
Data-driven discovery of non-Newtonian astronomy via learning non-Euclidean Hamiltonian

Oswin So*

Massachusetts Institute of Technology, USA
oswinso@mit.edu

Gongjie Li

Georgia Institute of Technology, USA
gongjie.li@physics.gatech.edu

Evangelos A. Theodorou

Georgia Institute of Technology, USA
evangelos.theodorou@gatech.edu

Molei Tao

Georgia Institute of Technology, USA
mtao@gatech.edu

Abstract

Incorporating the Hamiltonian structure of physical dynamics into deep learning models provides a powerful way to improve the interpretability and prediction accuracy. While previous works are mostly limited to the Euclidean spaces, their extension to the Lie group manifold is needed when rotations form a key component of the dynamics, such as the higher-order physics beyond simple point-mass dynamics for N -body celestial interactions. Moreover, the multiscale nature of these processes presents a challenge to existing methods as a long time horizon is required. By leveraging a symplectic Lie-group manifold preserving integrator, we present a method for data-driven discovery of non-Newtonian astronomy. Preliminary results show the importance of both these properties in training stability and prediction accuracy.

1 Introduction

Deep Neural Networks (DNN) have been demonstrated to be effective tools for learning dynamical systems from data. One important class of systems to be learned have dynamics described by physical laws, whose structure can be exploited by learning the Hamiltonian of the system instead of the vector field [1, 2]. An appropriately learned Hamiltonian can endow the learned system with properties such as superior long prediction accuracy [3] and applicability to chaotic systems [3–5].

To learn continuous dynamics from discrete data, one important step is to bridge the continuous and discrete times. Seminal work initially approximated the time derivative via finite differences and then matched it with a learned (Hamiltonian) vector field [1, 2]. Recent efforts avoid the inaccuracy of finite difference by numerically integrating the learned vector field. Especially relevant here is SRNN [6], which uses a symplectic integrator to ensure the learned dynamics is symplectic (a necessity for Hamiltonian systems). Although SRNN only demonstrated learning separable Hamiltonians, breakthrough in symplectic integration of arbitrary Hamiltonians [7] was used to extend SRNN [8]. Further efforts on improving the time integration error have also been made [9–11]. Meanwhile, alternative approaches based on learning a symplectic map instead of the Hamiltonian also demonstrated efficacy [3, 12], although these approaches have not been extended to non-Euclidean problems.




Figure 1: One planet’s orbit around a star: rigid body correction results in a *precession*, *i.e.* slow rotation of the orbital axis. Our method, ‘Lie T2’, learns V from data and predicts a trajectory that matches the ground truth with the rigid body potential.

*Work was done while Oswin was at Georgia Tech.

In fact, one relatively under-explored area is learning Hamiltonian dynamics on *manifolds* like the Lie group manifold family². One important member of this family is $\text{SO}(n)$, which describes isometries in \mathbb{R}^n and is important for, e.g., dynamical astronomy. The evolution of celestial bodies correspond to a mechanical system, and the 2- and 3-body problems have been a staple problem in works on learning Hamiltonian (e.g., [1, 3, 6, 15]); however, the Newtonian (point-mass) gravity considered is already well understood. Practical problems in planetary dynamics are complicated by higher-order physics such as planet spin-orbit interaction, tidal dissipation, and general relativistic correction. While it is unclear what would be a perfect scientific model for these effects, planetary rotation is a necessary component to account for spin-orbit interaction and tidal forcing, creating an $\text{SO}(3)^{\otimes N}$ component of the configuration space. To learn these physics from data, we need to learn on the Lie group.

Rigid body dynamics also play important roles in other applications such as robotics. In a seminal work [16], Hamiltonian dynamics on $\text{SO}(3)$ are used to learn rigid body dynamics for a quadrotor. In that work, Runge-Kutta 4 integrator is used. Consequently, the method is applicable to short time-horizon (see Sec.3 and last paragraph of Sec.2).

For our problem of learning non-Newtonian astronomy, the time-horizon has to be long. Hence, we use a different approach by leveraging a Lie-group preserving symplectic integrator. Structure-preserving integration of dynamical systems on manifolds has been extensively studied in literature, for example for Lie groups [17–21] and more broadly, geometric integration [22–25].

In summary, we propose a deep learning methodology for performing data-driven discovery of non-Newtonian astronomy. By leveraging the use of a symplectic Lie-group manifold preserving integrator, we show how a non-Euclidean Hamiltonian can be learned for accurate prediction of non-Newtonian effects. Moreover, we provide insights that show the importance of both symplecticity and exact preservation of the Lie-group manifold in training stability.

2 Method

Given observations of a dynamically evolving system, our goal is to learn the *physics* that governs its evolution from the data. Denote by $(\mathbf{q}_{k,l}, \mathbf{R}_{k,l}, \mathbf{p}_{k,l}, \mathbf{\Pi}_{k,l})_{k=1}^K$, $l \in [L]$ a dataset of snapshots of L continuous-time trajectories of a system with N interacting rigid bodies. That is,

$$(\mathbf{q}_{k,l}, \mathbf{R}_{k,l}, \mathbf{p}_{k,l}, \mathbf{\Pi}_{k,l}) = (\mathbf{q}_l(k\Delta t), \mathbf{R}_l(k\Delta t), \mathbf{p}_l(k\Delta t), \mathbf{\Pi}_l(k\Delta t)),$$

where $\mathbf{q}_l(t), \mathbf{R}_l(t), \mathbf{p}_l(t), \mathbf{\Pi}_l(t)$ is a solution of some latent Hamiltonian ODE to be learned corresponding to mechanical dynamics on $T^*\text{SE}(3)^{\otimes N}$. Δt is a (possibly large) observation timestep, $\mathbf{R} \in \text{SO}(3)^{\otimes N}$ is the rotational configuration of the N rigid bodies, and $\mathbf{\Pi} \in \mathfrak{so}(3)^{\otimes N}$ denotes each’s angular momentum in their respective body frames.

Importantly, since the configuration space $\mathcal{Q} = \text{SE}(3)^{\otimes N}$ is not flat, the mechanical dynamics are not given by $\dot{\mathbf{q}} = \frac{\partial \mathcal{H}}{\partial \mathbf{p}}$, $\dot{\mathbf{p}} = -\frac{\partial \mathcal{H}}{\partial \mathbf{q}}$ for some Hamiltonian \mathcal{H} that depends on the generalized coordinates $\mathbf{q} \in \mathcal{Q}$ and generalized momentum $\mathbf{p} \in T_{\mathbf{q}}^* \mathcal{Q}$. Instead, the equations of motion can be derived via either Lagrange multipliers [20, 26] or a Lie group variational principle [20, 27], which will be

$$\dot{\mathbf{q}}_i = \mathbf{p}_i/m_i, \quad (1a) \quad \dot{\mathbf{R}}_i = \mathbf{R}_i \widehat{\mathbf{J}_i^{-1} \mathbf{\Pi}_i} \quad (1c)$$

$$\dot{\mathbf{p}}_i = -\frac{\partial V}{\partial \mathbf{p}_i} + F_{\mathbf{p}_i} \quad (1b) \quad \dot{\mathbf{\Pi}}_i = \mathbf{\Pi}_i \times \mathbf{J}_i^{-1} \mathbf{\Pi}_i - \left(\mathbf{R}_i^\top \frac{\partial V}{\partial \mathbf{R}_i} - \left(\frac{\partial V}{\partial \mathbf{R}_i} \right)^\top \mathbf{R}_i \right)^\vee + F_{\mathbf{\Pi}_i} \quad (1d)$$

assuming a physical Hamiltonian $\mathcal{H}(\mathbf{q}, \mathbf{R}, \mathbf{p}, \mathbf{\Pi}) = \sum_{i=1}^N \frac{1}{2} \mathbf{p}_i^\top \mathbf{p}_i/m_i + \sum_{i=1}^N \frac{1}{2} \mathbf{\Pi}_i^\top \mathbf{J}_i^{-1} \mathbf{\Pi}_i + V(\mathbf{q}, \mathbf{R})$ that sums total (translation and rotational) kinetic energy and interaction potential V , where m_i, \mathbf{J}_i denote the mass and inertial tensor of the i th body, and $F_{\mathbf{p}}, F_{\mathbf{\Pi}}$ are forcing terms to model nonconservative forces. $\mathbf{\Pi}_i \in \mathbb{R}^3$ is a vector, $\widehat{\cdot}$ is the map from \mathbb{R}^3 to Skew_3 and \vee is its inverse ([20] for more details). By learning the potential V , external forcing $F_{\mathbf{p}}$ and torque $F_{\mathbf{\Pi}}$, we can learn the physics of the system.

²We note extensions to include holonomic constraints in [13] and to handle contact in [14]).

2.1 Machine Learning Challenges Posed by Dynamical Astronomy

We study this setup because it helps answer scientific questions like: what physics governs the motions of celestial bodies, such as planets in a planetary system? The leading order physics is of course already well known, namely these bodies can be approximated by point masses that are interacting through a $1/r$ gravitational potential. However, planets are not point masses, and their rotations matter because they shape planetary climates [28, 29] and even feedback to their orbits [30]. This already starts to alter V even if one only considers classical gravity. For example, the gravitational potential V for interacting bodies of finite sizes should be $V(\mathbf{q}, \mathbf{R}) = \sum_{i < j} V_{i,j}$, where

$$V_{i,j}(\mathbf{q}, \mathbf{R}) = \int_{\mathcal{B}_i} \int_{\mathcal{B}_j} -\frac{\mathcal{G}\rho(\mathbf{x}_i)\rho(\mathbf{x}_j)}{\|\mathbf{q}_i + \mathbf{R}_i\mathbf{x}_i - \mathbf{q}_j - \mathbf{R}_j\mathbf{x}_j\|} d\mathbf{x}_i d\mathbf{x}_j = \underbrace{-\frac{\mathcal{G}m_i m_j}{\|\mathbf{q}_i - \mathbf{q}_j\|}}_{V_{i,j,\text{point}}} + \underbrace{O\left(\frac{1}{\|q_i - q_j\|^2}\right)}_{V_{i,j,\text{resid}}}. \quad (2)$$

Working with the full potential is complicated since \mathcal{B}_i is not known and the integral is not analytically known. Can we directly learn V_{resid} from time-series data?

Classical gravity (i.e. Newtonian physics) is not the only driver of planetary motion — tidal forces and general relativity (GR) matter too. The former provides a dissipation mechanism and plays critical roles in altering planetary orbits [31, 32]; the latter doesn't need much explanation and has been demonstrated by, e.g., Mercury's precessing orbit [33]. Tidal forces depend on celestial bodies' rotations [34] and thus is a function of both \mathbf{q}, \mathbf{R} . GR's effects cannot be fully characterized with classical coordinates $\mathbf{q}, \mathbf{R}, \mathbf{p}, \mathbf{\Pi}$, but post-Newtonian approximations based purely on these coordinates are popular [35]. Can we learn both purely from data if we did not have theories for either?

In addition to the scientific questions, there are also significant machine learning challenges:

Multiscale dynamics. Rigid-body correction (V_{resid}), tidal force, and GR correction are all much smaller forces compared to point-mass gravity. Consequently, their effects do not manifest until long time. Thus, one challenge for learning them is that the dynamical system exhibits different behaviors over **multiple timescales**. It is reasonable to require long time series data for the small effects to be learned; meanwhile, when observations are expensive to make, the observation time step Δt can be much longer than the smallest timescales. Can we still learn the physics in this case? We will leverage symplectic integrator and its mild growth of error over long time [7, 26] to provide a positive answer.

Respecting the Lie group manifold. However, even having a symplectic integrator is not enough because the position variable of the latent dynamics (i.e. truth) stays on $\text{SE}(3)^{\otimes N}$. If the integrated solution falls off this manifold such that $\mathbf{R}^\top \mathbf{R} = I$ no longer holds, it is not only incorrect but likely misleading for the learning of $V(\mathbf{q}, \mathbf{R})$. Popular integrators such as forward Euler, Runge-Kutta 4 (RK4) and Leapfrog [1, 6, 36] unfortunately do no maintain the manifold structure.

2.2 Learning with Lie Symplectic RNNs

Our method can be viewed as a Lie-group generalization of the seminal work of SRNN [6], where a good integrator that is both *symplectic* and *Lie-group preserving* is employed as a recurrent block.

Lie T_2 : A Symplectic Lie-Group Preserving Integrator. To construct an integrator that achieves both properties, we borrow from [20] the idea of Lie-group and symplecticity preserving splitting, and split our Hamiltonian as $\mathcal{H} = \mathcal{H}_{\text{KE}} + \mathcal{H}_{\text{PE}} + \mathcal{H}_{\text{asym}}$, which contains the axial-symmetric kinetic energy, potential energy and asymmetric kinetic energy correction terms. This enables computing the *exact* integrators $\phi_t^{[\text{KE}]}$, $\phi_t^{[\text{PE}]}$ and $\phi_t^{[\text{asym}]}$ (see App B for details). We then construct a 2nd-order symplectic integrator Lie T_2 by applying the Strang composition scheme. To account for non-conservative forces, the corresponding non-conservative momentum update $\phi^{[\text{force}]} : (\mathbf{p}, \mathbf{\Pi}) \leftarrow F(\mathbf{q}, \mathbf{R}, \mathbf{p}, \mathbf{\Pi})$ is inserted in the middle of the composition [20]. This gives $\phi_h^{[\text{Lie } T_2]}$ for stepsize h as

$$\phi_h^{\text{Lie } T_2} := \phi_{h/2}^{[\text{KE}]} \circ \phi_{h/2}^{[\text{PE}]} \circ \phi_{h/2}^{[\text{asym}]} \circ \phi_h^{[\text{force}]} \circ \phi_{h/2}^{[\text{asym}]} \circ \phi_{h/2}^{[\text{PE}]} \circ \phi_{h/2}^{[\text{KE}]} \quad (3)$$

While the implementation of our integrator in the case of dynamical astronomy is specific to the $\text{SE}(3)^{\otimes N}$ manifold, we note that the use of a splitting method to construct Lie-group preserving




Figure 2: Inputs are fed through a recurrent layer with Lie T_2 . Prediction error is used as a loss on θ .

integrators can be applied to construct symplectic integrators for any Hamiltonian system defined on a Lie group [26].

A Recurrent Architecture for Nonlinear Regression. Given the simplicity of V_{point} , we assume this is known and learn V_{resid} and F^θ with multi-layer perceptron (MLP) V_{resid}^θ and F^θ without assuming any pairwise structure (see App C for discussion). We then use $\phi^{\theta, \text{Lie } T_2}$ to integrate dynamics forward, where θ denotes the dependence on the networks. However, when the temporal spacing between observations Δt is large, using a single $\phi_{\Delta t}^{\theta, \text{Lie } T_2}$ will result in large errors for the fast timescale dynamics. Instead, we compose $\phi_h^{\theta, \text{Lie } T_2}$ H times as $(\hat{q}_{k+1, l}, \hat{p}_{k+1, l}) = \phi_h^{\theta, \text{Lie } T_2} \circ \dots \circ \phi_h^{\theta, \text{Lie } T_2}(q_{k, l}, p_{k, l})$, where $H = h/\Delta t \in \mathbb{Z}$ determines the integration stepsize h , and we use $q = (\mathbf{q}, \mathbf{R})$ and $p = (\mathbf{p}, \mathbf{\Pi})$ to denote the position and momentum related terms respectively. We perform training by minimizing the following empirical loss over random minibatches of size N_b

$$\mathcal{L}(\theta) := \frac{1}{N_b K} \sum_{l=1}^{N_b} \sum_{k=1}^K \left\{ \|q_{k, l} - \hat{q}_{k, l}^\theta\|_2^2 + \|p_{k, l} - \hat{p}_{k, l}^\theta\|_2^2 \right\} \quad (4)$$

Note that we do not assume access to the true derivatives $\dot{q}_{k, j}$ and $\dot{p}_{k, j}$ used in the loss function of some works [1, 37, 38]. The forward pass of the training process is summarized in Fig. 2 (see App C for details).

Benefit. Learning an accurate V^θ, F^θ requires accurate numerical simulation which also leads to a trainable model. Without preservation of the manifold structure, training can lead to ‘shortcuts’ outside the manifold that seemingly match the data but completely mislead the learning. Symplecticity also plays a vital role in controlling the long time integration error — under reasonable conditions, a p th-order symplectic integrator has linear $\mathcal{O}(\Delta t h^p)$ error bound, whereas a p th-order nonsymplectic one has an exponential $\mathcal{O}(e^{C\Delta t h^p})$ error bound [7, 26]. While these bounds do not matter for small Δt , they are significant for multiscale problems where Δt is macroscopic but h is microscopic. Consequently, improving error estimates for a nonsymplectic integrator by reducing h makes the RNN exponentially deep — this often renders training difficult [39] and is not desirable.

3 Results

We aim to answer two questions. **Q1** Can we learn multiscale physics? **Q2** How important are **symplecticity** (\mathcal{S}) and **Lie-group preservation** (\mathcal{L}) for learning? The closest baseline for our problem is work of [16], which learns short timescale rigid-body Hamiltonian dynamics for robotics. Placed in our framework, their work corresponds to using RK4 for the recurrent block, which is neither \mathcal{S} nor \mathcal{L} . Therefore, to investigate **Q2**, we vary the choice of integrator in our framework as follows: *Normal*: Explicit Euler, RK4. \mathcal{S} : Verlet. \mathcal{L} : Lie RK2(CF2) and Lie RK4(CF4) [21]. We leave the precise details to Apps C and D.

Toy Two-Body Problem. We consider an illustrative two-body problem to demonstrate the effects of V_{rigid} . In Fig. 1 ‘Point’ & ‘Rigid’ denote exact solutions for a point-mass and rigid-body potential, and ‘Lie T2’ the prediction of our method based on a V learned from data. Compared to ‘Point’, ‘Rigid’ induces an *apsidal precession* (rotation of the orbital axis) due to spin-orbit couplings. Our method successfully predicts this interaction and matches the trajectory of ‘Rigid’.




Figure 3: Results on TRAPPIST-1 with short data separation. **Top**: $SO(3)$ manifold error. **Bottom**: Hamiltonian error over the integrated trajectory. Only Lie T2 achieves low errors in both metrics.

Table 1: TRAPPIST-1 with long data separation evaluated after 500 integrator steps. All methods except for Lie T_2 diverge during training. $\|\Delta\mathbf{R}\|$ is computed as the geodesic between the predicted and true rotations (see App E for details).

| | $\ \Delta q\ _2$ | $\ \Delta R\ $ | $\ \Delta\dot{p}\ _2$ | $\ \Delta\dot{\Pi}\ _2$ |
|--|------------------|----------------|-----------------------|-------------------------|
| No Corrections | 8.6e-6 | 1.6 | 3.3 | 4.7e-2 |
| Euler, RK4, Verlet, Lie RK2, Lie RK4 | ∞ | ∞ | ∞ | ∞ |
| Lie T_2 (Ours) | 6.1e-8 | 4.0e-2 | 1.2e-1 | 3.1e-3 |

Table 2: TRAPPIST-1 with short data separation evaluated after 500 integrator steps.

| | $\ \Delta q\ _2$ | $\ \Delta R\ $ | $\ \Delta\frac{\partial V}{\partial \mathbf{q}}\ _2$ | $\ \Delta\frac{\partial V}{\partial \mathbf{R}}\ _2$ |
|------------------------------------|------------------|----------------|--|--|
| Euler | ∞ | ∞ | ∞ | ∞ |
| RK4 | 1.6e-5 | 3.9e-1 | 1.5e+3 | 7.9e-1 |
| S Verlet | 8.6e-8 | 2.0e+7 | 1.6e+1 | 4.1e-1 |
| L Lie RK2 | 5.3e-5 | 3.3e-1 | 2.6e+2 | 7.6e-1 |
| L Lie RK4 | 1.6e-5 | 3.5e-1 | 1.5e+3 | 8.4e-1 |
| S L Lie T_2 (Ours) | 8.0e-8 | 2.4e-1 | 1.6e+1 | 4.0e-1 |

We next test our method by learning the dynamics of the TRAPPIST-1 system [40] which consists of seven earth-sized planets and is notable for potential habitability for terrestrial forms of lives.

TRAPPIST-1, Large Δt . To answer **Q1**, we choose a large data timestep $\Delta t = 2.4e-3$ yr. The closest planet has an orbital period of $\sim 2\Delta t$ (4.1e-3 yr), while the rigid body correction, tidal force and GR correction act on much longer scales. Only Lie T_2 successfully trains. All other methods diverge during training (denoted by ∞) despite attempts at stabilization with techniques such as normalization (LayerNorm [41], GroupNorm [42]). Reducing h improves integration accuracy, but increases the RNN depth and makes training more unstable. We compare with the solution for point-mass potential only (No Correction). Our method reduces the error up to two orders of magnitude in measures of trajectory error and potential gradients (Table 1). We evaluate each method by computing the norm of the error of the predicted positions and forces. The geodesic on $SO(3)$ is used to compute errors between rotations $\|\Delta\mathbf{R}\|$. See App E for more details on how each metric is defined and computed.

TRAPPIST-1, Small Δt . To gain more insight on **Q2**, we shrink Δt until almost all methods can converge and only consider conservative forces (*i.e.* no tidal force or GR). The mean errors in the predicted trajectory and derivatives of the learned potential V after 500 integrator steps are shown in Table 2. Both **S** methods achieve small errors in position related terms. Verlet has a large rotational error since it does not integrate on the rotation manifold. **L** methods achieve lower rotational errors but are worse elsewhere. Lie T_2 being both **S****L** achieves the lowest error on both fronts.

4 Broader Impact

Our work presents an approach for learning multiscale, higher order physics on the Lie-group manifold in the context of non-Newtonian astronomy. This research, though directly applicable to astronomy, can also be applied to perform data-driven discovery multiscale phenomena on Lie groups in other fields.

References

- [1] Samuel Greydanus, Misko Dzamba, and Jason Yosinski. Hamiltonian neural networks. *Advances in neural information processing systems*, 32, 2019.
- [2] Tom Bertalan, Felix Dietrich, Igor Mezić, and Ioannis G Kevrekidis. On learning hamiltonian systems from data. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 29(12):121107, 2019.
- [3] Renyi Chen and Molei Tao. Data-driven prediction of general hamiltonian dynamics via learning exactly-symplectic maps. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 1717–1727. PMLR. URL <https://proceedings.mlr.press/v139/chen21r.html>.
- [4] Anshul Choudhary, John F Lindner, Elliott G Holliday, Scott T Miller, Sudeshna Sinha, and William L Ditto. Physics-enhanced neural networks learn order and chaos. *Physical Review E*, 101(6):062207, 2020.

- [5] Chen-Di Han, Bryan Glaz, Mulugeta Haile, and Ying-Cheng Lai. Adaptable hamiltonian neural networks. *Physical Review Research*, 3(2):023156, 2021.
- [6] Zhengdao Chen, Jianyu Zhang, Martin Arjovsky, and Léon Bottou. Symplectic recurrent neural networks. *arXiv preprint arXiv:1909.13334*, 2019.
- [7] Molei Tao. Explicit symplectic approximation of nonseparable hamiltonians: Algorithm and long time performance. *Physical Review E*, 94(4):043303, 2016.
- [8] Shiyong Xiong, Yunjin Tong, Xingzhe He, Shuqi Yang, Cheng Yang, and Bo Zhu. Nonseparable symplectic neural networks. *arXiv preprint arXiv:2010.12636*, 2020.
- [9] Daniel DiPietro, Shiyong Xiong, and Bo Zhu. Sparse symplectically integrated neural networks. *Advances in Neural Information Processing Systems*, 33:6074–6085, 2020.
- [10] Marco David and Florian Méhats. Symplectic learning for hamiltonian neural networks. *arXiv preprint arXiv:2106.11753*, 2021.
- [11] Frederik Baymler Mathiesen, Bin Yang, and Jilin Hu. Hyperverlet: A symplectic hypersolver for hamiltonian systems. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 4575–4582, 2022.
- [12] Pengzhan Jin, Zhen Zhang, Aiqing Zhu, Yifa Tang, and George Em Karniadakis. Sympnets: Intrinsic structure-preserving symplectic networks for identifying hamiltonian systems. *Neural Networks*, 132:166–179, 2020.
- [13] Marc Finzi, Ke Alexander Wang, and Andrew G Wilson. Simplifying hamiltonian and lagrangian neural networks via explicit constraints. *Advances in neural information processing systems*, 33:13880–13889, 2020.
- [14] Yaofeng Desmond Zhong, Biswadip Dey, and Amit Chakraborty. Extending lagrangian and hamiltonian neural networks with differentiable contact models. *Advances in Neural Information Processing Systems*, 34:21910–21922, 2021.
- [15] Yupu Lu, Shijie Lin, Guanqi Chen, and Jia Pan. Modlanets: Learning generalisable dynamics via modularity and physical inductive bias. In *International Conference on Machine Learning*, pages 14384–14397. PMLR, 2022.
- [16] Thai Duong and Nikolay Atanasov. Hamiltonian-based neural ode networks on the se (3) manifold for dynamics learning and control. *arXiv preprint arXiv:2106.12782*, 2021.
- [17] Arieh Iserles, Hans Z Munthe-Kaas, Syvert P Nørsett, and Antonella Zanna. Lie-group methods. *Acta numerica*, 9:215–365, 2000.
- [18] Nawaf Bou-Rabee and Jerrold E Marsden. Hamilton–pontryagin integrators on lie groups part i: Introduction and structure-preserving properties. *Foundations of computational mathematics*, 9(2):197–219, 2009.
- [19] Elena Celledoni, Håkon Marthinsen, and Brynjulf Owren. An introduction to lie group integrators–basics, new developments and applications. *Journal of Computational Physics*, 257: 1040–1061, 2014.
- [20] Renyi Chen, Gongjie Li, and Molei Tao. Grit: A package for structure-preserving simulations of gravitationally interacting rigid bodies. *The Astrophysical Journal*, 919(1):50, 2021.
- [21] Elena Celledoni, Ergys Çokaj, Andrea Leone, Davide Murari, and Brynjulf Owren. Lie group integrators for mechanical systems. *International Journal of Computer Mathematics*, 99(1): 58–88, 2022.
- [22] Ernst Haier, Christian Lubich, and Gerhard Wanner. *Geometric Numerical integration: structure-preserving algorithms for ordinary differential equations*. Springer, 2006.
- [23] Benedict Leimkuhler and Sebastian Reich. *Simulating hamiltonian dynamics*. Number 14. Cambridge university press, 2004.

- [24] Sergio Blanes and Fernando Casas. *A concise introduction to geometric numerical integration*. CRC press, 2017.
- [25] Jesus-Maria Sanz-Serna and Mari-Paz Calvo. *Numerical hamiltonian problems*. Courier Dover Publications, 2018.
- [26] E. Hairer, Christian Lubich, and Gerhard Wanner. *Geometric Numerical Integration: Structure-Preserving Algorithms for Ordinary Differential Equations*. Number 31 in Springer Series in Computational Mathematics. Springer, 2nd ed edition.
- [27] Taeyoung Lee, N Harris McClamroch, and Melvin Leok. A lie group variational integrator for the attitude dynamics of a rigid body with applications to the 3D pendulum. In *Proceedings of 2005 IEEE Conference on Control Applications, 2005. CCA 2005.*, pages 962–967. IEEE. doi: 10.1109/CCA.2005.1507254. URL <http://ieeexplore.ieee.org/document/1507254/>.
- [28] Billy Quarles, Gongjie Li, and Jack J Lissauer. Milankovitch cycles for a circumstellar earth-analogue within α centauri-like binaries. *Monthly Notices of the Royal Astronomical Society*, 509(2):2736–2757, 2022.
- [29] Renyi Chen, Gongjie Li, and Molei Tao. Low spin-axis variations of circumbinary planets. *Monthly Notices of the Royal Astronomical Society*, 515(4):5175–5184, 2022.
- [30] Sarah Millholland and Gregory Laughlin. Obliquity-driven sculpting of exoplanetary systems. *Nature Astronomy*, 3:424–433, March 2019. doi: 10.1038/s41550-019-0701-7.
- [31] Rosemary A. Mardling and D. N. C. Lin. Calculating the Tidal, Spin, and Dynamical Evolution of Extrasolar Planetary Systems. 573(2):829–844. ISSN 0004-637X, 1538-4357. doi: 10.1086/340752. URL <https://iopscience.iop.org/article/10.1086/340752>.
- [32] Smadar Naoz, Will M Farr, Yoram Lithwick, Frederic A Rasio, and Jean Teysandier. Hot jupiters from secular planet–planet interactions. *Nature*, 473(7346):187–189, 2011.
- [33] Gerald Maurice Clemence. The relativity effect in planetary motions. *Reviews of Modern Physics*, 19(4):361, 1947.
- [34] P Hut. Tidal evolution in close binary systems. *Astronomy and Astrophysics*, 99:126–140, 1981.
- [35] Luc Blanchet. Gravitational radiation from post-newtonian sources and inspiralling compact binaries. *Living reviews in relativity*, 17(1):1–187, 2014.
- [36] Nate Gruver, Marc Anton Finzi, Samuel Don Stanton, and Andrew Gordon Wilson. Deconstructing the inductive biases of hamiltonian neural networks. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=EDeVYpT42oS>.
- [37] Sam Greydanus and Andrew Sosanya. Dissipative hamiltonian neural networks: Learning dissipative and conservative dynamics separately. *arXiv preprint arXiv:2201.10085*, 2022.
- [38] Miles Cranmer, Sam Greydanus, Stephan Hoyer, Peter Battaglia, David Spergel, and Shirley Ho. Lagrangian neural networks. *arXiv preprint arXiv:2003.04630*, 2020.
- [39] Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. On the difficulty of training recurrent neural networks. In *International conference on machine learning*, pages 1310–1318. PMLR, 2013.
- [40] Michaël Gillon, Emmanuël Jehin, Susan M Lederer, Laetitia Delrez, Julien de Wit, Artem Burdanov, Valérie Van Grootel, Adam J Burgasser, Amaury HMJ Triaud, Cyrielle Opitom, et al. Temperate earth-sized planets transiting a nearby ultracool dwarf star. *Nature*, 533(7602): 221–224, 2016.
- [41] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- [42] Yuxin Wu and Kaiming He. Group normalization. In *Proceedings of the European conference on computer vision (ECCV)*, pages 3–19, 2018.

- [43] Jerrold E Marsden and Tudor S Ratiu. *Introduction to mechanics and symmetry: a basic exposition of classical mechanical systems*, volume 17. Springer Science & Business Media, 2013.
- [44] Darryl D Holm, Tanya Schmah, and Cristina Stoica. *Geometric mechanics and symmetry: from finite to infinite dimensions*, volume 12. Oxford University Press, 2009.
- [45] James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018. URL <http://github.com/google/jax>.
- [46] Tom Hennigan, Trevor Cai, Tamara Norman, and Igor Babuschkin. Haiku: Sonnet for JAX, 2020. URL <http://github.com/deepmind/dm-haiku>.
- [47] Stefan Elfving, Eiji Uchibe, and Kenji Doya. Sigmoid-weighted linear units for neural network function approximation in reinforcement learning. *Neural Networks*, 107:3–11, 2018.
- [48] Juho Lee, Yoonho Lee, Jungtaek Kim, Adam Kosiorek, Seungjin Choi, and Yee Whye Teh. Set transformer: A framework for attention-based permutation-invariant neural networks. In *International conference on machine learning*, pages 3744–3753. PMLR, 2019.
- [49] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- [50] Yaofeng Desmond Zhong, Biswadip Dey, and Amit Chakraborty. Symplectic ODE-Net: Learning Hamiltonian Dynamics with Control. URL <http://arxiv.org/abs/1909.12077>.
- [51] Yaofeng Desmond Zhong, Biswadip Dey, and Amit Chakraborty. Benchmarking energy-conserving neural networks for learning dynamics from data. In *Learning for Dynamics and Control*, pages 1218–1229. PMLR, 2021.
- [52] Elena Celledoni, Arne Marthinsen, and Brynjulf Owren. Commutator-free Lie group methods. 19(3):341–352. ISSN 0167739X. doi: 10.1016/S0167-739X(02)00161-9. URL <https://linkinghub.elsevier.com/retrieve/pii/S0167739X02001619>.

Checklist

The checklist follows the references. Please read the checklist guidelines carefully for information on how to answer these questions. For each question, change the default **[TODO]** to **[Yes]**, **[No]**, or **[N/A]**. You are strongly encouraged to include a **justification to your answer**, either by referencing the appropriate section of your paper or providing a brief inline description. For example:

- Did you include the license to the code and datasets? **[Yes]** See Section ??.
- Did you include the license to the code and datasets? **[No]** The code and the data are proprietary.
- Did you include the license to the code and datasets? **[N/A]**

Please do not modify the questions and only use the provided macros for your answers. Note that the Checklist section does not count towards the page limit. In your paper, please delete this instructions block and only keep the Checklist section heading above along with the questions/answers below.

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? **[Yes]**
 - (b) Did you describe the limitations of your work? **[Yes]** See App F. Also see App C for a discussion on how the assumption of no prior knowledge on the structure of the potential limits the scalability of our method in comparison to exploiting the pairwise nature of the gravitational potential.
 - (c) Did you discuss any potential negative societal impacts of your work? **[N/A]**
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? **[Yes]**
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? **[Yes]** See Section 2 for the assumptions of the underlying theoretical framework.
 - (b) Did you include complete proofs of all theoretical results? **[Yes]** See Section 2 and Apps A and B for the derivation of the equations.
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? **[No]**
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? **[Yes]** See App C.
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? **[No]**
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? **[Yes]** See App C.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? **[Yes]** We, we cite all existing code used in App C.
 - (b) Did you mention the license of the assets? **[Yes]** See footnotes in App C.
 - (c) Did you include any new assets either in the supplemental material or as a URL? **[N/A]**
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? **[N/A]**
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? **[N/A]**
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? **[N/A]**

- (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
- (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

A Rigid Body Equations of Motion

We present two derivations from Chen et al. [20]. The first derivation derives a constrained Hamiltonian system via Lagrange multipliers, while the second uses a variational principle for mechanics on Lie groups. We also note that the same equations can be derived via the use of the Port-Hamiltonian framework as in Duong and Atanasov [16], though they choose to express the linear momentum coordinates in the moving body frame rather than the inertial frame as is done in this work.

A.1 Constrained Hamiltonian System

Since the manifold $SE(3)^{\otimes N}$ is a product space of the individual manifolds $SE(3)$, we consider the latter for brevity and drop indices i for each body. Furthermore, $SE(3)$ can be seen as the product space of $\mathbb{R}^3 \times SO(3)$ for which the former is unconstrained, so we focus our attention on the latter. We can view \mathbf{R} to be in the embedded Euclidean space $\mathbf{R}^{3 \times 3} \leftarrow SO(3)$ and use $\mathbf{R} \in SO(3) := h(\mathbf{R}) = \{\mathbf{R}^\top \mathbf{R} - \mathbf{I}_{3 \times 3} = \mathbf{0}_{3 \times 3}\}$ as a holonomic constraint. Using Lagrange multipliers Λ [26] for the constraint h gives us the following Lagrangian

$$L(\mathbf{R}, \dot{\mathbf{R}}) = \frac{1}{2} \text{Tr} \left[\dot{\mathbf{R}} \mathbf{J}_d \dot{\mathbf{R}} \right] - V(\mathbf{R}) - \frac{1}{2} \text{Tr} \left[\Lambda^\top (\mathbf{R}^\top \mathbf{R} - \mathbf{I}_{3 \times 3}) \right] \quad (\text{A1})$$

where \mathbf{J}_d denotes the nonstandard moment of inertia [20, 27]

$$\mathbf{J}_d := \int_{\mathcal{B}} \rho(\mathbf{x}) \mathbf{x} \mathbf{x}^\top d\mathbf{x}, \quad \mathbf{J}_d = \text{Tr}[\mathbf{J}] \mathbf{I}_{3 \times 3} - \mathbf{J}, \quad \mathbf{J} = \text{Tr}[\mathbf{J}_d] \mathbf{I}_{3 \times 3} - \mathbf{J}_d, \quad (\text{A2})$$

and

$$\Lambda = \begin{bmatrix} \lambda_1 & \lambda_4 & \lambda_6 \\ \lambda_4 & \lambda_2 & \lambda_5 \\ \lambda_6 & \lambda_5 & \lambda_3 \end{bmatrix} \in \mathbb{R}^{3 \times 3} \quad (\text{A3})$$

is a 6-dimensional symmetric matrix of Lagrange multipliers. Performing the Legendre transform for (A1) gives us the conjugate rotational momentum \mathbf{P} as

$$\mathbf{P} = \frac{\partial L(\mathbf{R}, \dot{\mathbf{R}})}{\partial \dot{\mathbf{R}}} = \dot{\mathbf{R}} \mathbf{J}_d \quad (\text{A4})$$

and corresponding Hamiltonian

$$H(\mathbf{R}, \mathbf{P}) = \frac{1}{2} \text{Tr} \left[\mathbf{P} \mathbf{J}_d^{-1} \mathbf{P} \right] + V(\mathbf{R}) + \frac{1}{2} \text{Tr} \left[\Lambda^\top (\mathbf{R}^\top \mathbf{R} - \mathbf{I}_{3 \times 3}) \right] \quad (\text{A5})$$

The constraint for \mathbf{P} can be obtained by taking the time derivative of $h(\mathbf{R}) = 0$ according to Haier et al. [22], i.e., $\mathbf{J}_d^{-1} \dot{\mathbf{P}}^\top \mathbf{R} + \mathbf{R}^\top \mathbf{P} \mathbf{J}_d^{-1} = \mathbf{0}_{3 \times 3}$. Hence, our equations of motion so far looks like

$$\begin{cases} \dot{\mathbf{R}} = \frac{\partial H}{\partial \mathbf{P}} = \mathbf{P} \mathbf{J}_d^{-1} \\ \dot{\mathbf{P}} = -\frac{\partial H}{\partial \mathbf{R}} = -\frac{\partial V}{\partial \mathbf{R}} - \mathbf{R} \Lambda \end{cases} \quad (\text{A6})$$

on the manifold

$$\mathcal{M} := \{(\mathbf{R}, \mathbf{P}) | \mathbf{R}^\top \mathbf{R} = \mathbf{I}_{3 \times 3}, \quad \mathbf{J}_d^{-1} \mathbf{P}^\top \mathbf{R} + \mathbf{R}^\top \mathbf{P} \mathbf{J}_d^{-1} = \mathbf{0}_{3 \times 3}\} \quad (\text{A7})$$

Let $\hat{\Omega} = \mathbf{R}^\top \dot{\mathbf{R}}$ denote the body's angular velocity in the inertial frame, where \wedge denotes the cross-product operation $\hat{u}v = u \times v$. Then, it can written using P as

$$\hat{\Omega} = \mathbf{R}^\top \mathbf{P} \mathbf{J}_d^{-1} \quad (\text{A8})$$

Taking the time derivative of $\hat{\Omega}$ gives

$$\dot{\hat{\Omega}} = \mathbf{J}_d^{-1} \dot{\mathbf{P}}^\top \mathbf{P} \mathbf{J}_d^{-1} + \mathbf{R}^\top \left(-\frac{\partial V(\mathbf{R})}{\partial \mathbf{R}} - \mathbf{R} \Lambda \right) \mathbf{J}_d^{-1} \quad (\text{A9})$$

Since our goal is to represent the dynamics of Π , we convert the above to the body frame momentum using $\Pi = \mathbf{J} \Omega$. Using properties of the hat map [20, Appendix A.1] gives us

$$\hat{\Pi} = \hat{\mathbf{J}} \hat{\Omega} = \text{Tr}[\mathbf{J}_d] \hat{\Omega} - \widehat{\mathbf{J}_d \Omega} = \hat{\Omega} \mathbf{J}_d - \mathbf{J}_d \hat{\Omega}^\top. \quad (\text{A10})$$

Taking the time derivative of the above then gives

$$\dot{\hat{\Pi}} = \dot{\hat{\Omega}}\mathbf{J}_d - \mathbf{J}_d\dot{\hat{\Omega}}, \quad (\text{A11})$$

$$= (\mathbf{J}_d^{-1}\mathbf{P}^\top\mathbf{P} - \mathbf{P}^\top\mathbf{P}\mathbf{J}_d^{-1}) - \left(\mathbf{R}^\top \frac{\partial V(\mathbf{R})}{\partial \mathbf{R}} + \left(\frac{\partial V(\mathbf{R})}{\partial \mathbf{R}} \right)^\top \mathbf{R} \right) - (\mathbf{\Lambda} - \mathbf{\Lambda}^\top), \quad (\text{A12})$$

$$= (\mathbf{J}_d^{-1}\mathbf{P}^\top\mathbf{P} - \mathbf{P}^\top\mathbf{P}\mathbf{J}_d^{-1}) - \left(\mathbf{R}^\top \frac{\partial V(\mathbf{R})}{\partial \mathbf{R}} + \left(\frac{\partial V(\mathbf{R})}{\partial \mathbf{R}} \right)^\top \mathbf{R} \right), \quad (\text{A13})$$

where the symmetric $\mathbf{\Lambda}$ vanishes on the last line.

Since $\mathbf{P} = \mathbf{R}\hat{\Omega}\mathbf{J}_d$, applying properties of the hat map [20, Appendix A.1] further simplifies the above as

$$\dot{\hat{\Pi}} = \left(\hat{\Omega}^\top \hat{\Omega} \mathbf{J}_d - \mathbf{J}_d \hat{\Omega}^\top \hat{\Omega} \right) - \left(\mathbf{R}^\top \frac{\partial V(\mathbf{R})}{\partial \mathbf{R}} + \left(\frac{\partial V(\mathbf{R})}{\partial \mathbf{R}} \right)^\top \mathbf{R} \right), \quad (\text{A14})$$

$$= \widehat{\Omega \times \mathbf{J}_d \Omega} - \left(\mathbf{R}^\top \frac{\partial V(\mathbf{R})}{\partial \mathbf{R}} + \left(\frac{\partial V(\mathbf{R})}{\partial \mathbf{R}} \right)^\top \mathbf{R} \right). \quad (\text{A15})$$

Thus, applying the vee map and converting back to \mathbf{J} and $\mathbf{\Pi}$ gives

$$\dot{\mathbf{\Pi}} = \Omega \times \mathbf{J}_d \Omega - \left(\mathbf{R}^\top \frac{\partial V(\mathbf{R})}{\partial \mathbf{R}} + \left(\frac{\partial V(\mathbf{R})}{\partial \mathbf{R}} \right)^\top \mathbf{R} \right)^\vee, \quad (\text{A16})$$

$$= \Omega \times (\cdot) \text{Tr}[\mathbf{J}] - \mathbf{J} \Omega - \left(\mathbf{R}^\top \frac{\partial V(\mathbf{R})}{\partial \mathbf{R}} + \left(\frac{\partial V(\mathbf{R})}{\partial \mathbf{R}} \right)^\top \mathbf{R} \right)^\vee, \quad (\text{A17})$$

$$= -\Omega \times \mathbf{J} \Omega - \left(\mathbf{R}^\top \frac{\partial V(\mathbf{R})}{\partial \mathbf{R}} + \left(\frac{\partial V(\mathbf{R})}{\partial \mathbf{R}} \right)^\top \mathbf{R} \right)^\vee, \quad (\text{A18})$$

$$= \mathbf{\Pi} \times \mathbf{J}^{-1} \mathbf{\Pi} - \left(\mathbf{R}^\top \frac{\partial V(\mathbf{R})}{\partial \mathbf{R}} + \left(\frac{\partial V(\mathbf{R})}{\partial \mathbf{R}} \right)^\top \mathbf{R} \right)^\vee. \quad (\text{A19})$$

Hence, the rotational equations of motion on $\text{SO}(3)$ are

$$\begin{cases} \dot{\mathbf{R}} = \mathbf{R}\hat{\mathbf{J}}\mathbf{\Pi}, \\ \dot{\mathbf{\Pi}} = \mathbf{\Pi} \times \mathbf{J}^{-1} \mathbf{\Pi} - \left(\mathbf{R}^\top \frac{\partial V(\mathbf{R})}{\partial \mathbf{R}} - \left(\frac{\partial V(\mathbf{R})}{\partial \mathbf{R}} \right)^\top \mathbf{R} \right)^\vee. \end{cases} \quad (\text{A20})$$

A.2 Variational Principle for Mechanics on the Lie Group

We can also apply the Euler-Lagrange equations for Hamilton's variational principle on a Lie group, a topic that has been well studied (e.g., Marsden and Ratiu [43], Holm et al. [44]) [20], and we summarize the results for the special case of rigid bodies from the expository part of Lee et al. [27].

Denote the infinitesimally varied rotation by $\mathbf{R}_\epsilon = \mathbf{R} \exp(\epsilon \hat{\eta})$, with $\epsilon \in \mathbb{R}$ and $\eta \in \mathbb{R}^3$, where $\exp(\cdot)$ is the exponential map from $\mathfrak{so}(3)$ to $\text{SO}(3)$. The varied angular velocity $\hat{\Omega}_\epsilon$ is

$$\hat{\Omega}_\epsilon = \mathbf{R}_\epsilon^\top \dot{\mathbf{R}}_\epsilon = \exp(-\epsilon \hat{\eta}) \left(\dot{\mathbf{R}} \exp(\epsilon \hat{\eta}) + \mathbf{R} \exp(\epsilon \hat{\eta}) \epsilon \hat{\eta} \right), \quad (\text{A21})$$

$$= \exp(-\epsilon \hat{\eta}) \hat{\Omega} \exp(-\epsilon \hat{\eta}) + \exp(-\epsilon \hat{\eta}) \hat{\eta}, \quad (\text{A22})$$

$$= \hat{\Omega} + \epsilon \left(\hat{\eta} + \hat{\Omega} \hat{\eta} - \hat{\eta} \hat{\Omega} \right) + O(\epsilon^2). \quad (\text{A23})$$

Consider the action

$$S(\Omega, \mathbf{R}) = \int_{t_0}^{t_1} L(\Omega, \mathbf{R}) dt = \int_{t_0}^{t_1} \frac{1}{2} \text{Tr} \left[\hat{\Omega} \mathbf{J}_d \hat{\Omega} \right] - V(\mathbf{R}) dt. \quad (\text{A24})$$

Taking the variation S_ϵ of the action S , we have

$$\begin{aligned}
S_\epsilon(\boldsymbol{\Omega}, \mathbf{R}) &= S(\boldsymbol{\Omega}_\epsilon, \mathbf{R}_\epsilon), \\
&= S(\boldsymbol{\Omega}, \mathbf{R}) + \epsilon \int_{t_0}^{t_1} \left\{ \frac{1}{2} \text{Tr} \left[-\hat{\eta} \left(\mathbf{J}_d \hat{\boldsymbol{\Omega}} + \hat{\boldsymbol{\Omega}} \mathbf{J}_d \right) \right. \right. \\
&\quad \left. \left. + \hat{\eta} \hat{\boldsymbol{\Omega}} \left(\mathbf{J}_d \hat{\boldsymbol{\Omega}} + \hat{\boldsymbol{\Omega}} \mathbf{J}_d \right) - \hat{\eta} \left(\mathbf{J}_d \hat{\boldsymbol{\Omega}} + \hat{\boldsymbol{\Omega}} \mathbf{J}_d \hat{\boldsymbol{\Omega}} \right) \right] + \text{Tr} \left[\hat{\eta} \mathbf{R}^\top \frac{\partial V(\mathbf{R})}{\partial \mathbf{R}} \right] \right\} dt \\
&\quad + O(\epsilon^2).
\end{aligned} \tag{A25}$$

Using *Hamilton's Principle*, we have $\frac{d}{d\epsilon} \Big|_{\epsilon=0} S_\epsilon = 0$, i.e.,

$$\frac{1}{2} \int_{t_0}^{t_1} \text{Tr} \left[\hat{\eta} \left\{ \widehat{\mathbf{J}}\boldsymbol{\Omega} + \boldsymbol{\Omega} \times \widehat{\mathbf{J}}\boldsymbol{\Omega} + 2\mathbf{R}^\top \frac{\partial V(\mathbf{R})}{\partial \mathbf{R}} \right\} \right] dt = 0, \tag{A26}$$

for any $\eta \in \mathbb{R}^3$. Hence, $\widehat{\mathbf{J}}\boldsymbol{\Omega} + \boldsymbol{\Omega} \times \widehat{\mathbf{J}}\boldsymbol{\Omega} + 2\mathbf{R}^\top \frac{\partial V(\mathbf{R})}{\partial \mathbf{R}}$ must be skew-symmetric, giving us

$$\widehat{\mathbf{J}}\boldsymbol{\Omega} = -\boldsymbol{\Omega} \times \widehat{\mathbf{J}}\boldsymbol{\Omega} + \left(\frac{\partial V(\mathbf{R})}{\partial \mathbf{R}}^\top \mathbf{R} - \mathbf{R}^\top \frac{\partial V(\mathbf{R})}{\partial \mathbf{R}} \right). \tag{A27}$$

Thus, by definition of $\boldsymbol{\Pi}$ and applying the vee map, we recover the same update

$$\dot{\boldsymbol{\Pi}} = \boldsymbol{\Pi} \times \mathbf{J}^{-1} \boldsymbol{\Pi} + \left(\frac{\partial V(\mathbf{R})}{\partial \mathbf{R}}^\top \mathbf{R} - \mathbf{R}^\top \frac{\partial V(\mathbf{R})}{\partial \mathbf{R}} \right)^\vee. \tag{A28}$$

B Details of the Lie \mathbb{T}_2 splitting integrator

The full Hamiltonian H takes the form

$$\mathcal{H}(\mathbf{q}, \mathbf{R}, \mathbf{p}, \boldsymbol{\Pi}) = \sum_{i=1}^N \frac{1}{2} \mathbf{p}_i^\top \mathbf{p}_i / m_i + \frac{1}{2} \boldsymbol{\Pi}_i^\top \mathbf{J}_i^{-1} \boldsymbol{\Pi}_i + V(\mathbf{q}, \mathbf{R}) \tag{B1}$$

We borrow from Chen et al. [20] the idea of Lie-group and symplecticity preserving splitting and split our Hamiltonian as $\mathcal{H} = \mathcal{H}_{\text{KE}} + \mathcal{H}_{\text{PE}} + \mathcal{H}_{\text{asym}}$, where

$$\mathcal{H}_{\text{KE}} := \sum_{i=1}^N \frac{1}{2} \mathbf{p}_i^\top \mathbf{p}_i / m_i + \frac{1}{2} \boldsymbol{\Pi}_i^\top \mathbf{J}_{i,\text{sym}}^{-1} \boldsymbol{\Pi}_i \tag{B2}$$

$$\mathcal{H}_{\text{PE}} := V(\mathbf{q}, \mathbf{R}) \tag{B3}$$

$$\mathcal{H}_{\text{asym}} := \sum_{i=1}^N \frac{1}{2} \boldsymbol{\Pi}_i^\top \mathbf{J}_{i,\text{asym}}^{-1} \boldsymbol{\Pi}_i \tag{B4}$$

where we assume that \mathbf{J}_i is axis-symmetric, i.e.,

$$\mathbf{J}_i := \begin{bmatrix} J_i^{(1)} & 0 & 0 \\ 0 & J_i^{(2)} & 0 \\ 0 & 0 & J_i^{(3)} \end{bmatrix} \tag{B5}$$

and $\mathbf{J}_{i,\text{sym}}^{-1}$ and $\mathbf{J}_{i,\text{asym}}^{-1}$ denote the axial-symmetric and residual terms of \mathbf{J}_i^{-1} such that $\mathbf{J}_{i,\text{sym}}^{-1} + \mathbf{J}_{i,\text{asym}}^{-1} = \mathbf{J}_i^{-1}$, i.e.,

$$\mathbf{J}_{i,\text{sym}}^{-1} := \begin{bmatrix} 1/J_i^{(1)} & 0 & 0 \\ 0 & 1/J_i^{(1)} & 0 \\ 0 & 0 & 1/J_i^{(3)} \end{bmatrix}, \quad \mathbf{J}_{i,\text{asym}}^{-1} := \begin{bmatrix} 0 & 0 & 0 \\ 0 & 1/J_i^{(2)} & -1/J_i^{(1)} \\ 0 & 0 & 0 \end{bmatrix}. \tag{B6}$$

Then, each of $\mathcal{H}_{\text{KE}} + \mathcal{H}_{\text{PE}} + \mathcal{H}_{\text{asym}}$ can be integrated exactly.

Exact integration of \mathcal{H}_{KE} For \mathcal{H}_{KE} , we have the equations of motion (using (1a)-(1d) but with $V = 0$)

$$\begin{cases} \dot{\mathbf{q}}_i = \mathbf{p}_i/m_i & (\text{B7a}) \\ \dot{\mathbf{p}}_i = 0 & (\text{B7b}) \\ \dot{\mathbf{R}}_i = \mathbf{R}_i \widehat{\mathbf{J}_{i,\text{sym}}^{-1} \boldsymbol{\Pi}_i} & (\text{B7c}) \\ \dot{\boldsymbol{\Pi}}_i = \boldsymbol{\Pi}_i \times \mathbf{J}_{i,\text{sym}}^{-1} \boldsymbol{\Pi}_i & (\text{B7d}) \end{cases}$$

The equation for $\dot{\boldsymbol{\Pi}}_i$ (B7d) is the Euler equation for a free rigid body [20]. It is exactly solvable with a simple expression for axial-symmetric bodies, since in this case $\widehat{\boldsymbol{\Pi}}$ simplifies as

$$\dot{\boldsymbol{\Pi}}_i = \boldsymbol{\Pi}_i \times \mathbf{J}_{i,\text{sym}}^{-1} \boldsymbol{\Pi}_i = \begin{bmatrix} \left(1/J_i^{(3)} - 1/J_i^{(1)}\right) \Pi_{i,y} \Pi_{i,z} \\ -\left(1/J_i^{(3)} - 1/J_i^{(1)}\right) \Pi_{i,y} \Pi_{i,z} \\ 0 \end{bmatrix} \quad (\text{B8})$$

Consequently, $\Pi_{i,z}(t) = \Pi_{i,z}(0)$, meaning that we can express the above as the linear differential equation

$$\dot{\boldsymbol{\Pi}}_i(t) = \begin{bmatrix} 0 & \left(1/J_i^{(3)} - 1/J_i^{(1)}\right) \Pi_{i,z}(0) & 0 \\ -\left(1/J_i^{(3)} - 1/J_i^{(1)}\right) \Pi_{i,z}(0) & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \boldsymbol{\Pi}_i(0) \quad (\text{B9})$$

$$= -\theta t \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \widehat{\boldsymbol{\Pi}}_i(0) \quad (\text{B10})$$

where $\theta := \left(1/J_i^{(3)} - 1/J_i^{(1)}\right) \Pi_{i,z}(0)$ and has the solution

$$\boldsymbol{\Pi}_i(t) = \exp\left(-\theta t \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}\right) \boldsymbol{\Pi}_i(0) = R_z^\top(\theta t) \boldsymbol{\Pi}_i(0) \quad (\text{B11})$$

where R_z denotes the rotation matrix around the z axis. Taking the above back to (B7c) then gives us the solution for \mathbf{R}_i as well, giving the flow $\phi_h^{[\text{KE}]}$ of \mathcal{H}_{KE} as

$$\begin{cases} \mathbf{q}_i(h) = \mathbf{q}_i(0) + \mathbf{p}_i/m_i h, \\ \mathbf{p}_i(h) = \mathbf{p}_i(0), \\ \mathbf{R}_i(h) = \mathbf{R}_i(0) R_{\boldsymbol{\Pi}_i(0)} \left(\|\boldsymbol{\Pi}_i(0)\| h / J_i^{(1)} \right) R_z(\theta h), \\ \boldsymbol{\Pi}_i(h) = R_z^\top(\theta h) \boldsymbol{\Pi}_i(0). \end{cases} \quad (\text{B12})$$

Exact integration of \mathcal{H}_{PE} For \mathcal{H}_{PE} , the equations of motion are [20]

$$\begin{cases} \dot{\mathbf{q}}_i = 0, \dot{\mathbf{p}}_i = -\frac{\partial V}{\partial \mathbf{q}_i}, \\ \dot{\mathbf{R}}_i = 0, \dot{\boldsymbol{\Pi}}_i = -\left(\mathbf{R}_i^\top \frac{\partial V}{\partial \mathbf{R}_i} - \left(\frac{\partial V}{\partial \mathbf{R}_i} \right)^\top \mathbf{R}_i \right)^\vee. \end{cases} \quad (\text{B13})$$

Since \mathbf{q}_i and \mathbf{p}_i stay constant, \mathbf{p}_i and $\boldsymbol{\Pi}_i$ change at constant rates. Hence, the flow $\phi_h^{[\text{PE}]}$ of \mathcal{H}_{PE} is given by

$$\begin{cases} \mathbf{q}_i(h) = \mathbf{q}_i(0) \\ \mathbf{p}_i(h) = \mathbf{p}_i(0) - \frac{\partial V}{\partial \mathbf{q}_i} h, \\ \mathbf{R}_i(h) = \mathbf{R}_i(0), \\ \boldsymbol{\Pi}_i(h) = \boldsymbol{\Pi}_i(0) - \left(\mathbf{R}_i^\top \frac{\partial V}{\partial \mathbf{R}_i} - \left(\frac{\partial V}{\partial \mathbf{R}_i} \right)^\top \mathbf{R}_i \right)^\vee h. \end{cases} \quad (\text{B14})$$

Exact integration of $\mathcal{H}_{\text{asym}}$ Finally, the equations of motion of $\mathcal{H}_{\text{asym}}$ are given by (again adapting (1a)-(1d))

$$\begin{cases} \dot{\mathbf{q}}_i = 0, \dot{\mathbf{p}}_i = 0, \\ \dot{\mathbf{R}}_i = \mathbf{R}_i \widehat{\mathbf{J}_{i,\text{asym}}^{-1}} \mathbf{\Pi}_i, \dot{\mathbf{\Pi}}_i = \mathbf{\Pi}_i \times \mathbf{J}_{i,\text{asym}}^{-1} \mathbf{\Pi}_i. \end{cases} \quad (\text{B15})$$

which can be solved to obtain the flow $\phi_h^{[\text{asym}]}$ of $\mathcal{H}_{\text{asym}}$ as

$$\begin{cases} \mathbf{q}_i(h) = \mathbf{q}_i(0) \\ \mathbf{p}_i(h) = \mathbf{p}_i(0), \\ \mathbf{R}_i(h) = R_y(\delta \mathbf{\Pi}_i^{(2)} h) \mathbf{R}_i(0), \\ \mathbf{\Pi}_i(h) = R_y(-\delta \mathbf{\Pi}_i^{(2)} h) \mathbf{\Pi}_i(0). \end{cases} \quad (\text{B16})$$

where $\delta := 1/J_i^{(2)} - 1/J_i^{(1)}$.

Having obtained analytical solutions for each of the flows $\phi^{[\text{KE}]}$, $\phi^{[\text{PE}]}$, $\phi^{[\text{asym}]}$, we then combine them with the non-conservative momentum update from the non-conservative forcing terms $F_{\mathbf{p}_i}$ and $F_{\mathbf{\Pi}_i}$ with flow $\phi_h^{[\text{force}]}$

$$\mathbf{q}_i(h) = \mathbf{q}_i(0), \mathbf{p}_i(h) = \mathbf{p}_i(0) + F_{\mathbf{p}_i} h, \mathbf{R}_i(h) = \mathbf{R}_i(0), \mathbf{\Pi}_i(h) = \mathbf{\Pi}_i(0) + F_{\mathbf{\Pi}_i} h. \quad (\text{B17})$$

Consequently, the full Lie T_2 integrator is obtained by applying the Strang composition scheme to obtain

$$\phi_h^{\text{Lie } T_2} := \phi_{h/2}^{[\text{KE}]} \circ \phi_{h/2}^{[\text{PE}]} \circ \phi_{h/2}^{[\text{asym}]} \circ \phi_h^{[\text{force}]} \circ \phi_{h/2}^{[\text{asym}]} \circ \phi_{h/2}^{[\text{PE}]} \circ \phi_{h/2}^{[\text{KE}]} \quad (\text{B18})$$

C Training Details

We implement our method using Jax³ [45] and use the Haiku framework [46]⁴ for constructing the deep neural networks. In all our experiments, we use a multilayer perceptron (MLP) with 3 hidden layers each of size 256 with the SiLU activation [47] for all networks (V_{resid}^θ , $F_{\mathbf{p}}^\theta$, $F_{\mathbf{\Pi}}^\theta$). Each method is run until convergence. Other common training hyperparameters used are summarized in Table 3.

Structure of V_{resid}^θ Note that we **do not** assume prior knowledge on the pairwise structure of the potential function V_{resid} or of the forcing terms $F_{\mathbf{p}}$, $F_{\mathbf{\Pi}}$. More specifically, since the true Gravitational potential (2) only acts *pairwise* between rigid bodies, the rigid body correction potential also acts pairwise and has the structure

$$V_{\text{resid}}(\mathbf{q}, \mathbf{R}) = \sum_{i < j} V_{i,j,\text{resid}}(\mathbf{q}_i, \mathbf{q}_j, \mathbf{R}_i, \mathbf{R}_j). \quad (\text{C1})$$

For the forcing terms, the tidal forcing term also acts pairwise with coupled effects on \mathbf{p}_i and $\mathbf{\Pi}_i$ due the relationship between forces and torques, while the post-Newton general-relativity correction term acts on each planet individually and affects only \mathbf{p}_i . However, in this work, we assume that none of this prior knowledge is available and aim to learn everything purely from data. Hence, we choose to learn the high dimensional forms of V_{resid} , $F_{\mathbf{p}}$, $F_{\mathbf{\Pi}}$. The fact that we were able to obtain improvements despite the high dimensionality of the input space (8 planets each with $(\mathbf{q}_i, \mathbf{R}_i) \in \mathbb{R}^3 \times \mathbb{R}^{3 \times 3} \approx \mathbb{R}^{12}$, i.e., \mathbb{R}^{96} for V and \mathbb{R}^{144} for $F_{\mathbf{p}}$, $F_{\mathbf{\Pi}}$) is an indication of the generality of our approach. We believe that the proposed approach can be made more scalable and have better generalization if we do assume some knowledge about the structure of the forces at play and instead, e.g., learn $V_{\text{resid}}^\theta(\mathbf{q}_i, \mathbf{q}_j, \mathbf{R}_i, \mathbf{R}_j, \phi_i, \phi_j)$ instead of the more general $V_{\text{resid}}^\theta(\mathbf{q}, \mathbf{R})$ and similarly for the forcing terms, where ϕ_i, ϕ_j contains known or potentially learned information about physical properties about each rigid body which are needed. Here, since the ordering of i and j is not important, some type of permutation-invariant encoding such as [48] should be used to further improve generalization.

³<https://github.com/google/jax>. The repository is licensed under Apache-2.0.

⁴<https://github.com/deepmind/dm-haiku>. The repository is licensed under Apache-2.0.

Table 3: Common hyperparameters used for training in all experiments.

| Name | Value |
|---------------|------------|
| Batch size | 256 |
| Optimizer | AdamW [49] |
| Learning Rate | $4e-4$ |

Training All experiments are done on a single RTX 3090 locally. The length of each run is dependent on the complexity of the integrator. For example, explicit Euler is the simplest and has the fastest per-iteration time, while Lie RK4 has the slowest per-iteration time. The Lie T2 integrator has a per-iteration time between the two extremes with each experiment taking approximately three hours.

Data generation We use the GRIT⁵ simulator [20] for generating the dataset in all cases. The GRIT simulator uses a high order Lie-group symplectic integrator derived using a tailored splitting scheme and can handle tidal forces (between all pairs of bodies) and general relativistic effects (between each planet and the host star, assuming that the host star is much more massive than the other planets) using a post-Newtonian correction term. Each dataset consists of 32 different trajectories, which is then split in a 80-20 training-validation split. Each trajectory is then further subsampled to provide 128 tuples $(\mathbf{q}, \mathbf{R}, \mathbf{p}, \mathbf{\Pi})$ of datapoints. Each trajectory is generated by adding small multiplicative Gaussian noise to the coordinates of each body.

C.1 Toy Two-Body Problem

The parameters for this system were hand-picked to provide intuition on how the additional corrections differ from the pure point-mass potential that is widely used in other Hamiltonian learning literature.

C.2 TRAPPIST-1

The initial conditions for this system were taken directly from the TRAPPIST-1 example from GRIT.

D Details on Integrators used for Comparison in Section 3

To answer the question **Q2** of how important are **symplecticity** (\mathcal{S}) and **Lie-group preservation** (\mathcal{L}) for learning, we vary the choice of integrator in our experiments. The integrators used can be broadly split into four categories:

Neither Symplectic \mathcal{S} nor Lie-group preserving \mathcal{L} : This category contains the popular explicit Euler and Runge-Kutta 4 integrators, neither of which are symplectic or Lie-group preserving. The finite-difference scheme from Greydanus et al. [1], Greydanus and Sosanya [37] can be interpreted as an application of the explicit Euler integrator David and Méhats [10]. Explicit Euler is also used in the works of [11]. Runge-Kutta 4 is a popular fourth-order integrator due to its implementation simplicity and is used in Finzi et al. [13], Zhong et al. [14], Duong and Atanasov [16], Zhong et al. [50, 51].

Symplectic \mathcal{S} but not Lie-group preserving \mathcal{L} : This category contains the verlet integrator, which we define loosely in the current work as the integrator obtained by using classical splitting and Strang composition, but *without* performing exact integration on the manifold. In the Euclidean case where $\mathcal{H}(\mathbf{q}, \mathbf{p}) = \frac{1}{2}\mathbf{p}^\top\mathbf{p}/m + V(\mathbf{q})$, this takes the form

$$\mathbf{q}(h/2) = \mathbf{q}(0) + \frac{\mathbf{p}(0)}{m}h, \tag{D1}$$

$$\mathbf{p}(h) = \mathbf{p}(0) - \frac{\partial V}{\partial \mathbf{q}}h, \tag{D2}$$

$$\mathbf{q}(h) = \mathbf{q}(h/2) + \frac{\mathbf{p}(1)}{m}h, \tag{D3}$$

⁵<https://github.com/GRIT-RBSim/GRIT>. The repository is licensed under Apache-2.0.

and corresponds to the ‘‘leapfrog’’ proposed in SRNN Chen et al. [6]. In our setting where the space is not Euclidean, we interpret a ‘‘naive’’ implementation to look like the following

$$\begin{cases} \mathbf{q}_i(h/2) = \mathbf{q}_i(0) + \frac{\mathbf{p}_i(0)}{m}h, \\ \mathbf{R}_i(h/2) = \mathbf{R}_i(0) + \mathbf{R}_i(0) \mathbf{J}_i^{-1} \widehat{\boldsymbol{\Pi}}_i(0)h, \end{cases} \quad (\text{D4a})$$

$$\quad \quad \quad (\text{D4b})$$

$$\begin{cases} \mathbf{p}_i(h) = \mathbf{p}_i(0) - \frac{\partial V}{\partial \mathbf{q}_i}h, \\ \boldsymbol{\Pi}_i(h) = \boldsymbol{\Pi}_i(0) + \left(\boldsymbol{\Pi}_i(0) \times \mathbf{J}_i^{-1} \boldsymbol{\Pi}_i(0) - \left(\mathbf{R}_i^\top(0) \frac{\partial V}{\partial \mathbf{R}_i} - \left(\frac{\partial V}{\partial \mathbf{R}_i} \right)^\top \mathbf{R}_i(0) \right)^\vee \right) h, \end{cases} \quad (\text{D5a})$$

$$\quad \quad \quad (\text{D5b})$$

$$\begin{cases} \mathbf{q}_i(h) = \mathbf{q}_i(h/2) + \frac{\mathbf{p}_i(h)}{m}h, \\ \mathbf{R}_i(h) = \mathbf{R}_i(h/2) + \mathbf{R}_i(h) \mathbf{J}_i^{-1} \widehat{\boldsymbol{\Pi}}_i(h)h, \end{cases} \quad (\text{D6a})$$

$$\quad \quad \quad (\text{D6b})$$

Not Symplectic \mathcal{S} but is Lie-group preserving \mathcal{L} : This category contains the Lie RK2 (CF2) and Lie RK4 (CF4) commutator-free Lie-group preserving integrators from Celledoni et al. [52]. These integrators are Lie-group preserving but are not symplectic. Specifically, the flow of Lie RK2 is described by

$$\begin{cases} \begin{bmatrix} \mathbf{q}_i(h/2) \\ \mathbf{p}_i(h/2) \\ \boldsymbol{\Pi}_i(h/2) \end{bmatrix} = \begin{bmatrix} \mathbf{q}_i(0) \\ \mathbf{p}_i(0) \\ \boldsymbol{\Pi}_i(0) \end{bmatrix} + \begin{bmatrix} \dot{\mathbf{q}}_i(0) \\ \dot{\mathbf{p}}_i(0) \\ \dot{\boldsymbol{\Pi}}_i(0) \end{bmatrix} h, \\ F_1 = \mathbf{J}_i^{-1} \widehat{\boldsymbol{\Pi}}_i(0), \\ \mathbf{R}_i(h/2) = \exp(F_1 h/2) \mathbf{R}_i(0), \end{cases} \quad (\text{D7})$$

$$\begin{cases} \begin{bmatrix} \mathbf{q}_i(h) \\ \mathbf{p}_i(h) \\ \boldsymbol{\Pi}_i(h) \end{bmatrix} = \begin{bmatrix} \mathbf{q}_i(0) \\ \mathbf{p}_i(0) \\ \boldsymbol{\Pi}_i(0) \end{bmatrix} + \begin{bmatrix} \dot{\mathbf{q}}_i(h/2) \\ \dot{\mathbf{p}}_i(h/2) \\ \dot{\boldsymbol{\Pi}}_i(h/2) \end{bmatrix} h, \\ F_2 = \mathbf{J}_i^{-1} \widehat{\boldsymbol{\Pi}}_i(h/2), \\ \mathbf{R}_i(h) = \exp(F_2 h) \mathbf{R}_i(0). \end{cases} \quad (\text{D8})$$

For Lie RK4, the coordinates in Euclidean space (i.e., \mathbf{q} , \mathbf{p} , $\boldsymbol{\Pi}$) follow the normal RK4 integration in a similar fashion as above, while the flow of the \mathbf{R} coordinate is described by

$$\begin{cases} F_1 = \mathbf{J}_i^{-1} \widehat{\boldsymbol{\Pi}}_i(0), \\ \mathbf{R}_i^1 = \exp(F_1 h/2) \mathbf{R}_i(0), \end{cases} \quad (\text{D9})$$

$$\begin{cases} F_2 = \mathbf{J}_i^{-1} \widehat{\boldsymbol{\Pi}}_i^1, \\ \mathbf{R}_i^2 = \exp(F_2 h/2) \mathbf{R}_i(0), \end{cases} \quad (\text{D10})$$

$$\begin{cases} F_3 = \mathbf{J}_i^{-1} \widehat{\boldsymbol{\Pi}}_i^2, \\ \mathbf{R}_i^3 = \exp((F_3 - F_1/2)h) \mathbf{R}_i^2, \end{cases} \quad (\text{D11})$$

$$\begin{cases} F_4 = \mathbf{J}_i^{-1} \widehat{\boldsymbol{\Pi}}_i^3, \\ \mathbf{R}_i(h/2) = \exp\left(\frac{h}{12}(3F_1 + 2F_2 + 2F_3 - F_4)\right) \mathbf{R}_i(0), \\ \mathbf{R}_i(h) = \exp\left(\frac{h}{12}(-F_1 + 2F_2 + 2F_3 + 3F_4)\right) \mathbf{R}_i(h/2), \end{cases} \quad (\text{D12})$$

where the superscripts denote the intermediate outputs of each stage.

Both Symplectic \mathcal{S} and Lie-group preserving \mathcal{L} : This category contains our proposed Lie T₂ integrator which is both symplectic and Lie-group preserving using a splitting technique that allows for exact integration, borrowed from Chen et al. [20]. See App B for a detailed derivation.

E Definition of Evaluation Metrics

The errors $\|\Delta\mathbf{q}\|_2$, $\|\Delta\mathbf{R}\|$ in Table 1 and Table 2 are computed by predicting a trajectory with 500 integrator steps, and then computing

$$\Delta\mathbf{q}_{k,l} := \mathbf{q}_{k,l} - \hat{q}_{k,l}, \quad (\text{E1})$$

$$\Delta\mathbf{R}_{k,l} := \mathbf{R}_{k,l} - \hat{R}_{k,l}, \quad (\text{E2})$$

and the norm for $\mathbf{R} \in \text{SO}(3)^{\otimes N}$ is the geodesic computed as

$$\|\mathbf{R}\| := \sum_{i=1}^N \|\mathbf{R}_i\|, \quad \|\mathbf{R}_i\| := \cos^{-1} \left(\frac{\text{tr}(\mathbf{R}_i^\top \mathbf{R}_i) - 1}{2} \right) \quad (\text{E3})$$

The numbers shown in the table are taken to be the mean across all LK samples.

The errors $\|\Delta\dot{p}\|$ and $\|\Delta\dot{\Pi}\|$ in Table 1 are intended to measure how well the forces (both conservative and non-conservative) are learned. Consequently, these are computed along the dataset and not the predicted trajectory and are defined as

$$\Delta\dot{p}_k = \frac{\hat{p}((k+1)h) - p(kh)}{h} - \frac{p((k+1)h) - p(kh)}{h}, \quad (\text{E4})$$

$$\Delta\dot{\Pi}_k = \frac{\hat{\Pi}((k+1)h) - \Pi(kh)}{h} - \frac{\Pi((k+1)h) - \Pi(kh)}{h}, \quad (\text{E5})$$

where \hat{p} and $\hat{\Pi}$ in the equation above denote the one-step predictions using the learned integrator ϕ_h , *i.e.*

$$\hat{q}(h), \hat{p}(h), \hat{R}(h), \hat{\Pi}(h), := \phi_h(q(0), p(0), R(0), \Pi(0)). \quad (\text{E6})$$

Similarly, the errors $\left\| \Delta \frac{\partial V}{\partial q} \right\|_2$ and $\left\| \Delta \frac{\partial V}{\partial \mathbf{q}} \right\|_2$ in Table 2 measure how well the potential (and hence the conservative forces) have been learned. Hence, these are again computed along the dataset and not the predicted trajectory and are defined as

$$\left\| \Delta \frac{\partial V}{\partial \mathbf{q}} \right\|_2^2 := \left\| \frac{\partial V_{\text{resid}}^\theta}{\partial \mathbf{q}} - \frac{\partial V_{\text{resid}}}{\partial \mathbf{q}} \right\|_2^2, \quad (\text{E7})$$

$$\left\| \Delta \frac{\partial V}{\partial \mathbf{R}} \right\|_2^2 := \sum_{i=1}^N \left\| \left(\mathbf{R}_i^\top \frac{\partial V_{\text{resid}}^\theta}{\partial \mathbf{R}_i} - \left(\frac{\partial V_{\text{resid}}^\theta}{\partial \mathbf{R}_i} \right)^\top \mathbf{R}_i \right)^\vee - \left(\mathbf{R}_i^\top \frac{\partial V_{\text{resid}}}{\partial \mathbf{R}_i} - \left(\frac{\partial V_{\text{resid}}}{\partial \mathbf{R}_i} \right)^\top \mathbf{R}_i \right)^\vee \right\|_2^2, \quad (\text{E8})$$

where the additional manipulations for $\Delta \frac{\partial V}{\partial \mathbf{R}}$ denote the projection of $\mathbf{R}_i^\top \frac{\partial V}{\partial \mathbf{R}_i}$ on the skew-symmetric matrices. This is done because only the skew-symmetric part of $\mathbf{R}_i^\top \frac{\partial V}{\partial \mathbf{R}_i}$ is used in the equations of motion (1a)–(1d) and hence the learned symmetric part can be arbitrarily defined.

F Limitations and Future Directions

One big assumption we make in our work is that the masses and inertial tensors m_i, \mathbf{J}_i are known. This may be a restricting assumption when applying this method to learning physics for rigid bodies where this information is not available and must be learned jointly with the physics, an approach taken in many works on learning with Hamiltonian structure (e.g., [16, 51]). Moreover, as discussed in App C, our framework does not assume that prior knowledge on the structure (e.g., pairwise, independent) of the potential or forcing terms is known. If we make this assumption, then additional physical properties for each rigid body may need to be provided to fully specify the structured physics (e.g., time lag and tide constants for tidal forcing). Extending this framework to handle the learning of per-body physical properties at the same time is left for future work.

Another direction that we have not been able to explore in this work is the robustness of our method to noise in the dataset. In this work, we use a dataset generated using the GRIT [20] simulator (see App C). Consequently, the dataset used is clean and not corrupted by any noise. As noted in SRNN

Chen et al. [6], integrating for multiple consecutive timesteps may allow the network to better discern the true noiseless trajectory when computing the loss function. Moreover, the benefit from performing this multi-step training is integrator dependent, with improvements not observed when using the naive explicit Euler integration. Given that, 1. our configuration space now lies on a Lie-group manifold, and 2. the physics of our system are multiscale, exploring whether those insights are applicable to the problem considered in this work is a future interesting direction that will help inform practitioners wishing to apply this methodology to real world data.