

Learning Hamiltonian Dynamics at Scale: A Differential-Geometric Approach

Katharina Friedl¹ Noémie Jaquier¹ Alyx Liao^{1,2} Danica Kragic¹

Abstract

Embedding physical intuition into network architectures allows the learning of dynamics that enforce fundamental properties, such as energy conservation laws, thereby leading to physically-plausible predictions. Yet, scaling these models to high-dimensional dynamical systems remains a significant challenge. This paper introduces Reduced-order Hamiltonian Neural Network (RO-HNN), a novel physics-inspired neural network that combines the conservation laws of Hamiltonian mechanics with the scalability of model order reduction. RO-HNN is built on two core components: a novel geometrically-constrained symplectic autoencoder that learns a low-dimensional, structure-preserving symplectic submanifold, and a geometric Hamiltonian neural network that models the dynamics on the submanifold. Our experiments demonstrate that RO-HNN provides physically-consistent, stable, and generalizable predictions of complex high-dimensional dynamics, thereby effectively extending the scope of Hamiltonian neural networks to high-dimensional physical systems.

1. Introduction

Learning the unknown governing equations of dynamical systems is of fundamental importance to model physical processes. In this context, generic neural models lack built-in physical intuition, thus resulting in limited explainability and poor generalization beyond the data support. Embedding fundamental physical properties, such as conservation laws and boundary conditions, into neural networks has been shown to drastically improve their performance. Various models incorporate physical intuition as soft constraints via penalty terms in the loss function (Saqlain et al., 2023). This

often leads to suboptimal enforcement of physical properties and to stiff optimization (Wang et al., 2021), motivating the embedding of domain-specific priors as hard constraints in specialized neural architectures. This allowed recent methods to learn dynamics that preserve energy (Greydanus et al., 2019; Cranmer et al., 2020; Lutter & Peters, 2023), conserve mass and momentum (Jnini et al., 2025), and strictly enforce general conservation laws (Liu et al., 2024b), thereby improving performances, generalization, and stability while yielding physically-consistent predictions.

Hamiltonian mechanics, introduced by Hamilton (1834) as a reformulation of Lagrangian mechanics, describe the evolution of a broad range of dynamical systems in robotics (Duong & Atanasov, 2021), fluid dynamics (Salmon, 1988), quantum mechanics (Schrödinger, 1926), and biology (Duarte et al., 1998), among others. Hamiltonian systems evolve on a phase space with symplectic structure, naturally enforcing energy conservation (Abraham & Marsden, 1987). Compared to Lagrangian mechanics, Hamiltonian mechanics provide a first-order formulation of dynamics that describes a broader range of physical systems. Hamiltonian neural networks (HNNs) are gray-box models that embed the Hamiltonian structure as hard constraints in specialized deep learning architectures. HNNs either directly learn the Hamiltonian function, ensuring conservation laws by construction (Greydanus et al., 2019; Lutter & Peters, 2023), or learn symplectomorphisms that preserve the invariants of interest via symplectic flows (Jin et al., 2020). HNNs were enhanced by including dissipation (Zhong et al., 2020a) and contact (Zhong et al., 2021) models, and utilized for model-based control (Duong & Atanasov, 2021; Zhong et al., 2020b). While most HNNs consider Hamiltonians characterized by a canonical symplectic form — exhibited at least locally for all Hamiltonian systems — few works proposed architectures handling non-canonical forms (Chen et al., 2021) and more general Poisson systems (Jin et al., 2021; Šípková et al., 2023). Although HNNs yielded drastic performance improvements over generic black-box models, their application remains limited to low-dimensional systems with 2-5 dimensions.

Learning the dynamics of high-dimensional physical systems, such as robots, continua, or fluids, is arguably a difficult problem due to the increasing complexity and non-linearity of their governing equations. Several approaches

¹Department of Robotics, Perception, and Learning, KTH Royal Institute of Technology, Sweden ²Ecole Normale Supérieure Paris, France. Correspondence to: Katharina Friedl <kfriedl@kth.se>.

combine data-driven sparse identification of nonlinear dynamics (SINDy) and dimensionality reduction to discover high-dimensional governing equations (Brunton et al., 2016; Champion et al., 2019). However, they disregard the *a priori*-known structures of physical systems. Another body of works (Sanchez-Gonzalez et al., 2019; Thangamuthu et al., 2022; Rahma et al., 2026) represents the state of a physical system with particles identified with nodes in a graph and learn a Hamiltonian function via a graph neural network, thus leveraging message-passing among nodes to approximate complex dynamics. While effective for homogeneous systems, the generalization of these approaches to high-dimensional systems is limited by the assumption that all particles share similar dynamics and interaction patterns. In contrast, Sharma & Kramer (2024); Friedl et al. (2025) took inspiration from model order reduction (MOR) to learn high-dimensional Lagrangian dynamics. MOR addresses the complexity of nonlinear high-dimensional governing equations, so-called full-order model (FOM), by finding a reduced-order model (ROM), i.e., a computationally-cheaper yet accurate low-dimensional surrogate model (Schilders et al., 2008). While MOR techniques are typically intrusive, i.e., they assume entirely-known FOM dynamics, Sharma et al. (2024) presented a novel non-intrusive MOR-based approach that learns the parameters of a high-dimensional Lagrangian system in a linear structure-preserving subspace. In a similar line, Friedl et al. (2025) adopted a Riemannian perspective on the problem and introduced a physics-inspired neural architecture that jointly learns a non-linear embedded submanifold via a biorthogonal Autoencoder (AE) and its associated low-dimensional conservative dynamics via a geometric Lagrangian neural network (LNN). A different line of works leverage Koopman operator theory to model nonlinear, e.g., Hamiltonian, dynamics via a learned surrogate linear dynamic model embedded in a higher-dimensional latent space (Lusch et al., 2018; Zhang et al., 2024).

This paper proposes a novel physics-inspired geometric deep neural network to learn the dynamics of high-dimensional Hamiltonian systems. In contrast to previous works that learn dynamics from high-dimensional observations such as images (Greydanus et al., 2019; Chen et al., 2021; Botev et al., 2021), we consider systems with *intrinsically high-dimensional* phase spaces. Taking inspiration from (Sharma & Kramer, 2024; Friedl et al., 2025), we build on recent advances in Hamiltonian MOR (Peng & Mohseni, 2016; Buchfink et al., 2024) and adopt a differential geometric perspective to embed the high-dimensional Hamiltonian structure as hard constraints in our architecture. **Our first contribution** is a geometrically-constrained symplectic AE that learns a low-dimensional symplectic submanifold from trajectories of a high-dimensional Hamiltonian system. Unlike soft-constrained symplectic networks (Buchfink et al.,

2023), our AE guarantees the preservation of the symplectic structure of the FOM, including its conservation laws and stability properties (Lepri et al., 2024), with increased expressivity compared to linear and quadratic symplectic projections (Bendokat & Zimmermann, 2022; Sharma et al., 2023). **Our second contribution** is a geometric HNN that models conservative and dissipative Hamiltonian dynamics while accounting for the Riemannian geometry of its parameters, and resorts to symplectic integration (Tao, 2016) for accurate long-term dynamics simulation. **Our third contribution** is a reduced-order Hamiltonian neural network (RO-HNN) that jointly learns a low-dimensional symplectic submanifold with a geometrically-constrained symplectic AE and the dynamics parameters of the associated Hamiltonian function with a geometric HNN. We validate our approach on three high-dimensional Hamiltonian systems: a pendulum, a thin cloth, and a particle vortex. Our experiments demonstrate that, due to its embedded geometries, RO-HNN predicts accurate, stable, and physically-consistent trajectories, outperforming traditional HNNs and state-of-the-art reduction approaches. Source code is available at https://github.com/katfriedl/reduced_hamiltonians.

2. Background

We briefly review Hamiltonian dynamics, Hamiltonian MOR, and related neural networks. App. A provides preliminaries on Riemannian and symplectic geometry, and a summary of the geometric terminology and notations.

2.1. Hamiltonian Dynamics on Symplectic Manifolds

A symplectic manifold $(\mathcal{M}, \omega)$ is a $2n$ -dimensional smooth manifold $\mathcal{M}$ equipped with a symplectic form ω , i.e., a closed ($d\omega = 0$), non-degenerate, differential 2-form represented by a skew-symmetric matrix ω in coordinates. We slightly abuse notation, equivalently denoting symplectic manifolds as $(\mathcal{M}, \omega)$. A Hamiltonian system $(\mathcal{M}, \omega, \mathcal{H})$ is a dynamical system evolving on a symplectic manifold $(\mathcal{M}, \omega)$ according to a smooth Hamiltonian function $\mathcal{H} : \mathcal{M} \rightarrow \mathbb{R}$. The Hamiltonian vector field $\mathbf{X}_{\mathcal{H}} = \omega^{-1} d\mathcal{H}$ is uniquely defined and preserves $\mathcal{H}$. Trajectories $\gamma : \mathcal{I} \rightarrow \mathcal{M}$ of the system over a time-interval $\mathcal{I} = [t_0, t_f]$ are solutions of the initial value problem (IVP)

$$\frac{d}{dt}\gamma|_t = \mathbf{X}_{\mathcal{H}}|_{\gamma(t)} \in T_{\gamma(t)}\mathcal{M}, \text{ with } \gamma(t_0) = \gamma_0 \in \mathcal{M}. \quad (1)$$

A diffeomorphism $f : (\mathcal{M}, \omega) \rightarrow (\mathcal{N}, \eta)$ between symplectic manifolds is a symplectomorphism if it preserves the symplectic form, i.e., $f^*\eta = \omega$ with $f^*\eta$ denoting the pullback of η by f .

Following Darboux theorem, there exists a canonical chart (U, ϕ) , $\mathbf{x} \in U$ for each point $\mathbf{x} \in \mathcal{M}$ in which the sym-

plectic form is represented as $\omega = \mathbb{J}_{2n}^\top$ via the canonical Poisson tensor

$$\mathbb{J}_{2n} = \begin{pmatrix} \mathbf{0} & \mathbf{I}_n \\ -\mathbf{I}_n & \mathbf{0} \end{pmatrix}, \text{ for which } \mathbb{J}_{2n}^\top = \mathbb{J}_{2n}^{-1} = -\mathbb{J}_{2n}. \quad (2)$$

In other words, every symplectic manifold is locally symplectomorphic to $(\mathbb{R}^{2n}, \mathbb{J}_{2n}^\top)$. A system $(\mathbb{R}^{2n}, \mathbb{J}_{2n}^\top, \mathcal{H})$ is called a canonical Hamiltonian system.

In this paper, we consider Hamiltonian systems $(\mathcal{M}, \omega, \mathcal{H})$, on $\mathcal{M}$ globally valid canonical symplectic form $\omega = \mathbb{J}_{2n}^\top$. In this case, the phase space $\mathcal{M}$ can be modeled on the cotangent bundle $\mathcal{T}^*\mathcal{Q}$ of a smooth n -dimensional manifold $\mathcal{Q}$ (Weinstein, 1971) with canonical coordinates $(\mathbf{q}, \mathbf{p})$ with position $\mathbf{q} \in \mathcal{Q}$ and conjugate momenta $\mathbf{p} \in \mathcal{T}_q^*\mathcal{Q}$. The Hamiltonian vector field simplifies to

$$\begin{pmatrix} \dot{\mathbf{q}} \\ \dot{\mathbf{p}} \end{pmatrix} = \mathbf{X}_{\mathcal{H}} = \mathbb{J}_{2n} d\mathcal{H}^\top = \begin{pmatrix} \frac{\partial \mathcal{H}}{\partial \mathbf{p}} \\ -\frac{\partial \mathcal{H}}{\partial \mathbf{q}} \end{pmatrix}. \quad (3)$$

Moreover, the Hamiltonian system $(\mathcal{T}^*\mathcal{Q}, \mathbb{J}_{2n}^\top, \mathcal{H})$ relates to a Lagrangian function $\mathcal{L} : \mathcal{T}\mathcal{Q} \rightarrow \mathbb{R}$ via the Legendre transform, which takes $\mathcal{L}$ to $\mathcal{H} = \dot{\mathbf{q}}^\top \mathbf{p} - \mathcal{L}$ with $\mathbf{p} = \frac{\partial \mathcal{L}}{\partial \dot{\mathbf{q}}}$ and $\dot{\mathbf{q}} \in \mathcal{T}_q\mathcal{Q}$. Mechanical systems often display a quadratic kinetic energy structure, where the configuration manifold $\mathcal{Q}$ is a Riemannian manifold endowed with the kinetic-energy metric equal to the system's mass-inertia matrix $\mathbf{M}(\mathbf{q})$. In this case, the Hamiltonian function is given by the sum of the system's kinetic $T(\mathbf{q}, \mathbf{p})$ and potential $V(\mathbf{q})$ energies as $\mathcal{H} = T(\mathbf{q}, \mathbf{p}) + V(\mathbf{q}) = \frac{1}{2} \mathbf{p}^\top \mathbf{M}^{-1}(\mathbf{q}) \mathbf{p} + V(\mathbf{q})$ and the momenta is $\mathbf{p} = \mathbf{M}(\mathbf{q}) \dot{\mathbf{q}}$.

2.2. Structure-Preserving Hamiltonian MOR

Given the known parametrized dynamic equations of a high-dimensional system, i.e., a FOM, MOR aims to construct a low-dimensional surrogate dynamic model, i.e., a ROM, that accurately and efficiently approximates the FOM trajectories. Structure-preserving MOR preserves the underlying geometric structure of the FOM, ensuring that its properties, e.g. stability and energy conservation, are maintained in the ROM. For Hamiltonian systems $(\mathcal{M}, \omega, \mathcal{H})$, the symplectic structure is preserved by constructing a reduced Hamiltonian $(\tilde{\mathcal{M}}, \tilde{\omega}, \tilde{\mathcal{H}})$ with $\dim(\tilde{\mathcal{M}}) = d \ll \dim(\mathcal{M}) = n$, whose vector field $\tilde{\mathbf{X}}_{\tilde{\mathcal{H}}}$ approximates the set of solutions $S = \{\gamma(t) \in \mathcal{M} \mid t \in \mathcal{I}\} \subseteq \mathcal{M}$ of (1).

Following the geometric perspective of Buchfink et al. (2024), the reduced Hamiltonian $(\tilde{\mathcal{M}}, \tilde{\omega}, \tilde{\mathcal{H}})$ is derived by identifying the submanifold $\tilde{\mathcal{M}}$ via a smooth embedding $\varphi : \tilde{\mathcal{M}} \rightarrow \mathcal{M}$ such that

$$\tilde{\omega} = \varphi^* \omega = d\varphi^\top \omega d\varphi, \quad (4)$$

is non-degenerate. This implies that $(\tilde{\mathcal{M}}, \tilde{\omega})$ is a symplectic manifold and φ is a symplectomorphism (Buchfink

et al., 2024, Lemma 5.13). Note that structure-preserving Hamiltonian MOR typically considers a canonical FOM $(\mathbb{R}^{2n}, \mathbb{J}_{2n}^\top, \mathcal{H})$ reduced to a canonical ROM $(\mathbb{R}^{2d}, \mathbb{J}_{2d}^\top, \tilde{\mathcal{H}})$ (Peng & Mohseni, 2016; Sharma et al., 2023; Buchfink et al., 2023). The Hamiltonian structure is preserved by constructing $\tilde{\mathcal{H}}$ via the pullback of the embedding as $\tilde{\mathcal{H}} = \varphi^* \mathcal{H} = \mathcal{H} \circ \varphi$. Trajectories $\tilde{\gamma}(t)$ of the reduced-order system are then obtained from the ROM $\frac{d}{dt} \tilde{\gamma}|_t = \tilde{\mathbf{X}}_{\tilde{\mathcal{H}}}(\tilde{\gamma}(t)) \in \mathcal{T}_{\tilde{\gamma}(t)} \tilde{\mathcal{M}}$, with $\tilde{\mathbf{X}}_{\tilde{\mathcal{H}}} = \tilde{\omega}^{-1} d\tilde{\mathcal{H}}$. The reduced initial value $\tilde{\gamma}_0 = \rho(\gamma_0) \in \tilde{\mathcal{M}}$ is computed via the point reduction map $\rho : \mathcal{M} \rightarrow \tilde{\mathcal{M}}$ associated with φ , which must satisfy the projection properties

$$\rho \circ \varphi = \text{id}_{\tilde{\mathcal{M}}} \quad \text{and} \quad d\rho|_{\varphi(\tilde{\mathbf{x}})} \circ d\varphi|_{\tilde{\mathbf{x}}} = \text{id}_{\mathcal{T}_{\tilde{\mathbf{x}}} \tilde{\mathcal{M}}}, \quad (5)$$

$\forall \tilde{\mathbf{x}} \in \tilde{\mathcal{M}}$. Trajectories of the original system are finally obtained as the approximation $\gamma(t) \approx \varphi(\tilde{\gamma}(t))$.

The embedding φ and point reduction ρ are key as they determine the ROM trajectories. Accurately approximating the FOM requires minimizing the reconstruction error

$$\ell_{\text{rec}} = \frac{1}{N} \sum_{i=1}^N \|\varphi \circ \rho(\mathbf{x}_i) - \mathbf{x}_i\|^2. \quad (6)$$

Exact reconstruction requires $d\rho$ to be the symplectic inverse of $d\varphi$, i.e., $d\rho = d\varphi^+ = \tilde{\omega}^{-1} d\varphi^\top \omega$. In this paper, we introduce a geometrically-constrained AE that fulfills (4) and (5) by design.

2.3. Hamiltonian Neural Networks

While MOR reduces the dimensionality of systems with known dynamics, HNNs aim to learn the unknown dynamics of typically low-dimensional systems while ensuring energy conservation. Most HNNs assume canonical Hamiltonian systems or Hamiltonian systems with canonical symplectic form (2). In this paper, we build on two HNN variants that (1) learn the Hamiltonian function as a single network $\mathcal{H}_\theta(\mathbf{q}, \mathbf{p})$ (Greydanus et al., 2019), or (2) learn the kinetic and potential energy as two distinct networks, i.e., $\mathcal{H}_\theta(\mathbf{q}, \mathbf{p}) = T_{\theta_T}(\mathbf{q}, \mathbf{p}) + V_{\theta_V}(\mathbf{q})$ (Zhong et al., 2020b; Lutter & Peters, 2023). Given a set of N observations $\{\mathbf{q}_i, \mathbf{p}_i, \dot{\mathbf{q}}_i, \dot{\mathbf{p}}_i\}_{i=1}^N$, the networks are trained to minimize the prediction error of the Hamiltonian vector field via the loss

$$\ell_{\text{HNN}} = \left\| \frac{\partial \mathcal{H}_\theta}{\partial \mathbf{p}}(\mathbf{q}_i, \mathbf{p}_i) - \dot{\mathbf{q}}_i \right\|^2 + \left\| \frac{\partial \mathcal{H}_\theta}{\partial \mathbf{q}}(\mathbf{q}_i, \mathbf{p}_i) - \dot{\mathbf{p}}_i \right\|^2. \quad (7)$$

3. Geometric Reduced-order Hamiltonian Neural Networks

We present the geometric reduced-order Hamiltonian neural network (RO-HNN) that learns the unknown dynamics of high-dimensional Hamiltonian systems.




Figure 1. Flowchart of the forward dynamics of the geometric RO-HNN. The geometrically-constrained symplectic AE (in blue) is built as the cotangent lift of a constrained AE (top right). The geometric HNN (in brown) is composed of two SPD networks, one MLP, and a Strang symplectic integrator.

We focus on systems $(\mathcal{M}, \mathbb{J}_{2n}^T, \mathcal{H})$ evolving on a phase space $\mathcal{M}$ with canonical symplectic form $\mathbb{J}_{2n}^T$ for which the solutions $\gamma(t)$ of the FOM (1) can be accurately approximated by a substantially lower dimensional surrogate model. Our goal is to learn a reduced Hamiltonian system $(\tilde{\mathcal{M}}, \tilde{\omega}, \tilde{\mathcal{H}})$ via non-intrusive structure-preserving MOR, where we set $\tilde{\mathcal{M}}$ as a phase space with $\tilde{\omega} = \mathbb{J}_{2d}^T$. Given a set of high-dimensional observations $\{q_i, p_i\}_{i=1}^N$, we identify low-dimensional dynamics by jointly learning a reduced symplectic manifold $(\tilde{\mathcal{M}}, \mathbb{J}_{2d}^T)$ via a smooth embedding φ and a reduction ρ , and a latent Hamiltonian function $\tilde{\mathcal{H}}$.

The proposed RO-HNN ensures the preservation of the Hamiltonian structure by fulfilling three necessary conditions by design: (1) the embedding φ is a symplectomorphism, or equivalently

$$\tilde{\omega} = \mathbb{J}_{2d} = d\varphi^T \mathbb{J}_{2n} d\varphi; \quad (8)$$

(2) the embedding φ and reduction map ρ satisfy the projection properties (5); and (3) $\tilde{\mathcal{H}}$ is a valid Hamiltonian function, thus preserving the reduced energy $\tilde{\mathcal{E}} = \mathcal{E} \circ \varphi$. The RO-HNN fulfill (1)-(2) via a novel geometrically-constrained symplectic AE (Sec. 3.1), while (3) is guaranteed by a reduced-order geometric HNN (Sec. 3.2), whose trajectories are obtained via symplectic integration (Sec. 3.3). Accurate modeling of the high-dimensional dynamics is achieved by jointly training the AE and the HNN (Sec. 3.4). The proposed RO-HNN is illustrated in Fig. 1.

3.1. Geometrically-constrained Symplectic Autoencoder

Preserving the geometric structure of the original Hamiltonian FOM is crucial for the learned ROM to display similar dynamics. We introduce a geometrically-constrained symplectic AE that projects a high-dimensional Hamiltonian system $(\mathcal{M}, \mathbb{J}_{2n}^T, \mathcal{H})$ onto a low-dimensional nonlinear symplectic manifold $(\tilde{\mathcal{M}}, \mathbb{J}_{2d}^T)$ such that the reduced system strictly retains the Hamiltonian structure of the FOM. We parametrize the point reduction $\rho : \mathcal{M} \rightarrow \tilde{\mathcal{M}}$ and embedding $\varphi : \tilde{\mathcal{M}} \rightarrow \mathcal{M}$ as the encoder and decoder of an AE designed

to satisfy symplecticity (8) and projection properties (5) by construction. To do so, we leverage the cotangent bundle structure of the phase space $\mathcal{M} = \mathcal{T}^*\mathcal{Q}$.

Given a smooth embedding $\varphi_{\mathcal{Q}} : \tilde{\mathcal{Q}} \rightarrow \mathcal{Q}$ and associated point reduction $\rho_{\mathcal{Q}} : \mathcal{Q} \rightarrow \tilde{\mathcal{Q}}$ satisfying (5), we define the cotangent-lifted embedding φ and point reduction ρ in canonical coordinates as

$$\varphi(\tilde{q}, \tilde{p}) = \begin{pmatrix} \varphi_{\mathcal{Q}}(\tilde{q}) \\ d\rho_{\mathcal{Q}}|_{\varphi_{\mathcal{Q}}(\tilde{q})}^T \tilde{p} \end{pmatrix} \text{ and } \rho(q, p) = \begin{pmatrix} \rho_{\mathcal{Q}}(q) \\ d\varphi_{\mathcal{Q}}|_{\rho_{\mathcal{Q}}(q)}^T p \end{pmatrix}, \quad (9)$$

where the pullbacks $d\rho_{\mathcal{Q}}|_{\varphi_{\mathcal{Q}}(\tilde{q})}^T \tilde{p}$ and $d\varphi_{\mathcal{Q}}|_{\rho_{\mathcal{Q}}(q)}^T p$ are computed using the Jacobian of $\rho_{\mathcal{Q}}$ and $\varphi_{\mathcal{Q}}$.

Proposition 3.1. *The reduction map $\rho(q, p)$ (9) satisfies the projection properties (5).*

Proof. See App. D. $\square$

Proposition 3.2. *The embedding $\varphi(\tilde{q}, \tilde{p})$ (9) satisfies the symplecticity property (8).*

Proof. See App. D. $\square$

Note that $d\rho_{\mathcal{Q}} \mathbb{J}_{2n} d\rho^T = \mathbb{J}_{2d}$ is shown to hold on $\varphi(\tilde{\mathcal{M}})$ with similar arguments.

In practice, we learn the embedding $\varphi_{\mathcal{Q}}$ and point reduction $\rho_{\mathcal{Q}}$ via the constrained AE architecture from Otto et al. (2023), and compute their differentials analytically to construct the cotangent-lifted maps (9) (see App. E for details). The encoder and decoder are given as a composition of feed-forward layers $\rho_{\mathcal{Q}} = \rho_{\mathcal{Q}}^{(1)} \circ \dots \circ \rho_{\mathcal{Q}}^{(L)}$ and $\varphi_{\mathcal{Q}} = \varphi_{\mathcal{Q}}^{(L)} \circ \dots \circ \varphi_{\mathcal{Q}}^{(1)}$ with $\rho_{\mathcal{Q}}^{(l)} : \mathbb{R}^{n_l} \rightarrow \mathbb{R}^{n_{l-1}}$, $\varphi_{\mathcal{Q}}^{(l)} : \mathbb{R}^{n_{l-1}} \rightarrow \mathbb{R}^{n_l}$, and $n_{l-1} \leq n_l$. The key to fulfill the projection properties (5) is the construction of the layer pairs as

$$\begin{aligned} \rho_{\mathcal{Q}}^{(l)}(q^{(l)}) &= \sigma_- \left(\Psi_l^T (q^{(l)} - b_l) \right), \\ \varphi_{\mathcal{Q}}^{(l)}(\tilde{q}^{(l-1)}) &= \Phi_l \sigma_+ (\tilde{q}^{(l-1)}) + b_l, \end{aligned} \quad (10)$$

where (Φ_l, Ψ_l) and (σ_+, σ_-) are pairs of weight matrices and smooth activation functions such that $\Psi_l^T \Phi_l = I_{n_{l-1}}$ and $\sigma_- \circ \sigma_+ = \text{id}$, respectively, and b_l are bias vectors. Therefore, each layer pair (10) satisfies $\rho_{\mathcal{Q}}^{(l)} \circ \varphi_{\mathcal{Q}}^{(l)} = \text{id}_{\mathbb{R}^{n_{l-1}}}$ and the constrained AE fulfills (5). Following (Friedl et al., 2025), we ensure that the pairs of weight matrices adhere to the biorthogonality constraint $\Psi_l^T \Phi_l = I_d$ by accounting for the Riemannian geometry of biorthogonal matrices (see App. B.2 for a background). Specifically, we consider each pair (Φ_l, Ψ_l) as an element of the biorthogonal manifold $\mathcal{B}_{n_l, n_{l-1}} = \{(\Phi, \Psi) \in \mathbb{R}^{n_l \times n_{l-1}} \times \mathbb{R}^{n_l \times n_{l-1}} : \Psi^T \Phi = I_{n_{l-1}}\}$ and optimize them to minimize the reconstruction error (6) via Riemannian optimization (Absil et al., 2007; Boumal, 2023) (see App. G). Note that this Riemannian approach was shown to consistently outperform the overparametrization proposed by Otto et al. (2023), achieving lower

reconstruction errors (Friedl et al., 2025). The constraint $\sigma_- \circ \sigma_+ = \text{id}$ is met by utilizing the smooth, invertible activation functions defined in (Otto et al., 2023, Eq. 12), see also App. E.

As will be shown in Sec. 4, the resulting geometrically-constrained symplectic AE provides increased expressivity compared to linear and quadratic symplectic projection approaches (Peng & Mohseni, 2016; Sharma et al., 2023), while guaranteeing the symplectic structure of the latent space in contrast to weakly-symplectic AEs based on soft constraints (Buchfink et al., 2023). In the intrusive case, i.e., if the FOM is known, we can construct the reduced Hamiltonian function via the pullback of the cotangent-lifted embedding as $\tilde{\mathcal{H}} = \varphi^* \mathcal{H}$, which yields the Hamiltonian ROM. Instead, in this paper, we consider the case where the high-dimensional dynamics are unknown, and learn the reduced-order Hamiltonian function $\tilde{\mathcal{H}}$ with a geometric HNN, as explained next.

3.2. Conservative and Dissipative Hamiltonian ROMs

Conservative dynamics. We propose to learn the reduced Hamiltonian dynamics in the embedded symplectic submanifold $(\tilde{\mathcal{M}}, \mathbb{J}_{2d}^\top)$ via a HNN. For general systems, we encode the reduced-order Hamiltonian function as a single neural network $\tilde{\mathcal{H}}_\theta(\tilde{q}, \tilde{p})$ with parameters θ , akin to (Greydanus et al., 2019). However, additional prior knowledge on the structure of the Hamiltonian is often available. For instance, the Hamiltonian function of mechanical systems sums a quadratic kinetic energy and a potential term. Leveraging that the learned symplectic submanifold preserves the original system structure, we propose to model the reduced Hamiltonian function as $\tilde{\mathcal{H}}_\theta(\tilde{q}, \tilde{p}) = \frac{1}{2} \tilde{p}^\top \tilde{M}_{\theta_T}^{-1}(\tilde{q}) \tilde{p} + \tilde{V}_{\theta_V}(\tilde{q})$ via two neural networks $\tilde{M}_{\theta_T}^{-1}$ and $\tilde{V}_{\theta_V}$ with parameters $\theta = \{\theta_T, \theta_V\}$. Existing HNNs (Lutter & Peters, 2023; Zhong et al., 2020a) enforce the symmetric positive-definiteness of the inverse mass-inertia matrix via a Euclidean network encoding its Cholesky decomposition L , i.e., $M^{-1} = LL^\top$. However, as for LNNs (Friedl et al., 2025), this parametrization leads to flawed measures of distances in the space of symmetric positive-definite (SPD) matrices and ultimately results in inaccurate dynamics predictions. To overcome this issue, we parametrize $\tilde{M}_{\theta_T}^{-1}$ via the SPD network from Friedl et al. (2025) that accounts for the Riemannian geometry of the SPD manifold $\mathcal{S}_{++}^d$ (see Apps. A, C). The network $\tilde{M}_{\theta_T}^{-1}(q) = (g_{\text{Exp}} \circ g_{\mathbb{R}})(q)$ is composed of (1) a standard Euclidean multilayer perceptron (MLP) $g_{\mathbb{R}} : \mathbb{R}^d \rightarrow \mathbb{R}^{d(d+1)/2}$ that maps the input configuration to the elements of a symmetric matrix $U \in \text{Sym}^d$, and (2) an exponential map layer g_{Exp} that interprets U as an element of the tangent space $\mathcal{T}_P \mathcal{S}_{++}^d$, and maps it onto $\mathcal{S}_{++}^d$.

Dissipative dynamics. While classical Hamiltonian dy-

namics conserve energy, dissipation and external inputs often appear in real-world systems. Both can be modeled in HNNs by complementing the Hamiltonian vector field with a force field $X_{\mathcal{F}}$, so that the total vector field is $X = X_{\mathcal{H}} + X_{\mathcal{F}}$ (Sosanya & Greydanus, 2022; Zhong et al., 2020a). We propose to leverage the structure-preserving symplectic submanifold and model dissipation and external inputs as a reduced-order force field $\tilde{X}_{\mathcal{F}}$ on $(\tilde{\mathcal{M}}, \mathbb{J}_{2d}^\top)$. Specifically, we consider high-dimensional systems with observed external inputs τ_{ext} and viscous damping following a Rayleigh dissipative function $\mathcal{D}(q, \dot{q}) = \frac{1}{2} \dot{q}^\top D(q) \dot{q}$ with unknown positive semi-definite dissipation matrix $D(q)$. The resulting force field is $X_{\mathcal{F}} = \begin{pmatrix} 0 \\ \tau_{\text{ext}} + \tau_d \end{pmatrix}$ with damping force $\tau_d = \frac{\partial \mathcal{D}(q, \dot{q})}{\partial \dot{q}} = D(q) \dot{q}$.

Proposition 3.3. *The reduced vector field $\tilde{X} = \varphi^* X$ obtained via the pullback of the cotangent-lifted embedding φ (9) preserves the structure of the total vector field $X = X_{\mathcal{H}} + X_{\mathcal{F}}$ with $X_{\mathcal{F}} = \begin{pmatrix} 0 \\ \tau_{\text{ext}} + \tau_d \end{pmatrix}$.*

Proof. See App. D. $\square$

We propose to model the reduced Rayleigh dissipative function $\tilde{\mathcal{D}}_{\theta_D}(\tilde{q}, \dot{\tilde{q}}) = \frac{1}{2} \dot{\tilde{q}}^\top \tilde{D}_{\theta_D}(\tilde{q}) \dot{\tilde{q}}$ via a neural network $\tilde{D}_{\theta_D}$. Dissipative HNNs (Zhong et al., 2020a) constrain the dissipation matrix to be positive semi-definite through its Cholesky decomposition, i.e., $D = LL^\top$, thus overlooking its intrinsic geometric structure. Instead, we utilize a second SPD network $\tilde{D}_{\theta_D}(\tilde{q}) = (g_{\text{Exp}} \circ g_{\mathbb{R}})(q)$.

Note that the dissipative dynamics no longer preserve a symplectic structure. However, Proposition 3.3 shows that the RO-HNN latent space still preserves the structure of high-dimensional dissipative dynamical systems characterized by a Rayleigh dissipative function. It is worth noting that these systems are equivalent to Port-Hamiltonian systems with energy dissipation matrix $R(q, p) = \begin{pmatrix} 0 & 0 \\ 0 & D(q) \end{pmatrix}$ and input $G(q, p)u = \tau_{\text{ext}}$, and to contact Hamiltonians $\mathcal{H}_c(q, p, s)$ with $\dot{s} = \mathcal{D}$.

Predicting dynamics. The geometric HNN predicts conservative and dissipative dynamics as

$$\dot{\tilde{q}}_p = \frac{\partial \tilde{\mathcal{H}}_\theta}{\partial \tilde{p}} \text{ and } \dot{\tilde{p}}_p = -\frac{\partial \tilde{\mathcal{H}}_\theta}{\partial \tilde{q}} - \tilde{D}_{\theta_D}(\tilde{q}) \frac{\partial \tilde{\mathcal{H}}}{\partial \tilde{p}} + \tilde{\tau}_{\text{ext}}, \quad (11)$$

where the predictions of the dynamic model are denoted via the subscript p. Note that $\tilde{D}_{\theta_D} = 0$ and $\tilde{\tau}_{\text{ext}} = 0$ in the conservative case. The architecture is illustrated in Fig. 1.

Predicting system trajectories according to the learned reduced-order Hamiltonian dynamics involves (1) integrating the latent predictions $(\dot{\tilde{q}}_p, \dot{\tilde{p}}_p)$ (11), and (2) decoding the obtained reduced-order position and momentum $(\tilde{q}_p, \tilde{p}_p)$ into the high-dimensional coordinates of the original system with the lifted point embedding φ (9), i.e., $(\tilde{q}_p, \tilde{p}_p) = \varphi(\tilde{q}_p, \tilde{p}_p)$. In this paper, we propose to integrate the learned reduced-order Hamiltonian flow via symplectic integration, as explained next.

3.3. Trajectory Prediction via Symplectic Integration

Symplectic integrators are particularly well suited to integrate Hamiltonian dynamics as they preserve the geometric structure and invariants of the Hamiltonian flow (Leimkuhler & Reich, 2005). They were shown to be key to accurately integrate learned HNNs dynamics, thus preventing long-term drifting of numerical solutions (Chen et al., 2020; Xiong et al., 2021).

The Hamiltonian dynamics learned in Sec. 3.2 are non-separable, thus prohibiting the usage of standard explicit integration schemes, e.g., leapfrog (Leimkuhler & Reich, 2005). Instead, we integrate the reduced-order Hamiltonian flow (11) using the second-order symplectic integrator of (Tao, 2016) based on Strang splitting, akin to (Xiong et al., 2021). In a nutshell, the integrator considers an augmented Hamiltonian function

$$\bar{\mathcal{H}}(\mathbf{q}, \mathbf{p}, \mathbf{x}, \mathbf{y}) = \mathcal{H}(\mathbf{q}, \mathbf{y}) + \mathcal{H}(\mathbf{p}, \mathbf{x}) + \frac{1}{2}w(\|\mathbf{q}, \mathbf{x}\|^2 + \|\mathbf{p}, \mathbf{y}\|^2),$$

with extended phase space, for which high-order separable symplectic integrators with explicit updates can be constructed. A numerical integrator approximating $\bar{\mathcal{H}}$ is obtained by composing the obtained explicit flows, which we refer to as Strang-symplectic integrator. Additional details are provided in App. F.

3.4. Model Training

Finally, we propose to jointly learn the parameters $\{\Phi_l, \Psi_l, \mathbf{b}_l\}_{l=1}^L$ of the AE and $\{\theta_T, \theta_V, \theta_D\}$ of the latent geometric HNN. As the learned dynamics are expected to predict multiple steps, we consider a loss that numerically integrates the latent predictions $(\dot{\mathbf{q}}_p, \dot{\mathbf{p}}_p)$ (11) via H forward Strang-symplectic integration steps before decoding. Given sets of observations $\{\mathbf{q}_i(\mathcal{I}_i), \mathbf{p}_i(\mathcal{I}_i), \tau_i(\mathcal{I}_i)\}_{i=1}^N$ over intervals $\mathcal{I}_i = [t_i, t_i + H\Delta t]$ with constant integration time Δt , the resulting multi-step loss is

$$\begin{aligned} \ell_{\text{RO-HNN}} = & \frac{1}{HN} \sum_{i=1}^N \sum_{j=1}^H \underbrace{\|\tilde{\mathbf{q}}_i(t_{i,j}) - \mathbf{q}_i(t_{i,j})\|^2}_{\ell_{\text{AE}}} \\ & + \underbrace{\|\tilde{\mathbf{p}}_i(t_{i,j}) - \mathbf{p}_i(t_{i,j})\|^2}_{\ell_{\text{AE}}} + \gamma \|\boldsymbol{\theta}\|_2^2 \\ & + \underbrace{\lambda \|\tilde{\mathbf{q}}_{p,i}(t_{i,j}) - \tilde{\mathbf{q}}_i(t_{i,j})\|^2 + \lambda \|\tilde{\mathbf{p}}_{p,i}(t_{i,j}) - \tilde{\mathbf{p}}_i(t_{i,j})\|^2}_{\ell_{\text{HNN},d}} \\ & + \underbrace{\|\tilde{\mathbf{q}}_{p,i}(t_{i,j}) - \mathbf{q}_i(t_{i,j})\|^2 + \|\tilde{\mathbf{p}}_{p,i}(t_{i,j}) - \mathbf{p}_i(t_{i,j})\|^2}_{\ell_{\text{HNN},n}}, \end{aligned} \quad (12)$$

where $(\tilde{\mathbf{q}}, \tilde{\mathbf{p}}) = \rho \circ \varphi(\mathbf{q}, \mathbf{p})$, $(\tilde{\mathbf{q}}_{p,i}, \tilde{\mathbf{p}}_{p,i}) = \varphi(\mathbf{q}_{p,i}, \mathbf{p}_{p,i})$, $\tilde{\mathbf{q}}_{p,i}(t_{i,j}) = \int_{t_i}^{t_{i,j}} \dot{\mathbf{q}}_{p,i} dt$ and $\tilde{\mathbf{p}}_{p,i}(t_{i,j}) = \int_{t_i}^{t_{i,j}} \dot{\mathbf{p}}_{p,i} dt$ with $t_{i,j} = t_i + j\Delta t$, and $\lambda, \gamma \in \mathbb{R}_{>0}$ are scaling factors. Note that initial conditions to the latent IVPs are given by the encoded observations at the initial timestep $\tilde{\mathbf{q}}_{p,i}(t_i) = \tilde{\mathbf{q}}_i$ and $\tilde{\mathbf{p}}_{p,i}(t_i) = \tilde{\mathbf{p}}_i$. We optimize the network parameters via Riemannian Adam (Becigneul & Ganeva, 2019), see App. G.

4. Experiments

We evaluate the proposed RO-HNN to learn the dynamics of three simulated high-dimensional Hamiltonian systems: a 15-degrees-of-freedom (DoF) pendulum, a 90-DoF particle vortex, and a 600-DoF thin cloth. Our experiments showcase that RO-HNNs accurately predict long-term trajectories of high-dimensional Hamiltonian systems, highlighting the importance of embedding geometric inductive biases as hard constraints in the AE and HNN. Details about datasets, network architectures, and model training are provided in App. H. Additional results and runtimes are provided in Apps. I and J.

4.1. Coupled Pendulum (15 DoF)

We consider a 15-DoF augmented pendulum whose nonlinear dynamics are specified from the symplectomorphism of a latent 3-DoF pendulum augmented with a 12-DoF mass-spring mesh. As the mesh oscillations are small, the system dynamics are intentionally constructed to showcase a system that is approximately reducible to 3 dimensions.


Learning high-dimensional dynamics. We train a RO-HNN with latent dimension $d = 3$ and a conservative geometric HNN with Strang-symplectic integration on 3000 observations $\{\mathbf{q}_i, \mathbf{p}_i\}$ (see App. H.1 for details). We compare our RO-HNN with a 15-DoF geometric HNN that directly learns high-dimensional dynamic parameters, and with a Hamiltonian neural Koopman operator (HNKO) (Zhang et al., 2024) that learns a discrete linear predictor embedded in a 100-dimensional lifted space. For completeness, we also consider a 3-DoF geometric HNN trained on observations of the latent system. Notice that this model would not be deployable in practice as it requires ground truth information, i.e., latent observations, that would not be available.

Table 1 reports short- and long-term relative prediction errors over $H\Delta t = \{0.25, 5\}$ s. The RO-HNN outperforms the 15-DoF HNN and the HNKO, leading to significantly lower prediction errors. Due to the high dimensionality, the 15-DoF HNN was difficult to train and did not lead to stable long-term predictions. As also shown in Fig. 2, the HNKO learns stable, but inaccurate long-term predictions. In contrast, the RO-HNN achieves similar long-term predictions as the 3-DoF HNN, which is expected to perform best as trained directly on the low-dimensional system (see also Figs. 7-8 in App. I.1). This validates the RO-HNN ability to jointly learn a latent symplectic submanifold and associated dynamics. Table 6 in App. I.1 shows that the RO-HNN is also robust to observation noise, consistently outperforming the HNKO.

AE architecture. The quality of the learned symplectic submanifold is crucial for learning accurate dynamics, as

Table 1. Pendulum: Mean and standard deviation of prediction errors ($\downarrow$) over 10 test trajectories.

	$H\Delta t$ (s)	RO-HNN	15-DoF HNN	3-DoF HNN	HNKO
$\frac{\ \hat{q} - q\ }{\ q\ }$	0.25	$(1.66 \pm 1.38) \times 10^{-1}$	$(5.33 \pm 6.02) \times 10^{-1}$	$(1.22 \pm 0.92) \times 10^{-1}$	$(5.64 \pm 4.41) \times 10^{-1}$
	5	$(7.08 \pm 7.56) \times 10^{-1}$	—	$(5.44 \pm 6.93) \times 10^{-1}$	$(1.32 \pm 0.94) \times 10^0$
$\frac{\ \hat{p} - p\ }{\ p\ }$	0.25	$(5.33 \pm 5.23) \times 10^{-2}$	$(1.76 \pm 2.40) \times 10^{-1}$	$(2.50 \pm 2.96) \times 10^{-2}$	$(5.93 \pm 10.73) \times 10^{-1}$
	5	$(1.98 \pm 2.67) \times 10^{-1}$	—	$(1.85 \pm 3.94) \times 10^{-1}$	$(1.23 \pm 2.08) \times 10^0$

 Figure 2. Pendulum: Median and quartiles of relative error and reconstructed trajectories of the RO-HNN (—), 3-DoF HNN (—), and HNKO (—) vs. ground truth (—) for a horizon $H\Delta t = 5s$. The 15-DoF HNN diverges and is not shown.

 Table 2. Pendulum: Mean and standard deviation of reconstruction, prediction, and symplecticity errors ($\downarrow$) of intrusive symplectic dimensionality reduction approaches over 10 test trajectories.

	Linear SMG	Quadr. SMG	Weakly-sympl. AE	Geom. Sympl. AE (ours)
$\frac{\ \hat{q} - q\ }{\ q\ }$	$(2.21 \pm 1.17) \times 10^{-1}$	$(2.84 \pm 4.23) \times 10^0$	$(1.43 \pm 0.68) \times 10^{-1}$	$(8.84 \pm 6.22) \times 10^{-2}$
$\frac{\ \hat{p} - p\ }{\ p\ }$	$(4.43 \pm 3.99) \times 10^{-1}$	$(2.75 \pm 1.60) \times 10^{-1}$	$(1.57 \pm 1.55) \times 10^{-1}$	$(4.09 \pm 3.99) \times 10^{-2}$
$\frac{\ \hat{q}_p - q\ }{\ q\ }$	$(2.58 \pm 2.33) \times 10^{-1}$	$(3.53 \pm 5.18) \times 10^0$	$(7.10 \pm 7.02) \times 10^{-1}$	$(1.13 \pm 0.92) \times 10^{-1}$
$\frac{\ \hat{p}_p - p\ }{\ p\ }$	$(2.16 \pm 1.89) \times 10^{-1}$	$(8.68 \pm 1.04) \times 10^{-1}$	$(2.16 \pm 1.89) \times 10^{-1}$	$(4.68 \pm 4.32) \times 10^{-2}$
$\ \mathbb{J}_{2d} - d\varphi^\top \mathbb{J}_{2n} d\varphi\ $	0.0 ± 0.0	0.0 ± 0.0	$(1.67 \pm 0.35) \times 10^{-2}$	0.0 ± 0.0
$\ d\rho - d\varphi^\top\ $	0.0 ± 0.0	0.0 ± 0.0	$(5.32 \pm 1.45) \times 10^0$	$(9.53 \pm 5.35) \times 10^{-1}$

they may systematically deviate from the ground truth if the submanifold does not accurately capture the solution space of the high-dimensional system. We analyze the influence of the reduction method in the RO-HNN and compare the proposed geometrically-constrained symplectic AE with linear and quadratic symplectic manifold Galerkin (SMG) projections (Peng & Mohseni, 2016; Sharma et al., 2023) which preserve the symplectic structure by construction, and a weakly-symplectic AE (Buchfink et al., 2023) which encourages structure preservation via a penalty term in the loss (see App. H.1 for details). We train each approach on 3000 observations of the 15-DoF pendulum. Here, we consider an intrusive MOR setup and project the known FOM dynamics onto the learned submanifold to predict new trajectories ($H\Delta t = 0.25s$). Table 2 shows that, due to their increased expressivity, the AEs outperform the linear and quadratic projections, with the geometrically-constrained symplectic AEs achieving the lowest reconstruction and prediction errors (see also Fig. 9 in App. I.1). Note that only the geometrically-constrained AE yielded stable longer-term predictions (see Fig. 2). Moreover, only the weakly-symplectic AE results in an error on the symplecticity condition (8), which is expected as both SMG projections and geometrically-constrained symplectic AE fulfill it by construction. Both SMG projections also ensure that the differential $d\rho$ is the symplectic inverse of $d\varphi$ by design, while the geometrically-constrained AE leads to a lower error than the weakly-constrained one. Note that jointly training the AE with the geometric HNN in the RO-HNN is beneficial as it further reduces this error to $(7.42 \pm 1.21) \times 10^{-1}$.

Latent HNN architecture. We compare the performance of our geometric HNN to learn the low-dimensional dynamics of the latent 3-DoF pendulum against (1) a non-geometric variant that parametrizes the inverse mass-inertia

matrix via a Cholesky network, and (2) two HNNs encoded as a single black-box network $\mathcal{H}_\theta$, where we consider two MLPs of 64- and 256-neurons width. As shown in Fig. 3-left, the geometric HNN achieves the lowest reconstruction error, followed by the Cholesky HNN (see also App. I.1). This showcases the importance of considering both the quadratic energy structure of mechanical systems, and the geometry of their mass-inertia matrices.

Latent integrator. We compare the Strang symplectic integrator against a symplectic leapfrog that disregards that the Hamiltonian is non-separable, and a Runge-Kutta of order 4 that overlooks its symplectic structure. Fig. 3-middle, right show that the Strang symplectic integrator achieves the lowest reconstruction error and conserves energy best during integration (see also App. I.1).

4.2. Particle Vortex (90-DoF)

Next, we learn the dynamics of a particle vortex composed of $n = 90$ particles with uniform interaction strengths. As the particle vortex dynamics are purely determined via the logarithmic interaction, its Hamiltonian function does not separate into kinetic and potential energies. We train RO-HNNs with $d = \{6, 10\}$ and (1) a geometric HNN and (2) a black-box HNN $\mathcal{H}_\theta$, both with Strang-symplectic integration (see App. H.2 for details). Fig. 4 depicts the predicted particle positions and momenta for a prediction horizon of $H = 100$, showing that the RO-HNN accurately predicts the particle vortex dynamics and generalizes beyond the data support ($t > 10s$). As shown in Table 3 and Fig. 13 in App. I.2, the geometric HNNs outperform the black-box HNNs despite the lack of structure of the ground truth Hamiltonian. This suggests that the AE learns a symplectomorphism to a latent space where the Hamiltonian can be decomposed into two energy terms, thereby taking




Figure 3. Pendulum: Ablation of the latent HNN architecture (left) and latent integrator (middle, right) of the RO-HNN.




Figure 4. Particle vortex: Predicted (•, •, •) vs ground truth (•, •, •) positions. Times beyond 10s are out of the training data distribution.

advantage of the additional structure of the geometric HNN. To further validate the role of symplecticity-preserving components, we compare variants of the RO-HNN with non-symplectic AE for $d = \{6, 10\}$ -dimensional geometric and black-box HNNs $\tilde{\mathcal{H}}_\theta$. Compared to the symplectic AE variants, the RO-HNNs with non-symplectic AE degrade the predictive performance of models using a latent geometric HNN. This suggests that a symplectic latent space simplifies the convergence of structure-preserving geometric HNN. This hypothesis is supported by the fact that the predictive performance remain comparable between both AE variant when combined with a latent black-box HNNs. Moreover, the 6-dimensional models slightly outperform the 10-dimensional ones, showing that the choice of latent dimension trades off between the latent space expressivity and the limitations of HNNs in higher dimensions (see App. I.2 for more results). Finally, Table 3 also reports the prediction errors of a HNKO, which is significantly outperformed by all RO-HNN variants, confirming the expressivity of nonlinear latent dynamic models.

4.3. Cloth (600-DoF)

Next, we learn the dynamics of a high-deformable damped system, namely a simulated 600-DoF thin cloth falling onto spheres of different radius, akin to (Friedl et al., 2025). The system is intrinsically damped due to external dissipation forces τ_d . We train two RO-HNNs with $d = \{6, 10\}$ and a dissipative geometric HNN with Strang-symplectic integration on 20 trajectories of 3000 observations $\{q_i, p_i, \tau_i\}$ each, where $\tau = \tau_c$ are measured external constraint forces (see App. H.3). Fig. 5 depicts the predicted cloth configurations for a horizon $H\Delta t = 0.5$ s, showing that the RO-HNN accurately predicts the high-dimensional dissipative dynamics of the cloth, generalizing beyond the data support ($t > 0.3$ s). We observe that the 10-dimensional model preserves more details of the cloth than the 6-dimensional one, as also validated by the prediction errors reported in Ta-

ble 4. We compare the RO-HNN against a RO-LNN (Friedl et al., 2025) with latent dimension $d = 10$, for which dissipation forces are not learned but provided as ground truth. The last row of Fig. 5 and Table 4 show that the RO-HNN predictions are more accurate than the RO-LNN’s, despite that the former learns the dissipation forces via the latent damping matrix. Table 4 also compares the RO-HNN with a HNKO, which leads to noticeably worse predictions compared to all reduced-order models. We hypothesize this due to the fact that HNKO does not explicitly enforce the symplectic structure of its high-dimensional ($d = 80$) latent dynamics. App. I.3 reports additional results, ablations, and comparisons with RO-LNNs (see Fig. 17).

Latent damping. We compare the performance of the dissipative RO-HNN against (1) a conservative RO-HNN, where the dissipation forces τ_d are not learned but provided as ground truth in the external input $\tau = \tau_c + \tau_d$, and (2) a dissipative RO-HNN where the dissipation matrix is parametrized via Cholesky decomposition. Note that the mass-inertia matrix is parametrized via SPD networks in all cases. Table 4 shows that both dissipative RO-HNNs successfully learn the dissipation forces, achieving similar prediction errors as the conservative RO-HNN (see also Fig. 14 in App. I.3). The geometric HNN slightly outperforms its Cholesky counterpart, showing the importance of considering geometry. However, the effect is less pronounced as when learning the inverse mass-inertia matrix, which we attribute to the reduced influence of damping compared to inertia in the overall dynamics. Table 4 reports additional results for a non-symplectic AE combined with a geometric latent HNN, and a symplectic AE combined with a black-box latent HNN. Both variants operate in a $d = 10$ -dimensional latent space and receive damping forces as ground truth values. Both variants are significantly outperformed by the RO-HNNs, indicating that preserving symplectic structure in both the latent representation and dynamic model contribute to accurate long-horizon pre-

Table 3. Particle vortex: Mean and standard deviation of RO-HNN prediction errors ($\downarrow$) for different AEs and latent HNNs, and of HNKO over 10 test trajectories.

AE	HNN	d	$\frac{\ q_p - q\ }{\ q\ }$	$\frac{\ p_p - p\ }{\ p\ }$
Symplectic	Black-box	6	$(6.73 \pm 2.83) \times 10^{-1}$	$(6.28 \pm 2.18) \times 10^{-1}$
		10	$(7.34 \pm 3.03) \times 10^{-1}$	$(7.04 \pm 2.27) \times 10^{-1}$
	Geometric	6	$(4.00 \pm 2.01) \times 10^{-1}$	$(4.33 \pm 0.14) \times 10^{-1}$
		10	$(4.44 \pm 0.60) \times 10^{-1}$	$(4.60 \pm 0.32) \times 10^{-1}$
Non-Symplectic	Black-box	6	$(6.49 \pm 3.19) \times 10^{-1}$	$(6.72 \pm 3.70) \times 10^{-1}$
		10	$(7.24 \pm 1.09) \times 10^{-1}$	$(7.48 \pm 2.58) \times 10^{-1}$
	Geometric	6	$(6.14 \pm 2.26) \times 10^{-1}$	$(7.18 \pm 2.84) \times 10^{-1}$
		10	$(7.31 \pm 0.88) \times 10^{-1}$	$(6.99 \pm 0.65) \times 10^{-1}$
HNKO		800	$(1.60 \pm 0.56) \times 10^0$	$(1.52 \pm 0.85) \times 10^0$

Table 4. Cloth: Mean and standard deviation of RO-HNN reconstruction and prediction errors ($\downarrow$) for different parametrization of the latent dissipation matrix $\tilde{D}$, and different architecture variants and baselines over 10 test trajectories. Note that the reduced-order LNN (RO-LNN) does not predict momenta p .

AE	HNN	d	$\tilde{D}$	$\ \hat{q}_p - q\ /\ q\ $	$\ \hat{p}_p - p\ /\ p\ $	$\ \hat{q}_p - \hat{q}\ /\ \hat{q}\ $	$\ \hat{p}_p - \hat{p}\ /\ \hat{p}\ $
Symplectic	Geometric	6	Cholesky	$(4.21 \pm 1.07) \times 10^{-2}$	$(4.11 \pm 2.93) \times 10^{-1}$	$(3.45 \pm 1.13) \times 10^{-2}$	$(6.90 \pm 10.90) \times 10^{-2}$
			SPD	$(4.15 \pm 2.10) \times 10^{-2}$	$(3.58 \pm 3.03) \times 10^{-1}$	$(3.11 \pm 1.86) \times 10^{-2}$	$(8.42 \pm 10.40) \times 10^{-2}$
			Ground truth	$(3.18 \pm 0.74) \times 10^{-2}$	$(3.45 \pm 3.99) \times 10^{-1}$	$(2.58 \pm 0.94) \times 10^{-2}$	$(4.77 \pm 5.36) \times 10^{-2}$
		10	Cholesky	$(2.62 \pm 0.74) \times 10^{-2}$	$(3.39 \pm 3.01) \times 10^{-1}$	$(1.86 \pm 0.69) \times 10^{-2}$	$(3.90 \pm 4.36) \times 10^{-2}$
			SPD	$(3.21 \pm 1.25) \times 10^{-2}$	$(3.33 \pm 2.88) \times 10^{-1}$	$(3.45 \pm 1.13) \times 10^{-2}$	$(1.96 \pm 1.10) \times 10^{-2}$
			Ground truth	$(2.31 \pm 0.70) \times 10^{-2}$	$(3.44 \pm 3.05) \times 10^{-1}$	$(1.37 \pm 0.63) \times 10^{-2}$	$(4.26 \pm 4.58) \times 10^{-2}$
Symplectic	Black-box	10	Ground truth	$(1.95 \pm 3.72) \times 10^0$	$(2.01 \pm 7.59) \times 10^2$	$(1.76 \pm 3.15) \times 10^0$	$(6.83 \pm 8.26) \times 10^1$
Non-Symplectic	Geometric			$(3.14 \pm 3.48) \times 10^0$	$(7.89 \pm 19.04) \times 10^1$	$(2.63 \pm 2.59) \times 10^0$	$(1.37 \pm 3.21) \times 10^2$
RO-LNN		10	Ground truth	$(1.05 \pm 0.78) \times 10^1$	—	$(1.87 \pm 1.35) \times 10^{-1}$	—
HNKO		80	Ground truth	$(4.56 \pm 14.13) \times 10^2$	$(3.86 \pm 1.52) \times 10^2$	$(1.76 \pm 0.52) \times 10^{-1}$	$(4.83 \pm 7.44) \times 10^{-1}$




Figure 5. Predicted positions of the damped cloth with RO-HNNs with $d = \{6, 10\}$ and RO-LNN with $d = 10$ for a $625 \times$ longer horizon than during training. Times beyond 0.3s are out of the training data distribution.

diction. As the dissipative dynamics do not preserve the symplectic structure, we compare the Strang symplectic integrator, which assumes a symplectic structure, against a non-symplectic Runge-Kutta integrator of order 4. Fig. 20 in App. I.3 shows that, despite the dissipative structure, the Strang symplectic integrator outperforms the Runge-Kutta one. We hypothesize that this is due to the fact that the evolution of this dissipative system is mostly governed by its Hamiltonian function, especially over the short timesteps taken by the integrators.

5. Conclusions

This paper proposed a novel physics-inspired neural network, RO-HNN, for learning the unknown dynamics of

high-dimensional Hamiltonian systems from data. Our model provides physically-consistent, accurate, and stable predictions that generalize beyond the data support. To achieve this, it systematically integrates geometric inductive bias by defining structure-preserving symplectic embeddings, considering the geometry of the dynamics parameters within the model and for optimization, and leveraging structure-preserving symplectic integrators. We showed that the structural incorporation of these priors in the architecture is essential to learn high-dimensional dynamics, as Euclidean and soft-constrained approaches consistently underperformed. Our experiments on systems ranging from 15 to 600 DoFs demonstrate that RO-HNNs efficiently scale to high-dimensional systems through their reduced descriptions. While these results highlight the scalability of RO-HNN, a remaining challenge is managing the trade-off inherent to the latent dimension: Although larger latent spaces can capture more complex dynamics, they increase the difficulty of training of the latent HNN. Future work will extend RO-HNN to Hamiltonian systems with non-canonical symplectic forms. To do so, we plan to leverage Darboux theorem to explore the development of local RO-HNNs. We will also generalize the RO-HNN to learn more general dynamics of Port-Hamiltonian and contact Hamiltonian systems. Finally, we will investigate model-based control strategies within the RO-HNN latent space, akin to the controllers developed by Friedl et al. (2026) for Lagrangian systems.

Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

Acknowledgements

This work was supported by ERC AdV grant BIRD 884807, Swedish Research Council, and the Knut and Alice Wallenberg Foundation, including the Wallenberg AI, Autonomous Systems and Software Program (WASP). The computations were enabled by the Berzelius resource provided by the Knut and Alice Wallenberg Foundation at the National Supercomputer Centre.

References

- Abraham, R. and Marsden, J. E. *Foundations of Mechanics*. Addison-Wesley Publishing Company, Inc., second edition, 1987.
- Absil, P.-A., Mahony, R., and Sepulchre, R. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, 2007. URL <https://press.princeton.edu/absil>.
- Becigneul, G. and Ganea, O.-E. Riemannian adaptive optimization methods. In *Intl. Conf. on Learning Representations (ICLR)*, 2019. URL <https://openreview.net/forum?id=rlei qi09K7>.
- Bendokat, T. and Zimmermann, R. Geometric optimization for structure-preserving model reduction of Hamiltonian systems. *IFAC-PapersOnLine*, 55(20):457–462, 2022. doi: 10.1016/j.ifacol.2022.09.137.
- Botev, A., Jaegle, A., Wirnsberger, P., Hennes, D., and Higgins, I. Which priors matter? Benchmarking models for learning latent dynamics. In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, volume 1, 2021. URL https://datasets-benchmarks-proceedings.neurips.cc/paper_files/paper/2021/file/f033ab37c30201f73f142449d037028d-Paper-round1.pdf.
- Boumal, N. *An introduction to optimization on smooth manifolds*. Cambridge University Press, 2023. URL <http://www.nicolasboumal.net/book>.
- Brunton, S. L., Proctor, J. L., and Kutz, J. N. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 113(15):3932–3937, 2016. doi: 10.1073/pnas.1517384113.
- Buchfink, P., Glas, S., and Haasdonk, B. Symplectic model reduction of Hamiltonian systems on nonlinear manifolds and approximation with weakly symplectic autoencoder. *SIAM Journal of Scientific Computing*, 45:A289–A311, 2023. doi: 10.1137/21m1466657.
- Buchfink, P., Glas, S., Haasdonk, B., and Unger, B. Model reduction on manifolds: A differential geometric framework. *Physica D: Nonlinear Phenomena*, 468:134299, 2024. doi: 10.1016/j.physd.2024.134299.
- Champion, K., Lusch, B., Kutz, J. N., and Brunton, S. L. Data-driven discovery of coordinates and governing equations. *Proceedings of the National Academy of Sciences*, 116(45):22445–22451, 2019. doi: 10.1073/pnas.1906995116.
- Chen, Y., Matsubara, T., and Yaguchi, T. Neural symplectic form: Learning Hamiltonian equations on general coordinate systems. In *Neural Information Processing Systems (NeurIPS)*, 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/file/8b519f198dd26772e3e82874826b04aa-Paper.pdf.
- Chen, Z., Zhang, J., Arjovsky, M., and Bottou, L. Symplectic recurrent neural networks. In *Intl. Conf. on Learning Representations (ICLR)*, 2020. URL <https://openreview.net/forum?id=BkgYPREtPr>.
- Cranmer, M. D., Greydanus, S., Hoyer, S., Battaglia, P. W., Spergel, D. N., and Ho, S. Lagrangian neural networks. In *ICLR Deep Differential Equations Workshop*, 2020. URL <https://arxiv.org/abs/2003.04630>.
- Duarte, P., Fernandes, R. L., and Oliva, W. M. Dynamics of the attractor in the lotka–volterra equations. *Journal of Differential Equations*, 149(1):143–189, 1998. doi: 10.1006/jdeq.1998.3443.
- Duong, T. and Atanasov, N. Hamiltonian-based neural ODE networks on the SE(3) manifold for dynamics learning and control. In *Robotics: Science and Systems (R:SS)*, 2021. doi: 10.15607/RSS.2021.XVII.086.
- Feragen, A. and Fuster, A. Geometries and interpolations for symmetric positive definite matrices. In Schultz, T., Özarslan, E., and Hotz, I. (eds.), *Modeling, analysis, and visualization of anisotropy*, Mathematics and Visualization, pp. 85–113. Springer, 2017. doi: 10.1007/978-3-319-61358-1_5.

- Friedl, K., Jaquier, N., Lundell, J., Asfour, T., and Kragic, D. A Riemannian framework for learning reduced-order Lagrangian dynamics. In *Intl. Conf. on Learning Representations (ICLR)*, 2025. URL <https://openreview.net/forum?id=RoN6M3i7gJ>.
- Friedl, K., Jaquier, N., Asfour, T., Lundell, J., and Kragic, D. Reduced-order control and geometric structure of learned lagrangian latent dynamics. *arXiv preprint arXiv:2602.08963*, 2026. doi: 10.48550/arXiv.2602.08963.
- Greydanus, S., Dzamba, M., and Yosinski, J. Hamiltonian neural networks. In *Neural Information Processing Systems (NeurIPS)*, volume 32, 2019. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/26cd8ecadce0d4efd6cc8a8725cbd1f8-Paper.pdf.
- Hamilton, W. R. On a general method in dynamics. *Philosophical Transactions of the Royal Society*, pp. 247–308, 1834. doi: 10.1098/rstl.1834.0017.
- Jin, P., Zhang, Z., Zhu, A., Tang, Y., and Karniadakis, G. E. Sympnets: Intrinsic structure-preserving symplectic networks for identifying Hamiltonian systems. *Neural Networks*, 132:166–179, 2020. doi: 10.1016/j.neunet.2020.08.017.
- Jin, P., Zhang, Z., Kevrekidis, I. G., and Karniadakis, G. E. Learning poisson systems and trajectories of autonomous systems via poisson neural networks. *IEEE Trans. on Neural Networks and Learning Systems*, 34:8271–8283, 2021. doi: 10.1109/tnnls.2022.3148734.
- Jnini, A., Breschi, L., and Vella, F. Riemann tensor neural networks: Learning conservative systems with physics-constrained networks. In *Intl. Conf. on Machine Learning (ICML)*, 2025. URL <https://openreview.net/forum?id=cPMhMoJLAX&noteId=I8LbqaDsgw>.
- Kochurov, M., Karimov, R., and Kozlukov, S. Geoopt: Riemannian optimization in PyTorch. *arXiv:2005.02819*, 2020. URL <https://github.com/geoopt/geoopt>.
- Lee, J. M. *Introduction to smooth manifolds*. Springer, 2013. doi: 10.1007/978-1-4419-9982-5.
- Leimkuhler, B. and Reich, S. *Simulating Hamiltonian Dynamics*. Cambridge University Press, 2005. doi: 10.1017/CBO9780511614118.
- Lepri, M., Bacciu, D., and Santina, C. D. Neural autoencoder-based structure-preserving model order reduction and control design for high-dimensional physical systems. *IEEE Control Systems Letters*, 8:133–138, 2024. doi: 10.1109/LCSYS.2023.3344286.
- Liu, J., Borja, P., and Santina, C. D. Physics-informed neural networks to model and control robots: A theoretical and experimental investigation. *Advanced Intelligent Systems*, 6(5), 2024a. doi: 10.1002/aisy.202300385.
- Liu, N., Fan, Y., Zeng, X., Klöwer, M., Zhang, L., and Yu, Y. Harnessing the power of neural operators with automatically encoded conservation laws. In *Intl. Conf. on Machine Learning (ICML)*, 2024b. URL <https://proceedings.mlr.press/v235/liu24p.html>.
- Lusch, B., Kutz, J. N., and Brunton, S. L. Deep learning for universal linear embeddings of nonlinear dynamics. *Nature Communications*, 9(1):4950, 2018. doi: 10.1038/s41467-018-07210-0.
- Lutter, M. and Peters, J. Combining physics and deep learning to learn continuous-time dynamics models. *Intl. Journal of Robotics Research*, 42(3):83–107, 2023. doi: 10.1177/02783649231169492.
- Otto, S. E., Macchio, G. R., and Rowley, C. W. Learning nonlinear projections for reduced-order modeling of dynamical systems using constrained autoencoders. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 33(11), 2023. doi: 10.1063/5.0169688.
- Peng, L. and Mohseni, K. Symplectic model reduction of Hamiltonian systems. *SIAM Journal on Scientific Computing*, 38(1):A1–A27, 2016. doi: 10.1137/140978922.
- Pennec, X., Fillard, P., and Ayache, N. A Riemannian framework for tensor computing. *International Journal of Computer Vision*, 66(1):41–66, 2006. doi: 10.1007/s11263-005-3222-z.
- Rahma, A., Datar, C., Cukarska, A., and Dietrich, F. Rapid training of hamiltonian graph networks using random features. In *Intl. Conf. on Learning Representations (ICLR)*, 2026. URL <https://openreview.net/forum?id=2T4KxtDquV>.
- Salmon, R. Hamiltonian fluid mechanics. *Annual Review of Fluid Mechanics*, 20:225–256, 1988. doi: 10.1146/annurev.fl.20.010188.001301.
- Sanchez-Gonzalez, A., Bapst, V., Cranmer, K., and Battaglia, P. Hamiltonian graph networks with ode integrators. In *NeurIPS Workshop on Machine Learning and the Physical Sciences*, 2019.
- Saqlain, S., Zhu, W., Charalampidis, E. G., and Kevrekidis, P. G. Discovering governing equations in discrete systems using pinns. *Communications in Nonlinear Science and Numerical Simulation*, 126:107498, 2023. doi: <https://doi.org/10.1016/j.cnsns.2023.107498>.

- Schilders, W. H. A., Vorst, H. A., and Rommes, J. *Model Order Reduction: Theory, Research Aspects and Applications*. Springer Verlag, 2008. doi: 10.1007/978-3-540-78841-6.
- Schrödinger, E. Quantisierung als eigenwertproblem. *Annalen der Physik*, 384(4):361–376, 1926. doi: 10.1002/andp.19263840404.
- Sharma, H. and Kramer, B. Preserving lagrangian structure in data-driven reduced-order modeling of large-scale dynamical systems. *Physica D: Nonlinear Phenomena*, 462: 134128, 2024. doi: 10.1016/j.physd.2024.134128.
- Sharma, H., Mu, H., Buchfink, P., Geelen, R., Glas, S., and Kramer, B. Symplectic model reduction of Hamiltonian systems using data-driven quadratic manifolds. *Computer Methods in Applied Mechanics and Engineering*, 417: 116402, 2023. doi: 10.1016/j.cma.2023.116402.
- Sharma, H., Najera-Flores, D. A., Todd, M. D., and Kramer, B. Lagrangian operator inference enhanced with structure-preserving machine learning for nonintrusive model reduction of mechanical systems. *Computer Methods in Applied Mechanics and Engineering*, 423:116865, 2024. doi: 10.1016/j.cma.2024.116865.
- Sosanya, A. and Greydanus, S. Dissipative Hamiltonian neural networks: Learning dissipative and conservative dynamics separately. *arXiv preprint arXiv:2201.10085*, 2022. doi: 10.48550/arxiv.2201.10085.
- Tao, M. Explicit symplectic approximation of nonseparable Hamiltonians: Algorithm and long time performance. *Phys. Rev. E*, 94:043303, 2016. doi: 10.1103/PhysRevE.94.043303.
- Thangamuthu, A., Kumar, G., Bishnoi, S., Bhattoo, R., Krishnan, N. M. A., and Ranu, S. Unraveling the performance of physics-informed graph neural networks for dynamical systems. In *Neural Information Processing Systems (NeurIPS)*, 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/17b598fda495256bef6785c2b76c3217-Paper-Datasets_and_Benchmarks.pdf.
- Todorov, E., Erez, T., and Tassa, Y. Mujoco: A physics engine for model-based control. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pp. 5026–5033, 2012. doi: 10.1109/IROS.2012.6386109.
- Wang, S., Teng, Y., and Perdikaris, P. Understanding and mitigating gradient flow pathologies in physics-informed neural networks. *SIAM Journal on Scientific Computing*, 43(5):A3055–A3081, 2021. doi: 10.1137/20m1318043.
- Weinstein, A. Symplectic manifolds and their Lagrangian submanifolds. *Advances in Mathematics*, 6(3):329–346, 1971. doi: [https://doi.org/10.1016/0001-8708\(71\)90020-X](https://doi.org/10.1016/0001-8708(71)90020-X).
- Xiong, S., Tong, Y., He, X., Yang, S., Yang, C., and Zhu, B. Nonseparable symplectic neural networks. In *Intl. Conf. on Learning Representations (ICLR)*, 2021. URL <https://openreview.net/forum?id=B5VvQrI49Pa>.
- Zhang, J., Zhu, Q., and Lin, W. Learning Hamiltonian neural Koopman operator and simultaneously sustaining and discovering conservation laws. *Physical Review Research*, 6(1), 2024. doi: 10.1103/PhysRevResearch.6.L012031.
- Zhong, Y. D., Dey, B., and Chakraborty, A. Dissipative SymODEN: Encoding Hamiltonian dynamics with dissipation and control into deep learning. In *ICLR Workshop on Integration of Deep Neural Models and Differential Equations (DeepDiffEq)*, 2020a. URL <https://openreview.net/pdf?id=knjWfNx6CN>.
- Zhong, Y. D., Dey, B., and Chakraborty, A. Symplectic ODE-net: Learning Hamiltonian dynamics with control. In *Intl. Conf. on Learning Representations (ICLR)*, 2020b. URL <https://openreview.net/forum?id=ryxmb1rKDS>.
- Zhong, Y. D., Dey, B., and Chakraborty, A. Extending Lagrangian and Hamiltonian neural networks with differentiable contact models. In *Neural Information Processing Systems (NeurIPS)*, volume 34, pp. 21910–21922, 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/file/b7a8486459730bea9569414ef76cf03f-Paper.pdf.
- Šípka, M., Pavelka, M., Esen, O., and Grmela, M. Direct poisson neural networks: learning non-symplectic mechanical systems. *Journal of Physics A: Mathematical and Theoretical*, 56(49):495201, 2023. doi: 10.1088/1751-8121/ad0803.

A. Riemannian and Symplectic Geometry

In this section, we provide a short background on Riemannian and symplectic geometry, which compose the theoretical backbone of the RO-HNN. We refer the interested reader to (Abraham & Marsden, 1987; Lee, 2013) for more details. Table 5 briefly summarizes the main geometric concepts and their notation.

Table 5. Main geometric terminology and notation.

Symbol	Meaning
$\mathcal{M}$	smooth manifold, $\dim(\mathcal{M}) = 2n$
ω	symplectic form
$(\mathcal{M}, \omega)$	symplectic manifold, i.e. smooth manifold $\mathcal{M}$ equipped with symplectic form ω
$\mathcal{H} : \mathcal{M} \rightarrow \mathbb{R}$	Hamiltonian functional defined on a symplectic manifold
$(\mathcal{M}, \omega, \mathcal{H})$	Hamiltonian system, defined by a Hamiltonian $\mathcal{H}$ evolving on a symplectic manifold $(\mathcal{M}, \omega)$
$\mathcal{X}_{\mathcal{H}}$	Hamiltonian vector field, governing system evolution of a Hamiltonian system
$\mathbb{J}_{2n}$	canonical symplectic form in coordinates
$(\mathcal{M}, \mathbb{J}_{2n}, \mathcal{H})$	canonical Hamiltonian system, i.e., Hamiltonian system on symplectic manifold with canonical symplectic form
$q \in \mathcal{Q}$	position on configuration manifold, $\dim(\mathcal{Q}) = n$
$\mathcal{T}_q \mathcal{Q}$	tangent space of configuration manifold $\mathcal{Q}$ at position q
$\mathcal{T}\mathcal{Q}$	tangent bundle of configuration manifold $\mathcal{Q}$
$\mathcal{T}_q^* \mathcal{Q}$	cotangent space of configuration manifold $\mathcal{Q}$ at position q
$\mathcal{T}^* \mathcal{Q}$	cotangent bundle of configuration manifold $\mathcal{Q}$
$p \in \mathcal{T}^* \mathcal{Q}$	momenta, elements of the cotangent bundle
$\varphi_{\mathcal{Q}}$	embedding of a low-dimensional submanifold
$\rho_{\mathcal{Q}}$	reduction onto a low-dimensional submanifold, left inverse of $\varphi_{\mathcal{Q}}$ on its domain.
φ	embedding of a low-dimensional symplectic submanifold
ρ	reduction onto a low-dimensional symplectic submanifold, left inverse of φ on its domain.
$\varphi^* \mathcal{H} = \mathcal{H} \circ \varphi$	pullback of the Hamiltonian $\mathcal{H}$ by function φ via pre-composition
$\varphi^* \omega = d\varphi^T \omega d\varphi$	pullback of the symplectic form ω by function φ
df	differential of a smooth map f , corresponding in coordinates to the Jacobian, or for scalar-valued functions, to the transpose gradient
$\vdots$	reduced, low-dimensional, submanifold quantities
symplectomorphism	map between two symplectic manifolds that preserves the symplectic form
cotangent lift	extending a map from a configuration manifold to a cotangent bundle

Riemannian and symplectic manifolds are smooth manifolds with special structures. A smooth manifold $\mathcal{M}$ of dimension n can be intuitively conceptualized as a manifold that is locally, but not globally, similar to the Euclidean space $\mathbb{R}^n$. The smooth structure of $\mathcal{M}$ allows the definition of derivative of curves on the manifold, which are tangent vectors. The set of all tangent vectors at a point $x \in \mathcal{M}$ defines the tangent space $\mathcal{T}_x \mathcal{M}$ which is a n -dimensional vector space. Tangent vectors can be represented on an ordered basis of $\mathcal{T}_x \mathcal{M}$ as $v = v^i \frac{\partial}{\partial x^i} |_{\mathbf{x}}$. The tangent bundle $\mathcal{T}\mathcal{M}$ is the disjoint union of all tangent spaces on $\mathcal{M}$ and is $2n$ -dimensional smooth manifold.

The cotangent space $\mathcal{T}_x^* \mathcal{M}$ at $x \in \mathcal{M}$ is the dual of the tangent space $\mathcal{T}_x \mathcal{M}$, i.e., $\mathcal{T}_x^* \mathcal{M} = \{\lambda | \lambda : \mathcal{T}_x \mathcal{M} \rightarrow \mathbb{R} \text{ linear}\}$. Cotangent vectors can be represented on an ordered basis of $\mathcal{T}_x^* \mathcal{M}$ as $\lambda = \lambda_i dx^i |_{\mathbf{x}}$. The cotangent bundle $\mathcal{T}^* \mathcal{M}$ is the disjoint union of all cotangent spaces on $\mathcal{M}$ and is $2n$ -dimensional smooth manifold, similarly as the tangent bundle.

A smooth mapping f between two smooth manifolds $\tilde{\mathcal{M}}$ and $\mathcal{M}$ with $\dim(\tilde{\mathcal{M}}) = d \ll \dim(\mathcal{M}) = n$ is an immersion if the differential $df|_{\tilde{\mathbf{x}}} : \mathcal{T}_{\tilde{\mathbf{x}}} \tilde{\mathcal{M}} \rightarrow \mathcal{T}_{f(\tilde{\mathbf{x}})} \mathcal{M}$ is injective at each $\tilde{\mathbf{x}} \in \tilde{\mathcal{M}}$. An embedding is an immersion that is also a homeomorphism onto its image, i.e., it is an injective and structure-preserving map. In this case, $\tilde{\mathcal{M}}$ is an embedded submanifold of $\mathcal{M}$. The pullback of a function $h : \mathcal{M} \rightarrow \mathbb{R}$ by a smooth mapping $f : \mathcal{N} \rightarrow \mathcal{M}$ between two smooth manifolds $\mathcal{N}$ and $\mathcal{M}$ is a smooth function $f^* h$ with

$$f^* h(\mathbf{x}) = h(f(\mathbf{x})) = (h \circ f)(\mathbf{x}). \quad (13)$$

A.1. Riemannian Geometry

A Riemannian manifold $(\mathcal{M}, g)$ is a smooth manifold $\mathcal{M}$ endowed with a Riemannian metric g , i.e., a smoothly-varying inner product $g_x : \mathcal{T}_x \mathcal{M} \times \mathcal{T}_x \mathcal{M} \rightarrow \mathbb{R}$. In coordinates, a Riemannian metric is represented by a SPD matrix. The Riemannian metric defines the notion of distance on the manifold, as well as the so-called geodesics, which are length-minimizing curves on the manifold.

Learning and optimization methods involving Riemannian data typically take advantage of their Euclidean tangent spaces to operate. Specifically, the exponential map $\text{Exp}_x : \mathcal{T}_x \mathcal{M} \rightarrow \mathcal{M}$ and logarithmic map $\text{Log}_x : \mathcal{M} \rightarrow \mathcal{T}_x \mathcal{M}$, derived from the Riemannian metric, allows us to map back and forth between the Euclidean tangent space and the manifold. Moreover,

the parallel transport $\text{PT}_{x \rightarrow y} : \mathcal{T}_x \mathcal{M} \rightarrow \mathcal{T}_y \mathcal{M}$ move tangent vectors across tangent spaces such that their inner product is conserved.

A Lagrangian system $(\mathcal{M}, g, \mathcal{L})$ is a dynamical system evolving on a Riemannian manifold $(\mathcal{M}, g)$ according to a smooth Lagrangian function $\mathcal{L} : \mathcal{T}\mathcal{M} \rightarrow \mathbb{R}$.

A.2. Symplectic Geometry

A symplectic manifold $(\mathcal{M}, \omega)$ is a $2n$ -dimensional smooth manifold $\mathcal{M}$ equipped with a symplectic form ω , i.e., a closed, non-degenerate, differential 2-form $g_x : \mathcal{T}_x \mathcal{M} \times \mathcal{T}_x \mathcal{M} \rightarrow \mathbb{R}$, which satisfies

$$\omega(\mathbf{u}, \mathbf{v}) = -\omega(\mathbf{v}, \mathbf{u}), \quad \omega(\mathbf{u}, \mathbf{v}) \forall \mathbf{v} \Rightarrow \mathbf{u} = \mathbf{0}, \quad \text{and} \quad d\omega = 0 \quad (14)$$

for all $\mathbf{u}, \mathbf{v} \in \mathcal{T}_x \mathcal{M}$. In coordinates, a symplectic form is represented by a skew-symmetric matrix ω . We slightly abuse notation, equivalently denoting symplectic manifolds as $(\mathcal{M}, \omega)$. Notice that the non-degeneracy of ω implies that all symplectic manifolds are of even dimension.

A diffeomorphism $f : (\mathcal{M}, \omega) \rightarrow (\mathcal{N}, \eta)$ between symplectic manifolds is a symplectomorphism if it preserves the symplectic form, i.e., $f^* \eta = \omega$ with $f^* \eta$ denoting the pullback of η by f . The Hamiltonian flow $\phi_t : (\mathcal{M}, \omega) \rightarrow (\mathcal{M}, \omega)$ induced by $\mathbf{X}_{\mathcal{H}}$ is a symplectomorphism, as it maps points $\mathbf{x} \in \mathcal{M}$ along the integral curves of the manifold thus preserving the symplectic form.

Following Darboux' theorem, there exists a canonical chart (U, ϕ) , $\mathbf{x} \in U$ for each point $\mathbf{x} \in \mathcal{M}$ in which the symplectic form is represented as $\omega = \mathbb{J}_{2n}^\top$ via the canonical Poisson tensor

$$\mathbb{J}_{2n} = \begin{pmatrix} \mathbf{0} & \mathbf{I}_n \\ -\mathbf{I}_n & \mathbf{0} \end{pmatrix}, \quad \text{for which} \quad \mathbb{J}_{2n}^\top = \mathbb{J}_{2n}^{-1} = -\mathbb{J}_{2n}.$$

In other words, every symplectic manifold is locally symplectomorphic to $(\mathbb{R}^{2n}, \mathbb{J}_{2n}^\top)$. A system $(\mathbb{R}^{2n}, \mathbb{J}_{2n}^\top, \mathcal{H})$ is called a canonical Hamiltonian system. Moreover, the cotangent bundle $\mathcal{T}^* \mathcal{M}$ any n -dimensional smooth manifold $\mathcal{M}$ carries a canonical symplectic structure, making it a symplectic manifold $(\mathcal{T}^* \mathcal{M}, \mathbb{J}_{2n})$.

A Hamiltonian system $(\mathcal{M}, \omega, \mathcal{H})$ is a dynamical system evolving on a symplectic manifold $(\mathcal{M}, \omega)$ according to a smooth Hamiltonian function $\mathcal{H} : \mathcal{M} \rightarrow \mathbb{R}$.

B. Riemannian Manifolds of Interest

This section provides a brief overview of the Riemannian manifolds of interest for this paper, namely the manifold of SPD matrices $\mathcal{S}_{++}^n$ (App. B.1), and the biorthogonal manifold $\mathcal{B}_{n,d}$ (App. B.2).

B.1. The Manifold of SPD Matrices

We denote the set of $n \times n$ symmetric matrices as $\text{Sym}^n = \{\mathbf{S} \in \mathbb{R}^{n \times n} | \mathbf{S} = \mathbf{S}^\top\}$. The set of SPD matrices $\mathcal{S}_{++}^n = \{\mathbf{\Sigma} \in \text{Sym}^n | \mathbf{\Sigma} \succ \mathbf{0}\}$ forms a smooth manifold of dimension $\dim(\mathcal{S}_{++}^n) = \frac{n(n+1)}{2}$, which can be represented as the interior of a convex cone embedded in Sym^n . The tangent space $\mathcal{T}_{\mathbf{\Sigma}} \mathcal{S}_{++}^n$ at a point $\mathbf{\Sigma} \in \mathcal{S}_{++}^n$ is identified with Sym^n .

The SPD manifold can be endowed with various Riemannian metrics, resulting in different theoretical properties and closed-form operations. We utilize the widely-used affine-invariant metric (Pennec et al., 2006), which places symmetric matrices with non-positive eigenvalues at infinite distance from any SPD matrix and prevents the well-known swelling effect (Feragen & Fuster, 2017). The affine-invariant metric defines the inner product $g : \mathcal{T}_{\mathbf{\Sigma}} \mathcal{S}_{++}^n \times \mathcal{T}_{\mathbf{\Sigma}} \mathcal{S}_{++}^n \rightarrow \mathbb{R}$ given two matrices $\mathbf{T}_1, \mathbf{T}_2 \in \mathcal{T}_{\mathbf{\Sigma}} \mathcal{S}_{++}^n$, as

$$\langle \mathbf{T}_1, \mathbf{T}_2 \rangle_{\mathbf{\Sigma}} = \text{tr}(\mathbf{\Sigma}^{-\frac{1}{2}} \mathbf{T}_1 \mathbf{\Sigma}^{-1} \mathbf{T}_2 \mathbf{\Sigma}^{-\frac{1}{2}}). \quad (15)$$

The corresponding geodesic distance, exponential map, logarithmic maps, and parallel transport are computed in closed

form as

$$d_{\mathcal{M}}(\Lambda, \Sigma) = \|\log(\Sigma^{-\frac{1}{2}} \Lambda \Sigma^{-\frac{1}{2}})\|_F, \quad (16)$$

$$\text{Exp}_{\Sigma}(S) = \Sigma^{\frac{1}{2}} \exp(\Sigma^{-\frac{1}{2}} S \Sigma^{-\frac{1}{2}}) \Sigma^{\frac{1}{2}}, \quad (17)$$

$$\text{Log}_{\Sigma}(\Lambda) = \Sigma^{\frac{1}{2}} \log(\Sigma^{-\frac{1}{2}} \Lambda \Sigma^{-\frac{1}{2}}) \Sigma^{\frac{1}{2}}, \quad (18)$$

$$\text{PT}_{\Sigma \rightarrow \Lambda}(T) = A_{\Sigma \rightarrow \Lambda} T A_{\Sigma \rightarrow \Lambda}^T, \quad (19)$$

where $\exp(\cdot)$ and $\log(\cdot)$ denote the matrix exponential and logarithm functions, and $A_{\Sigma \rightarrow \Lambda} = \Lambda^{\frac{1}{2}} \Sigma^{-\frac{1}{2}}$. These operations are key for the SPD networks encoding the mass-inertia and damping matrices in geometric HNNs (see Sec. 3.2), and for the on-manifold parameter optimization of SPD parameters of the network when training the model (see Sec. 3.4).

B.2. The Biorthogonal Manifold

The biorthogonal manifold is the smooth manifold $\mathcal{B}_{n,d} = \{(\Phi, \Psi) \in \mathbb{R}^{n \times d} \times \mathbb{R}^{n \times d} \mid \Psi^T \Phi = I_d\}$ formed by pairs of full-row-rank matrices $\Phi, \Psi \in \mathbb{R}^{n \times d}$, with $n \geq d \geq 1$ satisfying the biorthogonality condition $\Psi^T \Phi = I$ (Otto et al., 2023). The biorthogonal matrix manifold $\mathcal{B}_{n,d}$ is an embedded submanifold of the Euclidean product space $\mathbb{R}^{n \times d} \times \mathbb{R}^{n \times d}$ with dimension $\dim(\mathcal{B}_{n,d}) = 2nd - d^2$. The tangent space at a point $(\Phi, \Psi) \in \mathcal{B}_{n,d}$ is given by

$$\mathcal{T}_{(\Phi, \Psi)} \mathcal{B}_{n,d} = \{(V, W) \in \mathbb{R}^{n \times d} \times \mathbb{R}^{n \times d} : W^T \Phi + \Psi^T V = 0\}. \quad (20)$$

A pair of matrices $(X, Y) \in \mathbb{R}^{n \times d} \times \mathbb{R}^{n \times d}$ can be projected onto the tangent space $\mathcal{T}_{(\Phi, \Psi)} \mathcal{B}_{n,d}$ via the projection operation $\text{Proj}_{(\Phi, \Psi)} : \mathbb{R}^{n \times d} \times \mathbb{R}^{n \times d} \rightarrow \mathcal{T}_{(\Phi, \Psi)} \mathcal{B}_{n,d}$ defined as

$$\text{Proj}_{(\Phi, \Psi)}(X, Y) = (X - \Psi A, Y - \Phi A^T), \quad (21)$$

where A is a solution to the Sylvester equation $A(\Phi^T \Phi) + (\Psi^T \Psi)A = Y^T \Phi + \Psi^T X$.

When optimizing the parameters of the geometrically-constrained symplectic AE presented in Sec. 3.1, it is crucial to account for the biorthogonal geometry of the pairs of weight matrices (Friedl et al., 2025). Therefore, we train the model by optimizing pairs of weight matrices via Riemannian optimization on the biorthogonal manifold (see Sec. 3.4). Riemannian optimization algorithms utilize the exponential map and the parallel transport operations, which are difficult to obtain in closed form for the biorthogonal manifold. Therefore, we leverage a first-order approximation of the exponential map, i.e., a retraction map $R_{(\Phi, \Psi)} : \mathcal{T}_{(\Phi, \Psi)} \mathcal{B}_{n,d} \rightarrow \mathcal{B}_{n,d}$, defined as

$$R_{(\Phi, \Psi)}(V, W) = \left((\Phi + V) ((\Psi + W)^T (\Phi + V))^{-1}, (\Psi + W) \right). \quad (22)$$

Moreover, we use a first-order approximation of the parallel transport operation defined via the successive application of retraction and projection as

$$\text{PT}_{(\Phi_1, \Psi_1) \rightarrow (\Phi_2, \Psi_2)} = \text{Proj}_{(\Phi_2, \Psi_2)} \circ R_{(\Phi_1, \Psi_1)}. \quad (23)$$

C. SPD Network

As explained in Sec. 3.2, we learn reduced Hamiltonian dynamics in the embedded symplectic submanifold via a latent geometric HNN that parametrizes the inverse mass-inertia and damping matrices via SPD networks that account for their intrinsic geometry. We use a SPD network introduced in (Friedl et al., 2025) composed of (1) Euclidean layers $g_{\mathbb{R}}$, and (2) an exponential map layer g_{Exp} , which we detail next.

Euclidean Layers $g_{\mathbb{R}}$. The SPD network leverages classical fully-connected layers to model functions that return elements on the tangent space of a manifold. The output of the l -th Euclidean layer $x^{(l)}$ is given by

$$x^{(l)} = \sigma(A_l x^{(l-1)} + b_l), \quad (24)$$

with $A_l \in \mathbb{R}^{n_l \times n_{l-1}}$ and $b_l \in \mathbb{R}^{n_l}$ the weight matrix and bias of the layer l , and σ a nonlinear activation function of choice.

Exponential Map Layer g_{Exp} . The exponential map layer is used to map layer inputs $\mathbf{X}^{(l-1)} \in \text{Sym}^n$ from the tangent space onto the manifold $\mathcal{S}_{++}^n$. The layer output is given by

$$\mathbf{X}^{(l)} = \text{Exp}_{\mathbf{P}}(\mathbf{X}^{(l-1)}), \quad (25)$$

with $\mathbf{P} \in \mathcal{S}_{++}^n$ denoting the basepoint of the considered tangent space. Following the results of the ablation conducted in (Friedl et al., 2025), we define $\mathbf{P}$ as equal to the identity matrix $\mathbf{I}$, so that the layer input is assumed to lie in the tangent space at the origin of the cone.

Note that Friedl et al. (2025) additionally consider SPD layers mapping SPD matrices to SPD matrices, analogous to fully-connected Euclidean layers. However, the SPD networks with additional SPD layers were shown to achieve similar performances as those employing solely Euclidean and exponential-map layers. Therefore, we do not integrate such layers in the SPD networks of the RO-HNN.

D. Proofs of Propositions of Sec. 3

Proposition D.1. *The reduction map $\rho(\mathbf{q}, \mathbf{p})$ (9) satisfies the projection properties (5).*

Proof. It is clear that the cotangent-lifted map ρ fulfills (5) as $\rho_{\mathcal{Q}}$ satisfies (5) by assumption. $\square$

Proposition D.2. *The embedding $\varphi(\tilde{\mathbf{q}}, \tilde{\mathbf{p}})$ (9) satisfies the symplecticity property (8).*

Proof. Proving the statement is equivalent to show that the differential $d\varphi = \begin{pmatrix} d\varphi_{\mathcal{Q}} & \mathbf{0} \\ \frac{\partial(d\rho_{\mathcal{Q}}^{\top} \tilde{\mathbf{p}})}{\partial \tilde{\mathbf{q}}} & d\rho_{\mathcal{Q}}^{\top} \end{pmatrix}$ belongs to the symplectic Stiefel manifold $\text{Sp}(2n, 2d) = \{\mathbf{U} \in \mathbb{R}^{2n \times 2d} \mid \mathbf{U}^{\top} \mathbb{J}_{2n} \mathbf{U} = \mathbb{J}_{2d}\}$. A block matrix $\mathbf{U} = \begin{pmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{C} & \mathbf{D} \end{pmatrix}$ belongs to $\text{Sp}(2n, 2d)$ if its block elements satisfy the condition

$$\mathbf{U}^{\top} \mathbb{J}_{2n} \mathbf{U} = \begin{pmatrix} \mathbf{A}^{\top} & \mathbf{C}^{\top} \\ \mathbf{B}^{\top} & \mathbf{D}^{\top} \end{pmatrix} \begin{pmatrix} \mathbf{0} & -\mathbf{I}_n \\ \mathbf{I}_n & \mathbf{0} \end{pmatrix} \begin{pmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{C} & \mathbf{D} \end{pmatrix} = \begin{pmatrix} \mathbf{C}^{\top} \mathbf{A} - \mathbf{A}^{\top} \mathbf{C} & \mathbf{C}^{\top} \mathbf{B} - \mathbf{A}^{\top} \mathbf{D} \\ \mathbf{D}^{\top} \mathbf{A} - \mathbf{B}^{\top} \mathbf{C} & \mathbf{D}^{\top} \mathbf{B} - \mathbf{B}^{\top} \mathbf{D} \end{pmatrix} = \begin{pmatrix} \mathbf{0} & -\mathbf{I}_d \\ \mathbf{I}_d & \mathbf{0} \end{pmatrix},$$

i.e., the differential $d\varphi$ must satisfy

$$d\varphi^{\top} \mathbb{J} d\varphi = \begin{pmatrix} \mathbf{C}^{\top} d\varphi_{\mathcal{Q}} - d\varphi_{\mathcal{Q}}^{\top} \mathbf{C} & -d\varphi_{\mathcal{Q}}^{\top} d\rho_{\mathcal{Q}}^{\top} \\ d\rho_{\mathcal{Q}} d\varphi_{\mathcal{Q}} & \mathbf{0} \end{pmatrix} = \begin{pmatrix} \mathbf{0} & -\mathbf{I}_d \\ \mathbf{I}_d & \mathbf{0} \end{pmatrix} \quad \text{with} \quad \mathbf{C} = \frac{\partial(d\rho_{\mathcal{Q}}|_{\varphi_{\mathcal{Q}}(\tilde{\mathbf{q}})}^{\top} \tilde{\mathbf{p}})}{\partial \tilde{\mathbf{q}}}. \quad (26)$$

By assumption, $\rho_{\mathcal{Q}}$ fulfills the projection properties (5), so that $d\rho_{\mathcal{Q}} d\varphi_{\mathcal{Q}} = d\varphi_{\mathcal{Q}}^{\top} d\rho_{\mathcal{Q}}^{\top} = \mathbf{I}_d$ holds by construction. It remains to prove $\mathbf{C}^{\top} d\varphi_{\mathcal{Q}} - d\varphi_{\mathcal{Q}}^{\top} \mathbf{C} = \mathbf{0}$. We denote the elements of the canonical and reduced canonical coordinates as q^i, p_i and $\tilde{q}^{\alpha}, \tilde{p}_{\alpha}$, respectively. By definition, we have $(d\varphi_{\mathcal{Q}})_i^{\alpha} = \frac{\partial q^i}{\partial \tilde{q}^{\alpha}}$ and $(d\rho_{\mathcal{Q}})_i^{\alpha} = \frac{\partial \tilde{q}^{\alpha}}{\partial q^i}$ and the projection properties hold by assumption, i.e., $(d\rho_{\mathcal{Q}})_i^{\alpha} (d\varphi_{\mathcal{Q}})_i^{\beta} = \delta_{\beta}^{\alpha} \forall \tilde{q} \in \tilde{\mathcal{Q}}$, with $\delta_{\beta}^{\alpha} = 1$ if $\alpha = \beta$ and $\delta_{\beta}^{\alpha} = 0$ otherwise. Therefore, we have $p_i = (d\rho_{\mathcal{Q}})_i^{\alpha} \tilde{p}_{\alpha}$, and

$$C_{i\gamma} = \frac{\partial p_i}{\partial \tilde{q}^{\gamma}} = \frac{\partial}{\partial \tilde{q}^{\gamma}} ((d\rho_{\mathcal{Q}})_i^{\alpha} \tilde{p}_{\alpha}) = \frac{\partial (d\rho_{\mathcal{Q}})_i^{\alpha}}{\partial \tilde{q}^{\gamma}} \frac{\partial q^j}{\partial \tilde{q}^{\gamma}} \tilde{p}_{\alpha} = \frac{\partial (d\rho_{\mathcal{Q}})_i^{\alpha}}{\partial q^j} (d\varphi_{\mathcal{Q}})_i^j \tilde{p}_{\alpha}.$$

We aim to show that $\mathbf{C}^{\top} d\varphi_{\mathcal{Q}}$ is symmetric, i.e., $(d\varphi_{\mathcal{Q}})_i^j C_{i\gamma} = (d\varphi_{\mathcal{Q}})_i^j C_{i\beta}$. Using the projection properties, we can write $\tilde{p}_{\beta} = (d\varphi_{\mathcal{Q}})_i^j (d\rho_{\mathcal{Q}})_i^{\alpha} \tilde{p}_{\alpha}$. Differentiating with respect to $\tilde{q}^{\gamma}$ yields

$$\begin{aligned} 0 &= \frac{\partial}{\partial \tilde{q}^{\gamma}} ((d\varphi_{\mathcal{Q}})_i^j (d\rho_{\mathcal{Q}})_i^{\alpha} \tilde{p}_{\alpha}) = \frac{(d\varphi_{\mathcal{Q}})_i^j}{\partial \tilde{q}^{\gamma}} (d\rho_{\mathcal{Q}})_i^{\alpha} \tilde{p}_{\alpha} + (d\varphi_{\mathcal{Q}})_i^j \frac{(d\rho_{\mathcal{Q}})_i^{\alpha}}{\partial q^j} \frac{\partial q^j}{\partial \tilde{q}^{\gamma}} \tilde{p}_{\alpha} \\ &= \frac{(d\varphi_{\mathcal{Q}})_i^j}{\partial \tilde{q}^{\gamma}} p_i + (d\varphi_{\mathcal{Q}})_i^j C_{i\gamma} = \frac{\partial q^i}{\partial \tilde{q}^{\beta} \partial \tilde{q}^{\gamma}} p_i + (d\varphi_{\mathcal{Q}})_i^j C_{i\gamma}. \end{aligned}$$

As the Hessian in the first term is symmetric, the equality implies the symmetry of the second term, i.e.,

$$(d\varphi_{\mathcal{Q}})_i^j C_{i\gamma} = (d\varphi_{\mathcal{Q}})_i^j C_{i\beta},$$

and thus (26) holds. $\square$

Note that $d\rho\mathbb{J}_{2n}d\rho^\top = \mathbb{J}_{2d}$ is shown to hold on $\varphi(\check{\mathcal{M}})$ with similar arguments. Moreover, a similar proof is presented by Sharma et al. (2023) in the context of quadratic symplectic projections.

Proposition D.3. *The reduced vector field $\check{X} = \varphi^* X$ obtained via the pullback of the cotangent-lifted embedding φ (9) preserves the structure of the total vector field $X = X_{\mathcal{H}} + X_{\mathcal{F}}$ with $X_{\mathcal{F}} = \begin{pmatrix} 0 \\ \tau_{\text{ext}} + \tau_d \end{pmatrix}$.*

Proof. The reduced vector field decomposes as $\check{X} = \varphi^* X = \varphi^* X_{\mathcal{H}} + \varphi^* X_{\mathcal{F}} = \check{X}_{\check{\mathcal{H}}} + \check{X}_{\check{\mathcal{F}}}$ with $\check{X}_{\check{\mathcal{H}}} = \check{\omega}^{-1} d\check{\mathcal{H}}$ (see Sec. 2.2) and $\check{X}_{\check{\mathcal{F}}}$ is obtained by pulling back the external and damping terms. Since generalized forces belong to the cotangent bundle $\mathcal{T}^*Q$, they are embedded and reduced via the cotangent-lifted maps (9) as $\varphi(\check{q}, \check{\tau})$ and $\rho(q, \tau)$, leading to the reduced external inputs $\check{\tau}_{\text{ext}} = d\varphi_Q|_{\check{q}}^\top \tau_{\text{ext}}$. The reduced Rayleigh dissipative function is obtained via the pullback of the tangent-lifted embedding $\varphi_{\mathcal{T}Q}(\check{q}, \dot{\check{q}}) = (\varphi_Q(\check{q})^\top, (d\varphi_Q|_{\check{q}} \dot{\check{q}})^\top)^\top$ as $\check{D} = \varphi_{\mathcal{T}Q}^* D = \frac{1}{2} \dot{\check{q}}^\top \check{D}(\check{q}) \dot{\check{q}}$ with positive semi-definite reduced damping matrix $\check{D}(\check{q}) = d\varphi_Q^\top D(q) d\varphi_Q$. The reduced damping force is then $\check{\tau}_d = \frac{\partial \check{D}(\check{q}, \dot{\check{q}})}{\partial \dot{\check{q}}} = \check{D}(\check{q}) \dot{\check{q}}$. Therefore, the reduced force field is $\check{X}_{\check{\mathcal{F}}} = \begin{pmatrix} 0 \\ \check{\tau}_{\text{ext}} + \check{\tau}_d \end{pmatrix}$. $\square$

E. Additional Details on the Geometrically-constrained Symplectic Autoencoder

E.1. Constrained Autoencoder

The geometrically-constrained symplectic AE presented in Sec. 3.1 builds on the constrained AE architecture introduced in (Otto et al., 2023). Specifically, we learn the embedding φ_Q and associated point reduction ρ_Q via a constrained AE with layer pairs (10), and compute their differential to construct the tangent-lifted maps (9), as explained in Sec. 3.1. To guarantee the projection properties, the constrained AE architecture from (Otto et al., 2023) leverages pairs of biorthogonal weight matrices, which are described in Sec. 3.1, and pairs of invertible activation functions, which we introduce next.

The nonlinear activation functions σ_- and σ_+ employed in the encoder and decoder network must satisfy $\sigma_- \circ \sigma_+ = \text{id}$. To do so, they are defined as

$$\sigma_{\pm}(x_i) = \frac{bx_i}{a} \mp \frac{\sqrt{2}}{a \sin(\alpha)} \pm \frac{1}{a} \sqrt{\left(\frac{2x_i}{\sin(\alpha) \cos(\alpha)} \mp \frac{\sqrt{2}}{\cos(\alpha)} \right)^2 + 2a}, \quad (27)$$

with

$$\begin{cases} a &= \csc^2(\alpha) - \sec^2(\alpha), \\ b &= \csc^2(\alpha) + \sec^2(\alpha). \end{cases} \quad (28)$$

The activations then resemble smooth, rotation-symmetric versions of the common leaky ReLU activations. The parameter $0 < \alpha < \frac{\pi}{4}$ sets the slope of the activation functions. Throughout our experiments, we set $\alpha = \frac{\pi}{8}$.

Otto et al. (2023) proposed to incorporate the biorthogonality of the weight matrices by considering an overparametrization of the biorthogonal weights along with a soft constraint in the form of additional penalty losses. However, this approach does not guarantee the biorthogonality condition, in contrast to the Riemannian approach we use in this paper. Moreover, as shown in (Friedl et al., 2025), the overparametrized model leads to higher reconstruction errors compared to constrained AE trained on the biorthogonal manifold.

E.2. Computation of the Cotangent-lifted Maps

We construct the cotangent-lifted maps (9) by differentiating the outputs of the encoder ρ_Q and decoder φ_Q networks with respect to their inputs. To avoid the computational cost related to the automatically-differentiated transposed Jacobian-vector product, our implementation computes layer-wise analytical derivatives and obtains the full differentials via the chain rule. The derivatives of the nonlinear activations $\sigma_{\pm}$ are given by

$$\sigma'_{\pm}(x_i) = \frac{d}{dx_i} \sigma_{\pm}(x_i) = \frac{b}{a} \pm \frac{2}{a \sin(\alpha) \cos(\alpha)} \frac{\frac{2x_i}{\sin(\alpha) \cos(\alpha)} \mp \frac{\sqrt{2}}{\cos(\alpha)}}{\sqrt{\left(\frac{2x_i}{\sin(\alpha) \cos(\alpha)} \mp \frac{\sqrt{2}}{\cos(\alpha)} \right)^2 + 2a}}, \quad (29)$$

thus fulfilling the inverse-derivative property $\sigma'_-(\sigma_+(x_i)) \sigma'_+(x_i) = 1$ by construction.

The pullbacks $d\rho_Q|_{\varphi_Q(\tilde{q})}^\top \tilde{p}$ and $d\varphi_Q|_{\rho_Q(q)}^\top p$ are computed analytically as a composition of transposed layer derivatives $d\rho_Q^\top = d\rho_Q^{(L)\top} \circ \dots \circ d\rho_Q^{(1)\top}$ and $d\varphi_Q = d\varphi_Q^{(1)} \circ \dots \circ d\varphi_Q^{(L)}$, with $d\rho_Q^{(l)} \in \mathbb{R}^{n_{l-1} \times n_l}$ and $d\varphi_Q^{(l)} \in \mathbb{R}^{n_l \times n_{l-1}}$. From the definition of the layer pairs (10), the transpose of the layer derivatives are given as

$$d\rho_Q^{(l)}|_{q^{(l-1)}}^\top = \Psi_l \text{diag}(\sigma'_-(q^{(l-1)})) \quad \text{and} \quad d\varphi_Q^{(l)}|_{\tilde{q}^{(l-1)}}^\top = \text{diag}(\sigma'_+(\tilde{q}^{(l-1)})) \Phi_l^\top, \quad (30)$$

$$\text{with } \text{diag}(v) = \begin{pmatrix} v_1 & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & v_d \end{pmatrix}.$$

The computation of the layer derivatives requires storing the intermediate reduced and reconstructed positions $\tilde{q}^{(l-1)}$ and $q^{(l-1)}$ for each layer, which are obtained during the forward pass through the position encoder ρ_Q and decoder φ_Q . During the momentum forward pass, we store each intermediate $p^{(l-1)} = d\rho_Q^{(l)\top} p^{(l)}$ and $\tilde{p}^{(l)} = d\varphi_Q^{(l)\top} p^{(l-1)}$. This allows the computational cost of one momentum forward pass to roughly be equal to that of one forward pass of the position projection, scaling constantly through the matrix-multiplication of weights with system dimensionality $\dim(Q)$. We provide wall-clock evaluation times of our geometrically-constrained symplectic AE on the 600-DoF cloth dataset in App. H.3.

F. Strang-Symplectic Integrator

As explained in Sec. 3.3, we integrate the learned reduced-order Hamiltonian flow (11) using the second-order symplectic integrator of (Tao, 2016), which we refer to as Strang-symplectic integrator.

The Strang-symplectic integrator approximates the flow of a non-separable Hamiltonian function $\mathcal{H}(q, p)$ by considering an augmented Hamiltonian function

$$\bar{\mathcal{H}}(q, p, x, y) = \mathcal{H}_A(q, y) + \mathcal{H}_B(p, x) + w\mathcal{H}_C(q, p, x, y), \quad (31)$$

in an extended phase space, where $\mathcal{H}_A(q, y)$ and $\mathcal{H}_B(p, x)$ are two copies of the original system with mixed-up positions and momenta, and $\mathcal{H}_C = \frac{1}{2}(\|q, x\|^2 + \|p, y\|^2)$ is an artificial restraint with parameter w controlling the binding of $\mathcal{H}_A(q, y)$ and $\mathcal{H}_B(p, x)$. The dynamics of the augmented Hamiltonian $\bar{\mathcal{H}}$ are

$$\dot{q} = \frac{\partial}{\partial p} \bar{\mathcal{H}}(q, p, x, y) = \frac{\partial}{\partial p} H(x, p) + w(p - y) \quad (32)$$

$$\dot{p} = \frac{\partial}{\partial q} \bar{\mathcal{H}}(q, p, x, y) = \frac{\partial}{\partial q} H(q, y) - w(q - x) \quad (33)$$

$$\dot{x} = \frac{\partial}{\partial y} \bar{\mathcal{H}}(q, p, x, y) = \frac{\partial}{\partial y} H(q, y) + w(y - p) \quad (34)$$

$$\dot{y} = \frac{\partial}{\partial p} \bar{\mathcal{H}}(q, p, x, y) = \frac{\partial}{\partial p} H(x, p) - w(x - q) \quad (35)$$

and leads to the same exact IVP solutions as the original function $\mathcal{H}(q, p)$. High-order symplectic integrators can be construct for each of the component of the augmented Hamiltonian $\bar{\mathcal{H}}$ as

$$\phi_{\mathcal{H}_A}^\delta = \begin{pmatrix} q \\ p - \delta \frac{\partial}{\partial q} H(q, y) \\ x + \delta \frac{\partial}{\partial y} H(q, y) \\ y \end{pmatrix}, \quad \phi_{\mathcal{H}_B}^\delta = \begin{pmatrix} q + \delta \frac{\partial}{\partial p} H(x, p) \\ p \\ x \\ y - \delta \frac{\partial}{\partial x} H(x, p) \end{pmatrix}, \quad (36)$$

$$\phi_{w\mathcal{H}_C}^\delta = \frac{1}{2} \left(\begin{pmatrix} q+x \\ p+y \end{pmatrix} + R(\delta) \begin{pmatrix} q-x \\ p-y \end{pmatrix} \right), \quad \text{with } R(\delta) = \begin{pmatrix} \cos(2w\delta)\mathbf{I} & \sin(2w\delta)\mathbf{I} \\ -\sin(2w\delta)\mathbf{I} & \cos(2w\delta)\mathbf{I} \end{pmatrix}. \quad (37)$$

Tao (2016) proposed to construct a numerical symplectic integrator that approximates the flow of $\bar{\mathcal{H}}$ by composing these maps according to Strang splitting as

$$\phi_{\bar{\mathcal{H}}} = \phi_{\mathcal{H}_A}^{\delta/2} \circ \phi_{\mathcal{H}_B}^{\delta/2} \circ \phi_{w\mathcal{H}_C}^{\delta/2} \circ \phi_{\mathcal{H}_B}^{\delta/2} \circ \phi_{\mathcal{H}_A}^{\delta/2}. \quad (38)$$

The obtained Strang-symplectic integrator preserves the symplectic volume like the exact Hamiltonian flow.

The scalar parameter $w \in \mathbb{R}$, binding the two augmented Hamiltonians during the integration process, is obtained as optimization parameter during training. To enforce $w \leq 0$, we do not learn w directly. Instead, we learn it using the SoftPlus function with a small numerical offset for stability as $\log(1 + e^{\theta_w}) + 10^{-4}$, as part of the HNN network parameters $\theta_w \in \theta$.

G. Network Training via Riemannian Optimization

Training a neural network corresponds to finding a solution to an optimization problem

$$\min_{\mathbf{x} \in \mathcal{M}} \ell(\mathbf{x}), \quad (39)$$

where ℓ is the loss we aim to minimize, and $\mathbf{x} \in \mathcal{M}$ is the optimization variable, a.k.a the network parameters. For the RO-HNN, we train the network by minimizing the loss $\ell_{\text{RO-HNN}}$ (12). In this case, $\mathcal{M}$ is defined as a product of Euclidean, SPD, and biorthogonal manifolds to jointly optimize the parameters $\{\Phi_l, \Psi_l, \mathbf{b}_l\}_{l=1}^L$ of the AE and $\{\theta_{\tilde{T}}, \theta_{\tilde{V}}, \theta_{\tilde{D}}\}$ of the latent geometric HNN. To account for the curvature of the non-Euclidean parameter spaces, we leverage Riemannian optimization (Absil et al., 2007; Boumal, 2023) to optimize the RO-HNN loss $\ell_{\text{RO-HNN}}$ (12).

Conceptually, each iteration step in a first-order (stochastic) Riemannian optimization method consists of the three following successive operations:

$$\boldsymbol{\eta}_t \leftarrow h(\text{grad } \ell(\mathbf{x}_t), \boldsymbol{\tau}_{t-1}), \quad \mathbf{x}_{t+1} \leftarrow \text{Exp}_{\mathbf{x}_t}(-\alpha_t \boldsymbol{\eta}_t), \quad \boldsymbol{\tau}_t \leftarrow \text{PT}_{\mathbf{x}_t \rightarrow \mathbf{x}_{t+1}}(\boldsymbol{\eta}_t). \quad (40)$$

where (1) given the current parameter estimate $\mathbf{x}_t$, a search direction $\boldsymbol{\eta}_t \in \mathcal{T}_{\mathbf{x}_t} \mathcal{M}$ is computed based on a function h (determined by the choice of the optimization method) of the Riemannian gradient $\text{grad } \ell$, and of $\boldsymbol{\tau}_{t-1}$, which corresponds to the parallel-transport of the previous search direction on to the new estimate’s tangent space $\mathcal{T}_{\mathbf{x}_t} \mathcal{M}$; (2) the estimate $\mathbf{x}_t$ is updated by projecting the search direction $\boldsymbol{\eta}_t$ scaled by a learning rate α_t onto the manifold via the exponential map, (3) the current search direction is parallel-transported to the tangent space of the updated estimate to prepare for the next iteration. In this paper, we use the Riemannian Adam (Becigneul & Ganea, 2019) implemented in Geoopt (Kochurov et al., 2020) to optimize the RO-HNN parameters. The relevant manifold operations for the optimization procedure are given in closed-form in App. B.

H. Additional Details on Experiments

This section presents additional details on the experimental setup of Sec. 4.

H.1. Coupled Pendulum of Section 4.1

H.1.1. DATASET

System. Our first set of experiments is conducted on the augmented pendulum, a nonlinear conservative system with $n = 15$ -DoF. The pendulum dynamics are specified from the symplectomorphism of a latent Hamiltonian system composed of two independent subsystems: a 3-DoF planar pendulum, and a 12-DoF planar oscillating mass-spring mesh, see Fig. 6. The pendulum dynamics evolve on a slower timescale and with larger amplitude than the mesh oscillations. Consequently, a surrogate model based solely on the pendulum would capture the dominant behavior of the full system, i.e. the system is well-reducible with a Hamiltonian ROM. As we have access to the ground truth dynamics of the 15-DoF pendulum, this scenario allows for various ablations on the network architecture.

We simulate both subsystems in MUJOCO (Todorov et al., 2012). The pendulum’s links $i = \{1, 2, 3\}$ are modeled as capsules of radius $r_i = 0.025$ m, length $l_i = 0.5$ m, and mass $m_i = 0.5$ kg, connected via hinge joints. The initial configurations and velocities for each DoF are randomly sampled from the intervals $q_{\text{pend},i}(t = 0) \in [-30, 30]^\circ$ and $\dot{q}_{\text{pend},i}(t = 0) \in [-23, 23]^\circ \text{ s}^{-1}$. The mass-spring mesh consists of 6 masses $m_j = 0.005$ kg, equally spaced in a 3×2 grid along the x - and z -axes of the simulation environment. Each mass is connected to its immediate neighbors, and the top three masses are each additionally connected to a fixed anchor point above the grid, via springs of resting length $s_j = 0.5$ m and linear stiffness constants $k_j = 0.01 \text{ N m}^{-1}$. Initial displacements and velocities for each DoF are randomly sampled from the intervals $q_{\text{ms},j}(t = 0) \in [-1, 1] \times 10^{-2} \text{ m}$ and $\dot{q}_{\text{ms},j}(t = 0) \in [-2, 2] \times 10^{-3} \text{ m s}^{-1}$.

Data generation. Each simulation is recorded for $T = 5$ s at a timestep of $\Delta t = 10^{-2}$ s, yielding $N = 30$ training trajectories $\mathcal{D}_{\text{pend}} = \{\{q_{\text{pend},n,k}, p_{\text{pend},n,k}\}_{k=1}^K\}_{n=1}^N$ and $\mathcal{D}_{\text{ms}} = \{\{q_{\text{ms},n,k}, p_{\text{ms},n,k}\}_{k=1}^K\}_{n=1}^N$ with $K = 500$ observations




Figure 6. Illustration of the latent system used to obtain the dynamics of a 15-DoF augmented pendulum via a symplectomorphism. The latent system consists of an independent 3-DoF planar pendulum and a 12-DoF planar mass-spring mesh.

each. To form the full 15-dimensional dataset, the position and momentum vectors of the pendulum and mass-spring mesh are concatenated as $\mathbf{q}_{\text{aug}} = (\mathbf{q}_{\text{pend}}^\top \quad \mathbf{q}_{\text{ms}}^\top)^\top$ and $\mathbf{p}_{\text{aug}} = (\mathbf{p}_{\text{pend}}^\top \quad \mathbf{p}_{\text{ms}}^\top)^\top$ to obtain a dataset $\mathcal{D}_{\text{aug}} = \{(\mathbf{q}_{\text{aug},n,k}, \mathbf{p}_{\text{aug},n,k})_{k=1}^K\}_{n=1}^N$ with $(\mathbf{q}_{\text{aug},n,k}, \mathbf{p}_{\text{aug},n,k}) \in \mathcal{T}^*(\mathcal{Q}_{\text{pend}} \times \mathcal{Q}_{\text{ms}})$.

To ensure that the reducibility of the augmented dataset is not purely of numerical nature, we transform the observed dynamics of the latent system onto more complex ones via a symplectomorphism $h : (\mathcal{T}^*(\mathcal{Q}_{\text{pend}} \times \mathcal{Q}_{\text{ms}}), \mathbb{J}_{2n}) \rightarrow (\mathcal{T}^*\mathcal{Q}, \mathbb{J}_{2n})$ and obtain the final dataset $\mathcal{D} = \{h(\mathbf{q}_{\text{aug},n,k}, \mathbf{p}_{\text{aug},n,k})_{k=1}^K\}_{n=1}^N$. Practically, the symplectomorphism h is defined via the cotangent-lifted embedding φ of a map $\rho_{\mathcal{Q}} : \mathcal{Q}_{\text{pend}} \times \mathcal{Q}_{\text{ms}} \rightarrow \mathcal{Q}$, that we parametrize as a 3-layer encoder of the constrained AE from Sec. 3.1 and E. With $l = \{1, 2, 3\}$ layers of constant layer and latent dimension $n_l = n_0 = 15$, weights of each layer Ψ_l initialized as random orthogonal matrices $\mathbf{O} \in \mathbb{R}^{n_l \times n_l}$ sampled from the Haar distribution, and zero biases $\mathbf{b}_l = \mathbf{0}$. Notice that due to the constant dimension through the AE-layers, with decoder weights set to $\Phi_l = \Psi_l = \mathbf{O}$, the position decoder of the constrained AE returns an analytic inverse, and its cotangent lift h^{-1} .

The testing dataset is constructed in the same manner for $N = 10$ trajectories.

H.1.2. MODEL TRAINING

For the experiments of Sec. 4.1, we train a geometric RO-HNN composed of a geometrically-constrained symplectic AE and a latent geometric HNN. As described in Sec. 3.1, the geometrically-constrained symplectic AE is built from the cotangent lift of a constrained AE composed of layer pairs $\rho_{\mathcal{Q}}^{(l)} : \mathbb{R}^{n_l} \rightarrow \mathbb{R}^{n_{l-1}}$ and $\varphi_{\mathcal{Q}}^{(l)} : \mathbb{R}^{n_{l-1}} \rightarrow \mathbb{R}^{n_l}$ as defined in (10) (see Sec. 3.1 and App. E). We use $l = \{1, 2, 3\}$ pairwise biorthogonal encoder and decoder layers of sizes $n_l = \{6, 12, 15\}$ with latent space dimension $n_0 = 3$. The biorthogonal weight matrices are initialized by sampling a random orthogonal matrix $\mathbf{O} \in \mathbb{R}^{n_l \times n_l}$ from the Haar distribution and setting $\Phi = \Psi = \mathbf{O}_{[:,n_{l-1}]}$, where $\mathbf{O}_{[:,n_{l-1}]}$ are the first n_{l-1} column entries of $\mathbf{O}$. Bias vectors are initialized as $\mathbf{b}_l = \mathbf{0}$. For the latent geometric HNN, we parametrize the potential energy network $\tilde{V}_{\theta_{\tilde{V}}}$ and the Euclidean part $g_{\mathbb{R}}$ of the inverse mass-inertia network $\tilde{M}_{\theta_{\tilde{T}}}^{-1}$ each with $L_{\tilde{V}} = L_{\tilde{T},\mathbb{R}} = 2$ hidden Euclidean layers of 32 neurons and SoftPlus activation functions. We fix the basepoint of the exponential map layer g_{Exp} to the origin $\mathbf{P} = \mathbf{I}$. Weights are initialized by sampling from a Xavier normal distribution with gain $\sqrt{2}$ and bias vector entries set to 1. We train the model on the joint loss (12) with scaling factor $\lambda = 1$ for the latent loss on 3000 uniformly sampled random points from the dataset $\mathcal{D}$ with Strang-symplectic integration (see Sec. 3.3) over a training horizon of $H_{\mathcal{D}} = 12$ timesteps. We use a learning rate of 1.5×10^{-2} for the AE parameters and 7×10^{-4} for the HNN parameters. We train the model with Riemannian Adam (Becigneul & Ganea, 2019) until convergence at 3000 epochs.

AE baselines. In Sec. 4.1, we compare the geometrically-constrained symplectic AE with linear and quadratic symplectic manifold Galerkin (SMG) projections (Peng & Mohseni, 2016; Sharma et al., 2023), and a weakly-symplectic AE (Buchfink et al., 2023). We implement the linear and quadratic SMG projections onto a 3-dimensional symplectic submanifold following (Sharma et al., 2023). We compute the reduction parameters based on a singular value decomposition computed from 3000 randomly sampled training datapoints in $\mathcal{D}$.

For the weakly-symplectic AE (Buchfink et al., 2023), we train two independent constrained AEs for position and momentum reduction and embedding, i.e., $\rho_{\mathcal{Q}}^{(l)} : \mathbb{R}^{n_l} \rightarrow \mathbb{R}^{n_{l-1}}$, $\varphi_{\mathcal{Q}}^{(l)} : \mathbb{R}^{n_{l-1}} \rightarrow \mathbb{R}^{n_l}$, and $\rho_{\mathcal{P}}^{(l)} : \mathbb{R}^{n_l} \rightarrow \mathbb{R}^{n_{l-1}}$, $\varphi_{\mathcal{P}}^{(l)} : \mathbb{R}^{n_{l-1}} \rightarrow \mathbb{R}^{n_l}$, and

compute the embedding and reduction for the symplectic manifold as

$$\varphi(\check{\mathbf{q}}, \check{\mathbf{p}}) = \begin{pmatrix} \varphi_{\mathcal{Q}} \\ \varphi_{\mathcal{P}} \end{pmatrix} \quad \text{and} \quad \rho(\mathbf{q}, \mathbf{p}) = \begin{pmatrix} \rho_{\mathcal{Q}} \\ \rho_{\mathcal{P}} \end{pmatrix}. \quad (41)$$

Note that this architecture also fulfills the projection properties (5) by construction, as the other reduction approaches. However, it does not satisfy the symplecticity property (8). To enforce this property, Buchfink et al. (2023) introduces a symplecticity loss

$$\ell_{\text{symp}} = \frac{1}{N} \sum_{i=1}^N \|\mathbb{J}_{2d} - d\varphi^{\top} \mathbb{J}_{2n} d\varphi\|_{\text{F}}^2. \quad (42)$$

The weakly-symplectic AE is trained by minimizing the sum of the reconstruction loss ℓ_{AE} from (12) and the symplecticity loss (42).

For the geometrically-constrained symplectic AE, we consider the same architecture as in the RO-HNN described above.

All AE architectures consist of $l = \{1, 2, 3\}$ biorthogonal encoder and decoder layers with $n_l = \{6, 12, 15\}$ with latent space dimension $n_0 = 3$. We train both AE on 3000 samples from the dataset $\mathcal{D}$ with Riemannian Adam with a learning rate of 1.5×10^{-2} until convergence at 3000 epochs.

HNN baselines. In Sec. 4.1, we also ablate the choice of latent HNN and integrator. To isolate the HNN performance, we consider the low-dimensional dataset $\mathcal{D}_{\text{pend}}$ of the 3-DoF pendulum and no reduction. For the Cholesky HNN where the inverse mass-inertia matrix is parametrized via a Cholesky network, we implement shared parameters for the inverse mass-matrix and potential energy networks, i.e., $\theta_{\text{T}} \cap \theta_{\text{V}}$, following (Lutter & Peters, 2023). The MLP consists of 2 hidden Euclidean SoftPlus layers of 64 neurons, while separate output layers return the potential energy and the Cholesky decomposition. For the black-box HNNs, we use a single fully-connected MLP to model a Hamiltonian function $\mathcal{H}_{\theta}$. We conduct experiments with two black-box HNNs of 2 hidden layers with a width of 64, and 256 neurons, respectively. In all cases, the weights are initialized by sampling from a Xavier normal distribution with gain $\sqrt{2}$, and the bias vector entries are initialized to 1.

We train all architectures on 3000 datapoints of the dataset $\mathcal{D}_{\text{pend}}$ with Riemannian Adam optimizer on the HNN term $\ell_{\text{HNN},d}$ of the loss (12) over a training horizon of $H_{\mathcal{D}} = 12$ timesteps. For the ablation of the HNN architecture, we use the Strang-symplectic integrator. The geometric HNN and Cholesky networks are trained until convergence at 2500 epochs with learning rate set to 7×10^{-4} . The black-box HNNs are trained at a learning rate of 2×10^{-3} for 3000 epochs.

For the ablation of the integrator, we use the geometric HNN and compare the Strang-symplectic integrator with an explicit Euler integrator, a Runge-Kutta integrator of order 4, and a symplectic leapfrog integrator.

HNKO baseline. In Sec. 4.1, we compare the RO-HNN against the HNKO proposed by Zhang et al. (2024). Moreover, in App. I.1, we evaluate the performance of the RO-HNN against the HNKO under noisy observations. Hyperparameters are selected and refined empirically following the supplementary material and code provided by (Zhang et al., 2024).

The HNKO first maps the $2n = 30$ -dimensional observations into a $d = 100$ -dimensional lifted latent space via a fully-connected neural network of 6 hidden layers with Tanh activations. The latent dynamics are then propagated on a 50-dimensional sphere via the special orthogonal Hamiltonian Koopman operator, implemented by a constrained linear, bias-free layer with 100-dimensional input and output. The predicted states are mapped back onto the original space with a fully-connected neural network with 3 hidden layers and Tanh activations. The overall model is trained on 3000 randomly sampled datapoints of the dataset with $\mathcal{D}_{\text{pend}}$, using the Adam optimizer until convergence at 15000 epochs. For a fair comparison and for stable predictions over longer horizons, we adjusted the loss on the latent Koopman predictions, referred to as $\mathcal{L}_{\text{koop}}$ in (Zhang et al., 2024), to sum over a training horizon of $H_{\mathcal{D}} = 12$ timesteps, similar to our latent loss term $\ell_{\text{HNN},d}$ in (12).

H.2. Particle Vortex of Section 4.2

H.2.1. DATASET

System. In Sec. 4.2, we learn the dynamics of an $n = 90$ -dimensional particle vortex, consisting of $j = \{1, \dots, N\}$ particles with phase-space coordinates $\mathbf{x}_j = (q_j, p_j)^{\top}$ and uniform interaction strengths $\Gamma_j = 1$. The particle vortex dynamics are

governed by the Hamiltonian

$$\mathcal{H}(\mathbf{q}, \mathbf{p}) = - \sum_{j < k} \log |\mathbf{x}_j - \mathbf{x}_k|, \quad (43)$$

that models the interaction between each $j \neq k$ pair of particles (Xiong et al., 2021). Note that, as the particle vortex dynamics are purely determined via the logarithmic interaction, its Hamiltonian function does not separate into kinetic and potential energy, in contrast to mechanical systems such as the pendulum and the cloth.

Data generation. We generate a training dataset $\mathcal{D}_{\text{pv}} = \{\{\mathbf{q}_{n,k}, \mathbf{p}_{n,k}\}_{k=1}^K\}_{n=1}^N$ by simulating $N = 20$ trajectories of the conservative system over the time interval $\mathcal{I} = [0, 10.0]$ s with timestep $\Delta t = 10^{-3}$ s and Strang-symplectic solver with weight parameter $w = 0.1$, resulting in $K = 10000$ steps per trajectory. For each trajectory, initial conditions are randomly sampled to mimic clustered vortex distributions. The particles are evenly split among $j = \{1, 2, 3\}$ clusters. For each cluster, we randomly sample a center $\mathbf{c}_j$ within a radius of $R = 6$ m from the origin. Then, a cluster radius is sampled uniformly from $r_j \in [0.1, 2]$ m, and particles within a cluster are positioned following a Gaussian distribution $\mathcal{N} \sim (\mathbf{c}_j, r_j^2 \mathbf{I})$ around the center $\mathbf{c}_j$. For the testing dataset, we generate $N = 10$ trajectories via the same distribution of initial conditions, but simulating the system over a time interval of $\mathcal{I} = [0, 15.0]$ s.

H.2.2. MODEL TRAINING

The results presented in Sec. 4.2 are obtained via RO-HNNs composed of a geometrically-constrained symplectic AE and a latent geometric HNN. We conduct experiments with two RO-HNN with latent space dimensions $d = 6$ and $d = 10$. The constrained AE is composed of $l = \{1, 2, 3, 4\}$ pairwise biorthogonal encoder and decoder layers of sizes $n_l = \{32, 48, 64, 90\}$. The biorthogonal weight matrices are initialized by sampling a random orthogonal matrix $\mathbf{O} \in \mathbb{R}^{n_l \times n_l}$ from the Haar distribution and setting $\Phi = \Psi = \mathbf{O}_{[:, :n_{l-1}]}$, where $\mathbf{O}_{[:, :n_{l-1}]}$ are the first n_{l-1} column entries of $\mathbf{O}$. Bias vectors are initialized as $\mathbf{b}_l = \mathbf{0}$. The latent Hamiltonian network $\tilde{V}_{\theta_{\tilde{V}}}$ is parametrized by 2 hidden Euclidean layers of 32 neurons with SoftPlus activation functions. All weights are initialized by sampling from a Xavier normal distribution with gain $\sqrt{2}$, and all bias vector entries are initialized to 1.

We train the model on the joint loss (12) on 3000 random samples from the dataset $\mathcal{D}$ with Strang-symplectic integration over a training horizon $H_{\mathcal{D}} = 8$ timesteps. For better convergence, we scale the loss term $\ell_{\text{HNN},d}$ via a scalar factor $\lambda = 10^3$. The parameters are optimized via Riemannian Adam (Becigneul & Ganeva, 2019) until convergence at 3000 epochs with a learning rate of 1.5×10^{-2} for the AE parameters and 7×10^{-4} for the HNN parameters.

In Sec. 4.2, we consider a comparison with a RO-HNN with a latent black-box HNN $\tilde{\mathcal{H}}_{\theta}$ composed of 2 layers of 64 neurons. We set the learning rate to 2×10^{-3} . The remaining of the RO-HNN architecture and training pipeline are unchanged.

Non-symplectic AE baseline. In Table 3, we compare the RO-HNN with variants that use a non-symplectic AE instead of the proposed geometrically-constrained symplectic AE. We consider black-box and geometric HNNs in latent spaces of dimensions $d = \{6, 10\}$ learned by a non-symplectic projection-constrained AE. We use a single vanilla constrained AE with latent space dimension $n_0 = \{12, 20\}$ with 4 pairwise biorthogonal encoder-decoder layers of sizes $n_l = \{64, 96, 128, 180\}$. We assume the first $\{6, 10\}$ output dimensions of the latent space to correspond to reduced position $\tilde{\mathbf{q}}$ and the last $\{6, 10\}$ to correspond to the reduced momentum $\tilde{\mathbf{p}}$, which are used as inputs for the latent dissipative HNN. The training procedure remains the same as for the other RO-HNNs described above.

HNKO baseline. For the results in Table 3, the HNKO first maps the $2n = 180$ -dimensional particle observations into a $d = 800$ -dimensional latent space via a fully-connected neural network of 6 hidden layers with Tanh activations. The Hamiltonian latent dynamics are modeled on the 400-dimensional sphere. The resulting state predictions are reconstructed by a fully-connected neural network with 3 hidden layers with Tanh activations. The overall model is trained on 3000 randomly sampled datapoints, using the Adam optimizer until convergence at 15000 epochs. We train on a horizon of $H_{\mathcal{D}} = 8$ timesteps, as outlined in H.1.2.

H.3. Cloth of Section 4.3

H.3.1. DATASET

System. Our last set of experiments is conducted on a deformable thin cloth modeled in MUJOCO as a flexible composite object with $i = \{1, \dots, 200\}$ masses $m_i = 0.1$ kg, equally spaced over a width of 0.1 m and length of 0.2 m. Generalized coordinates are given by the Cartesian positions $\mathbf{q}_i = (x_i, y_i, z_i)^T$ of each mass' center of mass in the world frame. The

viscous damping coefficient is uniformly set to $d_i = 0.01 \text{ N s m}^{-1}$.

Data generation. Each trajectory captures the cloth falling on a sphere from a height of 0.12 m in the center above the origin of the sphere. To vary scenarios, the radius of the sphere is randomly-sampled from $r \in [0.02, 0.12] \text{ m}$. The state evolution is simulated with timestep $\Delta t = 10^{-4} \text{ s}$ over a time interval $\mathcal{I} = [0, 0.3] \text{ s}$, resulting in $K = 3000$ samples per trajectory. We generate $N = 20$ trajectories for a training dataset $\mathcal{D}_{\text{cloth}} = \{\{\mathbf{q}_{n,k}, \mathbf{p}_{n,k}, \boldsymbol{\tau}_{n,k}\}_{k=1}^K\}_{n=1}^N$, and $N = 10$ testing trajectories over a longer time interval $\mathcal{I} = [0, 0.5] \text{ s}$. When learning the damping force via a dissipative HNN, the generalized force vector consists of external constraint forces, i.e., $\boldsymbol{\tau} = \boldsymbol{\tau}_c$. The ablation of Sec. 4.3 compares the dissipative geometric HNN against a conservative HNN for which all external forces are provided. In this case, the training dataset is composed of generalized force vector $\boldsymbol{\tau} = \boldsymbol{\tau}_d + \boldsymbol{\tau}_c$ that contains both the damping forces $\boldsymbol{\tau}_d$ and the constraint forces $\boldsymbol{\tau}_c$.

H.3.2. MODEL TRAINING

For the RO-HNN experiments in Sec. 4.3, we train a RO-HNN composed of a geometrically-constrained symplectic AE and a latent dissipative geometric HNN. The underlying constrained AE has $l = \{1, 2, 3, 4\}$ pairwise biorthogonal encoder and decoder layers of sizes $n_l = \{32, 64, 128, 600\}$ with latent space dimension $n_0 = 6$ or $n_0 = 10$. The biorthogonal weight matrices are initialized by sampling a random orthogonal matrix $\mathbf{O} \in \mathbb{R}^{n_l \times n_l}$ from the Haar distribution and setting $\Phi_l = \Psi_l = \mathbf{O}_{[:,n_l-1]}$, where $\mathbf{O}_{[:,n_l-1]}$ are the first n_l-1 column entries of $\mathbf{O}$. Bias vectors are initialized as $\mathbf{b}_l = \mathbf{0}$.

The latent potential energy network $\check{V}_{\theta_V}$ is parametrized with $L_{\check{V}} = L_{\check{T}, \mathbb{R}} = 2$ hidden Euclidean layers of 32 neurons. The Euclidean part $g_{\mathbb{R}}$ of the inverse mass-inertia network $\check{M}_{\theta_T}^{-1}$ and of the damping-matrix network $\check{D}_{\theta_D}$ are composed of 2 hidden layers with 32 neurons. For both networks, we fix the basepoint of the exponential map layer g_{Exp} to the origin $\mathbf{P} = \mathbf{I}$. All activation functions are SoftPlus, all weights are initialized by sampling from a Xavier normal distribution with gain $\sqrt{2}$, and all bias vector entries are initialized to 1.

We train the model on the joint loss (12) on 3000 samples from the dataset $\mathcal{D}$ with Strang-symplectic integration over a training horizon $H_{\mathcal{D}} = 8$ timesteps. The scaling constant on the latent loss term $\ell_{\text{HNN},d}$ is set to $\lambda = 10^4$. We train the RO-HNN with Riemannian Adam (Becigneul & Ganea, 2019) until convergence at 3000 epochs with a learning rate of 1.5×10^{-2} for the AE parameters and 7×10^{-4} for the HNN parameters.

In Fig. 17, we compare the geometric RO-HNN with a black-box variant where the latent HNN is encoded as a single black-box network $\mathcal{H}_{\theta}$ corresponding to a fully-connected MLP of 2 hidden layers with a width of 256 neurons. The HNN weights are initialized by sampling from a Xavier normal distribution with gain $\sqrt{2}$, and the bias vector entries are initialized to 1. This black-box RO-HNN is training with the same parameters as the geometric RO-HNN, except for the learning rate of the HNN parameters, which we set as 2×10^{-3} .

Sequentially-trained baseline. To assess the effectiveness of the proposed joint training procedure, we compare the jointly-trained RO-HNN with a variant that sequentially trains first the geometrically-constrained AE, and second the latent HNN. As convergence is difficult to achieve when training only the latent HNN on a fully-trained representation of the AE, we first train only the AE by optimizing ℓ_{AE} for 3000 epochs with a learning rate of 1.5×10^{-2} . Subsequently, we jointly optimize the AE and latent loss (12). We train the networks jointly within the RO-HNN with Riemannian Adam (Becigneul & Ganea, 2019) until convergence at learning rates 1.5×10^{-2} for the AE parameters and 7×10^{-4} for the HNN parameters.

Projection and AE baselines. In App. I.3 (see Fig. 18), we compare the ability of a latent HNN to learn accurate dynamics using different reduction methods to obtain the symplectic embedding φ and corresponding reduction ρ . We compare the RO-HNN with geometrically-constrained symplectic AE with linear and quadratic symplectic manifold Galerkin (SMG) projections (Peng & Mohseni, 2016; Sharma et al., 2023), and a weakly-symplectic AE (Buchfink et al., 2023). We compute the linear and quadratic SMG projections onto latent spaces of symplectic submanifolds of three different dimensionalities $d = \{2, 6, 10\}$, following (Sharma et al., 2023), via 3000 training datapoints. In both cases, we then train a latent HNN on the terms $\ell_{\text{HNN},n}$ and $\ell_{\text{HNN},d}$ of the joint loss equation (12) on 3000 samples from the dataset $\mathcal{D}$ with Strang-symplectic integration over a training horizon $H_{\mathcal{D}} = 8$ timesteps. The model is trained with Riemannian Adam (Becigneul & Ganea, 2019) until convergence at 3000 epochs with a learning rate of 7×10^{-4} . Note that this essentially corresponds to a scenario with pre-trained symplectic submanifolds, as the parameter optimization for the linear and quadratic embedding maps happens once in the beginning.

The weakly-symplectic AE consists of two independent constrained AEs for position and momentum. We use 4 layers of size $n_l = \{32, 64, 128, 600\}$ with varying latent space dimension. We train the network jointly on the sum of the losses (12) and (42) via Riemannian Adam (Becigneul & Ganea, 2019) until convergence at 3000 epochs with a learning rate of 1.5×10^{-2} .

for the AE parameters and 7×10^{-4} for the HNN parameters.

Non-symplectic AE baseline. In App. I.3 (see Fig. 20), we ablate geometrically-constrained symplectic AE of the RO-HNN in a dissipative scenario. To do so, we train a RO-HNN that utilizes a non-symplectic projection-constrained AE instead of geometrically-constrained symplectic AE. Specifically, we use a single vanilla constrained AE with latent space dimension $n_0 = 20$ with 4 pairwise biorthogonal encoder-decoder layers of sizes $n_l = \{64, 128, 256, 1200\}$. Note that we assume the first 10 output dimensions of the latent space to correspond to reduced position $\tilde{q}$ and the last 10 to correspond to the reduced momentum $\tilde{p}$, which are used as inputs for the latent dissipative HNN.

Comparison against RO-LNN. In Sec. 4.3 and App. I.3, we compare the RO-HNNs against RO-LNNs (Friedl et al., 2025). The RO-LNNs are trained on the dataset $\mathcal{D}_{\text{cloth, vel}} = \{\{\mathbf{q}_{n,k}, \dot{\mathbf{q}}_{n,k}, \boldsymbol{\tau}_{n,k}\}_{k=1}^K\}_{n=1}^N$ obtained from the dataset described in App. H.3.1 by transforming the momentum data into velocities via $\dot{\mathbf{q}} = \mathbf{M}(\mathbf{q})^{-1}\mathbf{p}$.

We construct the RO-LNN following the procedure described in (Friedl et al., 2025) and use the same architecture for the constrained AE as for the RO-HNN. Specifically, we consider a latent space of dimension $n_0 = 10$ and use $l = \{1, 2, 3, 4\}$ pairwise biorthogonal encoder and decoder layers of sizes $n_l = \{32, 64, 128, 600\}$. The kinetic and potential energy networks of the latent geometric LNN consist of 2 hidden Euclidean layers with 64 neurons and SoftPlus activation functions, initialized as for the RO-HNN. Notice that, for the RO-LNN, the dissipation forces $\boldsymbol{\tau}_d$ are not learned but provided as ground truth in the external input $\boldsymbol{\tau} = \boldsymbol{\tau}_c + \boldsymbol{\tau}_d$.

The RO-LNN is trained on 3000 samples from $\mathcal{D}_{\text{cloth, vel}}$ with a Runge-Kutta integrator of order 4 over a training horizon $H_{\mathcal{D}} = 8$ timesteps. We train the RO-HNN with Riemannian Adam (Becigneul & Ganeva, 2019) until convergence at 3000 epochs with a learning rate of 5×10^{-2} for the AE parameters, 2×10^{-4} for the HNN parameters, and a regularization $\gamma = 2 \times 10^{-5}$ for 3000 epochs.

In our comparison, we also consider a black-box version of the RO-LNN, hereinafter referred to as black-box RO-LNN, where the latent LNN is encoded as a single black-box network $\mathcal{L}_{\theta}$ representing the Lagrangian function, which we model via a single fully-connected MLP of 2 hidden layers with a width of 256 neurons. The network weights are initialized by sampling from a Xavier normal distribution with gain $\sqrt{2}$, and the bias vector entries are initialized to 1. The black-box RO-LNN is trained with the same parameters as the original RO-LNN, except for the learning rate for the LNN parameters, which is set as of 2×10^{-3} .

Comparison against HNKO. For the prediction errors reported in Tab. 4, we train an HNKO baseline in a $d = 80$ -dimensional latent space using a fully connected neural network with 6 hidden layers and Tanh activations. The Hamiltonian latent dynamics are modeled on a 40-dimensional sphere. State reconstruction is performed by a fully-connected neural network with 3 hidden layers and Tanh activations. We train the network on 3000 training samples from $\mathcal{D}_{\text{cloth}}$, using Adam for 15000 epochs on a horizon of $H_{\mathcal{D}} = 8$ timesteps, adjusting the original loss to multi-step prediction as described in H.1.2.

I. Additional Experimental Results

This section presents additional results, complementing those presented in Sec. 4.

I.1. Coupled Pendulum of Section 4.1

This section presents additional results on learning the Hamiltonian dynamics of a 15-DoF coupled pendulum.

Learning high-dimensional dynamics. Fig. 7 complements Fig. 2 by depicting the predicted long-term (5s) positions and momenta. For the ease of visualization, we change the prediction coordinates and plot the first 3-DoF corresponding to the latent pendulum. We observe that RO-HNN leads to accurate long-term predictions similar to those of the 3-DoF HNN. Unlike the 15-dimensional full-order HNN, the HNKO yields stable, but inaccurate long-term predictions, exhibiting significantly higher deviation from the ground truth trajectory than the predictions of the RO-HNN.

Fig. 8 depicts the original 15-DoF trajectories projected into the 3-DoF latent space learned by the RO-HNN, along with its latent dynamic predictions over the full prediction horizon of 5s. We observe that the latent predictions are accurate and match the projected original trajectories. As expected, the learned latent space does not coincide with the phase-space of the original pendulum due to the nonlinear dimensionality reduction conducted via the AE, but displays comparable frequencies and amplitudes.




Figure 7. Reconstructed trajectories of the RO-HNN (—), 3-DoF HNN (—), and HNKO (—) compared to ground truth (—). The 15-DoF HNN leads to unstable long-term predictions and is not depicted.




Figure 8. Trajectories of the original 15-DoF pendulum projected into the latent space of the RO-HNN (—), and corresponding dynamic predictions obtained via the latent HNN (—). As expected, they do not coincide directly with the trajectories of the underlying 3-DoF pendulum representation (—).




Figure 9. Prediction errors ($\downarrow$) of intrusive symplectic dimensionality reduction approaches over 10 test pendulum trajectories.




Figure 10. *Left*: Ablation of the latent HNN architecture on a doubled training set size $|\mathcal{D}| = 6000$ compared to Fig. 10-left. *Middle, right*: Ablation of the latent integrator of the geometric RO-HNN at $|\mathcal{D}| = 3000$ for learning the dynamics of a 15-DoF pendulum. Errors are obtained via short-term prediction horizons $H\Delta t = 0.25$ s.

AE architecture. Figure 9 accompanies and validates the results of Table 2 by displaying the median and quartiles of the prediction errors obtained by different symplectic dimensionality reduction methods in the intrusive MOR scenario of Sec. 4.1.

Latent HNN architecture. Here we further evaluate the impact of HNN architecture. We compare the performance of our geometric HNN to learn the low-dimensional dynamics of the latent 3-DoF pendulum against (1) a non-geometric variant that parametrizes the inverse mass-inertia matrix via a Cholesky network, and (2) two HNNs encoded as a single black-box network $\mathcal{H}_\theta$, where we consider two MLPs of 64- and 256-neurons width. Compared to Sec. 4.1, we consider a doubled amount of training datapoints with 6000 random samples. As shown in Fig. 10-left, the geometric HNN still achieves the lowest reconstruction error, with differences compared to the black-box HNN increased compared to the smaller dataset of Fig. 3-left. This showcases the importance of considering both the quadratic energy structure of mechanical systems, and the geometry of their mass-inertia matrices, for both enhanced performance and data efficiency.

Latent integrator. We compare the Strang symplectic integrator against (1) a symplectic leapfrog integrator that disregards that the Hamiltonian is non-separable, (2) a Runge-Kutta integrator of order 4 that overlooks its symplectic structure, and (3) an explicit Euler integrator that also overlooks the symplectic structure. Compared to Sec. 4.1 (see Fig. 3-middle,right), we consider shorter prediction horizons, feeding the model with ground truth initial conditions every $H\Delta t = 0.25$ s, since the explicit Euler integrator did not lead to stable long-term predictions for $H\Delta t = 5$ s. Figs. 10-middle, right show that the networks trained via the Strang-symplectic integrator achieve the lowest reconstruction error and conserves energy best during integration, showcasing the importance of considering the symplectic structure of the system during numerical integration for stable predictions on short- and long-term time horizons.

Training under noisy observations. To assess the robustness of the RO-HNNs, we evaluate its performance under noisy observations and compare it against the HNKOs baseline, which is reported to be robust to noise in high-dimensional systems.

We generate noisy training data $\{\mathbf{q}_i + \boldsymbol{\epsilon}_{q,i}, \mathbf{p}_i + \boldsymbol{\epsilon}_{p,i}\}$ corrupted with zero-mean Gaussian noise $\boldsymbol{\epsilon}_{q,i} \sim \mathcal{N}(\mathbf{0}, \sigma_q^2 \mathbf{I})$ and $\boldsymbol{\epsilon}_{p,i} \sim \mathcal{N}(\mathbf{0}, \sigma_p^2 \mathbf{I})$. The noise level c_{noise} determines the standard deviations, which is also proportional to the maximum entry of the position and momentum, i.e., $\sigma_q = c_{\text{noise}} \max_{j,k} |q_{j,k}|$ and $\sigma_p = c_{\text{noise}} \max_{j,k} |p_{j,k}|$.

Table 6 reports the prediction errors on a testing dataset of 10 noise-free trajectories over time horizons $H\Delta t = \{0.25, 5\}$ s

Table 6. Mean and standard deviation of prediction errors ($\downarrow$) over $N = 10$ noise-free test pendulum trajectories, comparing the performances of the RO-HNNs and HNKOs trained on noisy observations.

	$H\Delta t$ (s)	$c_{\text{noise}} = 0$		$c_{\text{noise}} = 5\%$		$c_{\text{noise}} = 10\%$	
		RO-HNN	HNKO	RO-HNN	HNKO	RO-HNN	HNKO
$\frac{\ \hat{q}_t - q\ }{\ q\ }$	0.25	$(1.66 \pm 1.38) \times 10^{-1}$	$(5.64 \pm 4.41) \times 10^{-1}$	$(2.44 \pm 1.94) \times 10^{-1}$	$(9.76 \pm 10.48) \times 10^{-1}$	$(3.23 \pm 2.86) \times 10^{-1}$	$(1.02 \pm 0.78) \times 10^0$
	5	$(7.08 \pm 7.56) \times 10^{-1}$	$(1.32 \pm 0.94) \times 10^0$	$(8.40 \pm 8.17) \times 10^{-1}$	$(1.98 \pm 1.84) \times 10^0$	$(9.14 \pm 10.02) \times 10^{-1}$	$(4.06 \pm 3.47) \times 10^0$
$\frac{\ \hat{p}_t - p\ }{\ p\ }$	0.25	$(5.33 \pm 5.23) \times 10^{-2}$	$(5.93 \pm 10.73) \times 10^{-1}$	$(7.02 \pm 7.17) \times 10^{-2}$	$(4.56 \pm 8.34) \times 10^{-1}$	$(9.64 \pm 10.30) \times 10^{-2}$	$(9.60 \pm 17.23) \times 10^{-1}$
	5	$(1.98 \pm 2.67) \times 10^{-1}$	$(1.23 \pm 2.08) \times 10^0$	$(3.34 \pm 4.97) \times 10^{-1}$	$(1.62 \pm 2.97) \times 10^0$	$(3.61 \pm 4.94) \times 10^{-1}$	$(3.13 \pm 5.58) \times 10^0$

Table 7. Mean and standard deviation of testing loss ($\downarrow$) over different training set size $\mathcal{D}$. Errors are computed over $N = 10$ test trajectories and 10 different RO-HNNs that started training from different random seeds.

$ \mathcal{D} $	1000	3000	6000
$\ell_{\text{RO-HNN}}$	$(9.84 \pm 4.64) \times 10^{-3}$	$(1.73 \pm 0.33) \times 10^{-3}$	$(1.67 \pm 0.21) \times 10^{-3}$

for three noise level $c_{\text{noise}} = \{0, 0.05, 0.1\}$. The noiseless results are repeated from Table 1 for completeness. As expected, the performance of both models decreases with increasing noise magnitude. In each scenario, the RO-HNN outperforms the HNKO baseline. Note that the RO-HNN trained at $c_{\text{noise}} = 0.1$ outperforms the HNKO trained without noise, demonstrating the enhanced accuracy and robustness of the RO-HNN to noisy observations.

Number of training samples. Table 7 assesses the influence of the number of available training samples on the RO-HNN convergence. It reports the mean and standard deviation of the loss (12) across $N = 10$ testing trajectories and 10 individual RO-HNNs trained from different random seeds. We observe decreasing prediction errors with increasing number of training samples, before a stabilization around a high number of training data.

I.2. Particle Vortex (90-DoF) of Section 4.2

This section presents additional results on learning the Hamiltonian dynamics of a 90-DoF particle vortex.

Fig. 11 depicts the predicted positions and momenta of the particles along with the ground truth in the high-dimensional state space for RO-HNNs with latent dimension $d = \{6, 10\}$. Fig. 12 depicts the predicted positions and momenta of the particles in the reduced phase space of the AE along with the projected ground truth. We observe that both models accurately predict the particle vortex dynamics, with the $d = 6$ -dimensional model slightly outperforming the 10-dimensional one (see also Table 3 and Fig. 13). This shows that the choice of latent dimension is a trade off between the latent space expressivity and the limitations of HNNs in higher dimensions. In general, we observed that errors initially decrease as the latent dimension increases, suggesting that higher-dimensional latent spaces better capture the original high-dimensional dynamics. The errors then increase beyond a certain latent dimension, indicating that the latent HNN becomes harder to train.

Fig. 13 accompanies and validates the results of Table 3 by displaying the median and quartiles of the prediction errors obtained by different latent HNNs in combination with the symplectic AE.




Figure 11. Predicted ($\bullet, \bullet, \bullet$) vs ground truth ($\bullet, \bullet, \bullet$) positions of the particle vortex. The dynamics are learned with RO-HNN with $d = 6$ and $d = 10$. Times beyond 10s are out of the training data distribution.




Figure 12. Predicted (•, •, •) vs ground truth (•, •, •) reduced positions of the particle vortex in the latent space of the RO-HNN with $d = 6$ and $d = 10$. Times beyond 10s are out of the training data distribution.




Figure 13. RO-HNN prediction errors ($\downarrow$) for black-box and geometric HNNs in latent spaces learned by a symplectic AE with latent dimensions $d = 6$ (■) and $d = 10$ (■) over 10 particle vortex trajectories.

I.3. Cloth (600-DoF) of Section 4.3

This section presents additional results on learning the Hamiltonian dynamics of a 600-DoF thin cloth falling on a sphere.

Learning high-dimensional dynamics with dissipation. As shown in Table 4, the 10-dimensional model slightly outperforms the 6-dimensional one, modeling more details of the cloth, as observed in Fig. 5.

Fig. 14 accompanies Table 4 by visualizing the median and quartiles of the RO-HNN ($d = 10$) reconstructed and latent prediction errors over time for two different parametrization of the reduced dissipation matrix $\tilde{D}$. Both dissipative RO-HNN perform similarly to the conservative RO-HNN, showing that the RO-HNN can successfully predict dissipative dynamics in a stable manner, including beyond the training time horizon.

Fig. 15 shows the predictions of the RO-HNNs with different parametrizations of the dissipation matrix $\tilde{D}$ for selected dimensions of a test trajectory. This shows that the dissipative RO-HNNs successfully learn the dissipation forces, achieving similar prediction errors as the conservative models Fig. 16 displays the predicted latent energy to be compared with the ground-truth energy projected in the symplectic latent space. Overall, our results demonstrate the ability of the RO-HNN to infer long-term predictions of dissipative systems.

Comparison against RO-LNN. Fig. 17 compares the latent prediction and reconstructed prediction errors of the geometric and black-box RO-HNN and RO-LNN over time.

Our results show that the geometric RO-HNN outperforms the RO-LNN, leading to more accurate predictions even as the




Figure 14. Median and quartiles of the latent and reconstructed prediction errors of 10-dimensional RO-HNNs with latent dissipation matrix parametrized with a SPD network (—), a Cholesky network (—), and ground truth values (—). The gray-shaded area indicates the time horizon beyond the training data support.




Figure 15. Predicted cloth positions and momenta for 6-dimensional RO-HNNs with latent dissipation matrix parametrized with a SPD network (—), a Cholesky network (---), and ground truth values (····), and 10-dimensional RO-HNNs with latent dissipation matrix parametrized with a SPD network (—), a Cholesky network (---), and ground truth values (····). The grey-shaded areas indicates interval beyond the data support.




Figure 16. Top: Ground truth (—) and predicted latent energies for 6-dimensional RO-HNNs with latent dissipation matrix parametrized with a SPD network (—), a Cholesky network (---), and ground truth values (····), and 10-dimensional RO-HNNs with latent dissipation matrix parametrized with a SPD network (—), a Cholesky network (---), and ground truth values (····). Bottom: Energy errors for the same models. The grey-shaded areas indicate intervals beyond the data support, for which the ground truth is extrapolated from the last observation.




Figure 17. Median and quartiles of the latent and reconstructed prediction errors of 10-dimensional geometric RO-HNN (—), geometric RO-LNN (—), a RO-HNN with latent black-box HNN (····), and a RO-LNN with latent black-box LNN (····).




Figure 18. Mean and standard deviation of the relative reconstruction (left), latent prediction (middle), and reconstructed prediction (right) errors over 10 cloth trajectories with $H\Delta t = 0.0025$ s. Our RO-HNN with geometrically-constrained symplectic AE (—) is compared against linear SMG reduction (—), quadratic SMG reduction (—), a weakly symplectic AE (—), and a sequentially-trained RO-HNN with pretrained geometrically-constrained symplectic AE (—). The pretrained AE (—) is depicted for completeness. Notice that the linear SMG and quadratic SMG projections led to diverging dynamics for $d > 2$ and $d > 6$, respectively, for which results are not depicted.

RO-HNN learns the dissipation forces via the latent damping matrix, whereas the RO-LNN receives their ground truth measurements. Moreover, Fig. 17 shows that the geometric RO-HNN and RO-LNN featuring geometric latent HNN and LNN outperform their black-box counterparts, showcasing the importance of considering the quadratic energy structure of mechanical systems in both network types.

We hypothesize that the improved accuracy of the RO-HNN compared to the RO-LNN can be attributed to (1) the first-order dynamic formulation stemming from Hamiltonian mechanics, which is easier to learn and optimize than the second-order Lagrangian formulation, (2) the Strang-symplectic integrator which is specifically designed for Hamiltonian systems, in contrast to the Runge Kutta integrators typically used in the case of continuous-time Lagrangians. This aligns with the discussions in (Liu et al., 2024a), which showed that, by using position and momentum observations, HNNs learn mass-inertia matrices that are close to the physical solutions, while LNNs only learn one of the solutions satisfying the Euler-Lagrange equations.

Latent dimension and training ablation. We compare the performance of our dissipative RO-HNN across several latent dimensions $d = \{2, 6, 10\}$ with jointly-trained geometrically-constrained symplectic AE and latent geometric HNN against sequentially-trained architectures. Specifically, we consider (1) linear and (2) quadratic symplectic manifold Galerkin (SMG) projections (Peng & Mohseni, 2016; Sharma et al., 2023), (3) a weakly-symplectic AE trained jointly with a latent geometric HNN, and (4) a RO-HNN with pretrained geometrically-constrained AE. Fig. 18 shows that our jointly-trained RO-HNN significantly outperforms all baselines for all dimensions, leading to reduced relative reconstruction, latent prediction, and reconstructed prediction errors. This showcases (1) the higher expressivity of the AEs compared to linear and quadratic projection methods, (2) the importance of structurally-embedding the symplecticity condition, unlike the weakly-symplectic AE, and (3) the importance of joint training, allowing the RO-HNN to jointly learn a symplectic submanifold and the associated dynamics.

Finally, we compare the performance of the dissipative RO-HNN against (1) a conservative RO-HNN, where the dissipation forces τ_d are not learned but provided as ground truth in the external input $\tau = \tau_c + \tau_d$, and (2) a dissipative RO-HNN where the dissipation matrix is parametrized via Cholesky decomposition for latent dimensions $d = \{2, 6, 10\}$. The mass-inertia matrix is parametrized via SPD networks in all cases. Fig. 19 shows the obtained latent prediction and reconstructed prediction errors. Both dissipative HNNs achieve errors close to the conservative HNN where the ground truth dissipative forces are provided, with the geometric HNN slightly outperforming its Cholesky counterpart. However, the effect is less pronounced as when learning the inverse mass-inertia matrix, which we attribute to the reduced influence of damping compared to inertia in the overall dynamics.




Figure 19. Mean and standard deviation of the latent prediction (*left*) and reconstructed prediction (*right*) errors for different parametrization of the latent dissipation matrix $\tilde{D}$ over 10 test cloth trajectories. We compare our SPD network (—) against a Cholesky network (—), and the ground truth parametrization (—).




Figure 20. Prediction errors ($\downarrow$) of RO-HNNs with geometrically-constrained symplectic or vanilla constrained AE with Strang symplectic integrator (■) or Runge-Kutta integrator of order 4 (●) over 10 testing cloth trajectories.

Ablation of the symplectic architecture to learn dissipative dynamics. As discussed in Sec. 3.2, the dissipative dynamics do not preserve a symplectic structure. Our proposed RO-HNN features two main symplecticity-preserving components, namely the geometrically-constrained symplectic AE and the Strang-symplectic integrator, which we ablate here. We consider two variations of each components, i.e., (1) our geometrically-constrained symplectic AE and a vanilla projection-constrained AE, and (2) the Strang-symplectic integrator and a Runge-Kutta integrator of order 4. Throughout all experiments, we set the latent dimension $d = 10$. Fig. 20 shows the obtained reconstructed prediction errors. We observe that our geometrically-constrained symplectic AE leads to significantly lower median errors than vanilla projection-constrained AE independently of the choice of integrator, showcasing the benefit of preserving the structure of FOM vector field in the ROM (see Proposition 3.3). Moreover, despite the dissipative structure, the RO-HNN obtained with the Strang-symplectic integrator outperforms the non-symplectic Runge-Kutta integrator. We hypothesize that this is due to the fact that the evolution of this dissipative system is mostly governed by its Hamiltonian function, especially over the short timesteps taken by the integrators.

J. Runtimes

This section compares the runtimes of different models. All experiments were performed locally on a MacBook Pro with M3 CPU.

Speedup of simulation via the RO-HNN. We aim at providing an idea of the computational effort of the RO-HNN compared to the evaluation of the FOMs. To do so, we symbolically derive the Hamiltonian equations of motion with known physical quantities of the 15-DoF coupled pendulum of Sec. 4.1 and obtain the equations of motion of the 600-DoF cloth from Sec. 4.3 from Mujoco. We compare the wall-clock time of the evaluation of these two FOMs against the respective RO-HNN averaged over 10 single trajectory roll-outs from the same initial conditions as in the testing dataset. We consider two different integrators, namely an Euler forward and the Strang symplectic integrator, and different step size Δt on a time-horizon of $H\Delta t = 5s$ for the pendulum and $H\Delta t = 0.5s$ for the cloth. The evaluation times for the coupled pendulum and the cloth are given in Table 8, with the corresponding relative position prediction errors depicted in Fig. 21. We observe a prominent reduction of the evaluation time for the RO-HNNs compared to the FOMs. This reduction is exacerbated for higher-dimensional systems, e.g., the cloth, where the evaluation time of the RO-HNN remains significantly lower than that of the FOM evaluated with the computationally-cheaper forward Euler integrator or increased step size. Moreover, as shown in Fig. 21, the RO-HNNs exhibit lower prediction errors than their FOM counterparts in addition to a reduced computational

Table 8. Evaluation wall clock times for different ODE-solvers on analytic FOM compared to RO-HNN. Runtimes are averaged over 10 forward passes and given in s.

	pendulum ($n = 15, d = 3$)			cloth ($n = 600, d = 10$)		
	$\Delta t = 10^{-2}$ s	$\Delta t = 10^{-1}$ s		$\Delta t = 10^{-4}$ s	$\Delta t = 10^{-3}$ s	
	Strang	Euler	Strang	Strang	Euler	Strang
FOM, n -DoF	1.29	0.75	0.18	255.24	77.62	109.74
RO-HNN, d -DoF	0.79	—	—	16.01	—	—




Figure 21. Position prediction errors introduced by various simulation speedup methods: Euler-forward integration (—), Strang integration (—) at larger stepsize, and RO-HNN (—).

complexity. This showcases that the RO-HNN not only enable accurate learning of unknown high-dimensional Hamiltonian dynamics, but also the computationally-efficient and accurate evaluation of known systems via surrogate dynamics.

Comparison of runtimes of HNNs and RO-HNNs. Table 9 reports the averaged runtimes for the forward pass of the differently-sized network architectures considered in Sec. 4.1. The reported times correspond to the wall clock time of one forward pass of a batch of 10 initial conditions, predicted over $H = 10$ timesteps with the Strang-symplectic integrator. We observe that the RO-HNN speeds up the forward dynamics computation compared to the HNN, highlighting the computational advantages of ROMs compared to FOMs. Moreover, the black-box HNN is computationally more efficient than the geometric HNN at the expense of prediction accuracy.

Runtimes of the lifted AE. Table 10 reports the average wall-clock times for one forward pass of a batch of 100 states for several projection scenarios on the 600-DoF cloth dataset. We consider reduction to a latent space of dimension $d = 10$ via 4 layers of size $n_l = \{32, 64, 128, 600\}$. We consider (1) a position-level constrained AE, where only position projections $\tilde{q} = \varphi_Q \circ \rho_Q(q)$ are computed via the encoder and decoder layers, (2) a geometric symplectic AE, whose lifted mappings (9) are computed analytically as a composition of layer derivatives (see App. E) to project both positions and momenta via $(\tilde{q}, \tilde{p}) = \varphi \circ \rho(q, p)$, (3) a geometric symplectic AE whose lifted maps are computed via automatic differentiation using Pytorch’s `autograd.vjp` function, and (4) a naive constrained AE that jointly projects the positions and momenta in $2d = 20$ -dimensional latent space via doubled layer dimensions, i.e., $n_l = \{64, 128, 256, 1200\}$. The first two variants were evaluated under `torch.no_grad()`, reflecting a realistic scenario for evaluation of forward dynamics in the RO-HNN.

The runtimes reported in Table 10 show that the analytic computation of the lifted mappings is significantly faster than the automatic-differentiation-based implementation. This is expected, as our analytic implementation avoids the construction of a backward graph. It is worth emphasizing that the geometric AE with analytic lifts requires significantly less than twice the runtime of the position-level constrained AE. Therefore, using a single cotangent-lifted AE that jointly projects positions and momenta is computationally more advantageous than training two separate AEs for separate projections.

Table 9. Evaluation wall clock times for different network architectures on the 15-DoF pendulum. Runtimes are averaged over 10 forward passes and given in ms.

15-DoF		3-DoF	
Geometric HNN	Geometric RO-HNN	Geometric HNN	Black-box HNN
100.25	26.34	18.10	8.04

Table 10. Evaluation wall clock times of different variants for the AE on the 600-DoF cloth dataset. Runtimes are averaged over 10 forward passes and given in ms.

Position-level AE $\varphi_Q \circ \rho_Q(q)$	Geometric AE with analytic lift $\varphi \circ \rho(q, p)$	Geometric AE with <code>autodiff.vjp</code> lift $\varphi \circ \rho(q, p)$	AE with naive lift $\varphi_Q \circ \rho_Q(q, p)$
16.92	20.93	54.24	19.03