

Comparing Physics-Informed and Physics-Enforced Neural Networks for the Chaotic Three-Body Problem

Mithun Meenakshi Subramanian^{1,✉}, D. B. Shanmugam¹

¹. Data Science, SRM IST Ramapuram, Chennai, IND

Received: April 23, 2026 | Review began: July 14, 2026 | Review ended: August 31, 2026 | Published: September 21, 2026

© Copyright 2026

This is an open access article distributed under the terms of the Creative Commons Attribution License CC-BY 4.0., which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Abstract

The three-body gravitational problem is a canonical benchmark for chaotic dynamical systems, admitting no general closed-form solution and exhibiting sensitive dependence on initial conditions. We present a controlled comparative benchmark of four neural network architectures - a Deep Neural Network (DNN), a Physics-Informed Neural Network (PINN), a Hamiltonian Neural Network (HNN), and a Lagrangian Neural Network (LNN) - for trajectory prediction in the planar equal-mass three-body problem. All models are trained on 1,000 high-precision trajectories generated using the DOP853 integrator and evaluated on trajectory accuracy, long-horizon generalisation, and energy conservation. To the best of our knowledge, this is the first systematic evaluation of HNNs and LNNs as trajectory predictors for the gravitational three-body problem. The DNN achieves the lowest in-horizon error (mean absolute error: 0.3166), while the LNN surpasses all models in long-horizon generalisation and energy conservation. The PINN achieves the best ordinary differential equation compliance but the worst trajectory accuracy, demonstrating the challenge of loss balancing in chaotic systems. These results, obtained under deliberately controlled but limited conditions (1,000 trajectories, equal-mass planar case), provide initial evidence that Lagrangian structural bias offers a measurable advantage for this benchmark, motivating larger-scale evaluation as a next step.

Categories: Machine Learning (ML), Modeling

Keywords: three-body problem, physics-informed neural networks, hamiltonian neural networks, lagrangian neural networks, chaotic dynamical systems, scientific machine learning

Introduction

The three-body problem - the prediction of the trajectories of three mutually gravitating bodies from their initial conditions - has resisted closed-form analytical resolution since Newton implicitly posed it in 1687 [1]. Euler [2] and Lagrange [3] identified special equilibrium configurations, and Poincaré's 1890 memoir [4] established the problem as a paradigmatic example of chaos: arbitrarily small perturbations produce exponentially diverging trajectories. Sundman's [5] convergent power-series solution is computationally impractical because of its slow convergence. Moore [6] and Chenciner and Montgomery [7] later demonstrated richer periodic solutions, although the general problem remains intractable.

Breen et al. [8] demonstrated that a feedforward Deep Neural Network (DNN) can predict three-body outcomes up to 10^8 times faster than traditional integrators while maintaining comparable short-horizon accuracy. The sensitivity of such trajectories to initial conditions is characterised by positive Lyapunov exponents [9]. Ground-truth trajectories in this work are generated using the DOP853 integrator [10]. Physics-Informed Neural Networks (PINNs) embed governing differential equations as a soft loss penalty [11], whereas Hamiltonian Neural Networks (HNNs) [12] and Lagrangian Neural Networks (LNNs) [13] encode Hamiltonian or Lagrangian mechanics directly into the network architecture, guaranteeing energy

How to cite this article:

Subramanian M, Shanmugam D B (September 21, 2026) Comparing Physics-Informed and Physics-Enforced Neural Networks for the Chaotic Three-Body Problem. Cureus J Comput Sci 3 : es44389-026-00285-7. DOI <https://doi.org/10.7759/s44389-026-00285-7>

conservation by construction. Physics-informed approaches have been broadly reviewed in the literature [14]. Symplectic and graph-based architectures [15–22] offer related guarantees at greater complexity. Reservoir computing approaches have demonstrated model-free prediction of chaotic attractors [23,24]. Santos Pereira et al. [25] applied a Physics-Informed ResNet to the three-body problem, achieving far lower Ordinary Differential Equation (ODE) residuals than a plain DNN but at the cost of trajectory accuracy, building on the residual-connection architecture of He et al. [26] and early neural-network-based solvers for differential equations [27]. Boekholt et al.'s [28] Brutus integrator provides the high-precision ground truth on which this line of work is trained and validated. Choudhary et al. [29] demonstrated that HNNs can function in the chaotic Hénon-Heiles system.

Machine learning has also been applied to closely related three-body and hierarchical triple prediction tasks over the past three years, indicating that this remains an actively developing area rather than a settled one. Lalande and Trani [30] trained a convolutional neural network to classify the long-term dynamical stability of hierarchical triple systems, outperforming classical semi-analytic stability criteria. Li et al. [31] used machine learning to classify capture into, or escape from, mean-motion resonances in the restricted three-body problem and, separately, introduced a large-step neural network that learns a symplectic, structure-preserving map of the dynamics directly, permitting much larger integration steps than a conventional integrator [32]. Vynatheya et al. [33] combined a generalised algebraic Hill-stability criterion with a machine-learning classifier for hierarchical triple-star stability, benchmarking directly against the Mardling-Aarseth criterion, while Georgakarakos et al. [34] cross-validated new empirical circumbinary-planet stability formulae using a machine-learning model against real Kepler- and TESS-discovered systems. Attia and Sibony [35] most recently combined a modified secular code with a systematic exploration of a seven-dimensional parameter space of hierarchical triple black-hole systems. Transformer-based self-attention has likewise been extended beyond trajectory prediction to the certification of the global stability of dynamical systems through symbolic Lyapunov-function discovery [36], and classical variational periodic-orbit methods have recently been hybridised with evolutionary and stochastic optimisation techniques [37].

Despite this body of work, HNNs and LNNs have not been systematically benchmarked on the full planar three-body problem: the original HNN paper [12] included only preliminary experiments, LNNs [13] have not been evaluated on gravitational systems, and Santos Pereira et al. [25] explicitly identified both as open directions for future work. Learned three-body prediction is itself a broader family of approaches than any single paper has tested together: plain feedforward networks [8], residual physics-informed variants [25], graph-based simulators that generalise across particle counts at the cost of any conservation guarantee [22], and symbolic-regression methods that recover closed-form force laws from trained networks [38] each occupy a different point on the accuracy-conservation-generalisation trade-off. More recently, transformer-based self-attention architectures have been applied to post-Newtonian three-body trajectory prediction, reporting state-of-the-art accuracy for relativistic corrections [39], although they have not yet been evaluated within a Newtonian, conservation-law-controlled comparison of the kind performed here.

Notably, none of the seven most recent machine-learning contributions to three-body and hierarchical-triple prediction [30–36] targets the specific comparison performed in this paper. Four address classification rather than trajectory-prediction tasks: Lalande and Trani [30], Vynatheya et al. [33], and Georgakarakos et al. [34] predict a binary or scalar stability outcome from orbital parameters rather than a full state trajectory, while Li et al. [31] classified resonance capture rather than rolling out dynamics. Li et al. [32] and Ciccarelli et al. [37] preserve structure (symplectic or variational) but do so for map learning and periodic-orbit discovery, respectively, neither of which is evaluated against a Hamiltonian- or Lagrangian energy-conservation baseline. Attia and Sibony [35] explore a parameter space using a secular-plus-learned hybrid rather than a first-principles neural architecture, and Alfarano et al. [36] apply self-attention to certifying stability in general dynamical systems, not to gravitational trajectory rollout. Consequently, even accounting for this recent activity, no prior work has compared four architectures, DNN, PINN, HNN, and LNN, using identical training data, an identical evaluation protocol, and a matched parameter count where architecturally possible. This paper closes that gap with the first controlled, quantitative comparison of these four models under a unified evaluation framework.

How to cite this article:

Subramanian M, Shanmugam D B (September 21, 2026) Comparing Physics-Informed and Physics-Enforced Neural Networks for the Chaotic Three-Body Problem. *Cureus J Comput Sci* 3 : es44389-026-00285-7. DOI <https://doi.org/10.7759/s44389-026-00285-7>

Materials And Methods

Problem and dataset

We consider the planar equal-mass three-body problem with unit masses $m_1 = m_2 = m_3 = 1$ and gravitational constant $G = 1$, following the approach of Breen et al. [8] and Santos Pereira et al. [25]. The equations of motion for each body $i \in \{1, 2, 3\}$ are as follows:

$$\ddot{\mathbf{r}}_i = \sum_{j \neq i} \frac{\mathbf{r}_j - \mathbf{r}_i}{|\mathbf{r}_j - \mathbf{r}_i|^3}, \quad i \in \{1, 2, 3\}$$

where $\mathbf{r}_i \in \mathbb{R}^3$ is the position vector of body i , and the double dot denotes the second derivative with respect to time t . The denominator $|\mathbf{r}_j - \mathbf{r}_i|^3$ encodes the inverse-square gravitational force law; the singularity as $|\mathbf{r}_j - \mathbf{r}_i| \rightarrow 0$ corresponds to close-encounter events and is the primary source of numerical and learned-model difficulty.

Under the planar constraint (all y -coordinates fixed to zero), the full system state at time t is the 12-dimensional vector:

$$\mathbf{s}(t) = (x_1, z_1, x_2, z_2, x_3, z_3, \dot{x}_1, \dot{z}_1, \dot{x}_2, \dot{z}_2, \dot{x}_3, \dot{z}_3) \in \mathbb{R}^{12}$$

where x_i and z_i are the planar coordinates of body i , and $\dot{x}_i$ and $\dot{z}_i$ are the corresponding velocities. The total mechanical energy of the system, used as the ground-truth conservation quantity throughout the evaluation, is

$$H = \frac{1}{2} \sum_{i=1}^3 |\dot{\mathbf{r}}_i|^2 - \sum_{1 \leq i < j \leq 3} \frac{1}{|\mathbf{r}_i - \mathbf{r}_j|}$$

where the first term is the total kinetic energy and the second term is the Newtonian gravitational potential energy between each pair of bodies.

Dataset generation

The dataset follows the initial condition sampling algorithm described above. Body 1 is fixed at position $(1, 0, 0)$. Body 2 is placed by sampling an angle $\theta \sim \mathcal{U}(0, \pi/2)$ and a scale s , drawn uniformly between 0 and 1, with components $x_2 = -s \cdot \min(0.5, \cos\theta)$ and $z_2 = s \cdot \sin\theta$. Body 3 is determined by the centre-of-mass constraint $\mathbf{p}_3 = -(\mathbf{p}_1 + \mathbf{p}_2)$, which fixes the total centre of mass at the origin. All initial velocities are zero.

Ground-truth trajectories are generated using SciPy's `solve_ivp` routine with the DOP853 eighth-order Dormand-Prince Runge-Kutta integrator, configured with relative tolerance $\text{rtol} = 10^{-8}$, absolute tolerance $\text{atol} = 10^{-10}$, and a maximum step size of 0.01 to prevent stalling during stiff close-encounter events. Each trajectory spans from $t = 0$ to $t = 10$ and is sampled at 129 uniformly spaced output points, yielding a time step of $\Delta t \approx 0.078$. The dataset was generated with a fixed random seed of 42 for full reproducibility and stored in compressed NumPy format (.npz) with arrays t (shape 129), Y (shape $1,000 \times 129 \times 18$), y_0 (shape $1,000 \times 18$), and `failed_indices` [40]. Initial conditions that cause the integrator to fail - primarily near-collision configurations - are discarded. The final dataset comprises 1,000 successfully integrated trajectories, split 80/10/10 into training (800), validation (100), and test (100) sets. Although only 12 of these variables are dynamically independent under the planar constraint, the archived arrays retain the full three-dimensional ambient representation (9 positions and 9 velocities, with all y -components identically zero) for compatibility with the general-purpose numerical integrator. Downstream models operate on the reduced 12-dimensional planar state.

Model architectures

All four models share a unified training protocol: the Adam optimiser [41] with a learning rate of 10^{-3} and weight decay of 10^{-5} , a ReduceLROnPlateau scheduler (decay factor 0.5 upon validation plateau), a batch size of 2,048, and a maximum of 2,000 training epochs. All experiments were conducted using a fixed random seed.

Deep Neural Network

How to cite this article:

Subramanian M, Shanmugam D B (September 21, 2026) Comparing Physics-Informed and Physics-Enforced Neural Networks for the Chaotic Three-Body Problem. *Cureus J Comput Sci* 3 : es44389-026-00285-7. DOI <https://doi.org/10.7759/s44389-026-00285-7>

The DNN is a 10-layer multilayer perceptron (128 units per layer, ReLU activations) following Breen et al. [8]. The input is a 7-dimensional vector of initial positions $(\mathbf{x1}, z1, \mathbf{x2}, z2, \mathbf{x3}, z3)$ and a scalar query time t_n ; the output is the 12-dimensional state $\mathbf{s}(t_n)$ predicted directly (non-autoregressively). The training loss is

$$\mathcal{L}_{\text{DNN}} = \frac{1}{N} \sum_{n=1}^N \|\hat{\mathbf{s}}(t_n) - \mathbf{s}(t_n)\|_1$$

where $\hat{\mathbf{s}}(t_n)$ is the predicted state, $\mathbf{s}(t_n)$ is the ground-truth state, and N is the number of training samples. Total parameters: 149,636.

Physics-Informed Neural Network

The PINN is a 5-block residual network (128 units, ReLU) following the PI-ResNet design of Santos Pereira et al. [25]. The total training loss combines a data term and a physics residual term evaluated at 10,000 collocation points per batch:

$$\mathcal{L}_{\text{PINN}} = \mathcal{L}_{\text{data}} + \lambda_{\text{phys}} \cdot \mathcal{L}_{\text{phys}}, \quad \lambda_{\text{phys}} = 0.1$$

where $\mathcal{L}_{\text{phys}}$ is the mean absolute error (MAE) on labelled trajectory states, $\mathcal{L}_{\text{phys}}$ is the mean squared ODE residual (the discrepancy between the network's predicted accelerations and the gravitational accelerations implied by its predicted positions via Newton's law), and λ_{phys} is a fixed physics loss weight. Gradient-norm clipping is applied to manage loss-landscape volatility in chaotic regions. Total parameters: 166,148.

Hamiltonian Neural Network

The HNN has a separable Hamiltonian decomposition, following Greydanus et al. [12]. The network takes the full planar state $(\mathbf{q}, \mathbf{p}) \in \mathbb{R}^{12}$ - positions and momenta (equal to velocities under unit masses) - and outputs a learned scalar $\hat{H}(\mathbf{q}, \mathbf{p})$. Equations of motion are derived entirely via automatic differentiation through Hamilton's canonical equations:

$$\dot{\mathbf{q}}_k = \frac{\partial \hat{H}}{\partial \mathbf{p}_k}, \quad \dot{\mathbf{p}}_k = -\frac{\partial \hat{H}}{\partial \mathbf{q}_k}$$

where $\mathbf{q}_k$ and $\mathbf{p}_k$ denote the k -th generalised coordinate and conjugate momentum, respectively. The quantity $\hat{H}$ is exactly conserved along every predicted trajectory by construction. The training loss is

$$\mathcal{L}_{\text{HNN}} = \lambda_{\text{deriv}} \cdot \mathcal{L}_{\text{deriv}} + \lambda_H \cdot \mathcal{L}_H + \lambda_{\text{sup}} \cdot \mathcal{L}_{\text{sup}}$$

where $\mathcal{L}_{\text{deriv}}$ is the MSE between the canonical equations' predictions and the observed time derivatives $\dot{\mathbf{s}}(t_n)$; $\mathcal{L}_H$ penalises deviation of $\hat{H}$ from the analytically computed total energy; $\mathcal{L}_{\text{sup}}$ provides direct Hamiltonian supervision; and $\lambda_{\text{deriv}} = \lambda_H = \lambda_{\text{sup}} = 0.1$. Architecture: separable subnetwork (3 layers $\times$ 64 units, tanh) embedded in a wrapper (6 layers $\times$ 256 units, tanh). Trajectory rollout uses fourth-order Runge-Kutta (RK4) integration of the learned canonical vector field $\mathbf{J} \nabla \hat{H}$. Total parameters: 348,422.

Lagrangian Neural Network

The LNN follows Cranmer et al. [13] with a separable Lagrangian decomposition. The network takes generalised coordinates and velocities $(\mathbf{q}, \dot{\mathbf{q}}) \in \mathbb{R}^{12}$ and outputs a learned scalar $\hat{L}(\mathbf{q}, \dot{\mathbf{q}})$. Generalised accelerations are obtained by solving the Euler-Lagrange equations differentially:

$$\ddot{\mathbf{q}} = \mathbf{M}(\mathbf{q}, \dot{\mathbf{q}})^{-1} \left[\frac{\partial \hat{L}}{\partial \mathbf{q}} - \frac{\partial^2 \hat{L}}{\partial \dot{\mathbf{q}} \partial \mathbf{q}} \dot{\mathbf{q}} \right]$$

How to cite this article:

where $\mathbf{M}(\mathbf{q}, \dot{\mathbf{q}}) = \frac{\partial^2 \hat{\mathcal{L}}}{\partial \dot{\mathbf{q}}^2}$ is the mass matrix - the Hessian of the learned Lagrangian with respect to generalised velocities - computed via second-order automatic differentiation using PyTorch's `torch.func.hessian` and `torch.func.jacrev`. The training loss is

$$\mathcal{L}_{\text{LNN}} = \lambda_{\text{phys}} \cdot \mathcal{L}_{\text{acc}} + \lambda_{\text{EL}} \cdot \mathcal{L}_{\text{EL}} + \lambda_{\text{lag}} \cdot \mathcal{L}_{\text{lag}}$$

where $\mathcal{L}_{\text{acc}}$ is the MSE between Euler-Lagrange-predicted and ground-truth accelerations; $\mathcal{L}_{\text{EL}}$ penalises violations of the Euler-Lagrange equations at training states; $\mathcal{L}_{\text{lag}}$ provides direct Lagrangian supervision; and $\lambda_{\text{phys}} = 0.1$, $\lambda_{\text{EL}} = 0.01$, and $\lambda_{\text{lag}} = 0.1$. The architecture is identical to that of the HNN (3×64 tanh subnetwork + 6×256 tanh wrapper), enabling a controlled equal-capacity comparison. Rollout uses RK4 integration of the Euler-Lagrange accelerations. The total number of trainable parameters is 348,422.

Evaluation metrics

Tier A (All Models)

The primary trajectory-accuracy metrics are the MAE and root mean squared error (RMSE) computed on the 12-dimensional state vector over the full rollout horizon. The relative rollout error (RR) normalises trajectory deviation by the magnitude of the true state, providing a chaos-robust metric:

$$\text{RR}(t) = \frac{\|\hat{\mathbf{s}}(t) - \mathbf{s}(t)\|}{\|\mathbf{s}(t)\|}$$

Long-horizon rollout error is evaluated at $T^* = 15$ (50% beyond the training horizon of $T = 10$), with ground truth extended by re-running DOP853 on the same initial conditions. Energy drift is measured as

$$\Delta E = \max_{t \in [0, T]} |H(\hat{\mathbf{s}}(t)) - H(\hat{\mathbf{s}}(0))|$$

where H is the analytical three-body energy formula applied to the predicted state.

Tier B (Model-Specific)

The ODE residual measures how closely each model's predicted states satisfy the three-body equations of motion, regardless of the ground-truth proximity. For the LNN, the Euler-Lagrange residual and acceleration MSE are additionally reported. For the HNN, the Hamiltonian vector-field residual and Hamiltonian MSE are reported.

Results And Discussion

Experimental setup

All experiments were implemented in Python (3.12.6) using PyTorch (2.0+) as the primary deep-learning framework. Dataset generation and numerical integration used SciPy (1.9+) [10] and NumPy (1.23+). Experiments were run on a system equipped with an NVIDIA RTX 3060 GPU (12 GB VRAM), an AMD Ryzen 7 7445HS processor, and 16 GB of system RAM. The second-order automatic differentiation required by the LNN was computed using PyTorch's `torch.func.hessian` and `torch.func.jacrev` APIs, which necessitated GPU acceleration to remain tractable. All four models were trained under the unified protocol described in the Model Architectures section (Adam optimiser, learning rate 10^{-3} , weight decay = 10^{-5} , batch size 2,048, up to 2,000 epochs). Model checkpoints were saved as PyTorch state dictionaries alongside `config.json` files recording all hyperparameters, enabling full reproducibility. The dataset and source code are publicly available (see Data Availability).

In-horizon trajectory accuracy

The DNN achieves the lowest in-horizon error across all metrics (MAE: 0.3166, RMSE: 1.2463, relative RR: 0.4555), confirming the findings of Breen et al. [8] even on the smaller 1,000-trajectory dataset. As shown in Table 1, error distributions are strongly right-skewed across all models; medians are well below means, indicating that a small subset of

How to cite this article:

Subramanian M, Shanmugam D B (September 21, 2026) Comparing Physics-Informed and Physics-Enforced Neural Networks for the Chaotic Three-Body Problem. *Cureus J Comput Sci* 3 : es44389-026-00285-7. DOI <https://doi.org/10.7759/s44389-026-00285-7>

difficult near-collision trajectories dominates the mean [9]. The DNN's trajectory predictions, shown in Figure 1, demonstrate close correspondence with the DOP853 ground truth across diverse trajectory types, ranging from tightly bound configurations (Trajectory 99) to weakly interacting systems (Trajectory 166).

The LNN ranks second on all accuracy metrics (MAE: 0.3929), marginally outperforming the HNN (MAE: 0.4016) at identical parameter counts (348,422). The PINN ranks last on every accuracy metric (Table 1): its mean rollout error of 0.9530 approaches 1.0, showing that the physics loss ($\lambda = 0.1$, evaluated at 10,000 collocation points per batch) dominated training at the expense of trajectory-fitting accuracy. This contrasts with Pereira et al. [25], who reported that a physics-informed residual network achieved a mean-squared ODE residual of 3.29×10^{-7} compared with 3.0×10^{-2} for their non-physics-informed baseline, representing roughly five orders of magnitude of improvement, while the data-fitting MAE was only marginally worse for the physics-informed variant (0.252 vs. 0.243) on a 30,000-trajectory dataset. Our PINN shows the same qualitative trend, namely substantially improved ODE residuals (2.44 vs. 35.29 for the DNN) at the cost of data accuracy, but the cost is considerably larger in our controlled, smaller (1,000-trajectory) setting (MAE: 0.585 vs. 0.317 for the DNN, an 85% relative increase, compared with Pereira et al.'s [25] approximately 4% increase). This is consistent with the loss-balancing problem being sensitive to training-set scale: with an order of magnitude less data, the physics-residual term appears to dominate optimisation more severely, amplifying rather than contradicting Pereira et al.'s [25] findings.

Long-horizon generalisation and energy conservation

The long-horizon evaluation ($T^* = 15$, 50% beyond the training horizon of $T = 10$) produces the benchmark's most striking result, as summarised in Table 2. The LNN achieves the best mean (1.0842) and median (0.9199) long-horizon errors, narrowly surpassing the DNN (1.0866 and 0.9468, respectively), which led all in-horizon metrics. Within the scope of this benchmark, these results provide evidence that the Euler-Lagrange inductive bias confers a genuine generalisation advantage on the three-body problem. The LNN and HNN trajectory predictions are shown in Figures 2-3, respectively; the LNN visually retains qualitatively correct orbital structure further into the rollout horizon than the HNN.

The HNN's long-horizon performance (1.2560) is comparable to that of the PINN (1.2392) despite enforcing exact conservation of its learned Hamiltonian, $\hat{H}$. Because $\hat{H}$ only approximates the true energy E , the gap between the two remains large near close-encounter singularities. The PINN's trajectory predictions (Figure 4) show visible divergence from the ground truth across multiple test trajectories, consistent with its high rollout error. The LNN achieves the best median energy drift (29.17) and the second-lowest mean ODE residual (4.90), confirming that Lagrangian structural constraints provide the best overall balance between physical fidelity and trajectory accuracy (Table 2). The PINN achieves the lowest ODE residual (mean: 2.44, median: 2.30), confirming effective physics enforcement, but incurs a consistent accuracy penalty on every Tier A metric, providing a direct demonstration of the loss-balancing problem in chaotic systems identified by Pereira et al. [25] and Karniadakis et al. [14].

Figure 5 shows the mean absolute energy error over time for all four models across the test set, presented on both linear and logarithmic scales. All models exhibit similar baseline energy-conservation behaviour during the non-encounter phases of the trajectories, with sharp spikes in energy error coinciding with close-encounter events, the gravitational singularities that dominate prediction error across all architectures. On the logarithmic scale, the LNN and HNN show marginally lower energy-error floors between encounter events, consistent with their architectural conservation properties, whereas the DNN and PINN show comparatively higher variance throughout.

How to cite this article:

Model	MAE Mean	MAE Median	RMSE Mean	RMSE Median	RR Mean	RR Median
DNN	0.3166	0.1946	1.2463	0.8083	0.4555	0.4502
PINN	0.5852	0.4517	1.9577	1.4677	0.9530	0.9887
LNN	0.3929	0.2192	1.4308	0.8227	0.6229	0.5994
HNN	0.4016	0.2386	1.4636	0.8498	0.6754	0.6303

TABLE 1: Trajectory accuracy metrics on held-out test set (lower is better)

DNN, Deep Neural Network; HNN, Hamiltonian Neural Network; LNN, Lagrangian Neural Network; MAE, Mean Absolute Error; PINN, Physics-Informed Neural Network; RMSE, Root Mean Square Error; RR, Relative Rollout

Model	LH Mean	LH Median	E-Drift (Median)	ODE Res. Mean	ODE Res. Median	Params
DNN	1.0866	0.9468	30.345	35.291	11.458	149K
PINN	1.2392	1.0821	30.389	2.441	2.302	166K
LNN	1.0842	0.9199	29.172	4.902	3.336	348K
HNN	1.2560	1.0443	29.370	14.479	2.526	348K

TABLE 2: Long-horizon error ($T^* = 15$), energy drift (median), and ODE residual (lower is better for all)

DNN, Deep Neural Network; E-Drift, Energy Drift; HNN, Hamiltonian Neural Network; LH, Long Horizon; LNN, Lagrangian Neural Network; ODE, Ordinary Differential Equation; PINN, Physics-Informed Neural Network

Figures 1-5 show representative trajectory predictions for each model against the DOP853 ground truth, along with the mean energy drift over time across the test set. Solid lines denote the ground truth, whereas dashed lines denote model predictions. Colours represent individual bodies.

How to cite this article:




FIGURE 1: DNN trajectory predictions vs ground truth

DNN, Deep Neural Network

How to cite this article:

Subramanian M, Shanmugam D B (September 21, 2026) Comparing Physics-Informed and Physics-Enforced Neural Networks for the Chaotic Three-Body Problem. Cureus J Comput Sci 3 : es44389-026-00285-7. DOI <https://doi.org/10.7759/s44389-026-00285-7>

LNN: Ground Truth vs Prediction (●=start)




FIGURE 2: LNN trajectory predictions vs ground truth

LNN, Lagrangian Neural Network

How to cite this article:

Subramanian M, Shanmugam D B (September 21, 2026) Comparing Physics-Informed and Physics-Enforced Neural Networks for the Chaotic Three-Body Problem. Cureus J Comput Sci 3 : es44389-026-00285-7. DOI <https://doi.org/10.7759/s44389-026-00285-7>




FIGURE 3: HNN trajectory predictions vs ground truth

HNN, Hamiltonian Neural Network

How to cite this article:

Subramanian M, Shanmugam D B (September 21, 2026) Comparing Physics-Informed and Physics-Enforced Neural Networks for the Chaotic Three-Body Problem. Cureus J Comput Sci 3 : es44389-026-00285-7. DOI <https://doi.org/10.7759/s44389-026-00285-7>




FIGURE 4: PINN trajectory predictions vs ground truth

PINN, Physics-Informed Neural Network




FIGURE 5: Mean absolute energy error over time for all four models ($\pm$ std across test set)

DNN, Deep Neural Network; HNN, Hamiltonian Neural Network; LNN, Lagrangian Neural Network; PINN, Physics-Informed Neural Network

How to cite this article:

Subramanian M, Shanmugam D B (September 21, 2026) Comparing Physics-Informed and Physics-Enforced Neural Networks for the Chaotic Three-Body Problem. Cureus J Comput Sci 3 : es44389-026-00285-7. DOI <https://doi.org/10.7759/s44389-026-00285-7>

Discussion

These findings suggest several directions for future work. Training on larger datasets approaching the scale of Pereira et al. [25] would determine whether the accuracy gap between DNNs and structure-preserving models narrows with additional data. Sequence-to-sequence architectures [42] offer a natural extension by conditioning each predicted state on prior trajectory history. Transformer-based architectures [43], which treat each state vector as a token, are particularly promising; self-attention over trajectory history could appropriately weight recent close-encounter events while remaining compatible with Hamiltonian or Lagrangian structural constraints.

Limitations

This benchmark has three main limitations. First, the dataset (1,000 trajectories) is an order of magnitude smaller than the 30,000-trajectory dataset used by Pereira et al. [25]; the discussion above suggests that this scale difference explains the larger PINN accuracy penalty observed here. Second, this benchmark is restricted to the equal-mass planar three-body problem; none of the four architectures was tested for generalisation across mass ratios, particle counts, or three-dimensional motion, a limitation that graph-based architectures are specifically designed to address [22]. Third, all evaluation was performed against DOP853 synthetic ground truth; no comparison with observational or real astronomical trajectory data was performed.

Conclusions

We presented the first controlled comparison of DNN, PINN, HNN, and LNN on the planar gravitational three-body problem. The DNN is the strongest in-horizon predictor, whereas the LNN is the best overall model, surpassing the DNN in long-horizon generalisation and energy conservation while remaining competitive in-horizon - evidence, within this controlled benchmark, that the Euler-Lagrange inductive bias yields a measurable advantage on this canonical chaotic system. The PINN produces the most physically consistent predictions (lowest ODE residual) but the worst trajectories, suggesting that dynamic loss scheduling may be necessary for chaotic systems. The HNN enforces conservation by construction but does not close the gap with the LNN at equal model capacity, indicating that the Lagrangian formulation is better suited to gravitational problems at the scales evaluated here. We emphasise that these findings are benchmark-specific: they establish that the effect is present and measurable under controlled conditions, not that it generalises beyond the scope tested here. Future work will extend this comparison to datasets approaching the 30,000-trajectory scale, and to unequal mass ratios and three-dimensional configurations, in order to determine whether the DNN-LNN trade-off and the LNN's long-horizon advantage persist beyond the narrow setting evaluated in this paper.

Additional Information

Author Contributions

All authors have reviewed the final version to be published and agreed to be accountable for all aspects of the work.

Concept and design: Mithun Meenakshi Subramanian, D. B. Shanmugam

Acquisition, analysis, or interpretation of data: Mithun Meenakshi Subramanian, D. B. Shanmugam

Drafting of the manuscript: Mithun Meenakshi Subramanian, D. B. Shanmugam

Critical review of the manuscript for important intellectual content: Mithun Meenakshi Subramanian, D. B. Shanmugam

Supervision: D. B. Shanmugam

Disclosures

Human subjects: All authors have confirmed that this study did not involve human participants or tissue.

How to cite this article:

Subramanian M, Shanmugam D B (September 21, 2026) Comparing Physics-Informed and Physics-Enforced Neural Networks for the Chaotic Three-Body Problem. *Cureus J Comput Sci* 3 : es44389-026-00285-7. DOI <https://doi.org/10.7759/s44389-026-00285-7>

Animal subjects: All authors have confirmed that this study did not involve animal subjects or tissue.

Conflicts of interest: In compliance with [the ICMJE uniform disclosure form](#), all authors declare the following:

- **Payment/services info:** All authors have declared that no financial support was received from any organization for the submitted work.
- **Financial relationships:** All authors have declared that they have no financial relationships at present or within the previous three years with any organizations that might have an interest in the submitted work.
- **Other relationships:** All authors have declared that there are no other relationships or activities that could appear to have influenced the submitted work.

Data Availability Statements

The dataset and source code supporting this study are openly available on GitHub at [<https://github.com/emmcquared/Comparing-Physics-Informed-and-Physics-Enforced-Neural-Networks-for-the-Chaotic-Three-Body-Problem>] and archived on Zenodo at [<https://doi.org/10.5281/zenodo.21193210>]

References

1. Newton I: [Philosophiæ naturalis principia mathematica \[Mathematical Principles of Natural Philosophy\]](#). Jussu Societatis Regiæ ac Typis Joseph Streater, London; 1687.
2. Euler L: [De motu rectilineo trium corporum se mutuo attrahentium \[On the Rectilinear Motion of Three Mutually Attracting Bodies\]](#). Novi Commentarii academiae scientiarum Petropolitanae. 1767, 11:144-151.
3. Lagrange J-L: [Essai sur le Problème des Trois Corps \[Essay on the Three-Body Problem\]](#). Recueil des pièces qui ont remporté les prix de l'Académie royale des sciences, Tome IX. Claude Jombert, Paris; 1772.
4. Poincaré H: [The Three-Body Problem and the Equations of Dynamics: Poincaré's Foundational Work on Dynamical Systems Theory](#). Springer, Cham; 2017. [10.1007/978-3-319-52899-1](#)
5. Sundman KF: [Mémoire sur le problème des trois corps \[On the Three-Body Problem and the Equations of Dynamics\]](#). Acta Mathematica. 1913, 36:105-179. [10.1007/BF02422379](#)
6. Moore C: [Braids in classical dynamics](#). Physical Review Letters. 1993, 70:3675-3679. [10.1103/physrevlett.70.3675](#)
7. Chenciner A, Montgomery R: [A remarkable periodic solution of the three-body problem in the case of equal masses](#). Annals of Mathematics. 2000, 152:881-901. [10.2307/2661357](#)
8. Breen PG, Foley CN, Boekholt T, Portegies Zwart S: [Newton versus the machine: solving the chaotic three-body problem using deep neural networks](#). Monthly Notices of the Royal Astronomical Society. 2020, 494:2465-2470. [10.1093/mnras/staa713](#)
9. Oseledets VI: [A multiplicative ergodic theorem. Lyapunov characteristic numbers for dynamical systems](#). Transactions of the Moscow Mathematical Society. 1968, 19:197-231.
10. Virtanen P, Gommers R, Oliphant TE, et al.: [SciPy 1.0: fundamental algorithms for scientific computing in Python](#). Nature Methods. 2020, 17:261-272. [10.1038/s41592-019-0686-2](#)
11. Raissi M, Perdikaris P, Karniadakis GE: [Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations](#). Journal of Computational Physics. 2019, 378:686-707. [10.1016/j.jcp.2018.10.045](#)
12. Greydanus S, Dzamba M, Yosinski J: [Hamiltonian neural networks](#). Proceedings of the 33rd International Conference on Neural Information Processing Systems. 2019, 15379-15389.
13. Cranmer M, Greydanus S, Hoyer S, Battaglia P, Spergel D, Ho S: [Lagrangian neural networks](#). ICLR Workshop on Deep Differential Equations. 2020, 1-6.
14. Karniadakis GE, Kevrekidis IG, Lu L, Perdikaris P, Wang S, Yang L: [Physics-informed machine learning](#). Nature Reviews Physics. 2021, 3:422-440. [10.1038/s42254-021-00314-5](#)

How to cite this article:

Subramanian M, Shanmugam D B (September 21, 2026) Comparing Physics-Informed and Physics-Enforced Neural Networks for the Chaotic Three-Body Problem. Cureus J Comput Sci 3 : es44389-026-00285-7. DOI <https://doi.org/10.7759/s44389-026-00285-7>

15. Chen Z, Zhang J, Arjovsky M, et al.: [Symplectic recurrent neural networks](#). International Conference on Learning Representations. 2020, 1-23.
16. Zhong YD, Dey B, Chakraborty A: [Symplectic ODE-Net: Learning Hamiltonian dynamics with control](#). International Conference on Learning Representations. 2020, 1-17.
17. Jin P, Zhang Z, Zhu A, Tang Y, Karniadakis GE: [SympNets: Intrinsic structure-preserving symplectic networks for identifying Hamiltonian systems](#). Neural Networks. 2020, 132:166-179. [10.1016/j.neunet.2020.08.017](#)
18. Hairer E, Wanner G, Lubich C: [Geometric Numerical Integration: Structure-Preserving Algorithms for Ordinary Differential Equations](#). Springer, Berlin, Heidelberg; 2006. [10.1007/3-540-30666-8](#)
19. Cai MX, Portegies Zwart S, Podareanu D: [Neural symplectic integrator with Hamiltonian inductive bias for the gravitational N-body problem \[PREPRINT\]](#). arXiv. 2021, [10.48550/arXiv.2111.15631](#)
20. Saz Ulibarrena V, Horn P, Portegies Zwart S, Sellentin E, Koren B, Cai MX: [A hybrid approach for solving the gravitational N-body problem with Artificial Neural Networks](#). Journal of Computational Physics. 2024, 496:112596. [10.1016/j.jcp.2023.112596](#)
21. Battaglia PW, Pascanu R, Lai M, Rezende DJ, Kavukcuoglu K: [Interaction networks for learning about objects, relations and physics](#). NIPS'16: Proceedings of the 30th International Conference on Neural Information Processing Systems. 2016, 4509-4517.
22. Sanchez-Gonzalez A, Godwin J, Pfaff T, Ying R, Leskovec J, Battaglia PW: [Learning to simulate complex physics with graph networks](#). ICML'20: Proceedings of the 37th International Conference on Machine Learning. 2020, 8459-8468.
23. Pathak J, Lu Z, Hunt BR, Girvan M, Ott E: [Using machine learning to replicate chaotic attractors and calculate Lyapunov exponents from data](#). Chaos. 2017, 27:121102. [10.1063/1.5010300](#)
24. Pathak J, Hunt B, Girvan M, Lu Z, Ott E: [Model-free prediction of large spatiotemporally chaotic systems from data: A reservoir computing approach](#). Physical Review Letters. 2018, 120:024102. [10.1103/physrevlett.120.024102](#)
25. Santos Pereira M, Tripa L, Lima N, Caldas F, Soares C: [Advancing solutions for the three-body problem through physics-informed neural networks](#). Proceedings of the 75th International Astronautical Congress. 2024, 1-14.
26. He K, Zhang X, Ren S, Sun J: [Deep residual learning for image recognition](#). 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV, USA. 2016, 770-778. [10.1109/CVPR.2016.90](#)
27. Lagaris IE, Likas A, Fotiadis DI: [Artificial neural networks for solving ordinary and partial differential equations](#). IEEE Transactions on Neural Networks. 1998, 9:987-1000. [10.1109/72.712178](#)
28. Boekholt T, Portegies Zwart S: [On the reliability of N-body simulations](#). Computational Astrophysics and Cosmology. 2015, 2:2. [10.1186/s40668-014-0005-3](#)
29. Choudhary A, Lindner JF, Holliday EG, Miller ST, Sinha S, Ditto WL: [Physics-enhanced neural networks learn order and chaos](#). Physical Review E. 2020, 101:062207. [10.1103/physreve.101.062207](#)
30. Lalande F, Trani AA: [Predicting the stability of hierarchical triple systems with convolutional neural networks](#). The Astrophysical Journal. 2022, 938:18. [10.3847/1538-4357/ac8eab](#)
31. Li X, Li J, Xia ZJ, Georgakarakos N: [Machine-learning prediction for mean motion resonance behaviour - The planar case](#). Monthly Notices of the Royal Astronomical Society. 2022, 511:2218-2228. [10.1093/mnras/stac166](#)
32. Li X, Li J, Xia ZJ, Georgakarakos N: [Large-step neural network for learning the symplectic evolution from partitioned data](#). Monthly Notices of the Royal Astronomical Society. 2023, 524:1374-1385. [10.1093/mnras/stad1948](#)
33. Vynatheya N, Mardling RA, Trani AA, Bellinger EP: [Algebraic and machine learning approach to hierarchical triple-star stability](#). Monthly Notices of the Royal Astronomical Society. 2022, 516:4146-4155. [10.1093/mnras/stac2540](#)
34. Georgakarakos N, Eggel S, Ali-Dib M, Dobbs-Dixon I: [Empirical stability criteria for 3D hierarchical triple systems. I. Circumbinary planets](#). The Astronomical Journal. 2024, 168:224. [10.3847/1538-3881/ad7a78](#)
35. Attia M, Sibony Y: [Exploring the parameter space of hierarchical triple black hole systems](#). Astronomy & Astrophysics. 2025, 700:A35. [10.1051/0004-6361/202555165](#)
36. Alfarano A, Charton F, Hayat A: [Global Lyapunov functions: a long-standing open problem in mathematics, with symbolic transformers \[PREPRINT\]](#). arXiv. 2024, [10.48550/arXiv.2410.08304](#)
37. Ciccirelli R, Introna M, Terracini S, Vasile M: [New solutions for the symmetrical N-body problem through variational approach and optimisation techniques](#). Aerotecnica Missili & Spazio. 2026, 105:153-170. [10.1007/s42496-025-00265-5](#)

How to cite this article:

Subramanian M, Shanmugam D B (September 21, 2026) Comparing Physics-Informed and Physics-Enforced Neural Networks for the Chaotic Three-Body Problem. Cureus J Comput Sci 3 : es44389-026-00285-7. DOI <https://doi.org/10.7759/s44389-026-00285-7>

38. Cranmer M, Sanchez-Gonzalez A, Battaglia P, Xu R, Cranmer K, Spergel D, Ho S: [Discovering symbolic models from deep learning with inductive biases](#). NIPS'20: Proceedings of the 34th International Conference on Neural Information Processing Systems. 2020, 17429-17442.
39. SiTu H, Wu XJ, Yang B, Lin W: [Self-attention-enhanced deep neural networks for simulating post-Newtonian dynamics of three-body systems](#). The Astrophysical Journal Supplement Series. 2025, 281:42. [10.3847/1538-4365/ae173f](#)
40. [emmcesquared/Comparing-Physics-Informed-and-Physics-Enforced-Neural-Networks-for-the-Chaotic-Three-Body-Problem: submission](#). Zenodo. (2026). <https://doi.org/10.5281/zenodo.21193210>.
41. Kingma DP, Ba J: [Adam: A method for stochastic optimization](#). International Conference on Learning Representations. 2015, 1-15.
42. Sutskever I, Vinyals O, Le QV: [Sequence to sequence learning with neural networks](#). NIPS'14: Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 2. 2014, 3104-3112.
43. Vaswani A, Shazeer N, Parmar N, et al.: [Attention is all you need](#). NIPS'17: Proceedings of the 31st International Conference on Neural Information Processing Systems. 2017, 6000-6010.

How to cite this article:

Subramanian M, Shanmugam D B (September 21, 2026) Comparing Physics-Informed and Physics-Enforced Neural Networks for the Chaotic Three-Body Problem. Cureus J Comput Sci 3 : es44389-026-00285-7. DOI <https://doi.org/10.7759/s44389-026-00285-7>