified.

5.4 First-Passage and Event Validation

The transport classification depends on the accurate detection and ordering of terminal events along each nonlinear manifold trajectory. Event detection was therefore validated independently of the discrete outcome label ultimately assigned to a trajectory.

For the vertical transport section

$$h(\mathbf{X}) = x - x_{\text{sec}} = 0, \quad (122)$$

an accepted section event was required to satisfy both the section condition and the prescribed crossing-direction and transversality criteria. The residual of a detected crossing was monitored through

$$E_{\text{sec}} = |x(t_{\text{sec}}) - x_{\text{sec}}|. \quad (123)$$

Because

$$\nabla h = [1 \ 0 \ 0 \ 0]^T, \quad (124)$$

the local transversality diagnostic reduces to

$$\gamma = |\nabla h^T \mathbf{f}(\mathbf{X})| = |v_x|. \quad (125)$$

Candidate section events failing the prescribed transversality or crossing-direction requirement were not accepted as first-passage crossings.

Dedicated boundary calculations provided an additional numerical check on the event detector. Across the analyzed boundary neighborhoods, the largest accepted section residual was approximately

$$E_{\text{sec,max}} \simeq 6.1 \times 10^{-15}, \quad (126)$$

while the smallest recorded transversality magnitude among the accepted crossings was approximately

$$\gamma_{\text{min}} \simeq 4.74 \times 10^{-1}. \quad (127)$$

The accepted events were therefore well resolved relative to the section condition and were not associated with numerically grazing crossings in these validation calculations.

Section crossing was not the only possible terminal event. Propagation could instead terminate upon entry into one of the prescribed primary-proximity regions. If neither an accepted section crossing nor a primary-proximity event occurred, integration continued until the fixed finite-time horizon

$$T_{\text{max}} = 160. \quad (128)$$

The corresponding classification, NO_CROSSING_BY_160, therefore denotes only the absence of an accepted crossing within the specified integration interval; it does not imply that the trajectory can never cross the section at a later time.

When multiple event conditions were possible, the earliest admissible terminal event determined the physical classification. Numerical integration failures, invalid event evaluations, or unresolved cases were treated separately and were not silently mapped

onto SECTION_CROSSING, PRIMARY_PROXIMITY, or NO_CROSSING_BY_160. No integration failures occurred in the dedicated boundary-scaling validation set used to assess these event conditions.

Together, the small section residuals, finite transversality margins, explicit event ordering, and separate treatment of numerical failures support the use of the first-passage classifier as a reproducible finite-time outcome definition.

5.5 Pullback-Metric Validation

The pullback metric relies on a local linearization of the manifold-seed map

$$\mathbf{X} = \mathbf{X}(\mathbf{p}), \quad \mathbf{p} = [\theta \quad A_x]^T. \quad (129)$$

For a sufficiently small parameter perturbation $\delta\mathbf{p}$, the corresponding change in the physical seed state is approximated by

$$\delta\mathbf{X} \approx J(\mathbf{p}) \delta\mathbf{p}, \quad J = \frac{\partial\mathbf{X}}{\partial\mathbf{p}}. \quad (130)$$

The induced local metric was computed as

$$G_\eta = J^T W_\eta J, \quad (131)$$

using the same state weighting and normalization employed in the robustness calculation. Numerical validation addressed both the algebraic properties of G_η and the accuracy of the local linearization from which it is derived.

At the representative validation points, the computed metric tensors were symmetric to numerical precision and had strictly positive eigenvalues. Thus, the tested matrices defined positive-definite local quadratic forms on the two-dimensional parameter family. The matrices were nevertheless anisotropic, with condition numbers ranging from approximately

$$7 \times 10^3 \lesssim \kappa(G_\eta) \lesssim 7 \times 10^4 \quad (132)$$

over the representative metric-validation samples. This anisotropy was retained rather than regularized away, since it arises from the local mapping between parameter perturbations and weighted physical state-space perturbations.

The local approximation was tested directly by comparing its predicted weighted displacement with the displacement obtained by reevaluating the nonlinear seed map. For a perturbation $\delta\mathbf{p}$, the metric prediction is

$$d_{\text{pred}} = \sqrt{\delta\mathbf{p}^T G_\eta(\mathbf{p}) \delta\mathbf{p}}, \quad (133)$$

while the corresponding nonlinear displacement was evaluated from the two seed states using the same weighted state-space norm,

$$d_{\text{true}} = \|\mathbf{X}(\mathbf{p} + \delta\mathbf{p}) - \mathbf{X}(\mathbf{p})\|_{W_\eta}. \quad (134)$$

Their relative discrepancy was measured as

$$E_{\text{metric}} = \frac{|d_{\text{pred}} - d_{\text{true}}|}{d_{\text{true}}}. \quad (135)$$

At the largest perturbation scale used in the local validation tests, representative relative errors were of order

$$E_{\text{metric}} \sim 10^{-5}, \quad (136)$$

with observed values approximately between 1.3×10^{-5} and 2.6×10^{-5} for the sampled cases. Successive reductions of the perturbation magnitude produced corresponding reductions in the discrepancy, with the error decreasing approximately linearly as the perturbation scale was halved. This behavior is consistent with the expected convergence of the first-order local approximation in Eq. (130).

These tests validate the numerical construction of the local pullback metric and its interpretation as the differential geometry induced by the prescribed manifold-seed family. They do not establish that distance under this metric is, by itself, a predictor of a finite-time transport outcome change; that question is examined separately using the computed outcome boundaries.

5.6 Boundary Resolution and Computational Reproducibility

In the final atlas calculation, 484 coarse outcome-changing edges were identified and subjected to adaptive refinement. The resulting set contained 484 numerical boundary samples. For the resolved brackets, the minimum and median recorded resolution were both approximately

$$\delta_B \simeq 1.95 \times 10^{-6}, \quad (137)$$

whereas the largest recorded value was approximately

$$\delta_{B,\text{max}} \simeq 9.82 \times 10^{-3}. \quad (138)$$

The large spread in resolution indicates that the boundary was not uniformly localized throughout the parameter domain. Two cells were additionally identified as having complex local outcome geometry and were retained explicitly rather than being reduced to a single simple binary transition. Subsequent distance-to-boundary calculations therefore inherit the finite and locally nonuniform resolution of this numerical boundary representation.

Because boundary refinement requires additional nonlinear propagations beyond the base parameter

grid, the production calculation was also designed to be restartable. Completed base-atlas evaluations and boundary-refinement states were stored using configuration-specific, deterministic identifiers. During refinement, the current bracket endpoints, midpoint evaluations, refinement depth, and completion state were persisted so that an interrupted calculation could resume from the last completed operation rather than restart the edge refinement.

The restart mechanism was tested by deliberately interrupting a refinement calculation after a completed midpoint evaluation and then resuming from the stored state. The completed midpoint was not recomputed, and no duplicate completed propagations were observed after restart. Checkpoint updates were written through a temporary file, synchronized to persistent storage, and atomically replaced, reducing the possibility that an interruption could leave a partially written checkpoint interpreted as valid.

The final production atlas was executed serially. This choice followed a representative process-parallel benchmark in which additional worker processes did not reduce execution time for the propagation workload. Serial execution also provided a simple deterministic ordering for checkpointed evaluations. The completed base atlas contained

$$160 \times 33 = 5280 \quad (139)$$

parameter samples, all generated from the same validated periodic-orbit family and finite-time classification pipeline.

These checks do not remove the finite-resolution limitations of the computed boundary. Rather, they establish that those limitations are recorded explicitly and that interruption, restart, and execution order do not silently alter already completed trajectory classifications. Boundary-resolution effects are therefore retained as part of the interpretation of the subsequent robustness analysis.

6 Transport Robustness Framework

The finite-time outcome atlas partitions the prescribed family of manifold-seeded initial conditions according to their subsequent transport behavior. This partition permits a geometric notion of robustness based on proximity to a change in finite-time outcome. The construction developed here is deliberately restricted to the two-dimensional family used to generate the transport atlas; it is not intended to define robustness with respect to arbitrary perturbations of the full CR3BP state.

6.1 Family-Constrained Perturbation Space

Let

$$\mathbf{p} = \begin{bmatrix} \theta \\ A_x \end{bmatrix} \in \mathcal{D} \quad (140)$$

denote a point in the parameter domain

$$\mathcal{D} = [0, 2\pi) \times [0.006, 0.014]. \quad (141)$$

The phase coordinate θ selects a location on a planar Lyapunov orbit, while A_x selects the member of the continued periodic-orbit family. Together with the prescribed manifold branch and perturbation convention, these coordinates determine a physical manifold-seeded initial state through the map

$$\mathbf{X} : \mathcal{D} \rightarrow \mathbb{R}^4, \quad \mathbf{p} \mapsto \mathbf{X}(\mathbf{p}). \quad (142)$$

A perturbation considered in the present robustness analysis therefore takes the form

$$\mathbf{p} \rightarrow \mathbf{p} + \delta\mathbf{p}, \quad (143)$$

which induces the physical state perturbation

$$\delta\mathbf{X} = \mathbf{X}(\mathbf{p} + \delta\mathbf{p}) - \mathbf{X}(\mathbf{p}). \quad (144)$$

The allowed state perturbations are consequently constrained by the geometry of the seed map $\mathbf{X}(\mathbf{p})$. They form a two-dimensional family embedded in the four-dimensional planar CR3BP state space rather than an unrestricted neighborhood of $\mathbf{X}(\mathbf{p})$ in $\mathbb{R}^4$.

For sufficiently small $\delta\mathbf{p}$, this map is locally approximated by

$$\delta\mathbf{X} = J(\mathbf{p}) \delta\mathbf{p} + \mathcal{O}(\|\delta\mathbf{p}\|^2), \quad (145)$$

where

$$J(\mathbf{p}) = \frac{\partial\mathbf{X}}{\partial\mathbf{p}} \in \mathbb{R}^{4 \times 2}. \quad (146)$$

This local relation is the basis for the pullback metric introduced in Sec. 6.3. It also makes clear why raw Euclidean distance in (θ, A_x) is not adopted as the robustness measure: the two coordinates have different meanings and scales, and equal coordinate displacements need not produce equal perturbations of the physical initial state.

The robustness quantities defined below are therefore *family-constrained*. A large distance from an outcome boundary within $\mathcal{D}$ does not exclude the existence of a much smaller outcome-changing perturbation in a direction transverse to the two-dimensional seed family. No claim of robustness to arbitrary position and velocity errors is implied.

6.2 Finite-Time Outcome Boundary

The first-passage classifier defines a discrete outcome map

$$\mathcal{O} : \mathcal{D} \rightarrow \mathcal{C}, \quad (147)$$

where

$$\mathcal{C} = \{\text{SECTION_CROSSING}, \text{PRIMARY_PROXIMITY}, \text{NO_CROSSING_BY_160}\}. \quad (148)$$

Thus, each parameter point $\mathbf{p} \in \mathcal{D}$ is associated with one finite-time transport outcome under the propagation and event rules defined in Sec. 4.5.

For each outcome $c \in \mathcal{C}$, define the corresponding subset of the parameter domain as

$$\mathcal{D}_c = \{\mathbf{p} \in \mathcal{D} : \mathcal{O}(\mathbf{p}) = c\}. \quad (149)$$

A finite-time outcome boundary occurs where arbitrarily small parameter-space neighborhoods contain points assigned to different outcome classes. The conceptual boundary set may therefore be written as

$$\mathcal{B} = \bigcup_{c \in \mathcal{C}} \partial \mathcal{D}_c \cap \text{int}(\mathcal{D}), \quad (150)$$

where $\partial \mathcal{D}_c$ denotes the boundary of the outcome region associated with class c . Restricting the definition to the interior of $\mathcal{D}$ prevents the outer edge of the sampled parameter domain from being interpreted as an outcome-change boundary.

The set $\mathcal{B}$ is not available analytically. The numerical procedure instead constructs a finite approximation from neighboring samples whose classifications differ. Adaptive edge refinement then produces a collection of resolved boundary samples,

$$\hat{\mathcal{B}} = \{\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_{N_B}\}, \quad (151)$$

with each $\mathbf{b}_k$ accompanied by a finite refinement resolution. The distinction between $\mathcal{B}$ and $\hat{\mathcal{B}}$ is retained throughout the analysis: $\mathcal{B}$ denotes the conceptual finite-time outcome boundary, whereas $\hat{\mathcal{B}}$ is the numerically sampled set actually available for distance calculations.

The boundary is determined by the finite-time classifier and therefore depends on its defining choices, including the transport section, primary-proximity termination conditions, and integration horizon. In particular, the boundary separating NO_CROSSING_BY_160 from another class is a boundary of finite-time behavior rather than a statement about asymptotic escape or confinement.

Figure 10 illustrates the numerical construction of $\hat{\mathcal{B}}$. Once this boundary representation has been obtained, the local pullback geometry of Fig. 11 provides a physically induced measure of distance from a nominal family member to the sampled outcome boundary.

6.3 Metric-Weighted Distance to Outcome Change

Let $\mathbf{p} \in \mathcal{D}$ denote a nominal member of the manifold-seeded initial-condition family, and let $\mathbf{b} \in \hat{\mathcal{B}}$ denote a numerically resolved outcome-boundary sample. The coordinate displacement between the two points is

$$\Delta \mathbf{p} = \mathbf{b} - \mathbf{p}. \quad (152)$$

Rather than measuring this displacement using the Euclidean norm in (θ, A_x) , the local pullback metric at the nominal point is used. For a weighting choice η , the metric-weighted distance from $\mathbf{p}$ to $\mathbf{b}$ is defined as

$$d_\eta(\mathbf{p}, \mathbf{b}) = \sqrt{(\mathbf{b} - \mathbf{p})^T G_\eta(\mathbf{p})(\mathbf{b} - \mathbf{p})}. \quad (153)$$

The metric is evaluated at the nominal point $\mathbf{p}$; it is not assumed to remain constant along the finite displacement to $\mathbf{b}$. Equation (153) should therefore be interpreted as a local, nominal-point approximation to the weighted physical displacement associated with movement through the two-dimensional seed family.

For each nominal point, define the nearest sampled boundary point under this metric as

$$\mathbf{b}_\eta^*(\mathbf{p}) = \arg \min_{\mathbf{b} \in \hat{\mathcal{B}}} d_\eta(\mathbf{p}, \mathbf{b}). \quad (154)$$

The corresponding family-constrained boundary distance is

$$\rho_\eta(\mathbf{p}) = \min_{\mathbf{b} \in \hat{\mathcal{B}}} d_\eta(\mathbf{p}, \mathbf{b}) = d_\eta(\mathbf{p}, \mathbf{b}_\eta^*(\mathbf{p})). \quad (155)$$

Geometrically, Eq. (155) can be visualized by expanding local equal-distance contours centered on $\mathbf{p}$ until one first reaches a sample in $\hat{\mathcal{B}}$. Because the distance is induced by G_η , these contours are generally elliptical in the (θ, A_x) coordinates rather than circular, as illustrated in Fig. 12.

The use of $\hat{\mathcal{B}}$ rather than the unknown continuous boundary $\mathcal{B}$ is significant. The computed ρ_η depends on both the local metric approximation and the finite spatial resolution of the numerical boundary set. It is therefore more precisely a distance to the *resolved sampled boundary* than an exact distance to the underlying outcome boundary. Locally complex or insufficiently resolved boundary geometry can consequently limit the accuracy with which ρ_η approximates the corresponding continuous-boundary distance.

Furthermore, Eq. (155) does not by itself establish that every perturbation satisfying

$$\sqrt{\delta \mathbf{p}^T G_\eta(\mathbf{p}) \delta \mathbf{p}} < \rho_\eta(\mathbf{p}) \quad (156)$$

must preserve the nominal finite-time outcome. Such an interpretation would require the sampled boundary to capture the relevant local outcome geometry sufficiently well and the local metric approximation to remain adequate over the displacement being considered. Accordingly, ρ_η is treated here as a *candidate family-constrained boundary-distance diagnostic*, and its relationship to actual outcome-changing perturbations is tested empirically rather than assumed from the definition.

6.4 Metric Weighting and Sensitivity

The pullback metric G_η depends on the state-space weighting W_η introduced in Sec. 6.3. Consequently, the numerical value of ρ_η is not expected to be invariant under changes in η . This dependence is part of the definition of the diagnostic: different choices of state weighting assign different relative costs to the physical components of an induced seed-state perturbation.

To determine whether the inferred robustness structure is stable under reasonable changes in this weighting, the analysis was repeated for

$$\eta \in \{0.1, 0.3, 1, 3, 10\}. \quad (157)$$

For each value of η , the same outcome atlas and resolved boundary set $\hat{\mathcal{B}}$ were retained. Only the metric used to measure distance from a nominal parameter point to the boundary was changed. This isolates metric sensitivity from changes in the underlying trajectory classifications or boundary-refinement procedure.

Absolute values of ρ_η cannot be compared across weighting choices without accounting for the corresponding change in metric scale. The analysis therefore also examines whether different choices of η preserve the *relative ordering* of nominal points by their boundary distance. For two weighting choices η_i and η_j , this agreement is quantified using the Spearman rank correlation

$$r_s(\eta_i, \eta_j) = \text{corr}_{\text{rank}} [\rho_{\eta_i}(\mathbf{p}), \rho_{\eta_j}(\mathbf{p})], \quad (158)$$

evaluated over the common set of parameter points for which the boundary-distance diagnostic is defined.

Values of r_s near unity indicate that two metric choices produce similar rankings of the sampled family, even if the numerical magnitudes of ρ_η differ. Lower correlations indicate that the weighting choice changes which regions of the family are identified as relatively close to or far from the resolved outcome boundary.

This comparison provides a direct test of whether the spatial robustness pattern is largely preserved across the prescribed metric family or whether conclusions about relative robustness depend substantially on the choice of η . No metric-independence is assumed a priori.

6.5 Empirical Test of the Boundary-Distance Interpretation

The definition of ρ_η establishes a geometric distance to the resolved outcome boundary, but it does not guarantee that ρ_η coincides with the smallest perturbation capable of changing the observed finite-time transport outcome. This stronger interpretation was therefore examined directly at a set of representative nominal points.

For a nominal point $\mathbf{p}_0$, define the local metric magnitude of a parameter perturbation $\delta\mathbf{p}$ as

$$r_\eta(\delta\mathbf{p}; \mathbf{p}_0) = \sqrt{\delta\mathbf{p}^T G_\eta(\mathbf{p}_0) \delta\mathbf{p}}. \quad (159)$$

The first test examines perturbations whose metric magnitude is a prescribed fraction of the computed boundary distance,

$$r_\eta = \alpha \rho_\eta(\mathbf{p}_0), \quad 0 < \alpha < 1. \quad (160)$$

Each perturbed parameter point is mapped through the nonlinear seed construction and propagated using the same first-passage classifier as the nominal point. Its outcome is then compared directly with $\mathcal{O}(\mathbf{p}_0)$. If ρ_η behaved as a strict local outcome-preserving radius, every admissible perturbation satisfying Eq. (160) would retain the nominal outcome. Observed counterexamples therefore provide direct evidence against such a guaranteed-radius interpretation.

A complementary directional search was used to look for the smallest observed outcome-changing perturbation. For a collection of search directions $\mathbf{u}_k$ in parameter space, perturbations were constructed along rays from the nominal point. The directions were normalized with respect to the local metric so that

$$\mathbf{u}_k^T G_\eta(\mathbf{p}_0) \mathbf{u}_k = 1. \quad (161)$$

A perturbation along direction k can then be written as

$$\mathbf{p}(r, k) = \mathbf{p}_0 + r \mathbf{u}_k, \quad (162)$$

where r is the local metric displacement predicted at $\mathbf{p}_0$. Whenever an outcome change was bracketed along a search direction, radial refinement was used to localize the transition.

The minimum outcome-changing displacement observed by this finite search is denoted

$$\rho_\eta^{\text{obs}}(\mathbf{p}_0) = \min_k \{r : \mathcal{O}(\mathbf{p}_0 + r \mathbf{u}_k) \neq \mathcal{O}(\mathbf{p}_0)\}, \quad (163)$$

when at least one such transition is detected. If no outcome change is found within the prescribed directional search, no value of ρ_η^{obs} is assigned. Such a case indicates

only that the finite search did not observe a flip; it does not establish that no closer outcome-changing perturbation exists.

The comparison between ρ_η , inside- ρ_η perturbations, and ρ_η^{obs} separates three distinct questions: whether the resolved boundary is geometrically close under the chosen metric, whether perturbations smaller than that computed distance empirically preserve outcome, and whether a finite directional search can independently recover a nearby outcome transition. These tests are used to determine how strongly ρ_η can be interpreted as a transport-robustness quantity rather than assuming that interpretation from its geometric definition.

Accordingly, the strongest interpretation adopted in this work is conditional on the empirical tests. In the absence of consistent outcome preservation inside ρ_η and independently observed outcome changes near its predicted scale, ρ_η is retained as a family-constrained, finite-time boundary-distance diagnostic rather than a certified safe radius.

7 Results

7.1 Finite-Time Transport Outcome Atlas

The production outcome atlas was evaluated over the full parameter domain using 160 uniformly spaced phase samples and 33 members of the continued Lyapunov family, giving

$$N_{\text{atlas}} = 160 \times 33 = 5280 \quad (164)$$

classified initial conditions. Each point was propagated under the first-passage rules defined in Sec. 4.5, with a maximum integration time of $T_{\text{max}} = 160$.

All three finite-time outcome classes occurred in the resulting atlas. Of the 5280 initial conditions,

$$\begin{aligned} N_{\text{SECTION_CROSSING}} &= 4854, \\ N_{\text{PRIMARY_PROXIMITY}} &= 408, \\ N_{\text{NO_CROSSING_BY_160}} &= 18. \end{aligned} \quad (165)$$

These correspond to approximately 91.9%, 7.7%, and 0.34% of the sampled parameter domain, respectively. Section crossing therefore dominates the sampled family, but the minority outcomes form localized regions within the (θ, A_x) domain rather than disappearing at the production resolution.

Figure 14 shows that the outcome partition is strongly nonuniform in parameter space. The dominant SECTION_CROSSING class occupies most of the atlas, while PRIMARY_PROXIMITY appears in comparatively narrow phase-amplitude regions. The small NO_CROSSING_BY_160 population is confined to still more limited portions of the sampled domain.

The interfaces between these regions are the structures relevant to the robustness analysis. Their geometry is not represented by the class counts alone: even a minority outcome can generate an extended boundary when embedded as a narrow region within the dominant class. For this reason, the subsequent analysis uses the adaptively refined outcome-change interfaces rather than treating outcome frequency as a measure of robustness.

The atlas also demonstrates why the outcome labels are retained as finite-time classifications. In particular, NO_CROSSING_BY_160 states are known only not to have satisfied the section-crossing criterion within the prescribed integration horizon. No asymptotic confinement or eventual transport outcome is inferred from that label.

7.2 Resolved Outcome-Boundary Structure

The discrete atlas in Fig. 14 was searched for neighboring parameter samples assigned to different finite-time outcomes. This procedure identified 484 coarse outcome-changing edges across the sampled (θ, A_x) domain. Adaptive refinement of these edges produced 484 numerical boundary samples forming the resolved set $\hat{\mathcal{B}}$ used in the subsequent distance calculations.

Figure 15 shows that the resolved boundary follows several qualitatively different structures. Extended interfaces occur around the larger primary-proximity regions, while narrow branches and isolated minority-outcome features generate additional localized boundary segments. The resulting set is therefore not well represented by a single smooth curve through the parameter domain.

Boundary localization was also nonuniform. The minimum and median recorded refinement resolution were approximately 1.95×10^{-6} , whereas the largest recorded value was approximately 9.82×10^{-3} . Thus, most refined interfaces were localized substantially more tightly than the least-resolved portions of the boundary. The latter remain an explicit source of uncertainty for distances computed to $\hat{\mathcal{B}}$.

Two cells were identified as containing complex local outcome geometry. These cells were retained rather than forcing their classifications into a single binary interface. Their presence is consistent with the irregular structure visible in the underlying atlas and provides a specific example of why the numerical boundary should not be treated as an exact smooth separatrix.

The boundary set used below is consequently both finite and resolution-limited. Two refinement cases were flagged as having complex local outcome geometry and were retained explicitly rather than being forced into a single simple binary boundary representation.




Figure 14: Finite-time transport outcome atlas over the manifold-seed family parameterized by Lyapunov-orbit phase θ and amplitude A_x . Each of the 5280 sampled initial conditions is assigned according to the earliest admissible terminal event: section crossing, primary proximity, or no section crossing before $T_{\max} = 160$. The plotted classes represent finite-time outcomes under the prescribed event definitions and should not be interpreted as asymptotic basins.

7.3 Family-Constrained Boundary Distance

The resolved boundary set $\hat{\mathcal{B}}$ was combined with the local pullback metric to evaluate the family-constrained boundary distance ρ_η defined in Eq. (155). Figure 16 shows the resulting field for $\eta = 1$, used here as a representative intermediate weighting.

The computed field is strongly nonuniform across the (θ, A_x) domain. Small values of ρ_1 occur near the resolved interfaces identified in Fig. 15, whereas parameter regions farther from those interfaces generally exhibit larger values. The result therefore associates the discrete outcome partition with a continuous-valued geometric diagnostic evaluated on the sampled family, while retaining the local state-space scaling induced by the seed map.

The magnitude of this variation is illustrated by representative samples from the production calculation. At $(\theta, A_x) \approx (2.356, 0.0135)$, the computed value was

$$\rho_1 \approx 1.03 \times 10^{-3}, \quad (166)$$

whereas at $(\theta, A_x) \approx (1.335, 0.00825)$,

$$\rho_1 \approx 3.25 \times 10^{-2}. \quad (167)$$

The latter value is more than an order of magnitude larger under the same metric definition. Intermediate representative points likewise span a broad range of boundary distances, showing that the diagnostic does not reduce to a nearly uniform scaling across the parameter domain.

These values should not be interpreted as Euclidean distances in (θ, A_x) . The local tensor $G_1(\mathbf{p})$ maps pa-

rameter displacements into the weighted geometry of the physical seed-state family. Consequently, two nominal points with similar coordinate separations from $\hat{\mathcal{B}}$ can receive different values of ρ_1 when their local seed-map geometries differ.

The field also inherits the limitations of the resolved boundary. Distances near tightly refined boundary segments are supported by a more localized numerical representation than distances whose nearest sample lies in a poorly resolved or locally complex region. Thus, Fig. 16 should be read together with the boundary-resolution information in Fig. 15.

At this stage, the calculation establishes that ρ_η provides a spatially varying measure of proximity to the resolved finite-time outcome boundary within the prescribed seed family. Whether this geometric distance remains stable under changes in the state weighting, and whether its magnitude corresponds to observed outcome-changing perturbations, are separate empirical questions considered in the following sections.

7.4 Sensitivity to Metric Weighting

The boundary-distance calculation was repeated for $\eta \in \{0.1, 0.3, 1, 3, 10\}$ while retaining the same outcome atlas and resolved boundary set. Changing η therefore modifies the local geometry used to measure displacement without changing the underlying finite-time classifications.

The resulting rankings of ρ_η were compared using the Spearman correlation defined in Eq. (158). The complete pairwise comparison is shown in Fig. 17.




Figure 15: Adaptively resolved finite-time outcome boundary $\hat{\mathcal{B}}$ superimposed on the production outcome atlas. Boundary samples are obtained by refining neighboring parameter samples with different finite-time classifications. The resolved boundary is nonuniform in both geometry and numerical resolution; it represents the sampled approximation used in the family-constrained boundary-distance calculation, not an analytic separatrix.

The strongest agreement occurred between the two smallest neighboring weightings,

$$r_s(0.1, 0.3) \approx 0.9988, \quad (168)$$

indicating that these choices produce almost identical rankings of nominal points by boundary distance. Strong correlations were also obtained for several other neighboring choices, including

$$\begin{aligned} r_s(0.3, 1) &\approx 0.9795, \\ r_s(1, 3) &\approx 0.9481, \\ r_s(3, 10) &\approx 0.9521. \end{aligned} \quad (169)$$

The agreement decreases as the weighting choices become more widely separated. In particular,

$$r_s(0.1, 10) \approx 0.7944, \quad (170)$$

was the lowest pairwise correlation among the tested metrics. Other widely separated comparisons showed the same trend, with $r_s(0.1, 3) \approx 0.8877$ and $r_s(0.3, 10) \approx 0.8005$.

Thus, the broad ordering of the parameter family is substantially preserved across the tested metrics, but it is not invariant to the state weighting. Nearby choices of η produce highly consistent rankings, while sufficiently

different weightings can appreciably reorder which nominal points lie relatively close to or far from the resolved outcome boundary.

The dependence is also visible in the absolute boundary distances. For example, at $(\theta, A_x) \approx (2.356, 0.0135)$, the computed values decrease from

$$\rho_{0.1} \approx 9.10 \times 10^{-3} \quad (171)$$

to

$$\rho_{10} \approx 5.02 \times 10^{-4}. \quad (172)$$

This change in magnitude is expected because altering η changes the metric itself; absolute distances obtained under different metrics do not share a common scale. The rank comparison is therefore more informative for assessing whether the spatial pattern identified as relatively near or far from the boundary is stable under changes in weighting.

These results support the flag `METRIC_SENSITIVE`: the family-constrained boundary-distance structure is robust to small changes in the tested weighting but cannot be regarded as metric-independent.




Figure 16: Family-constrained boundary-distance diagnostic $\rho_\eta(\theta, A_x)$ for $\eta = 1$. Each value is the minimum local pullback-metric distance from a nominal manifold-seed family member to the numerically resolved finite-time outcome boundary $\hat{\mathcal{B}}$. Smaller values indicate proximity to a resolved outcome-change interface under the prescribed metric. The quantity is a local, family-constrained boundary-distance diagnostic and is not interpreted here as a guaranteed outcome-preserving radius.

7.5 Empirical Test of the Boundary-Distance Interpretation

The geometric construction of ρ_η does not by itself imply that perturbations smaller than the computed boundary distance preserve the nominal finite-time outcome. The empirical tests defined in Sec. 6.5 were therefore applied to representative points spanning different locations and boundary-distance scales within the parameter family.

For the first representative point, the inside-radius test evaluated 26 perturbations with local metric magnitudes below the computed ρ_η . Of these,

$$25/26 \quad (173)$$

retained the nominal finite-time outcome. One perturbation, however, changed classification to PRIMARY_PROXIMITY despite having a prescribed local metric magnitude of only

$$r_\eta = 0.5 \rho_\eta. \quad (174)$$

This single counterexample is sufficient to reject a strict outcome-preserving interpretation of the computed ρ_η at that nominal point. In particular, the numerical boundary distance cannot be treated as a certified radius within which every family-constrained perturbation is guaranteed to preserve the finite-time classification.

The remaining representative points did not provide corresponding inside-radius test records in the production

validation output. Accordingly, no preservation rate is inferred for those points, and the available inside-radius evidence is treated as incomplete rather than extrapolated across the parameter family.

A separate directional search was performed at all five representative points to look for independently observed outcome changes. Each search used eight parameter-space directions together with radial refinement when a transition could be bracketed. No outcome-changing perturbation was detected within the prescribed search for any of the five nominal points. Consequently,

$$\rho_\eta^{\text{obs}} \quad (175)$$

could not be assigned at any of these locations under the finite search design. These cases were recorded as NO_OBSERVED_FLIP rather than interpreted as evidence that no nearby outcome-changing perturbation exists.

The two tests therefore do not support interpreting ρ_η as a sharp empirical outcome-change threshold. The inside-radius counterexample shows that the sampled-boundary distance need not define an outcome-preserving neighborhood, while the absence of detected flips in the finite directional searches prevents an independent numerical identification of a transition distance near the predicted ρ_η scale.

These findings do not invalidate the pullback metric itself. The local metric construction was independently validated in Sec. 5.5, where nonlinear seed-state displacements agreed with the local metric prediction




Figure 17: Spearman rank correlation between family-constrained boundary distances computed using different state-weighting parameters η . Nearby weighting choices produce nearly identical rankings, whereas larger changes in η produce measurable reordering of the sampled parameter family.

under shrinking parameter perturbations. Instead, the discrepancy concerns the stronger step from a distance to the numerically resolved boundary $\hat{\mathcal{B}}$ to a certified or empirically sharp outcome-change radius.

The production evidence therefore supports the more limited interpretation adopted throughout this work: ρ_η is a family-constrained, finite-time distance to the resolved outcome boundary under a prescribed local metric. Its value can be used to compare geometric proximity to the resolved transport interfaces, but the present calculations do not establish it as a guaranteed safe radius or as the minimum perturbation required to change transport outcome.

7.6 Summary of Results

The production calculations resolve a structured finite-time transport partition over the two-dimensional manifold-seed family. Of the 5280 sampled initial conditions, 4854 produced an admissible section crossing, 408 terminated through primary proximity, and 18 did not cross the prescribed section before $T_{\max} = 160$. Adaptive refinement of outcome-changing neighboring samples produced 484 numerical boundary samples, together with two refinement cases flagged for complex local outcome geometry.

The pullback-metric construction converts proximity to this resolved boundary into the family-constrained diagnostic ρ_η . The resulting boundary-distance field varies strongly across the sampled parameter domain and re-

flects both the location of the finite-time outcome interfaces and the local geometry of the manifold-seed map. The underlying metric construction passed the local numerical validation tests, but the resulting distances remain dependent on the chosen state weighting.

This dependence is modest for nearby values of η but becomes appreciable across the full tested range. Spearman correlations between the ρ_η rankings ranged from approximately 0.794 to 0.999. Thus, the broad spatial ordering is substantially preserved under small changes in weighting, but no metric-independent ranking of transport robustness is obtained.

The direct empirical tests place a stronger restriction on the interpretation of ρ_η . At the representative point for which inside-radius perturbations were available, 25 of 26 perturbations preserved the nominal finite-time outcome, but one outcome change occurred at approximately $0.5\rho_\eta$. In addition, the finite directional searches at all five representative points produced no observed outcome flip from which an independent transition distance could be assigned.

Taken together, these results support ρ_η as a family-constrained geometric measure of distance to the numerically resolved finite-time outcome boundary. They do not support its stronger interpretation as a certified outcome-preserving radius or as an empirically established minimum perturbation required to change transport outcome. The implications of this distinction, including the roles of boundary resolution, local metric approximation, and the restricted perturbation family, are considered in Sec. 8.

8 Discussion

8.1 What the Boundary-Distance Diagnostic Measures

The calculations separate two properties that can appear similar when viewed only through an outcome atlas: geometric proximity to a resolved outcome boundary and robustness to an actual outcome-changing perturbation. The former is directly quantified by ρ_η ; the latter is not established by the present results.

Within the prescribed manifold-seed family, the finite-time outcome map contains structured interfaces between section-crossing, primary-proximity, and no-crossing classifications. The pullback metric provides a physically induced local geometry on this family, and ρ_η measures the distance from a nominal point to the sampled representation of those interfaces. In this sense, the construction successfully replaces raw coordinate distance in (θ, A_x) with a distance that accounts for how parameter perturbations map into changes of the physical initial state.

The numerical validation supports this geometric part




Figure 18: Empirical tests of the family-constrained boundary-distance interpretation at the representative validation points. Inside- ρ_η perturbations are compared with the nominal finite-time outcome, while directional searches test for independently observed outcome changes. One tested perturbation at $0.5\rho_\eta$ changed outcome despite lying inside the computed boundary distance. No directional search produced a bracketed outcome flip from which ρ_η^{obs} could be assigned.

of the construction. The pullback tensors were positive definite at the tested points, and the local metric prediction converged toward the nonlinear seed-state displacement as the parameter perturbation was reduced. Accordingly, the failure of the stronger robustness interpretation cannot be attributed simply to an internally inconsistent local metric.

The distinction appears when ρ_η is compared with actual outcome changes. The inside-radius counterexample in Fig. 18 demonstrates that proximity to the sampled boundary does not define a guaranteed outcome-preserving neighborhood. A perturbation with local metric magnitude approximately

$$r_\eta = 0.5\rho_\eta \quad (176)$$

changed the finite-time classification even though its magnitude was smaller than the computed distance to $\hat{\mathcal{B}}$. The implication

$$r_\eta < \rho_\eta \implies \mathcal{O}(\mathbf{p} + \delta\mathbf{p}) = \mathcal{O}(\mathbf{p}) \quad (177)$$

therefore does not hold generally for the numerical construction tested here.

The directional searches provide a complementary limitation. No outcome-changing perturbation was observed at any of the five representative points under the

prescribed finite search. Hence the experiments did not recover an independent value ρ_η^{obs} that could be compared directly with the boundary-distance prediction. This absence should not be interpreted as evidence for unlimited local robustness; it indicates only that the finite directional search did not locate a transition.

The combined evidence therefore supports a narrower interpretation. ρ_η is a local-metric distance to a finite numerical representation of an outcome boundary within a specified two-dimensional family. It identifies points that are geometrically near or far from the transport interfaces resolved by the computation, but the present experiments do not establish that it equals the minimum physical perturbation required to reach a different outcome.

This distinction is important for condition-like interpretations of the quantity. A true distance-to-outcome-change measure would require the nearest relevant outcome-changing state to be represented accurately by the admissible perturbation set and boundary construction. Here, the metric, numerical boundary, finite-time classifier, and restricted seed family each enter the definition. The resulting quantity is therefore conditional on those choices rather than an intrinsic scalar property of the underlying CR3BP trajectory.

8.2 Sources of Discrepancy Between Boundary Distance and Outcome Change

The counterexample in Fig. 18 shows that the computed distance to $\hat{\mathcal{B}}$ can exceed the magnitude of an observed outcome-changing perturbation. Several features of the construction can produce such a discrepancy. The present experiment does not isolate their individual contributions, but it identifies where the geometric boundary-distance interpretation differs from a certified minimum distance to outcome change.

First, $\hat{\mathcal{B}}$ is a finite numerical representation of the underlying outcome boundary $\mathcal{B}$. Boundary candidates are discovered from outcome changes between neighboring atlas samples. Consequently, an outcome-changing structure that does not produce a detectable class transition across the sampled grid can be absent from $\hat{\mathcal{B}}$. This issue is particularly relevant for narrow features, disconnected components, or structures occurring below the sampling scale. The complex refinement cases in Fig. 15 provide direct evidence that the outcome geometry is not uniformly reducible to isolated smooth binary interfaces.

Second, the numerical boundary has nonuniform finite resolution. Most refinement brackets were localized tightly, but the largest recorded bracket span was orders of magnitude larger than the minimum and median values. A nearest sampled boundary point is therefore not necessarily an equally accurate approximation to $\mathcal{B}$ everywhere in the parameter domain. Since ρ_η minimizes distance over $\hat{\mathcal{B}}$, errors in the location or completeness of this set propagate directly into the resulting diagnostic.

Third, the distance in Eq. (153) uses the metric tensor evaluated at the nominal point,

$$d_\eta(\mathbf{p}, \mathbf{b}) = \sqrt{(\mathbf{b} - \mathbf{p})^T G_\eta(\mathbf{p})(\mathbf{b} - \mathbf{p})}. \quad (178)$$

This is a local approximation. The validation in Sec. 5.5 shows that it accurately reproduces sufficiently small seed-state displacements, but it does not imply that G_η remains constant over a finite displacement extending from $\mathbf{p}$ to a boundary sample. If the seed-map geometry varies appreciably along that displacement, the nominal-point quadratic form need not equal the distance obtained by accounting for that variation along the path.

A more geometric alternative would treat G_η as a position-dependent metric and define the length of a path $\gamma(s)$ in parameter space by

$$L_\eta[\gamma] = \int_0^1 \sqrt{\dot{\gamma}(s)^T G_\eta(\gamma(s)) \dot{\gamma}(s)} ds, \quad (179)$$

with a corresponding geodesic distance obtained by minimizing L_η over admissible paths. Such a construction

was not evaluated in the present work, and it is therefore not known whether replacing the nominal-point approximation with a path-dependent distance would improve agreement with observed outcome changes.

Fourth, the robustness construction is restricted to the chosen two-dimensional seed family. This restriction is intentional, but it means that ρ_η cannot characterize the nearest outcome-changing perturbation in the complete planar CR3BP state space. Directions transverse to the family are absent from the optimization. The present result should therefore not be generalized to arbitrary spacecraft position and velocity uncertainty.

Finally, the outcome boundary itself is conditional on a finite-time classification rule. Changing the section, primary-proximity criterion, integration horizon, or event-direction requirement can alter the partition of $\mathcal{D}$ and hence the boundary to which distance is measured. In particular, NO_CROSSING_BY_160 represents the absence of a specified event within a finite horizon, not an asymptotic dynamical state.

These effects need not contribute equally to the observed discrepancy, and the present data do not identify a unique cause of the inside- ρ_η outcome change. The counterexample instead establishes the practical consequence: the combination of a sampled finite-time boundary and a nominal-point local metric is insufficient, in the tested configuration, to certify an outcome-preserving neighborhood.

8.3 Metric Dependence and the Meaning of Robustness

The sensitivity analysis in Fig. 17 shows that the ordering induced by ρ_η is not independent of the state-space weighting. This dependence is expected because W_η specifies which physical components of a seed-state perturbation contribute most strongly to distance. Changing η therefore changes the geometry in which the nearest resolved boundary is sought.

For nearby weighting choices, this dependence is weak. The rankings obtained with $\eta = 0.1$ and $\eta = 0.3$, for example, have a Spearman correlation of approximately

$$r_s(0.1, 0.3) = 0.999. \quad (180)$$

Similarly high correlations occur between several neighboring metric choices. The broad structure identified as relatively close to or far from the boundary is therefore stable under modest changes in the tested weighting.

This stability does not extend uniformly across the full range of η . The correlation between the two extreme tested choices falls to

$$r_s(0.1, 10) = 0.794. \quad (181)$$

The reduction is substantial enough that the weighting can change the relative ordering of nominal points even though the underlying outcome atlas and numerical boundary are identical. The flag `METRIC_SENSITIVE` therefore reflects a property of the robustness definition rather than a change in the underlying transport dynamics.

This result has a direct consequence for the interpretation of ρ_η . A statement that one manifold-seeded initial condition is “more robust” than another is incomplete unless the metric used to compare perturbations is specified. The present construction does not produce a unique scalar robustness ordering intrinsic to the CR3BP flow. It produces an ordering conditional on the selected family, outcome definition, boundary representation, and state-space weighting.

Such conditionality is not necessarily undesirable. In an applied uncertainty model, position and velocity errors need not be equally important or expressed on directly comparable scales. A metric can encode those distinctions explicitly. In that setting, metric dependence has a physical interpretation: robustness is measured with respect to a specified model of admissible perturbations rather than an arbitrary coordinate norm.

The present calculations, however, do not derive W_η from a measured navigation-error covariance, maneuver-execution model, or mission-specific uncertainty distribution. The tested values of η instead probe the sensitivity of the construction to changes in the prescribed weighting. Consequently, none of the tested values should be interpreted as a uniquely physical choice of metric.

A mission-specific extension could replace the parametric weighting with an uncertainty model defined from a covariance matrix or another physically motivated perturbation set. For example, if a positive-definite state covariance Σ were available, a local quadratic state-space measure could be based on

$$\|\delta\mathbf{X}\|_\Sigma^2 = \delta\mathbf{X}^T \Sigma^{-1} \delta\mathbf{X}, \quad (182)$$

which would induce the family metric

$$G_\Sigma(\mathbf{p}) = J(\mathbf{p})^T \Sigma^{-1} J(\mathbf{p}). \quad (183)$$

In such a formulation, distance would have a direct relationship to the assumed uncertainty geometry. This extension was not evaluated here, but it illustrates how the metric dependence observed in Fig. 17 could be converted from a free modeling choice into a mission-dependent physical specification.

The present results therefore argue against referring to ρ_η simply as “the” robustness of a trajectory. A more precise description is the family-constrained distance to

the resolved finite-time outcome boundary under a specified perturbation metric. That terminology retains the useful geometric information in the quantity without implying metric-independent or probabilistic robustness.

8.4 Relation to Condition Numbers and Distance to Ill-Posedness

The construction of ρ_η was motivated in part by the general relationship between sensitivity and proximity to a set at which a problem changes character or becomes ill posed. In classical numerical analysis, condition numbers can often be related to the inverse distance from a well-posed problem instance to an appropriate set of ill-posed instances. This viewpoint suggests an appealing analogy for transport: a trajectory family member lying close to an outcome-change boundary might be expected to be more sensitive to perturbation than one lying far from that boundary.

The present calculations show that this analogy must be applied carefully to finite-time CR3BP transport. The outcome map

$$\mathcal{O} : \mathcal{D} \rightarrow \mathcal{C} \quad (184)$$

is discrete and generally nonsmooth at the interfaces between outcome classes. A distance to such an interface is well defined once a metric and perturbation family are specified, but that distance is not automatically a condition number. In particular, the computation of ρ_η contains no derivative of a continuous output quantity whose relative response to perturbation is being measured.

One could formally construct an inverse-distance quantity,

$$\kappa_\eta^{(\mathcal{B})}(\mathbf{p}) = \frac{1}{\rho_\eta(\mathbf{p})}, \quad (185)$$

so that nominal points closer to the resolved boundary receive larger values. Equation (185), however, would only reparameterize the same geometric information contained in ρ_η . It would not by itself establish the perturbation-response properties normally associated with a condition number.

The empirical tests make this distinction concrete. If ρ_η represented the true minimum metric distance to an outcome-changing perturbation, its reciprocal could reasonably be interpreted as a condition-like measure of proximity to outcome ill-posedness within the prescribed family. The observed outcome change at approximately $0.5\rho_\eta$ demonstrates that the computed distance to $\hat{\mathcal{B}}$ does not satisfy that stronger property in the tested case. Conversely, the directional searches did not recover independent outcome-changing distances at the five representative points. The numerical evidence therefore does not

support identifying $1/\rho_\eta$ with an empirically validated transport condition number.

This negative result is informative because it identifies additional requirements for such an interpretation. The relevant outcome-change set would need to be represented with sufficient completeness and resolution, the perturbation metric would need a specified physical meaning, and the distance calculation would need to approximate the true minimum admissible displacement to that set. Direct perturbation tests would then be required to verify that the inferred distance corresponds to actual changes in transport behavior.

The results of the earlier boundary-scaling experiment point in the same direction. Across the 16 examined outcome boundaries, outcome switching dominated the local behavior, while passage time, residence time, and finite-amplification diagnostics did not exhibit a common monotonic precursor. The absence of a common sensitivity signature motivated the present shift from searching for a universal precursor to measuring the geometry of the outcome partition directly. The EXP-009 results show that the geometric formulation is more sharply defined, but they also show that distance to a numerically sampled boundary should not be promoted to a universal condition number without further validation.

A useful distinction is therefore between a *condition-like boundary diagnostic* and a *validated condition number*. The former can rank nominal points according to their metric proximity to a specified outcome-change set. The latter would require a demonstrated quantitative relationship between that distance and the minimum perturbation producing the relevant change in behavior. The present work establishes the former construction but not the latter.

8.5 Limitations and Scope

The conclusions of this study are intentionally restricted to the finite-time, family-constrained transport problem examined here. Several features of the computational design limit how broadly the resulting boundary geometry and robustness diagnostic can be interpreted.

First, the analysis is confined to the planar Earth-Moon CR3BP. The model neglects out-of-plane motion, solar and planetary perturbations, ephemeris variation, nonspherical gravity, radiation pressure, and finite-thrust or maneuver errors. These effects are absent by construction and the resolved outcome boundaries should therefore not be interpreted as operational boundaries for a physical cislunar mission.

Second, perturbations are restricted to the two-dimensional family parameterized by phase θ and Lyapunov-orbit amplitude A_x . This restriction makes the boundary-distance problem computationally tractable

and gives the perturbations a consistent dynamical origin, but it samples only a two-dimensional subset of the four-dimensional planar state space. A smaller outcome-changing perturbation may exist in a direction transverse to this family. Consequently, ρ_η is not a full-state robustness measure.

Third, the outcome partition is defined by a finite-time first-passage classifier. The classes SECTION_CROSSING, PRIMARY_PROXIMITY, and NO_CROSSING_BY_160 describe behavior relative to the prescribed events and integration horizon. They do not constitute a complete asymptotic classification of CR3BP trajectories. Changing the section location, proximity criterion, event direction, or $T_{\max}$ can alter the partition and its associated boundary.

Fourth, the numerical boundary $\hat{\mathcal{B}}$ is necessarily incomplete at sufficiently fine scales. Boundary discovery begins from neighboring samples on a 160×33 atlas, after which detected outcome-changing edges are refined adaptively. Refinement improves the localization of detected interfaces but cannot recover a boundary component that was never identified by the base sampling. Narrow islands, disconnected components, or fine interleaving of outcomes may therefore remain unresolved.

The two refinement cases flagged for complex local outcome geometry reinforce this limitation. They indicate that at least part of the resolved partition cannot be treated uncritically as a collection of isolated smooth binary interfaces. The present calculation retains these cases explicitly, but it does not attempt to determine whether finer structure continues below the achieved resolution.

Fifth, the distance calculation is local in its treatment of the pullback geometry. The metric tensor $G_\eta(\mathbf{p})$ is evaluated at the nominal point and applied to the displacement between that point and a sampled boundary location. Although the local metric implementation was numerically validated, this approximation does not account for variation of G_η along a finite path through the parameter domain. The computed quantity should therefore not be identified with a global geodesic distance on the seed family.

Sixth, the state weighting is prescribed rather than derived from a mission-specific uncertainty model. The observed dependence on η demonstrates that the resulting robustness ordering is partly conditional on this choice. The sensitivity study quantifies that dependence over the tested range but does not select a uniquely physical metric.

Finally, empirical validation of ρ_η as an outcome-change distance remains limited. Inside-radius perturbation records were available for only one of the five representative points, and the finite directional searches did not locate an outcome flip at any of the five points. The single observed inside-radius outcome change is sufficient to reject a guaranteed-radius interpretation

for the tested construction, but the present validation set is not large enough to characterize systematically how the discrepancy between ρ_η and the true nearest outcome-changing perturbation varies over the domain.

These restrictions define the scope of the conclusions rather than invalidate the numerical outcome atlas itself. Within the stated model, sampling, event definitions, and perturbation family, the calculations provide a reproducible finite-time outcome partition and a metric-dependent measure of proximity to its numerically resolved boundary. Claims beyond that scope require additional dynamical models, perturbation dimensions, boundary-resolution studies, and direct outcome-change validation.

8.6 Implications and Future Work

The present results suggest that outcome-boundary geometry can provide a useful description of finite-time transport sensitivity, but they also identify what would be required to turn that description into a stronger robustness framework. In particular, future work should distinguish between improving the representation of the outcome boundary and improving the metric used to measure distance to it.

A first priority is more complete boundary discovery. The current procedure refines interfaces only after an outcome change has been detected between neighboring atlas samples. A more aggressive adaptive scheme could recursively subdivide regions containing rapid changes in event time, close approaches to competing terminal events, or other indicators of unresolved outcome structure. Such a procedure would test whether the inside- ρ_η counterexample is associated with an outcome-boundary component that was missed by the original atlas.

Boundary convergence should then be assessed explicitly. Repeating the calculation with progressively finer base grids and refinement tolerances would allow the stability of $\hat{\mathcal{B}}$ and ρ_η to be measured as numerical resolution increases. A boundary-distance quantity intended for robustness analysis should converge, within a stated tolerance, under such refinement before it is interpreted physically.

A second extension concerns the geometry of the seed family. Rather than freezing the pullback tensor at the nominal point, future calculations could evaluate the path-dependent distance introduced in Eq. (179). The resulting geodesic problem would account for variation of the seed-map Jacobian across the parameter domain. Comparison between nominal-point and geodesic distances would determine whether the local quadratic approximation contributes materially to the discrepancy observed in the present validation.

The weighting itself can also be made physically spe-

cific. Replacing the exploratory family W_η with a metric derived from navigation covariance, injection uncertainty, or maneuver-execution statistics would connect the geometry directly to an assumed error model. Under a probabilistic uncertainty distribution, the analysis could then move beyond distance alone and estimate the probability that uncertainty mass intersects a competing outcome region. Such a probability would represent a different quantity from ρ_η and should be validated independently.

Most importantly, a future robustness calculation should search directly for the minimum outcome-changing perturbation. Within the two-dimensional family, this can be formulated as the constrained problem

$$r_\eta^*(\mathbf{p}) = \inf_{\delta\mathbf{p}} \left\{ \sqrt{\delta\mathbf{p}^T G_\eta(\mathbf{p}) \delta\mathbf{p}} : \mathcal{O}(\mathbf{p} + \delta\mathbf{p}) \neq \mathcal{O}(\mathbf{p}) \right\}. \quad (186)$$

Unlike the present nearest-sampled-boundary construction, Eq. (186) defines the target quantity in terms of an actual outcome change. A sufficiently reliable numerical solution could therefore provide an independent reference against which ρ_η is tested. Agreement between the two quantities under boundary and search refinement would be substantially stronger evidence for a distance-to-outcome-change interpretation.

The perturbation space could subsequently be expanded beyond (θ, A_x) . In the planar problem, optimization over the full four-dimensional initial state would test whether the selected seed family excludes closer outcome-changing directions. Extension to the spatial CR3BP would further introduce out-of-plane perturbations. Ultimately, the same methodology could be evaluated in higher-fidelity ephemeris models, where the persistence of the identified outcome structure would become a separate question.

These extensions also provide a route toward a more defensible condition-like transport quantity. Rather than defining a condition number solely as the reciprocal of a sampled boundary distance, one could first compute or approximate the minimum physically admissible outcome-changing perturbation and then study its variation over a trajectory family. Only after numerical convergence and direct perturbation validation would an inverse-distance interpretation have a strong basis as a transport condition measure.

The main implication of the present study is therefore methodological. Finite-time outcome boundaries contain useful information about transport sensitivity, but geometric proximity, numerical boundary resolution, uncertainty geometry, and actual outcome-changing perturbations must be treated as distinct objects. Keeping these quantities separate makes it possible to test, rather than assume, whether a boundary-distance diagnostic can support a stronger notion of trajectory robustness.

9 Conclusion

This study investigated whether finite-time transport robustness in the planar Earth–Moon CR3BP can be characterized geometrically through distance to changes in transport outcome. A validated family of L_1 planar Lyapunov orbits and their invariant-manifold seeds was used to construct a two-dimensional parameter family in phase θ and orbit amplitude A_x . Each seed was propagated under a first-passage classification, producing a finite-time outcome atlas containing section-crossing, primary-proximity, and no-crossing trajectories.

The production atlas contained 5280 initial conditions and revealed a structured partition of the sampled family. Adaptive refinement of outcome-changing neighboring samples produced 484 numerical boundary samples, including refinement cases with complex local outcome geometry. These results show that even within the restricted two-dimensional family, changes in finite-time transport behavior are organized by nontrivial interfaces in parameter space rather than by a single simple transition.

To measure proximity to these interfaces without using an arbitrary Euclidean norm in (θ, A_x) , the seed-state map was used to induce a local pullback metric on the parameter family. The resulting quantity ρ_η measures the metric distance from a nominal seed to the numerically resolved outcome boundary. Numerical tests confirmed the local consistency of the metric construction, while the resulting distance field varied substantially across the sampled family.

The experiments also establish important limits on the interpretation of this quantity. The ordering produced by ρ_η depends on the chosen state-space weighting, with Spearman correlations ranging from approximately 0.794 to 0.999 across the tested metrics. More importantly, one directly tested perturbation changed outcome at approximately $0.5\rho_\eta$. Finite directional searches at five representative points did not recover independent outcome-changing distances. The computed ρ_η therefore cannot be identified, on the basis of the present experiments, with either a guaranteed outcome-preserving radius or the minimum perturbation required to change transport outcome.

The resulting distinction is central to the interpretation of condition-like measures for nonlinear transport. Distance to a sampled outcome boundary is a useful geometric diagnostic, but its reciprocal does not become a validated transport condition number without an independently established relationship to actual outcome-changing perturbations. Boundary completeness, numerical resolution, the perturbation family, and the physical metric all enter that relationship.

Within these limits, the present framework pro-

vides a reproducible way to map finite-time transport outcomes and quantify family-constrained proximity to their resolved boundaries. A stronger robustness measure would require direct computation of minimum outcome-changing perturbations, convergence of the boundary representation under refinement, and a metric tied to a physically specified uncertainty model. The present results therefore support boundary distance as a diagnostic of transport geometry, while leaving the construction of a validated distance-to-outcome-change condition measure as a separate problem.

References

- [1] Analysis, design, and optimization of robust trajectories in cislunar environment for limited-capability spacecraft. *The Journal of the Astronautical Sciences*, 2023.
- [2] Jean-Baptiste Caillau and B. Daoud. Minimum time control of the restricted three-body problem. *SIAM Journal on Control and Optimization*, 50(6):3178–3202, 2012.
- [3] Jean-Baptiste Caillau, B. Daoud, and J. Gergaud. Minimum fuel control of the planar circular restricted three-body problem. *Celestial Mechanics and Dynamical Astronomy*, 114(1):137–150, 2012.
- [4] Sheila C. de Assis and Maisa O. Terra. Escape dynamics and fractal basin boundaries in the planar earth–moon system. *Celestial Mechanics and Dynamical Astronomy*, 120(2):105–130, 2014.
- [5] James W. Demmel. The geometry of ill-conditioning. *Journal of Complexity*, 3(2):201–229, 1987.
- [6] James W. Demmel. On condition numbers and the distance to the nearest ill-posed problem. *Numerische Mathematik*, 51(3):251–289, 1987.
- [7] Marco Giancotti, Mauro Pontani, and Paolo Teofilatto. Cylindrical isomorphic mapping applied to invariant manifold dynamics for earth–moon missions. *Celestial Mechanics and Dynamical Astronomy*, 120:249–268, 2014.
- [8] Gerard Gómez, Àngel Jorba, Josep Masdemont, and Carles Simó. *Dynamics and Mission Design Near Libration Points: Volume III: Advanced Methods for Collinear Points*. World Scientific, 2001.
- [9] Gerard Gómez, Jaume Llibre, Regina Martínez, and Carles Simó. *Dynamics and Mission Design Near Libration Points: Volume I: Fundamentals: The Case of Collinear Libration Points*. World Scientific, 2001.

- [10] Kathleen Connor Howell. Three-dimensional, periodic, ‘halo’ orbits. *Celestial Mechanics*, 32(1):53–71, 1984.
- [11] Wang Sang Koon, Martin W. Lo, Jerrold E. Marsden, and Shane D. Ross. Heteroclinic connections between periodic orbits and resonance transitions in celestial mechanics. *Chaos*, 10(2):427–469, 2000.
- [12] Wang Sang Koon, Martin W. Lo, Jerrold E. Marsden, and Shane D. Ross. *Dynamical Systems, the Three-Body Problem and Space Mission Design*. Marsden Books, 2022.
- [13] Mauro Pontani and Paolo Teofilatto. Polyhedral representation of invariant manifolds applied to orbit transfers in the earth–moon system. *Acta Astronautica*, 119:218–232, 2016.
- [14] Giuseppe Rega and Stefano Lenci. Dynamical integrity and control of nonlinear mechanical oscillators. *Journal of Vibration and Control*, 14(1–2), 2008.
- [15] David L. Richardson. Analytic construction of periodic orbits about the collinear points. *Celestial Mechanics*, 22(3):241–253, 1980.
- [16] Victor Szebehely. *Theory of Orbits: The Restricted Problem of Three Bodies*. Academic Press, New York, 1967.
- [17] Euaggelos E. Zotos. Classifying orbits in the restricted three-body problem. *Nonlinear Dynamics*, 82(3):1233–1250, 2015.
- [18] Euaggelos E. Zotos. Escape dynamics and fractal basin boundaries in the three-dimensional earth–moon system. *Astrophysics and Space Science*, 361, 2016.