

SympGal-NODE: Continuous-Time Symplectic Neural ODEs for Relativistic Stellar Swarms: Predicting Escape Times, Exit Trajectories, and Tidal Horizons

Quan Van

vanhaminhquan2406@gmail.com

Independent Researcher <https://orcid.org/0009-0002-8222-0076>

Research Article

Keywords: Escape Trajectory Prediction, Escape Time Horizon, Symplectic Integrators, Relativistic Hamiltonian Dynamics, Neural Ordinary Differential Equations, Equivariant Graph Neural Networks, Post-Newtonian Astrodynamics

Posted Date: October 14th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11355949/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.
[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

Highlights

SympGal-NODE: Continuous-Time Symplectic Neural ODEs for Relativistic Stellar Swarms: Predicting Escape Times, Exit Trajectories, and Tidal Horizons

Quan Van

- SympGal-NODE predicts continuous escape trajectories, exit corridors, and escape times (T_{escape}) of relativistic stellar swarms.
- Preserves canonical symplectic 2-form $d\mathbf{q} \wedge d\mathbf{p}$ and Liouville volume ($\det \mathbf{J} \equiv 1$), eliminating artificial cluster dissolution.
- Relativistic Hamiltonian formulation couples 1PN Schwarzschild curvature, J_2 oblateness, and solar radiation pressure.
- Hierarchical Dynamic Hill Routing scales to $N = 10,000,000$ bodies on commodity 4GB GPU via constant-memory streaming.
- Outperforms extensive 2025–2026 baselines on chaotic cluster evaporation and empirical NASA JPL Horizons ephemerides.

SympGal-NODE: Continuous-Time Symplectic Neural ODEs for Relativistic Stellar Swarms: Predicting Escape Times, Exit Trajectories, and Tidal Horizons

Quan Van^{a,*}

^a*Independent Researcher, Vietnam*

Abstract

Predicting the long-horizon collective dissolution, continuous escape trajectories, and evaporation timescales of massive gravitational swarms in relativistic environments is a foundational challenge in celestial mechanics and computational astrodynamics. Classical numerical integrators suffer from quadratic computational complexity and secular energy drift, whereas conventional deep learning models violate canonical phase-space geometry, producing artificial numerical dissipation that accelerates cluster dissolution and distorts chaotic escape corridors. In this work, we propose SympGal-NODE, an end-to-end structure-preserving geometric deep learning framework designed specifically to predict continuous-time escape trajectories and exit horizons of relativistic stellar swarms. SympGal-NODE integrates: (1) an exact first Post-Newtonian relativistic Hamiltonian formulation in the Solar System Barycenter frame with oblateness and radiation pressure perturbations; (2) an Equivariant Graph Neural Network equipped with Hierarchical Dynamic Hill Sphere edge routing that achieves near-linear computational scaling; (3) an unconditionally symplectic implicit midpoint neural integrator that rigorously preserves canonical symplectic geometry and Liouville volume conservation; and (4) a continuous-time multi-risk survival analysis head predicting cause-specific escape hazards, exit epochs, and directional Lagrange aperture trajectories. Benchmarked against extensive 2025–2026 baselines across swarm populations up to ten million bodies and validated on empirical NASA JPL Horizons ephemerides, SympGal-NODE strictly preserves phase-space volume, delivers superior escape time concordance, and achieves a 28.4-fold speedup over conventional symplectic integrators while operating within commodity hardware memory constraints.

Keywords: Escape Trajectory Prediction, Escape Time Horizon, Symplectic Integrators, Relativistic Hamiltonian Dynamics, Neural Ordinary Differential Equations, Equivariant Graph Neural Networks, Post-Newtonian Astrodynamics

PACS: 95.10.Ce, 95.75.Pq, 04.25.Nx

2020 MSC: 85-08, 68T07, 70F10, 37M15

*Corresponding author

Email address: vanhaminhquan2406@gmail.com (Quan Van)

1. Introduction & Literature Review

The dynamical evolution of massive gravitational swarms ($N \sim 10^3\text{--}10^6$) orbiting in the extreme gravitational environments of binary supermassive black holes, hierarchical stellar triples, and dense galactic nuclei represents one of the foundational open frontiers in computational astrophysics, celestial mechanics, and relativistic astrodynamics [1, 2]. Within these extreme environments, classical Newtonian gravity ceases to be a sufficient description. Relativistic phenomena—including first Post-Newtonian (1PN) Schwarzschild apsidal precession, frame-dragging gravitomagnetic Lense-Thirring torques, and primary oblateness perturbations (J_2)—interlock with multi-body resonances, inducing complex chaos, tidal stripping, and violent stellar ejections [3, 4, 5].

Accurately resolving these multi-scale trajectories is essential for tracking stellar streams, predicting extreme mass-ratio inspiral (EMRI) progenitors, and modeling deep-space satellite swarms in relativistic planetary systems. However, direct numerical integration of massive relativistic swarms faces fundamental algorithmic and physical hurdles:

1. **Curse of Quadratic Complexity:** Direct relativistic N -body force evaluation scales strictly as $\mathcal{O}(N^2)$ per timestep [6?]. For million-body clusters ($N = 10^6$), simulating millions of dynamical orbits requires millions of GPU-hours, rendering exhaustive parametric exploration intractable.
2. **Geometric Phase-Space Inconsistency in Classical Deep Learning:** Recent advances in Scientific Machine Learning (SciML) and continuous-time deep learning—notably Neural Ordinary Differential Equations (NODEs) [7], Physics-Informed Neural Networks (PINNs) [8, 9], and Fourier Neural Operators (FNO) [?]—have demonstrated remarkable predictive capabilities. Nevertheless, conventional deep surrogates parameterize arbitrary non-conservative vector fields $\dot{\mathbf{x}} = \mathbf{f}_\theta(\mathbf{x})$. When integrated over astronomical horizons, these models fail to preserve the canonical symplectic 2-form $\omega = d\mathbf{q} \wedge d\mathbf{p}$. In accordance with Liouville’s theorem, non-symplectic surrogates introduce spurious artificial dissipation or exponential runaway energy growth, causing unphysical spiral infall or artificial cluster evaporation.
3. **Chaotic Horizon Escape and Escape Trajectory Prediction:** In gravitational swarms, determining when individual stars escape across the Jacobi tidal horizon (R_{tidal}) and mapping their asymptotic exit trajectories into galactic tidal streams presents extreme sensitivity to initial conditions (Lyapunov chaos). Standard regression surrogates or discrete classification networks fail to predict continuous escape times (T_{escape}) or handle right-censored observation windows (particles remaining bound until simulation end T_{max}). Furthermore, when calculating long-range exit vectors through Lagrange saddle nozzles (L_1, L_2), non-symplectic errors corrupt the angular momentum distribution, leading to artificial, unphysical stellar stream morphologies.

1.1. Prior Work and the Research Gap

To enforce physical invariants, Hamiltonian Neural Networks (HNN) [10] and Lagrangian Neural Networks (LNN) [11] pioneered learning scalar potential energies $\mathcal{H}_\theta(\mathbf{q}, \mathbf{p})$ from which

conservative vector fields are derived via exact automatic differentiation:

$$\dot{\mathbf{q}} = \frac{\partial \mathcal{H}_\theta}{\partial \mathbf{p}}, \quad \dot{\mathbf{p}} = -\frac{\partial \mathcal{H}_\theta}{\partial \mathbf{q}}. \quad (1)$$

Subsequently, SympNets [12] and Variational Integrator Networks [13] formulated discrete symplectic mappings composed of upper and lower triangular shear operators. Simultaneously, Equivariant Graph Neural Networks (EGNN) [? 14] and Lie-symmetric neural representations [?] established exact $E(n)$ and $SO(3)$ geometric equivariance for particle systems.

Despite these advancements, an acute research gap persists at the intersection of relativistic astrodynamics and geometric deep learning:

- Existing Hamiltonian neural architectures are universally formulated for idealized, isolated Newtonian toy problems (e.g., pendulums, Keplerian two-body systems) and assume separable Hamiltonians $\mathcal{H}(\mathbf{q}, \mathbf{p}) = \frac{1}{2} \mathbf{p}^T \mathbf{M}^{-1} \mathbf{p} + V(\mathbf{q})$. They completely break down in relativistic regimes where the first Post-Newtonian (1PN) potential explicitly couples canonical positions and canonical momenta through cross-product angular momentum terms $V_{1PN}(\mathbf{q}, \mathbf{p}) \propto \|(\mathbf{q} - \mathbf{q}_j) \times (\mathbf{p} - \mathbf{p}_j)\|^2$, rendering the system strictly non-separable.
- Standard graph neural ODEs suffer from catastrophic over-smoothing and $\mathcal{O}(N^2)$ message passing when deployed on massive stellar swarms, lacking physical domain-aware dynamic edge pruning.
- No existing framework bridges continuous-time relativistic symplectic dynamics with right-censored competing-risk survival analysis to forecast chaotic tidal escape horizons ($\log_{10} T_{\text{escape}}$) directly from initial phase-space coordinates.

1.2. Contributions of This Work

To address these limitations, we introduce **SympGal-NODE** (**S**ymplectic **G**alactic **N**eural **O**rdinary **D**ifferential **E**quations). The core contributions of this paper are structured as follows:

1. **Rigorous Relativistic Hamiltonian Formulation:** We derive the complete multi-body non-separable Hamiltonian $\mathcal{H}_{\text{total}}(\mathbf{q}, \mathbf{p})$ in the Solar System Barycenter (SSB) frame, rigorously incorporating Newtonian point-mass interactions, 1PN Schwarzschild apsidal precession, primary oblateness perturbations (J_2), Solar Radiation Pressure (SRP), 1.5PN Lense-Thirring gravitomagnetic spin-orbit coupling, and Chandrasekhar dynamical friction via Strang operator splitting.
2. **Hierarchical Dynamic Hill Sphere (R_{SOI}) Edge Routing:** We design an equivariant message-passing scheme that exploits gravitational Sphere of Influence (SOI) hierarchies, decomposing swarm interactions into active primary wells and cluster mean-field backgrounds, reducing graph complexity from $\mathcal{O}(N^2)$ to $\mathcal{O}(N \log N)$.

3. **Structure-Preserving Neural Architecture & Implicit Midpoint Integration:** We propose SympGal-Net, combining an $\text{SO}(3)$ -invariant geometric encoder, a volume-preserving Poincaré jump operator ($\det \mathbf{J} \equiv 1$), and an unconditionally symplectic implicit midpoint neural integrator that exactly preserves the canonical symplectic 2-form for non-separable 1PN potentials.
4. **Multi-Risk Survival Head for Chaotic Escape Horizons:** We introduce a continuous-time competing-risk survival head that resolves right-censored tidal horizon crossings, outputting discrete hazard curves h_k , cumulative survival functions $S(t)$, and ejection horizons ($\log_{10} T_{\text{escape}}$).
5. **Large-Scale Benchmark Scaling and Empirical Ephemeris Validation:** We evaluate SympGal-NODE against eight competitive baselines across swarm populations up to $N = 1,000,000$, and validate the learned trajectory vectors against high-precision NASA JPL Horizons ephemerides, establishing superior stability, energy conservation, and runtime throughput.

2. Relativistic Physical Mechanics & Hamiltonian Formulation

2.1. Barycentric Coordinate Frame & Canonical Phase Space

We consider a relativistic stellar swarm composed of K primary gravitational sources (referred to as “Mother Stars” or massive central bodies, indexed by $j \in \{1, \dots, K\}$) with masses M_j , and N child stars or test particles (indexed by $i \in \{1, \dots, N\}$) with individual masses $m_i \ll M_j$. The dynamical state of each entity is described in a global, quasi-inertial Solar System Barycenter (SSB) reference frame by canonical coordinates:

$$\mathbf{z}_i = (\mathbf{q}_i, \mathbf{p}_i) \in \mathbb{R}^3 \times \mathbb{R}^3 \cong \mathbb{R}^6, \quad (2)$$

where $\mathbf{q}_i = [x_i, y_i, z_i]^T$ denotes the barycentric Cartesian position vector, and $\mathbf{p}_i = [p_{x,i}, p_{y,i}, p_{z,i}]^T$ denotes the canonical momentum per unit specific mass (such that $\mathbf{p}_i = \mathbf{v}_i$ in the non-relativistic limit).

To eliminate unphysical linear momentum drift and coordinate frame acceleration, the combined system strictly satisfies barycentric momentum and center-of-mass constraints:

$$\sum_{j=1}^K M_j \mathbf{q}_j + \sum_{i=1}^N m_i \mathbf{q}_i = \mathbf{0}, \quad \sum_{j=1}^K M_j \mathbf{p}_j + \sum_{i=1}^N m_i \mathbf{p}_i = \mathbf{0}. \quad (3)$$

2.2. The Total Relativistic Multi-Body Hamiltonian

The total governing Hamiltonian of a test particle i moving through the collective curved spacetime potential generated by the K primary sources and the diffuse cluster background is given by:

$$\mathcal{H}_{\text{total}}(\mathbf{q}_i, \mathbf{p}_i) = \mathcal{T}(\mathbf{p}_i) + \Phi_{\text{eff}}(\mathbf{q}_i, \mathbf{p}_i), \quad (4)$$

where $\mathcal{T}(\mathbf{p}_i) = \frac{1}{2} \|\mathbf{p}_i\|^2$ represents the specific kinetic energy, and Φ_{eff} represents the effective potential per unit mass, decomposed as:

$$\Phi_{\text{eff}}(\mathbf{q}_i, \mathbf{p}_i) = \Phi_{\text{Newt}}(\mathbf{q}_i) + V_{\text{1PN}}(\mathbf{q}_i, \mathbf{p}_i) + \Phi_{J_2}(\mathbf{q}_i) + \Phi_{\text{bg}}(\mathbf{q}_i) + \Phi_{\text{SRP}}(\mathbf{q}_i, t). \quad (5)$$

2.2.1. Newtonian Point-Mass Gravitational Potential

The direct Newtonian gravitational interaction exerted by the K primaries on particle i is parameterized with a Plummer softening parameter $\epsilon > 0$ to prevent numerical singularities during close encounters:

$$\Phi_{\text{Newt}}(\mathbf{q}_i) = - \sum_{j=1}^K \frac{GM_j}{(\|\mathbf{q}_i - \mathbf{q}_j\|^2 + \epsilon^2)^{1/2}}, \quad (6)$$

where G is the universal gravitational constant.

2.2.2. 1PN Post-Newtonian Relativistic Schwarzschild Potential

In the relativistic regime, test particles experience apsidal precession governed by the Einstein field equations. In the isotropic post-Newtonian formulation, the first Post-Newtonian (1PN) Schwarzschild interaction can be expressed through an effective momentum-dependent potential term [3, 4]:

$$V_{\text{1PN}}(\mathbf{q}_i, \mathbf{p}_i) = - \sum_{j=1}^K \frac{GM_j \|\mathbf{L}_{ij}\|^2}{c^2 (\|\mathbf{q}_i - \mathbf{q}_j\|^2 + \epsilon^2)^{3/2}}, \quad (7)$$

where c is the speed of light, and $\mathbf{L}_{ij}$ is the relative specific angular momentum vector between particle i and primary j :

$$\mathbf{L}_{ij} = (\mathbf{q}_i - \mathbf{q}_j) \times \left(\mathbf{p}_i - \frac{\mathbf{p}_j}{M_j} \right). \quad (8)$$

Notice that Equation (7) induces an explicit, non-separable coupling between $\mathbf{q}_i$ and $\mathbf{p}_i$. The resulting canonical equations of motion are given by:

$$\dot{\mathbf{q}}_i = \frac{\partial \mathcal{H}_{\text{total}}}{\partial \mathbf{p}_i} = \mathbf{p}_i + \frac{\partial V_{\text{1PN}}}{\partial \mathbf{p}_i}, \quad \dot{\mathbf{p}}_i = - \frac{\partial \mathcal{H}_{\text{total}}}{\partial \mathbf{q}_i} = - \nabla_{\mathbf{q}_i} \Phi_{\text{eff}}(\mathbf{q}_i, \mathbf{p}_i). \quad (9)$$

Because $\partial V_{\text{1PN}} / \partial \mathbf{p}_i \neq \mathbf{0}$, the velocity $\dot{\mathbf{q}}_i$ does not coincide trivially with $\mathbf{p}_i$, requiring specialized symplectic numerical treatment.

2.2.3. Oblateness Perturbation (J_2)

When the central primary bodies exhibit non-spherical mass distributions due to rotation, the gravitational quadrupole moment induces an oblateness potential perturbation Φ_{J_2} . For a primary body with equatorial radius R_j and dimensionless oblateness coefficient $J_{2,j}$ aligned with spin axis $\hat{\mathbf{s}}_j$:

$$\Phi_{J_2}(\mathbf{q}_i) = \sum_{j=1}^K \frac{GM_j R_j^2 J_{2,j}}{2 \|\mathbf{r}_{ij}\|^3} \left[3 \left(\frac{\mathbf{r}_{ij} \cdot \hat{\mathbf{s}}_j}{\|\mathbf{r}_{ij}\|} \right)^2 - 1 \right], \quad (10)$$

where $\mathbf{r}_{ij} = \mathbf{q}_i - \mathbf{q}_j$. This quadrupole term drives nodal regression and periapsis shift, interacting constructively with 1PN relativistic precession.

2.2.4. Solar Radiation Pressure (SRP) Model

For small celestial bodies, dust grains, and spacecraft swarms, solar radiation pressure (SRP) exerts an outward, non-conservative or time-dependent potential force directed away from luminous primaries:

$$\mathbf{F}_{\text{SRP},j}(\mathbf{q}_i) = \frac{P_{\text{rad},j} A_i (1 + \eta_i)}{m_i} \frac{\mathbf{r}_{ij}}{\|\mathbf{r}_{ij}\|^3}, \quad (11)$$

where $P_{\text{rad},j} = \mathcal{L}_j / (4\pi c)$ is the radiation pressure constant at unit distance associated with stellar luminosity $\mathcal{L}_j$, A_i/m_i is the area-to-mass ratio of particle i , and $\eta_i \in [0, 1]$ represents the surface reflectivity coefficient. For conservative approximations, $\mathbf{F}_{\text{SRP}}$ is incorporated via an effective repulsive scalar potential $\Phi_{\text{SRP}}(\mathbf{q}_i) = \sum_j \frac{C_{\text{SRP},ij}}{\|\mathbf{r}_{ij}\|}$, where $C_{\text{SRP},ij} = \frac{P_{\text{rad},j} A_i (1 + \eta_i)}{m_i}$.

2.2.5. Smooth Cluster Mean-Field Background

To capture the collective self-gravity of the unmodeled stellar cluster or galactic nucleus, we incorporate a smooth mean-field background potential $\Phi_{\text{bg}}(\mathbf{q}_i)$. Our framework supports both a spherical Plummer sphere and an axisymmetric flattened Miyamoto-Nagai disc [1]:

$$\Phi_{\text{bg}}(\mathbf{q}_i) = \begin{cases} -\frac{GM_0}{\sqrt{\|\mathbf{q}_i\|^2 + b_0^2}}, & \text{(Plummer)} \\ -\frac{GM_0}{\sqrt{x_i^2 + y_i^2 + \left(a_{\text{MN}} + \sqrt{z_i^2 + b_{\text{MN}}^2}\right)^2 + \epsilon^2}}, & \text{(Miyamoto-Nagai)} \end{cases} \quad (12)$$

where M_0 is the total background mass, b_0 is the Plummer scale radius, and $a_{\text{MN}}, b_{\text{MN}}$ parameterize the radial and vertical disc scalelengths.

2.3. 1.5PN Gravitomagnetic Lense-Thirring Coupling & Dynamical Friction

In addition to the conservative Hamiltonian flow, compact spinning primaries impart a 1.5PN gravitomagnetic frame-dragging acceleration on orbiting swarm particles via spin-orbit coupling [15]:

$$\mathbf{a}_{\text{LT},j}(\mathbf{q}_i, \mathbf{p}_i) = \frac{G}{c^2 \|\mathbf{r}_{ij}\|^3} \left[\frac{3(\mathbf{r}_{ij} \cdot \mathbf{S}_j)}{\|\mathbf{r}_{ij}\|^2} (\mathbf{r}_{ij} \times \mathbf{v}_{ij}) + 2(\mathbf{v}_{ij} \times \mathbf{S}_j) \right], \quad (13)$$

where $\mathbf{S}_j$ is the intrinsic spin angular momentum vector of primary j , and $\mathbf{v}_{ij} = \mathbf{p}_i - \mathbf{p}_j / M_j$.

Simultaneously, massive primaries traversing the background stellar density $\rho_{\text{bg}}(\mathbf{q}_j)$ undergo Chandrasekhar dynamical friction [16]:

$$\mathbf{a}_{\text{df},j} = -4\pi G^2 M_j \rho_{\text{bg}}(\mathbf{q}_j) \ln \Lambda \left[\text{erf}(X_j) - \frac{2X_j}{\sqrt{\pi}} e^{-X_j^2} \right] \frac{\mathbf{v}_j}{\|\mathbf{v}_j\|^3}, \quad (14)$$

where $\ln \Lambda$ is the Coulomb logarithm, $X_j = \|\mathbf{v}_j\| / (\sqrt{2}\sigma_v)$, and σ_v is the background velocity dispersion.

2.4. Symplectic Phase-Space Conservation & Liouville's Theorem

A time-continuous dynamical flow $\psi_t : (\mathbf{q}_0, \mathbf{p}_0) \mapsto (\mathbf{q}(t), \mathbf{p}(t))$ generated by any Hamiltonian system is a symplectomorphism. Formally, let $\mathbf{J}$ denote the standard $2d \times 2d$ symplectic matrix:

$$\mathbf{J} = \begin{bmatrix} \mathbf{0} & \mathbf{I}_d \\ -\mathbf{I}_d & \mathbf{0} \end{bmatrix}. \quad (15)$$

The Jacobian matrix $\mathbf{M} = \frac{\partial \psi_t}{\partial (\mathbf{q}_0, \mathbf{p}_0)}$ satisfies the canonical symplectic condition:

$$\mathbf{M}^T \mathbf{J} \mathbf{M} = \mathbf{J}, \quad \forall t \geq 0. \quad (16)$$

A direct consequence of Equation (16) is that the phase-space volume is strictly conserved:

$$\det(\mathbf{M}^T \mathbf{J} \mathbf{M}) = \det(\mathbf{M})^2 \det(\mathbf{J}) = \det(\mathbf{J}) \implies \det(\mathbf{M}) = +1. \quad (17)$$

This mathematical equality is the differential manifestation of **Liouville's Theorem**: the density of states in canonical phase space along a trajectory is constant in time ($\frac{d\rho}{dt} = 0$). Any numerical integrator or neural ODE surrogate that fails to satisfy Equation (16) introduces artificial numerical dissipation ($\det \mathbf{M} < 1$) or artificial explosion ($\det \mathbf{M} > 1$), corrupting orbital resonance locks and phase coherence over long temporal baselines.

2.5. Tidal Evaporation Horizons & Escape Criterion

In multi-body astrodynamics, particles are continually perturbed through close encounters and three-body sling-shot mechanisms. A particle i is defined to cross the **Tidal Horizon** (R_{tidal}) and undergo irreversible hyperbolic escape if two simultaneous conditions are met:

$$\|\mathbf{q}_i(t)\| \geq R_{\text{tidal}} \quad \text{and} \quad \mathcal{E}_i(t) = \frac{1}{2} \|\mathbf{p}_i(t)\|^2 + \Phi_{\text{eff}}(\mathbf{q}_i(t), \mathbf{p}_i(t)) > 0, \quad (18)$$

where $\mathcal{E}_i(t)$ represents the specific orbital energy. The continuous escape time is defined by the infimum:


$$T_{\text{escape},i} \triangleq \inf \{t > 0 \mid \|\mathbf{q}_i(t)\| \geq R_{\text{tidal}} \wedge \mathcal{E}_i(t) > 0\}. \quad (19)$$

If a particle remains bound within the tidal radius throughout the entire observational or simulation window $[0, T_{\text{max}}]$, its escape time is right-censored at T_{max} with event indicator $\delta_i = 0$. Otherwise, if the particle crosses the horizon, $\delta_i = 1$ and the logged escape horizon is measured in logarithmic decades: $\log_{10} T_{\text{escape},i}$. Conversely, if a particle passes within the physical collision radius of a primary ($\|\mathbf{q}_i - \mathbf{q}_j\| \leq R_{\text{coll}}$), it is flagged as a tidal collision capture event ($\delta_i = 2$).

3. SympGal-NODE Architecture & Dynamic Hill Routing

3.1. Hierarchical Dynamic Hill Sphere (R_{SOI}) Edge Routing

Direct graph neural networks computing dense pairwise interactions between N swarm particles scale as $\mathcal{O}(N^2)$, which becomes computationally prohibitive for $N \geq 10^5$. To resolve this bottleneck, we introduce **Hierarchical Dynamic Hill Sphere Routing**.



(a) Top-Down Ecliptic Plane Map (b) 3D Spacetime Curvature Funnel (c) Mollweide All-Sky Escape Map

Figure 1: Astrodynamical visualizations generated by the SympGal-NODE computational framework. (a) Top-down orbital trajectory projection of the relativistic central binary core and orbiting swarm test particles, color-coded by logarithmic escape epoch $\log_{10} T_{\text{escape}}$. (b) 3D curved spacetime effective potential surface (“rubber trampoline sheet” analogy) illustrating primary gravitational funnels and Lagrange saddle bottlenecks. (c) All-sky Mollweide projection (λ, β) mapping anisotropic escape corridors and directional horizon ejection vectors.

In celestial mechanics, the Sphere of Influence (SOI) or Hill sphere of primary body j with semi-major axis a_j relative to the central barycenter is given by:

$$R_{\text{SOI},j} = a_j \left(\frac{M_j}{3M_{\text{total}}} \right)^{1/3}. \quad (20)$$

We construct a dynamically routed hierarchical bipartite graph $\mathcal{G}_t = (\mathcal{V}_{\text{swarm}} \cup \mathcal{V}_{\text{mothers}}, \mathcal{E}_t)$ at each timestep t . Edges between swarm particle i and primary j are assigned via an adaptive radial gating function:

$$e_{ij} = \begin{cases} 1, & \text{if } \|\mathbf{q}_i - \mathbf{q}_j\| \leq \gamma R_{\text{SOI},j}, \\ 0, & \text{otherwise,} \end{cases} \quad (21)$$

where $\gamma \geq 1.5$ is a buffer margin ensuring smooth transitioning across orbital resonance boundaries. Swarm particles outside all individual primary Hill spheres interact exclusively through the cluster mean-field multipole expansion Φ_{bg} . This hierarchical routing guarantees that each swarm node maintains active directed edges only to relevant gravitational wells, reducing active edge cardinality from $|\mathcal{E}| = \mathcal{O}(K \cdot N)$ to $|\mathcal{E}_{\text{active}}| = \mathcal{O}(N)$, thereby enabling linear spatial and temporal scaling.

3.2. Equivariant Relativistic Feature Embeddings

To ensure that all learned latent representations respect physical coordinate rotation symmetries, we design an $\text{SO}(3)$ -invariant feature extractor. For each particle-mother pair (i, j) with relative separation $\mathbf{d}_{ij} = \mathbf{q}_i - \mathbf{q}_j$ and relative velocity $\mathbf{v}_{ij} = \mathbf{p}_i - \mathbf{p}_j/M_j$, we

construct a 9-dimensional scalar invariant feature vector:

$$\chi_{ij} = \begin{bmatrix} \|\mathbf{d}_{ij}\| \\ \|\mathbf{v}_{ij}\| \\ \mathbf{d}_{ij} \cdot \mathbf{v}_{ij} \\ \frac{\|\mathbf{d}_{ij}\| + \epsilon}{\|\mathbf{d}_{ij} \times \mathbf{v}_{ij}\|} \\ \Phi_{\text{1PN},ij} \\ \|\mathcal{T}_{\text{tidal},ij}\|_F \\ M_j \\ \|\mathbf{q}_i\| \\ \|\mathbf{p}_i\| \end{bmatrix} \in \mathbb{R}^9, \quad (22)$$

where $\|\mathcal{T}_{\text{tidal},ij}\|_F = \frac{\sqrt{6}GM_j}{\|\mathbf{d}_{ij}\|^3 + \epsilon^3}$ is the Frobenius norm of the local quadrupole tidal gravity gradient tensor $\mathcal{T}_{ab} = \partial^2 \Phi / \partial x_a \partial x_b$, and $\Phi_{\text{1PN},ij} = -\frac{GM_j \|\mathbf{d}_{ij} \times \mathbf{v}_{ij}\|^2}{c^2 (\|\mathbf{d}_{ij}\|^2 + \epsilon^2)^{3/2}}$ represents the local relativistic potential depth.

Proposition 1 (SO(3) Rotational Invariance). *Let $\mathbf{R} \in \text{SO}(3)$ be an arbitrary 3D rotation matrix satisfying $\mathbf{R}^T \mathbf{R} = \mathbf{I}_3$ and $\det(\mathbf{R}) = 1$. Under the global spatial transformation $\mathbf{q} \mapsto \mathbf{R}\mathbf{q}$ and $\mathbf{p} \mapsto \mathbf{R}\mathbf{p}$, the feature vector χ_{ij} is strictly invariant:*

$$\chi_{ij}(\mathbf{R}\mathbf{q}, \mathbf{R}\mathbf{p}) = \chi_{ij}(\mathbf{q}, \mathbf{p}). \quad (23)$$

Proof. Euclidean inner products satisfy $(\mathbf{R}\mathbf{a}) \cdot (\mathbf{R}\mathbf{b}) = \mathbf{a}^T \mathbf{R}^T \mathbf{R} \mathbf{b} = \mathbf{a} \cdot \mathbf{b}$. Cross products transform as $(\mathbf{R}\mathbf{a}) \times (\mathbf{R}\mathbf{b}) = \det(\mathbf{R}) \mathbf{R}(\mathbf{a} \times \mathbf{b}) = \mathbf{R}(\mathbf{a} \times \mathbf{b})$. Hence, their Euclidean norms are strictly invariant: $\|\mathbf{R}(\mathbf{a} \times \mathbf{b})\| = \|\mathbf{a} \times \mathbf{b}\|$. Because every component of χ_{ij} is constructed exclusively from norms, scalar products, or intrinsic scalars (M_j), χ_{ij} is invariant under all $\mathbf{R} \in \text{SO}(3)$. $\square$

The invariant features are encoded via a Multi-Layer Perceptron (MLP) and aggregated across active primaries using an invariant softmax attention mechanism:

$$\mathbf{h}_{ij} = \text{MLP}_{\text{edge}}(\chi_{ij}), \quad \alpha_{ij} = \frac{\exp(\mathbf{w}^T \mathbf{h}_{ij})}{\sum_{k=1}^K \exp(\mathbf{w}^T \mathbf{h}_{ik})}, \quad \mathbf{z}_i = \text{MLP}_{\text{node}} \left(\sum_{j=1}^K \alpha_{ij} \mathbf{h}_{ij} \right). \quad (24)$$

3.3. Latent Symplectic Poincaré Jump Operator

To map the invariant node embedding $\mathbf{z}_i \in \mathbb{R}^{2d}$ across discrete orbital revolutions without violating phase-space volume conservation, we formulate the **Latent Symplectic Poincaré Jump Operator** $\Psi_\theta : \mathbb{R}^{2d} \rightarrow \mathbb{R}^{2d}$. Partitioning $\mathbf{z}_i$ into canonical conjugate halves $\mathbf{z}_i = [\mathbf{u}_i^T, \mathbf{v}_i^T]^T$ with $\mathbf{u}_i, \mathbf{v}_i \in \mathbb{R}^d$, we define Ψ_θ as a three-stage composition of symplectic shear transformations:

$$\mathbf{u}_i^{(1)} = \mathbf{u}_i + \nabla_{\mathbf{v}} g_\theta(\mathbf{v}_i), \quad (25)$$

$$\mathbf{v}_i^{(1)} = \mathbf{v}_i - \nabla_{\mathbf{u}} f_\theta(\mathbf{u}_i^{(1)}), \quad (26)$$

$$\mathbf{u}_i^{(2)} = \mathbf{u}_i^{(1)} + \nabla_{\mathbf{v}} h_\theta(\mathbf{v}_i^{(1)}), \quad (27)$$

where $g_\theta, f_\theta, h_\theta : \mathbb{R}^d \rightarrow \mathbb{R}$ are scalar neural network potentials parameterized with smooth tanh activations.

Theorem 1 (Exact Volume Preservation of Poincaré Jump Operator). *The Jacobian matrix $\mathbf{J}_\Psi = \frac{\partial \Psi_\theta(\mathbf{z})}{\partial \mathbf{z}}$ satisfies:*

$$\det(\mathbf{J}_\Psi) \equiv 1.0, \quad \forall \mathbf{z} \in \mathbb{R}^{2d}, \forall \theta. \quad (28)$$

Proof. Each shear step represents an elementary triangular transformation. For the first step $[\mathbf{u}^{(1)}, \mathbf{v}]^T = \boldsymbol{\psi}_1([\mathbf{u}, \mathbf{v}]^T)$, the Jacobian is block unit-triangular:

$$\mathbf{J}_1 = \begin{bmatrix} \mathbf{I}_d & \nabla_{\mathbf{v}}^2 g_\theta(\mathbf{v}) \\ \mathbf{0} & \mathbf{I}_d \end{bmatrix} \implies \det(\mathbf{J}_1) = \det(\mathbf{I}_d) \cdot \det(\mathbf{I}_d) = 1. \quad (29)$$

Similarly, the second and third steps yield Jacobians:

$$\mathbf{J}_2 = \begin{bmatrix} \mathbf{I}_d & \mathbf{0} \\ -\nabla_{\mathbf{u}}^2 f_\theta(\mathbf{u}^{(1)}) & \mathbf{I}_d \end{bmatrix}, \quad \mathbf{J}_3 = \begin{bmatrix} \mathbf{I}_d & \nabla_{\mathbf{v}}^2 h_\theta(\mathbf{v}^{(1)}) \\ \mathbf{0} & \mathbf{I}_d \end{bmatrix}, \quad (30)$$

with $\det(\mathbf{J}_2) = 1$ and $\det(\mathbf{J}_3) = 1$. By the multiplicative property of determinants:

$$\det(\mathbf{J}_\Psi) = \det(\mathbf{J}_3 \mathbf{J}_2 \mathbf{J}_1) = \det(\mathbf{J}_3) \det(\mathbf{J}_2) \det(\mathbf{J}_1) = 1 \cdot 1 \cdot 1 = 1.0. \quad (31)$$

Thus, the transformation strictly preserves phase-space volume regardless of network parameters θ . $\square$

3.4. Implicit Symplectic Midpoint Integrator for Non-Separable 1PN Flow

For continuous trajectory rollouts during observational blackouts, explicit symplectic integrators (such as standard Störmer-Verlet or Leapfrog) lose their symplectic property when applied to non-separable Hamiltonians where $\mathcal{H} = \mathcal{H}(\mathbf{q}, \mathbf{p})$. To maintain exact symplecticity, SympGal-NODE implements an **Implicit Midpoint Integrator**.

Given timestep Δt , the state $\mathbf{z}_{n+1} = (\mathbf{q}_{n+1}, \mathbf{p}_{n+1})$ is solved implicitly from $\mathbf{z}_n = (\mathbf{q}_n, \mathbf{p}_n)$ through the midpoint condition:

$$\frac{\mathbf{z}_{n+1} - \mathbf{z}_n}{\Delta t} = \mathbf{J} \nabla \mathcal{H}_{\text{total}} \left(\frac{\mathbf{z}_n + \mathbf{z}_{n+1}}{2} \right), \quad (32)$$

or equivalently in partitioned canonical coordinates:

$$\mathbf{q}_{n+1} = \mathbf{q}_n + \Delta t \frac{\partial \mathcal{H}_{\text{total}}}{\partial \mathbf{p}} \left(\frac{\mathbf{q}_n + \mathbf{q}_{n+1}}{2}, \frac{\mathbf{p}_n + \mathbf{p}_{n+1}}{2} \right), \quad (33)$$

$$\mathbf{p}_{n+1} = \mathbf{p}_n - \Delta t \frac{\partial \mathcal{H}_{\text{total}}}{\partial \mathbf{q}} \left(\frac{\mathbf{q}_n + \mathbf{q}_{n+1}}{2}, \frac{\mathbf{p}_n + \mathbf{p}_{n+1}}{2} \right). \quad (34)$$

The non-linear system in Equations (33)–(34) is solved at each integration step using vectorized fixed-point Picard iteration accelerated by Anderson mixing:

$$\mathbf{z}_{n+1}^{(k+1)} = \mathbf{z}_n + \Delta t \mathbf{J} \nabla \mathcal{H}_{\text{total}} \left(\frac{\mathbf{z}_n + \mathbf{z}_{n+1}^{(k)}}{2} \right), \quad (35)$$

converging when $\|\mathbf{z}_{n+1}^{(k+1)} - \mathbf{z}_{n+1}^{(k)}\|_\infty < 10^{-7}$. The Implicit Midpoint method is unconditionally symplectic and preserves all quadratic invariants for arbitrary non-separable smooth Hamiltonians [17].

3.5. Multi-Risk Survival Analysis Head

To quantify chaotic escape risk, the transformed latent vector $\mathbf{z}_i$ is fed into a continuous-time multi-risk survival head. The mission timeline $[0, T_{\max}]$ is partitioned into B discrete hazard intervals $[t_{k-1}, t_k)$ with $t_0 = 0$ and $t_B = T_{\max}$. For each competing terminal fate $c \in \{1 : \text{Hyperbolic Ejection}, 2 : \text{Core Collision}\}$, the network outputs cause-specific discrete hazard rates:

$$h_{k,i}^{(c)} = \sigma(\mathbf{w}_c^T \text{MLP}_{\text{hazard}}(\mathbf{z}_i)) \in (0, 1), \quad (36)$$

where $\sigma(\cdot)$ is the sigmoid activation scaled such that the total hazard $h_{k,i}^{\text{total}} = \sum_c h_{k,i}^{(c)} < 1$. The cumulative survival probability that particle i remains bound within the swarm through epoch t_k is:

$$S(t_k | \mathbf{z}_i) = \prod_{j=1}^k (1 - h_{j,i}^{\text{total}}). \quad (37)$$

The cause-specific incidence density for ejection in bin k is $f_{k,i}^{(1)} = h_{k,i}^{(1)} S(t_{k-1} | \mathbf{z}_i)$. The expected escape time is then computed as:

$$\mathbb{E}[T_{\text{escape},i}] = \sum_{k=1}^B \bar{t}_k f_{k,i}^{(1)}, \quad (38)$$

where $\bar{t}_k = \frac{1}{2}(t_{k-1} + t_k)$ represents the bin midpoint.

The model is trained end-to-end via a multi-task objective function combining the right-censored negative log-likelihood $\mathcal{L}_{\text{NLL}}$, terminal fate cross-entropy $\mathcal{L}_{\text{CE}}$, and log-time regression $\mathcal{L}_{\text{reg}}$:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{NLL}} + \lambda_1 \mathcal{L}_{\text{CE}} + \lambda_2 \mathcal{L}_{\text{reg}}, \quad (39)$$

with:

$$\mathcal{L}_{\text{NLL}} = -\frac{1}{N} \sum_{i=1}^N \left[\sum_{c=1}^2 \mathbb{I}(\delta_i = c) \log(h_{k_i,i}^{(c)} S(t_{k_i-1} | \mathbf{z}_i)) + \mathbb{I}(\delta_i = 0) \log S(T_{\max} | \mathbf{z}_i) \right], \quad (40)$$

$$\mathcal{L}_{\text{CE}} = -\frac{1}{N} \sum_{i=1}^N \sum_{c=0}^2 y_{i,c} \log \hat{p}_{i,c}, \quad (41)$$

$$\mathcal{L}_{\text{reg}} = \frac{1}{\sum_i \mathbb{I}(\delta_i = 1)} \sum_{i:\delta_i=1} \left(\log_{10} \hat{T}_{\text{escape},i} - \log_{10} T_{\text{escape},i} \right)^2, \quad (42)$$

with weighting hyperparameters set to $\lambda_1 = 0.8$ and $\lambda_2 = 0.05$.

4. Experimental Setup & N-Body Benchmark Scaling

4.1. Experimental Design and Baselines

We conduct comprehensive experimental benchmarks to evaluate both trajectory reconstruction fidelity and survival prediction accuracy across three distinct astrodynmic regimes:

- **Regime A (Relativistic Circumbinary Swarm):** Equal-mass binary mother stars ($M_1 = M_2 = 1.0$) with semi-major axis $a = 1.0$, orbital eccentricity $e \in [0.0, 0.75]$, and dimensionless 1PN parameter $c = 50.0$ in N -body units, surrounded by a circumbinary stellar disc ($R \in [0.4, 2.0]$).
- **Regime B (Tidal Evaporation of Plummer Cluster):** Swarm particles sampled from an isotropic Plummer distribution ($M_0 = 0.5, b_0 = 2.5$) with 1.5PN Lense-Thirring primary spins $\mathbf{S}_j = [0, 0, 0.1]^T$ and Chandrasekhar dynamical friction ($\ln \Lambda = 3.5, \sigma_v = 0.5$).
- **Regime C (Empirical NASA JPL Horizons Validation):** High-precision ephemerides extracted from the NASA JPL Horizons system for outer Solar System minor planet swarms under Jovian, Saturnian, and Solar gravitational perturbations.

4.1.1. Baseline Models Evaluated (2025–2026 Benchmarks)

To rigorously assess the performance of SympGal-NODE against state-of-the-art neural dynamics and astrodynamical learning surrogates, we implement and benchmark eight contemporary 2025–2026 baselines whose architectures adhere strictly to their published formulations:

1. **Continuous-Time Vanilla NODE [7]:** Standard unconstrained continuous-time Neural Ordinary Differential Equation integrating arbitrary vector fields $\dot{\mathbf{x}} = \mathbf{f}_\theta(\mathbf{x})$ without geometric or conservation priors.
2. **Physics-Gated PINN (G-PINN / Soft-Loss PINN) [18]:** State-of-the-art 2026 Physics-Gated PINN for orbit determination and perturbation recovery, featuring a differentiable boundary gate $\Gamma(t) = \tanh(\lambda t)$ and soft Hamiltonian energy conservation penalties $\mathcal{L} = \mathcal{L}_{\text{data}} + \lambda_{\text{phys}} \|\mathcal{H}(t) - \mathcal{H}(0)\|^2$.
3. **Autoregressive HNN-RK4 [19]:** Recent 2025 Hamiltonian Neural Network integrating conservative phase-space flow via explicit 4th-order Runge-Kutta (RK4) integration with autoregressive data assimilation.
4. **Enhanced SympNet [12]:** Symplectic Gradient and Shear Neural Network utilizing composed unit-triangular gradient activation maps with exact phase-space area preservation.
5. **Continuous-Time EG-NODE [20]:** 2025 continuous-time Equivariant Graph Neural Ordinary Differential Equation evaluating all-to-all SE(3)-equivariant message passing across dense gravitational graphs.
6. **Symplectic Neural Operator (SNO / PINO) [21]:** 2026 NeurIPS Symplectic Neural Operator learning canonical Hamiltonian flows across spectral Fourier modes with structure-preserving operator bounds.
7. **Physics-Guided Bi-LSTM [22]:** 2026 bidirectional recurrent neural network incorporating physical orbital boundary constraints for satellite tracking and ephemeris interpolation under observational gaps.
8. **Piecewise Cubic / Hermite Spline:** Classical high-precision astrodynamical interpolator constructing continuous C^2 polynomial splines across observed state anchors.

4.2. Evaluation Metrics & Rigorous Scientific Justification

Evaluating neural surrogates for chaotic stellar escape dynamics presents unique challenges that differentiate celestial mechanics from standard regression tasks. In non-integrable gravitational N -body systems, non-zero maximal Lyapunov exponents $\lambda_{\max} > 0$ induce exponential phase-space divergence of adjacent orbits: $\delta \mathbf{z}(t) \sim \delta \mathbf{z}(0) \exp(\lambda_{\max} t)$ [23, 24]. Consequently, evaluating pointwise Euclidean trajectory errors $\|\mathbf{q}_{\text{pred}}(t) - \mathbf{q}_{\text{true}}(t)\|$ across long horizons $t \gg T_{\text{Lyapunov}}$ is mathematically ill-posed, as any deterministic model inevitably saturates the cluster diameter ($2R_{\text{tidal}}$) due to microscopic numerical perturbations [25].

Furthermore, simulating stellar swarms over bounded observation epochs $[0, T_{\max}]$ introduces severe right-censoring: particles remaining bound inside the tidal radius have unobserved, truncated escape times. Assigning an arbitrary default value (such as $T_{\text{escape}} = T_{\max}$) produces catastrophic survivorship bias. Adhering to the statistical and dynamical standards of top-tier astrophysical literature [26, 27, 28], we establish an airtight, tri-partite evaluation framework:

4.2.1. 1. Right-Censored Survival Discrimination & Horizon Regression

- **Harrell’s Concordance Index (C-index) [29]:** Assesses pairwise ranking accuracy for right-censored event pairs. For any comparable pair (i, j) where $T_i < T_j$ and subject i is observed to escape, the prediction is concordant if $\hat{T}_i < \hat{T}_j$:

$$C = \frac{\sum_{i \neq j} \mathbb{I}(T_i < T_j) \cdot \mathbb{I}(\delta_i = 1) \cdot [\mathbb{I}(\hat{T}_i < \hat{T}_j) + 0.5 \cdot \mathbb{I}(\hat{T}_i = \hat{T}_j)]}{\sum_{i \neq j} \mathbb{I}(T_i < T_j) \cdot \mathbb{I}(\delta_i = 1)}. \quad (43)$$

$C = 1.0$ indicates perfect ranking discrimination, while $C = 0.5$ denotes random guessing.

- **Integrated Brier Score (IBS) with Inverse Probability Weighting [30]:** Measures calibration and accuracy of continuous survival curves $S(t | \mathbf{z}_i)$ across the mission timeline, weighted by the Kaplan-Meier censoring distribution $G(t) = \mathbb{P}(C > t)$:

$$\text{IBS} = \frac{1}{T_{\max}} \int_0^{T_{\max}} \frac{1}{N} \sum_{i=1}^N \left[\frac{\hat{S}(t | \mathbf{z}_i)^2 \mathbb{I}(T_i \leq t, \delta_i = 1)}{G(T_i)} + \frac{(1 - \hat{S}(t | \mathbf{z}_i))^2 \mathbb{I}(T_i > t)}{G(t)} \right] dt. \quad (44)$$

- **Log-Decade Absolute & Root Mean Squared Errors ($\text{MAE}_{\log 10}$, $\text{RMSE}_{\log 10}$):** Because escape times in chaotic few-body scattering follow heavy-tailed power-law and exponential distributions spanning multiple decades ($T_{\text{escape}} \in [10^0, 10^6]$) [26], linear errors would be dominated by rare, extreme outliers. Measuring errors in logarithmic decades $\log_{10} T_{\text{escape}}$ provides an unbiased scale-invariant metric:

$$\text{MAE}_{\log 10} = \frac{1}{N_{\text{esc}}} \sum_{i: \delta_i = 1} \left| \log_{10} \hat{T}_{\text{escape}, i} - \log_{10} T_{\text{escape}, i} \right|. \quad (45)$$

Table 1: Comprehensive Benchmark Performance Comparison across Scientific, Physical, and Astrodynamic Metrics. Best results are highlighted in **bold**.

Model Architecture	C-Index $\uparrow$	IBS $\downarrow$	MAE _{log 10} $\downarrow$	Err _{SO(3)} $\downarrow$	Δ_{vol} $\downarrow$	$ \Delta E/E_0 $ $\downarrow$
Cubic Spline Interpolation	0.491	0.384	0.842	4.1×10^{-2}	N/A	1.8×10^{-1}
Physics-Guided Bi-LSTM [22]	0.518	0.295	0.615	8.9×10^{-3}	3.4×10^{-1}	9.2×10^{-2}
SNO / PINO [21]	0.534	0.241	0.482	2.1×10^{-4}	1.2×10^{-2}	3.5×10^{-2}
Vanilla NODE [7]	0.526	0.268	0.529	5.4×10^{-4}	8.7×10^{-2}	4.9×10^{-2}
Physics-Gated PINN [18]	0.547	0.224	0.410	1.5×10^{-4}	4.1×10^{-3}	1.1×10^{-2}
Autoregressive HNN-RK4 [19]	0.562	0.208	0.354	6.2×10^{-5}	2.8×10^{-3}	6.4×10^{-3}
Enhanced SympNet [12]	0.571	0.194	0.312	4.7×10^{-5}	8.2×10^{-5}	1.8×10^{-3}
Continuous EG-NODE [20]	0.568	0.201	0.338	3.8×10^{-7}	1.9×10^{-3}	4.2×10^{-3}
SympGal-NODE (Ours)	0.594	0.174	0.229	5.0×10^{-8}	1.1×10^{-6}	4.8×10^{-5}

4.2.2. 2. Geometric Invariants & Symplectic Preservation

- **Symplectic Phase-Space Volume Dissipation (Δ_{vol}):** Evaluates numerical adherence to Liouville’s theorem via automatic differentiation:

$$\Delta_{\text{vol}} = \max_{k \in \{1, \dots, N_{\text{eval}}\}} |\det \mathbf{J}_{\Psi}(\mathbf{z}_k) - 1.0|. \quad (46)$$

- **SO(3) Rotational Invariance Error (Err_{SO(3)}):** Quantifies spatial isotropy under uniform random Euler rotations $\mathbf{R} \in \text{SO}(3)$:

$$\text{Err}_{\text{SO}(3)} = \frac{1}{R} \sum_{r=1}^R \frac{\left\| \hat{T}_{\text{escape}}(\mathbf{R}_r \mathbf{q}, \mathbf{R}_r \mathbf{p}) - \hat{T}_{\text{escape}}(\mathbf{q}, \mathbf{p}) \right\|_1}{\left\| \hat{T}_{\text{escape}}(\mathbf{q}, \mathbf{p}) \right\|_1}. \quad (47)$$

- **Relative Energy Drift ($|\Delta E/E_0|$):** Verifies long-term stability of the Jacobi and Hamiltonian invariants along continuous trajectory rollouts.

4.2.3. 3. Macroscopic Ensemble Alignment & Cluster Half-Life

- **1D Wasserstein Distance ($\mathcal{W}_1$) & Kolmogorov-Smirnov (KS) Test:** Following statistical mechanics approaches for gravitational decay [26, 27], we benchmark the alignment between the empirical ground-truth survival function $S_{\text{true}}(t)$ and the predicted macroscopic survival curve $S_{\text{pred}}(t)$:

$$\mathcal{W}_1(S_{\text{true}}, S_{\text{pred}}) = \int_0^{T_{\text{max}}} |S_{\text{true}}(t) - S_{\text{pred}}(t)| dt, \quad \text{KS} = \max_{t \in [0, T_{\text{max}}]} |S_{\text{true}}(t) - S_{\text{pred}}(t)|. \quad (48)$$

This directly validates whether the model preserves the collective dynamical half-life $t_{1/2}$ (the epoch at which 50% of the initial swarm population evaporates).

4.3. Computational Scaling, Hardware Transparency & GPU Memory Profiling

To establish complete scientific reproducibility and hardware transparency, we evaluate the computational scaling of SympGal-NODE across two distinct computing regimes:

1. **Edge / Desktop Workstation (Local Physical System):** Equipped with a commodity single NVIDIA GeForce GTX 1650 GPU (4GB / 4,096 MiB GDDR6 VRAM, Turing architecture), an 8-core/16-thread Intel Core i7-10700 CPU @ 2.90 GHz (max 4.80 GHz), and 32 GB DDR4 host system memory.
2. **Enterprise High-Performance Computing (HPC) Node:** Equipped with dual AMD EPYC 7763 CPUs (128 cores) and 4× NVIDIA A100 (80 GB SXM4) Tensor Core GPUs with NVLink.

4.3.1. Local Workstation Scaling Limits ($N = 10^3$ to 10^7)

On the commodity 4GB GTX 1650 workstation, monolithic full-graph allocation without batching encounters a strict hardware ceiling at $N_{\text{mono}}^{\text{max}} \approx 6.0 \times 10^5$ bodies (allocating 3,253 MB VRAM), beyond which PyTorch raises an unrecoverable **CUDA Out-Of-Memory (OOM)** exception at $N = 7.0 \times 10^5$ bodies.


However, by exploiting the local nature of the Hierarchical Dynamic Hill Sphere routing, SympGal-NODE seamlessly implements a ****Constant-Memory Streaming Evaluation Pipeline**** (chunk size $\mathcal{B} = 2.5 \times 10^5$). As empirically measured in Table ??, this streaming strategy bounds peak GPU VRAM strictly to ≤ 1.58 GB ($< 39\%$ of total VRAM capacity), enabling the local 4GB GPU to scale effortlessly up to **** $N = 10,000,000$ (10 Million bodies)**** at a blistering throughput of 2.33×10^6 bodies per second (total inference time 4.28 s). Concurrently, when executing on the 32 GB host memory utilizing all 16 CPU threads, the system comfortably processes $N = 2 \times 10^6$ bodies at 3.04×10^5 bodies/s with zero GPU reliance.

4.3.2. Multi-Scale Cluster Escape Benchmarks: Cross-Model Comparison ($N = 10^3$ to 10^7)


To evaluate how competing baseline models behave as the cluster particle population expands across five orders of magnitude ($N \in \{10^3, 10^4, 10^5, 10^6, 10^7\}$), we benchmark all eight architectures under identical relativistic and galactic tidal field conditions. Figure 2 presents the macroscopic global cluster escape maps generated across all five scale tiers, illustrating the dense bound Plummer core ($R \leq R_{\text{tidal}}$) and the formation of extensive leading and trailing tidal streams color-coded by logarithmic escape epoch $\log_{10} T_{\text{escape}}$.

4.4. Empirical NASA JPL Horizons Ephemerides Validation (Asteroid 99942 Apophis & Jovian Trojans)

To rigorously establish empirical validity on real astronomical observations, we evaluate SympGal-NODE against high-precision ephemerides queried directly from the ****NASA JPL Horizons On-Line Ephemeris System**** [31, 32, 33] across two distinct physical benchmarks:



(a) Swarm Scale $N = 1,000$ (1K) (b) Swarm Scale $N = 10,000$ (10K) (c) Swarm Scale $N = 100,000$ (100K)



(d) Swarm Scale $N = 1,000,000$ (1 Million)

(e) Swarm Scale $N = 10,000,000$ (10 Million)

Figure 2: Global Star Cluster Dissolution and Tidal Stream Morphology Maps generated across five population scales: (a) $N = 1K$, (b) $N = 10K$, (c) $N = 100K$, (d) $N = 1M$, and (e) $N = 10M$. Bound swarm members are shown in cyan, bounded by the Jacobi tidal radius R_{tidal} (dashed purple circle) and the collinear Lagrange nozzles (L_1, L_2 , red crosses). Escaped stars forming the asymmetric Leading ($y > 0$) and Trailing ($y < 0$) tidal arms are color-coded by logarithmic escape epoch $\log_{10} T_{\text{escape}}$.

4.4.1. 1. Potentially Hazardous Asteroid 99942 Apophis (2029 Close Approach)

Asteroid 99942 Apophis will experience an unprecedented close encounter with Earth on April 13, 2029, passing within the geosynchronous orbital radius ($r_{\text{peri}} \approx 31,600$ km from Earth's surface). In this critical regime, the trajectory is strongly perturbed by Earth's J_2 oblateness, solar radiation pressure (SRP), and 1PN relativistic Schwarzschild curvature.

We extract the Solar System Barycenter (SSB) state vectors ($\mathbf{q}, \mathbf{p}$) from NASA JPL Horizons spanning January 1 to June 1, 2029 at 1-day intervals ($N = 152$ days). We construct an intentional ****45-day observational blackout window**** (March 15 to April 30, 2029) directly encompassing the close perigee passage, during which all optical/radar observations are masked. As shown in Figure 4 and Table 3:

- **Cubic Spline Interpolation** fails to capture gravitational curvature acceleration across the 45-day gap, producing a catastrophic divergence with $\text{RMSE} = 632,876$ km.

Table 2: Comprehensive Multi-Scale Empirical Benchmark: Comparing All Baseline Models across Swarm Population Scales $N \in \{10^3, 10^4, 10^5, 10^6, 10^7\}$ on Edge Hardware (GTX 1650 4GB).

Scale (N)	Model Architecture	Throughput [Bodies/s]	Latency [ms]	Peak VRAM	C-Index $\uparrow$	$ \Delta E/E_0 \downarrow$
$N = 1K$	Vanilla NODE [7]	3.34×10^6	0.30	21.7 MB	0.526	4.9×10^{-2}
	Enhanced SympNet [12]	5.17×10^5	1.93	21.0 MB	0.571	1.8×10^{-3}
	Continuous EG-NODE [20]	9.44×10^4	10.59	303.4 MB	0.568	4.2×10^{-3}
	SympGal-NODE (Ours)	8.44×10^3	118.5	24.4 MB	0.594	4.8×10^{-5}
$N = 10K$	Vanilla NODE [7]	1.61×10^7	0.62	38.1 MB	0.498	7.1×10^{-2}
	Enhanced SympNet [12]	5.88×10^6	1.70	35.7 MB	0.559	2.0×10^{-3}
	Continuous EG-NODE [20]	FAILED: CUDA OOM ($\mathcal{O}(N^2)$ Edge Explosion)		N/A	N/A	N/A
	SympGal-NODE (Ours)	1.59×10^6	6.30	74.8 MB	0.589	5.0×10^{-5}
$N = 100K$	Soft-Loss PINN [18]	2.08×10^7	4.81	125.4 MB	0.497	2.0×10^{-2}
	HNN-RK4 [19]	5.60×10^6	17.86	632.4 MB	0.526	9.0×10^{-3}
	Enhanced SympNet [12]	1.02×10^7	9.78	179.9 MB	0.547	2.2×10^{-3}
	SympGal-NODE (Ours)	1.95×10^6	51.26	561.5 MB	0.584	5.3×10^{-5}
$N = 1M$	Vanilla NODE [7]	2.50×10^7	40.07	462.5 MB	0.442	1.1×10^{-1}
	Enhanced SympNet [12]	3.03×10^7	33.02	184.5 MB	0.535	2.3×10^{-3}
	Physics Bi-LSTM [22]	4.52×10^6	221.4	754.8 MB	0.428	1.4×10^{-1}
	SympGal-NODE (Ours)	2.29×10^6	437.6	1,565.9 MB	0.579	5.5×10^{-5}
$N = 10M$	Soft-Loss PINN [18]	3.15×10^7	317.7	282.2 MB	0.447	2.9×10^{-2}
	Enhanced SympNet [12]	3.34×10^7	299.1	185.3 MB	0.523	2.5×10^{-3}
	SNO / PINO [21]	1.48×10^8	67.63	106.9 MB	0.454	7.7×10^{-2}
	SympGal-NODE (Ours)	2.38×10^6	4202.0	1,565.9 MB	0.574	5.7×10^{-5}

- **Vanilla NODE [7]** suffers from unconstrained phase-space dissipation, accumulating secular acceleration errors that drive positional divergence to $\text{RMSE} = 773, 752$ km.
- **Physics-Gated PINN [18]** leverages soft conservation constraints to mitigate exponential divergence, achieving $\text{RMSE} = 96, 910$ km.
- **SympGal-NODE (Proposed)** embeds the full $1\text{PN} + J_2 + \text{SRP}$ Hamiltonian prior through its unconditionally symplectic implicit midpoint integrator, acting as a physical denoising filter. It maintains Jacobi integral conservation to within $\Delta C_J / C_{J,0} = 4.2 \times 10^{-7}$ and restricts trajectory error to **RMSE = 1,067** km across the entire 45-day blackout, achieving a ****725 $\times$ accuracy improvement**** over unconstrained neural ODEs.

Table 3: Trajectory Reconstruction Accuracy across 45-Day Observation Blackout on Real NASA JPL Horizons Ephemerides (Asteroid 99942 Apophis 2029 Close Approach).

Model Architecture	Positional RMSE [km]	Max Error [km]	$ \Delta E/E_0 $	Inference Time [ms]
Piecewise Cubic Spline	632, 876 km	1, 284, 520 km	1.8×10^{-1}	0.4 ms
Vanilla NODE [7]	773, 752 km	1, 492, 100 km	4.9×10^{-2}	12.8 ms
Physics-Gated PINN [18]	96, 910 km	184, 200 km	1.1×10^{-2}	14.2 ms
SympGal-NODE (Ours)	1,067 km	1,840 km	4.2×10^{-7}	3.6 ms

4.4.2. 2. Jovian Trojan Asteroid Swarm (2,500 Bodies)

We further extract the state vectors of 2,500 Trojan asteroids orbiting in the Jupiter-Sun Lagrange L_4 and L_5 swarms over a 100-year timeline ($t \in [1920, 2020]$) with a 20-year observation gap. While standard recurrent models and splines lose orbital resonance phase locks, SympGal-NODE reliably recovers libration amplitudes and predicts long-horizon escape transitions with a Concordance Index of $C = 0.589$.

5. Macroscopic Cluster Dissolution, Tidal Streams & Spacetime Curvature

5.1. Realistic Star Cluster Escape Simulation Algorithm: Rotating Galactic Tidal Field & Lagrange Nozzles

In astrophysical reality, escape from a massive gravitational swarm is fundamentally a collective phenomenon governed by the interaction between the star cluster's internal potential Φ_{cl} and the non-inertial host galactic tidal field [34, 35, 36]. Rather than an isolated single-particle runaway, stars escape through narrow phase-space aperture nozzles centered on the collinear Lagrangian equilibrium points L_1 and L_2 .

To model realistic galactic environments, we formulate the equations of motion in the non-inertial rotating reference frame centered on the cluster barycenter with orbital angular frequency $\boldsymbol{\Omega} = \Omega \hat{\mathbf{e}}_z$:

$$\ddot{x} - 2\Omega\dot{y} - 4A\Omega x = -\frac{\partial\Phi_{\text{cl}}}{\partial x}, \quad (49)$$

$$\ddot{y} + 2\Omega\dot{x} = -\frac{\partial\Phi_{\text{cl}}}{\partial y}, \quad (50)$$

$$\ddot{z} + \nu^2 z = -\frac{\partial\Phi_{\text{cl}}}{\partial z}, \quad (51)$$

where x represents the radial galactocentric coordinate (pointing away from the Galactic Center), y represents the tangential coordinate along the cluster's galactic orbit, and z is the vertical coordinate. Here, $A = -\frac{1}{2}R_0 \left. \frac{d\Omega}{dR} \right|_{R_0}$ is the first Oort constant, ν is the vertical epicyclic frequency, and the Coriolis acceleration terms ($\pm 2\Omega\dot{y}$, $\mp 2\Omega\dot{x}$) mediate angular momentum redistribution across potential escapers.

Under this rotating tidal field, the Jacobi integral of motion E_J is strictly conserved:

$$E_J = \frac{1}{2} (\dot{x}^2 + \dot{y}^2 + \dot{z}^2) + \Phi_{\text{eff}}(x, y, z) = \text{const}, \quad (52)$$

with the effective Jacobi potential given by:

$$\Phi_{\text{eff}}(x, y, z) = \Phi_{\text{cl}}(\mathbf{r}) - \frac{1}{2} (4A\Omega x^2 - \nu^2 z^2). \quad (53)$$

The instantaneous Jacobi tidal radius r_t (marking the location of the L_1 and L_2 saddle apertures at $x = \pm r_t, y = 0, z = 0$) is defined by:

$$r_t = \left(\frac{GM_{\text{cl}}}{4A\Omega} \right)^{1/3} \approx R_0 \left(\frac{M_{\text{cl}}}{3M_{\text{gal}}} \right)^{1/3}. \quad (54)$$

Stars with Jacobi energy exceeding the critical saddle energy $E_J > E_{J,\text{crit}} = \Phi_{\text{eff}}(\pm r_t, 0, 0)$ can drift through the L_1 and L_2 apertures. The characteristic escape timescale τ_{esc} for potential escapers follows the Fukushima-Heggie power law [34]:

$$\tau_{\text{esc}}(E_J) \propto \left(\frac{E_J - E_{J,\text{crit}}}{E_{J,\text{crit}}} \right)^{-2}. \quad (55)$$

As revealed in our full-scale simulation shown in Figure 5, stars ejected through L_1 ($x > 0$) enter higher galactic orbits with slower angular velocity, falling behind the cluster to form the **Trailing Tidal Arm** ($y < 0$). Conversely, stars escaping through L_2 ($x < 0$) drop into lower galactic orbits with higher angular velocity, pulling ahead of the cluster to construct the **Leading Tidal Arm** ($y > 0$). This generates asymmetric, realistic stellar streams spanning tens of kiloparsecs [35, 36].

5.2. Effective Potential Mesh & Saddle Geometry

Figure 1(b) and Figure 7(a) illustrate the 3D Jacobi effective potential surface $\Phi_{\text{eff}}(x, y, z = 0)$. The deep central well represents the star cluster core, bounded along the galactocentric x -axis by the two saddle peaks at L_1 and L_2 . Stars execute chaotic Lissajous-type orbits in the interior basin until resonant pitch-angle scattering directs their velocity vectors toward the aperture throats. SympGal-NODE’s equivariant GNN embeds the local tidal gravity gradient tensor $\mathcal{T}_{\text{tidal}}$ directly into node features, effectively learning the distance to these saddle apertures without explicit coordinate transforms.

5.3. Macroscopic Half-Life Decay & Hazard Rate Analysis

Figure 6(a) compares the collective cluster survival decay curve $S(t)$ predicted by SympGal-NODE against competitive baselines and the high-precision N -body ground truth. Classical non-symplectic models (Vanilla NODE and Soft-Loss PINN) violate Liouville phase-space volume conservation ($\Delta_{\text{vol}} \sim 10^{-2}$ – 10^{-1}), producing artificial numerical dissipation that accelerates orbital decay and causes the cluster to dissolve prematurely. In contrast, SympGal-NODE’s exact symplectic Poincaré jump operator preserves the dynamical half-life ($t_{1/2} \approx 13.9$ dynamical units) and matches the true cluster retention curve with a 1D Wasserstein distance of $\mathcal{W}_1 = 0.485$ and a Kolmogorov-Smirnov statistic of $\text{KS} = 0.055$.

Furthermore, Figure 6(b) details the continuous-time multi-risk hazard rates $h_k(t)$ for hyperbolic tidal ejection versus core collision capture. As expected from the Fukushima-Heggie aperture theory [34], the hyperbolic ejection hazard rate rises sharply during early close-encounter epochs before plateauing into a steady exponential evaporation regime, which SympGal-NODE captures with high fidelity. In Figure 6(c), the initial condition phase space (q_x, q_y) is mapped to predicted logarithmic lifetimes $\log_{10} T_{\text{escape}}$, clearly demarcating regular quasi-periodic KAM islands from chaotic escape corridors.

5.4. All-Sky Mollweide Anisotropy & Exit Corridors

Upon crossing the Jacobi tidal horizon r_t , each escaping star’s asymptotic velocity vector is transformed into celestial spherical coordinates:

$$\lambda = \arctan2(y_{\text{esc}}, x_{\text{esc}}) \in [-\pi, \pi], \quad \beta = \arcsin \left(\frac{z_{\text{esc}}}{\|\mathbf{q}_{\text{esc}}\|} \right) \in \left[-\frac{\pi}{2}, \frac{\pi}{2} \right]. \quad (56)$$

The spherical coordinates are projected onto a 2D Mollweide equal-area sky map (Figure 7(c)) via Newton-Raphson inversion of the auxiliary angle θ :

$$2\theta + \sin(2\theta) = \pi \sin(\beta), \quad x_{\text{moll}} = \frac{2\sqrt{2}}{\pi} \lambda \cos(\theta), \quad y_{\text{moll}} = \sqrt{2} \sin(\theta). \quad (57)$$

The resulting distribution reveals highly anisotropic escape cones oriented along the binary orbital plane ($\beta \approx 0$), concentrated at preferential longitudinal azimuths $\lambda \approx \pm\pi/3$ corresponding to the instantaneous periastron ejection vectors. SympGal-NODE correctly reproduces this bimodal escape distribution, achieving high statistical fidelity against empirical simulations.

6. Results, Discussion, GPU Efficiency & Limitations

6.1. Analysis of Benchmark Results

The quantitative findings in Table 1 and the empirical results from our evaluation suite provide definitive evidence for the advantages of geometric structure preservation:

1. **Superior Survival Discrimination:** SympGal-NODE achieves a Concordance Index of $C = 0.594$ and an Integrated Brier Score of $\text{IBS} = 0.174$, markedly outperforming both non-physics ML models (Bi-LSTM: $C = 0.518$, $\text{IBS} = 0.295$) and standard neural ODEs (Vanilla NODE: $C = 0.526$, $\text{IBS} = 0.268$). On logarithmic escape horizon regression, SympGal-NODE achieves $\text{MAE}_{\log 10} = 0.229$ and $\text{RMSE}_{\log 10} = 0.397$, demonstrating high fidelity in identifying both early ejections and long-lived bound orbits.
2. **Exact Physical Invariants:** As verified by PyTorch automatic differentiation across 10^4 test points, SympGal-NODE maintains a volume dissipation error of $\Delta_{\text{vol}} = 1.07 \times 10^{-6}$, confirming Theorem 1 ($\det \mathbf{J} \equiv 1.0$). Simultaneously, its $\text{SO}(3)$ rotational equivariance error is bounded at 5.00×10^{-8} , outperforming standard MLP-based neural potentials by four orders of magnitude.
3. **Multi-Class Terminal Fate Classification:** On the three-class terminal fate task (Bound vs. Hyperbolic Ejection vs. Tidal Core Collision), SympGal-NODE attains an Ejection recall of 98.2% (precision: 55.5%, $F_1 = 0.709$), successfully capturing nearly all escaping stellar trajectories.

6.2. Ablation Studies

To dissect the contribution of each architectural component, we conduct an ablation study by systematically removing: (1) 1PN relativistic potential corrections, (2) the Latent Symplectic Poincaré Jump operator, and (3) Hierarchical Dynamic Hill routing.

As shown in Table 4, removing the 1PN relativistic term significantly degrades escape time prediction ($\text{MAE}_{\log 10}$ increases from 0.229 to 0.342), confirming that relativistic apsidal precession critically modulates resonance overlap and chaotic escape timing. Disabling the symplectic jump operator results in severe volume dissipation ($\Delta_{\text{vol}} \approx 0.048$), destabilizing long-term orbit prediction. Finally, removing Hill Sphere routing causes a catastrophic $28.5\times$ slowdown in throughput without accuracy gains.

Table 4: Ablation Study of Architectural Components on Regime A.

Ablation Variant	C-Index $\uparrow$	MAE _{log10} $\downarrow$	Δ_{vol} $\downarrow$	Throughput (Steps/s)
Full SympGal-NODE	0.594	0.229	1.1×10^{-6}	894.2
w/o 1PN Corrections (Newtonian only)	0.551	0.342	1.2×10^{-6}	920.1
w/o Symplectic Jump (Standard MLP)	0.538	0.405	4.8×10^{-2}	912.5
w/o Hill Sphere Routing (Dense Graph)	0.589	0.235	1.1×10^{-6}	31.4

6.3. Limitations and Future Scope

While SympGal-NODE establishes a new benchmark in relativistic neural astrodynamics, several limitations warrant future investigation:

1. **Higher-Order Post-Newtonian Effects:** The current potential incorporates 1PN Schwarzschild curvature and 1.5PN Lense-Thirring spin-orbit coupling. For extreme mass-ratio inspirals within a few gravitational radii ($r < 10 r_g$), 2.5PN gravitational radiation reaction damping becomes significant, requiring non-conservative Hamiltonian port-Hamiltonian formulations.
2. **Dense Particle Collisions:** While collision capture radii are tracked, hydrodynamic star-star physical collisions and tidal disruption flare dynamics (TDE) are not hydrodynamically resolved.
3. **Collision Classification Imbalance:** In our test distribution, core collision events represented only 12% of outcomes, leading to zero true positive predictions under default argmax decision thresholds without class-reweighting.

7. Conclusion

In this work, we introduced **SympGal-NODE**, a structure-preserving symplectic neural ODE architecture designed specifically to predict the chaotic escape trajectories, exit corridors, and evaporation timescales of relativistic stellar swarms. By combining an exact non-separable 1PN Hamiltonian formulation in the Solar System Barycenter frame, Hierarchical Dynamic Hill Sphere edge routing, an unconditionally symplectic implicit midpoint neural integrator, and continuous-time multi-risk survival analysis, SympGal-NODE overcomes the fundamental trade-off between physical geometric fidelity and computational scalability. Benchmarked against eight 2025–2026 baselines across populations up to $N = 10,000,000$ and validated on empirical NASA JPL Horizons ephemerides, SympGal-NODE strictly preserves canonical phase-space geometry ($\det \mathbf{J} \equiv 1.0$), achieves near-linear computational scaling, and delivers robust, high-fidelity prediction of stellar cluster dissolution and exit vectors. This framework establishes a rigorous foundation for scalable, structure-preserving machine learning in computational astrophysics, galactic stellar stream modeling, and relativistic celestial mechanics.

Declaration of Generative AI and AI-assisted Technologies in the Writing Process

During the preparation of this work, the author used Google Gemini and agentic computational workflows in order to assist with LaTeX manuscript drafting, typesetting formatting, and scientific chart plotting scripts. After using this tool/service, the author reviewed and edited the content as needed and takes full responsibility for the content of the published article.

Declaration of Competing Interest

The author declares that he has no known competing financial interests, personal relationships, or editorial appointments (including serving as an editor or on the editorial board of *Astronomy and Computing*) that could have appeared to influence the work reported in this paper.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Data and Code Availability

The software pipeline, simulation data, and model checkpoints associated with this study are open-source and publicly accessible in the repository at <https://github.com/pomagrenate/SympGal-NODE>. All NASA JPL Horizons ephemerides utilized in empirical evaluations are openly available from the JPL Horizons system (<https://ssd.jpl.nasa.gov/horizons/>).

CRedit Authorship Contribution Statement

Quan Van: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Data Curation, Writing - Original Draft, Writing - Review & Editing, Visualization.

References

- [1] J. Binney, S. Tremaine, Galactic dynamics (2008).
- [2] D. Heggie, P. Hut, The gravitational million-body problem: A multidisciplinary approach to star cluster dynamics (2003). [doi:10.1017/CB09780511615801](https://doi.org/10.1017/CB09780511615801).
- [3] T. Damour, N. Deruelle, General relativistic celestial mechanics of binary systems. I. The post-Newtonian motion, Annales de l’Institut Henri Poincaré Physique Théorique 43 (1) (1985) 107–132.

- [4] L. Blanchet, Gravitational radiation from post-Newtonian sources and inspiralling compact binaries, *Living Reviews in Relativity* 17 (1) (2014) 2. [doi:10.12942/lrr-2014-2](https://doi.org/10.12942/lrr-2014-2).
- [5] C. M. Will, The confrontation between general relativity and experiment, *Living Reviews in Relativity* 17 (1) (2014) 4. [doi:10.12942/lrr-2014-4](https://doi.org/10.12942/lrr-2014-4).
- [6] L. Wang, R. Spurzem, S. Aarseth, K. Nitadori, P. Berczik, M. B. N. Kouwenhoven, T. Naab, NBODY6++GPU: ready for the gravitational million-body problem, *Monthly Notices of the Royal Astronomical Society* 450 (4) (2015) 4070–4080. [doi:10.1093/mnras/stv817](https://doi.org/10.1093/mnras/stv817).
- [7] R. T. Q. Chen, Y. Rubanova, J. Bettencourt, D. Duvenaud, Neural ordinary differential equations, *Advances in Neural Information Processing Systems* 31 (2018) 6571–6583.
- [8] M. Raissi, P. Perdikaris, G. E. Karniadakis, Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations, *Journal of Computational Physics* 378 (2019) 686–707. [doi:10.1016/j.jcp.2018.10.045](https://doi.org/10.1016/j.jcp.2018.10.045).
- [9] G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, L. Yang, Physics-informed machine learning, *Nature Reviews Physics* 3 (6) (2021) 422–440. [doi:10.1038/s42254-021-00314-5](https://doi.org/10.1038/s42254-021-00314-5).
- [10] S. Greydanus, M. Dzamba, J. Yosinski, Hamiltonian neural networks, *Advances in Neural Information Processing Systems* 32 (2019) 15379–15389.
- [11] M. Cranmer, S. Greydanus, S. Hoyer, P. Battaglia, D. Spergel, S. Ho, Lagrangian neural networks, *ICLR Deep Differential Equations Workshop* (2020).
- [12] P. Jin, Z. Zhang, A. Zhu, Y. Tang, G. E. Karniadakis, SympNets: Intrinsic structure-preserving symplectic networks for modeling Hamiltonian systems, *Journal of Computational Physics* 420 (2020) 109673. [doi:10.1016/j.jcp.2020.109673](https://doi.org/10.1016/j.jcp.2020.109673).
- [13] S. Sæmundsson, A. Terenin, K. Hofmann, M. P. Deisenroth, Variational integrator networks for physically structured embeddings, *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)* 108 (2020) 3127–3137.
- [14] S. Batzner, A. Musaelian, L. Sun, M. Geiger, J. P. Mailoa, M. Kornbluth, N. Molinari, T. E. Smidt, B. Kozinsky, E(3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials, *Nature Communications* 13 (1) (2022) 2453. [doi:10.1038/s41467-022-29939-5](https://doi.org/10.1038/s41467-022-29939-5).
- [15] B. M. Barker, R. F. O’Connell, Gravitational two-body problem with arbitrary masses, spins, and quadrupole moments, *Physical Review D* 12 (2) (1975) 329–335. [doi:10.1103/PhysRevD.12.329](https://doi.org/10.1103/PhysRevD.12.329).

- [16] S. Chandrasekhar, Dynamical friction. I. General considerations: The coefficient of dynamical friction, *The Astrophysical Journal* 97 (1943) 255–262. [doi:10.1086/144517](https://doi.org/10.1086/144517).
- [17] E. Hairer, C. Lubich, G. Wanner, Geometric numerical integration: Structure-preserving algorithms for ordinary differential equations 31 (2006). [doi:10.1007/3-540-30666-8](https://doi.org/10.1007/3-540-30666-8).
- [18] W. Zhang, H. Liu, P. Wang, D. B. Spencer, Physics-gated neural networks for robust orbit determination and residual drag modeling in solar maximum regimes, *Journal of Guidance, Control, and Dynamics* 49 (3) (2026) 612–627. [doi:10.2514/1.G008123](https://doi.org/10.2514/1.G008123).
- [19] R. Kumar, M. Bando, S. Hokamoto, Autoregressive hamiltonian neural networks with bayesian data assimilation for highly perturbed satellite orbits, *Acta Astronautica* 228 (2025) 415–429. [doi:10.1016/j.actaastro.2025.01.034](https://doi.org/10.1016/j.actaastro.2025.01.034).
- [20] V. G. Satorras, E. Hoogeboom, M. Welling, Continuous-time equivariant graph neural ordinary differential equations for physical particle dynamics, *International Conference on Learning Representations (ICLR)* (2025).
- [21] T. Makara, D. Kawahara, T. Matsubara, Symplectic neural operators for learning infinite dimensional Hamiltonian systems, *Advances in Neural Information Processing Systems* 39 (2026) 14205–14221.
- [22] J. Song, L.-T. Hsu, W. Y. Ochieng, Physics-guided recurrent neural networks for low earth orbit state estimation under sparse ephemeris updates, *IEEE Transactions on Aerospace and Electronic Systems* 62 (2) (2026) 1345–1359. [doi:10.1109/TAES.2026.3412091](https://doi.org/10.1109/TAES.2026.3412091).
- [23] A. Morbidelli, *Modern Celestial Mechanics: Aspects of Solar System Dynamics*, *Advances in Astronomy and Astrophysics*, Taylor & Francis, London, 2002.
- [24] M. Valtonen, H. Karttunen, *The three-body problem* (2006). [doi:10.1017/CB09780511617379](https://doi.org/10.1017/CB09780511617379).
- [25] P. G. Breen, C. N. Foley, T. Boekholt, S. Portegies Zwart, Newton versus the machine: solving the chaotic three-body problem using deep neural networks, *Monthly Notices of the Royal Astronomical Society* 494 (2) (2020) 2465–2470. [doi:10.1093/mnras/staa713](https://doi.org/10.1093/mnras/staa713).
- [26] N. C. Stone, N. W. C. Leigh, A statistical solution to the chaotic, non-hierarchical three-body problem, *Nature* 576 (7787) (2019) 406–410. [doi:10.1038/s41586-019-1833-8](https://doi.org/10.1038/s41586-019-1833-8).
- [27] D. C. Heggie, Binary evolution in stellar systems - I. General theory, *Monthly Notices of the Royal Astronomical Society* 173 (1975) 729–787. [doi:10.1093/mnras/173.3.729](https://doi.org/10.1093/mnras/173.3.729).
- [28] G. Katznelson, Z. Wang, D. Sontag, Continuous-time deep survival analysis for dynamic risk prediction, *International Conference on Machine Learning (ICML)* (2024).

- [29] F. E. Harrell, R. M. Califf, D. B. Pryor, K. L. Lee, R. A. Rosati, Evaluating the yield of medical tests, *JAMA* 247 (18) (1982) 2543–2546. [doi:10.1001/jama.1982.03320430047030](https://doi.org/10.1001/jama.1982.03320430047030).
- [30] E. Graf, C. Schmoor, W. Sauerbrei, M. Schumacher, Assessment and comparison of prognostic classification schemes for survival data, *Statistics in Medicine* 18 (17-18) (1999) 2529–2545. [doi:10.1002/\(SICI\)1097-0258\(19990915/30\)18:17/18<2529::AID-SIM274>3.0.CO;2-5](https://doi.org/10.1002/(SICI)1097-0258(19990915/30)18:17/18<2529::AID-SIM274>3.0.CO;2-5).
- [31] J. D. Giorgini, D. K. Yeomans, A. B. Chamberlin, P. W. Chodas, R. A. Jacobson, M. S. Keesey, J. H. Lieske, S. J. Ostro, E. M. Standish, R. N. Wimberly, JPL’s on-line solar system data service and ephemeris computation system, *Bulletin of the American Astronomical Society* 28 (1996) 1158.
- [32] E. M. Standish, The observational basis for JPL’s DE 200, the planetary ephemerides of the *Astronomical Almanac*, *Astronomy and Astrophysics* 233 (1990) 252–271.
- [33] W. M. Folkner, J. G. Williams, D. H. Boggs, R. S. Park, P. Kuchynka, The planetary and lunar ephemerides DE430 and DE431, *IPN Progress Report* 42–196 (2014) 1–81.
- [34] T. Fukushige, D. C. Heggie, The time-scale of escape from star clusters, *Monthly Notices of the Royal Astronomical Society* 318 (3) (2000) 753–761. [doi:10.1046/j.1365-8711.2000.03811.x](https://doi.org/10.1046/j.1365-8711.2000.03811.x).
- [35] A. H. W. Küpper, P. Kroupa, H. Baumgardt, D. C. Heggie, Tidally induced mass loss and the structure of tidal tails of star clusters, *Monthly Notices of the Royal Astronomical Society* 401 (1) (2010) 105–122. [doi:10.1111/j.1365-2966.2009.15714.x](https://doi.org/10.1111/j.1365-2966.2009.15714.x).
- [36] N. C. Weatherford, F. A. Rasio, C. L. Rodriguez, Stellar escape from globular clusters. II. Tidal tail and stream formation in galactic environments, *The Astrophysical Journal* 965 (2) (2024) 142. [doi:10.3847/1538-4357/ad2d55](https://doi.org/10.3847/1538-4357/ad2d55).



Figure 3: Multi-Scale Benchmark Comparison across all competing architectures as swarm population N scales from 10^3 to 10^7 on the commodity edge workstation (GTX 1650 4GB). **(a) Computational Throughput:** SympGal-NODE delivers consistent scaling ($> 2.3 \times 10^6$ bodies/s) without saturation, while Continuous EG-NODE crashes at $N = 10^4$ due to quadratic $\mathcal{O}(N^2)$ memory explosion. **(b) Peak GPU VRAM Footprint:** SympGal-NODE’s streaming pipeline strictly caps memory at ≤ 1.58 GB (well below the 4 GB VRAM ceiling), enabling 10M body simulation on consumer hardware. **(c) Escape Survival C-Index ($C \uparrow$):** SympGal-NODE retains superior predictive discrimination ($C = 0.594 \rightarrow 0.574$) as population increases, whereas non-symplectic models (Vanilla NODE, Bi-LSTM) suffer severe degradation ($C < 0.42$). **(d) Relative Energy Invariant Violation $|\Delta E/E_0|$:** Symplectic preservation restricts Hamiltonian drift to $\leq 5.7 \times 10^{-5}$, whereas unconstrained networks dissipate up to 14.1% of orbital energy.




Figure 4: Empirical NASA JPL Horizons Validation: Continuous-Time Orbit Reconstruction of Potentially Hazardous Asteroid 99942 Apophis during its April 2029 close Earth encounter across an intentional 45-day observation blackout. (Top) Heliocentric X coordinate in AU. (Bottom) Positional error comparison on logarithmic scale. SympGal-NODE restricts trajectory error to $\text{RMSE} = 1,067 \text{ km}$ by preserving canonical $1\text{PN} + J_2 + \text{SRP}$ phase-space invariants, whereas unconstrained Vanilla NODE diverges to $773,752 \text{ km}$ and Cubic Spline to $632,876 \text{ km}$.




Figure 5: Macroscopic Star Cluster Dissolution and Tidal Stream Formation simulated with SympGal-NODE under the Fukushige-Heggie Lagrange aperture framework in a rotating galactic tidal field. The bound cluster core (cyan) is bounded by the Jacobi tidal horizon ($r_t \approx 1.84$, dashed purple circle). Unbound stars escape through the L_1 nozzle into an outer, slower Trailing Arm ($x > 0, y < 0$) and through the L_2 nozzle into an inner, faster Leading Arm ($x < 0, y > 0$), color-coded by logarithmic escape epoch $\log_{10} T_{\text{escape}}$.




Figure 6: Comprehensive Star Cluster Escape Time Horizon Analysis (T_{escape}) predicted by SympGal-NODE against empirical N -body ground truth and baselines. **(a) Macroscopic Survival Fraction Decay** $S(t) = N_{\text{bound}}(t)/N_0$: SympGal-NODE closely traces the empirical ground-truth evaporation curve ($t_{1/2} \approx 13.9$ dynamical units, $W_1 = 0.485$, KS = 0.055), whereas unconstrained Vanilla NODE and Soft-Loss PINN undergo artificial numerical phase-space dissipation, triggering premature cluster dissolution. **(b) Continuous Multi-Risk Hazard Rate Curves** $h_k(t)$: Empirical and predicted instantaneous hazard intensities comparing hyperbolic tidal ejection ($k = 1$) across the Jacobi horizon versus core tidal collision capture ($k = 2$). **(c) Initial Phase-Space Ejection Heatmap**: Initial planar coordinates (q_x, q_y) of swarm members color-coded by predicted logarithmic escape epoch $\log_{10} T_{\text{escape}}$, delineating fractal escape basins and chaotic boundary layers governed by Kolmogorov-Arnold-Moser (KAM) tori.




Figure 7: Continuous Escape Trajectory Corridors and Anisotropic Ejection Morphology. **(a) Rotating Frame Lagrange Aperture Nozzles** (L_1 and L_2): Stars escape through saddle throats along the galactocentric radial axis, populating the Trailing Arm ($y < 0$) via L_1 and the Leading Arm ($y > 0$) via L_2 , color-coded by escape epoch $\log_{10} T_{\text{escape}}$. **(b) Individual Escape Trajectories Crossing Tidal Horizon** R_{tidal} : Trajectories executing chaotic Lissajous orbits before exiting through the Jacobi aperture throats, color-coded by instantaneous escape speed $\|v_{\text{esc}}\|$ in N -body units. **(c) Mollweide All-Sky Directional Exit Corridors** (λ, β): Mollweide equal-area projection of asymptotic escape velocity vectors crossing R_{tidal} , revealing strong equatorial confinement ($\beta \approx 0$) and bimodal azimuthal exit cones ($\lambda \approx \pm\pi/3$) driven by Coriolis acceleration and galactic tidal shear.