

Port-Hamiltonian neural networks for learning explicit time-dependent dynamical systemsShaan A. Desai ^{1,*} Marios Mattheakis ² David Sondak,² Pavlos Protopapas,² and Stephen J. Roberts ¹¹*Machine Learning Research Group, University of Oxford Eagle House, Oxford OX26ED, United Kingdom*²*John A. Paulson School of Engineering and Applied Sciences, Harvard University Cambridge, Massachusetts 02138, USA*

(Received 8 June 2021; accepted 14 September 2021; published 29 September 2021)

Accurately learning the temporal behavior of dynamical systems requires models with well-chosen learning biases. Recent innovations embed the Hamiltonian and Lagrangian formalisms into neural networks and demonstrate a significant improvement over other approaches in predicting trajectories of physical systems. These methods generally tackle autonomous systems that depend implicitly on time or systems for which a control signal is known *a priori*. Despite this success, many real world dynamical systems are nonautonomous, driven by time-dependent forces and experience energy dissipation. In this study, we address the challenge of learning from such nonautonomous systems by embedding the port-Hamiltonian formalism into neural networks, a versatile framework that can capture energy dissipation and time-dependent control forces. We show that the proposed port-Hamiltonian neural network can efficiently learn the dynamics of nonlinear physical systems of practical interest and accurately recover the underlying stationary Hamiltonian, time-dependent force, and dissipative coefficient. A promising outcome of our network is its ability to learn and predict chaotic systems such as the Duffing equation, for which the trajectories are typically hard to learn.

DOI: [10.1103/PhysRevE.104.034312](https://doi.org/10.1103/PhysRevE.104.034312)**I. INTRODUCTION**

Neural networks (NNs), as universal function approximators [1], have shown resounding success across a host of domains including image segmentation [2], machine translation [3], and material property predictions [4,5]. However, their performance in learning and generalizing the long-term behavior of dynamic systems governed by known physical laws from state data has often been limited [6,7]. New research in scientific machine learning, a field that tackles scientific problems with domain-specific machine learning methods, is paving a way to address this challenge. Concretely, it has been shown that physically informed learning biases embedded in networks, such as Hamiltonian mechanics [6,8] and Lagrangians [9,10], Ordinary differential equations (ODEs) [11], physics-informed networks [12,13], generative networks [14], and graph neural networks [15,16] can significantly improve learning and generalization over vanilla neural networks in complex physical domains. The performance uplift arises primarily because learning biases are able to constrain networks to learn physically meaningful representations from data that are crucial to generalization.

Despite extensive research on learning biases, there is yet no method that accounts for nonautonomous systems, i.e., systems with explicit time dependence. Nonautonomous dynamics feature prominently in settings with externally driven or controlled time-dependent forces as well as in systems with energy dissipation, for example, interacting materials, forced

oscillators and charge-discharge cycles. Defining a network to accurately learn and predict the dynamics of such systems from position and momentum data is therefore of critical practical interest. We address this challenge by embedding the port-Hamiltonian formalism [17–21] into neural networks. We show that the structure of this formulation can be used to uncover the underlying Hamiltonian, force, and damping terms given position and momentum data and as such, can be used to accurately predict the long-range trajectories of many forced or damped systems. We extensively benchmark our network on a range of tasks including a simple mass-spring system with damping and external force, a Duffing system in both the nonchaotic and chaotic regimes, and a relativistic Duffing system. Our proposed network consistently outperforms other approaches while accurately recovering both the driving force and the damping coefficient. Furthermore, using minimal data, we show that our network can visually recover the Poincaré section of the Duffing system in a chaotic regime, emphasizing how our network can be used to identify and understand chaotic trajectories.

II. BACKGROUND**A. Hamiltonian neural networks**

Recently, the authors of Ref. [6] demonstrated that the dynamics of an energy conserving autonomous system can be accurately learned by guiding a neural network to predict a Hamiltonian, an important representation of a dynamical system. Considering a dynamical system of M objects, the Hamiltonian \mathcal{H} is a scalar function of a position vector $\mathbf{q}(t) = [q_1(t), q_2(t), \dots, q_M(t)]$ and momentum vector $\mathbf{p}(t) = [p_1(t), p_2(t), \dots, p_M(t)]$ that obeys Hamilton's

*Also at John A. Paulson School of Engineering and Applied Sciences, Harvard University, Cambridge, Massachusetts, USA; shaan@robots.ox.ac.uk

equations,

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

where $\dot{\mathbf{q}} = \frac{d\mathbf{q}}{dt}$ and $\dot{\mathbf{p}} = \frac{d\mathbf{p}}{dt}$.

Using the Hamiltonian formalism, Ref. [6] showed that a NN with parameters θ can be used to learn a Hamiltonian $\mathcal{H}_\theta(\mathbf{q}, \mathbf{p})$ given \mathbf{q} and \mathbf{p} as inputs to the network. The time derivatives are recovered from Eq. (1) by differentiating \mathcal{H}_θ with respect to its inputs using automatic differentiation. The resulting system has the form,

$$\dot{\mathbf{x}} = s(\mathbf{x}), \quad (2)$$

where $\mathbf{x} = (\mathbf{q}, \mathbf{p})$ and s is determined by Eq. (1) after differentiating the trained NN. Equation (2) can be discretized using a time integrator to determine the trajectory of an initial state. Moreover, since s is the symplectic gradient of the Hamiltonian, energy conservation is embedded into the method by construction. Given these advantages, Hamiltonian neural networks (HNNs) outperform traditional approaches that directly predict the state time derivatives from the input state. However, this formulation does not readily generalize to damped or forced time-varying systems.

B. Port-Hamiltonian framework

The port-Hamiltonian [17–21] is a well-studied formalism that generalizes Hamilton’s equations to incorporate energy dissipation and an external control input to a dynamical system. Hamilton’s equations in the port-Hamiltonian framework are represented as:

$$\begin{bmatrix} \dot{\mathbf{q}} \\ \dot{\mathbf{p}} \end{bmatrix} = \left(\begin{bmatrix} \mathbf{0} & \mathbf{I} \\ -\mathbf{I} & \mathbf{0} \end{bmatrix} + \mathbf{D}(\mathbf{q}) \right) \begin{bmatrix} \frac{\partial \mathcal{H}}{\partial \mathbf{q}} \\ \frac{\partial \mathcal{H}}{\partial \mathbf{p}} \end{bmatrix} + \begin{bmatrix} \mathbf{0} \\ \mathbf{G}(\mathbf{q}) \end{bmatrix} \mathbf{u}, \quad (3)$$

where $\mathbf{D}(\mathbf{q})$ is a damping matrix, \mathbf{u} is a temporal control input, $\mathbf{G}(\mathbf{q})$ is a nonlinear scaling of the position vector, \mathbf{I} is the identity matrix, and $\mathbf{0}$ the zero matrix. The damping matrix is semipositive definite. This general formalism readily reduces to the standard Hamiltonian system when $\mathbf{D} = \mathbf{0}$ and $\mathbf{u} = \mathbf{0}$. The port Hamiltonian has been used in control applications where explicit knowledge of the control term \mathbf{u} is known and was recently shown to reveal promising results in NNs [22]. Note that in Ref. [22], the vector $[\mathbf{q}, \mathbf{p}, \mathbf{u}]$ is provided as input to the network. However, in many applications the control force \mathbf{u} is unknown and it is therefore of interest to uncover the underlying forcing term from the data of the state vector, where no explicit knowledge of the control input and damping term is available.

C. Related work

While Hamiltonian mechanics presents one way to address learning dynamical systems, numerous recent methods highlight how incorporating other physically informed inductive biases into neural networks can improve learning. Functional priors, for example, embed the full functional form of an equation into the NN. Physics-informed neural networks (PINNs) [12] and Hamiltonian networks [8] are two such approaches that look at directly embedding the equations of motion into the loss function. While PINNs are data-driven

approaches that rely on AUTOGRAD [23] to compute partial derivatives of a hidden state, Hamiltonian networks are data-independent approaches that prespecify the full functional form of a system-specific Hamiltonian in the loss function.

Many recent methods have also sought to embed integrators into the training process. Indeed doing so induces an effectively continuous depth neural network able to perform large time-step predictions. NeuralODE [11] presents one way to tackle back-propagating through this continuous depth network more efficiently.

More recent work has looked at generalizing this approach to different and more complex data structures and topologies that standard NeuralODEs cannot represent [24]. Other work, such as Zhu *et al.* [25], theoretically shows the importance of using symplectic integrators over Runge-Kutta methods to evolve Hamiltonian systems in NNs.

In Ref. [15], the authors detailed how a graph neural network, designed to capture the relational structure of systems, can be used to learn dynamics of interacting particles. This work has been exploited in numerous advances [9,26,27] and emphasizes how a relational inductive bias can significantly improve learning.

In Refs. [9] and [6] it has been shown that by learning the Hamiltonian or the Lagrangian of a system, it is possible to accurately predict the temporal dynamics and conserve energy. The work of Ref. [10] also showed that by exploiting the Euler-Lagrange equation, it is possible to predict a controlled double pendulum, a system pertinent for controlled robots. A recent advance shows that Hamiltonian and Lagrangian NNs can be drastically improved if they are optimized over Cartesian coordinates with holonomic constraints [28].

Despite the significant breakthroughs, there is no existing method that investigates explicit time dependence and damping in dynamical systems, two elements that are often found in real world problems. As such, we outline a technique to address this challenge.

III. METHOD

A. Theory

In this section we introduce port-Hamiltonian neural networks (pHNNs). We begin by illustrating how a NN takes the form of the port-Hamiltonian formulation of Eq. (3) with two modifications. First, our approach exploits the fact that many damped systems consist of a nonzero, state-independent damping term in the bottom right quadrant, so we replace the damping matrix $\mathbf{D}(\mathbf{q})$ with a state-independent matrix for which only the bottom right term is nonzero and represented by \mathbf{N} . Second, in order to generalize to time-dependent forcing, we replace $\mathbf{G}(\mathbf{q})\mathbf{u}$ with the force field $\mathbf{F}(t)$. The resulting representation,

$$\begin{bmatrix} \dot{\mathbf{q}} \\ \dot{\mathbf{p}} \end{bmatrix} = \left(\begin{bmatrix} \mathbf{0} & \mathbf{I} \\ -\mathbf{I} & \mathbf{0} \end{bmatrix} + \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{N} \end{bmatrix} \right) \begin{bmatrix} \frac{\partial \mathcal{H}}{\partial \mathbf{q}} \\ \frac{\partial \mathcal{H}}{\partial \mathbf{p}} \end{bmatrix} + \begin{bmatrix} \mathbf{0} \\ \mathbf{F}(t) \end{bmatrix}, \quad (4)$$

is general enough to handle many well-known forced systems but also specific enough to tackle learning in physical domains of practical importance such as Duffing equation.

The architecture for the proposed pHNN model and a summary of comparable existing approaches are shown in

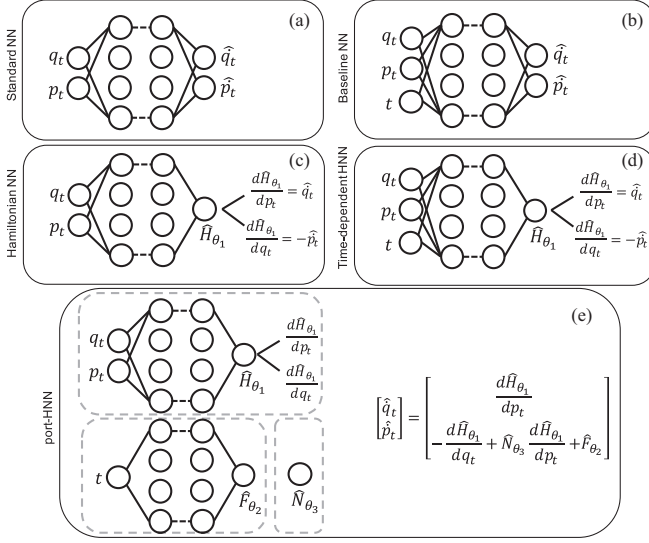


FIG. 1. Architectures used in this study to learn dynamical systems. The naive extension of a standard feed forward NN [outlined in (a)] to incorporate time as an additional variable is shown in (b) and considered as the baseline network. The standard HNN in (c) is extended to receive time as an input and demonstrated by (d). Our innovation is presented in (e), which exploits port Hamiltonians and explicitly learns the force F_{θ_2} , the damping term N_{θ_3} and the Hamiltonian \mathcal{H}_{θ_1} to predict the state-time derivatives.

Fig. 1. A standard feed forward NN in Fig. 1(a) and an HNN in Fig. 1(c) [6] take q and p as inputs and are trained to yield the time derivatives of the input, with the HNN learning an intermediate Hamiltonian and employing back propagation to compute the final output. A natural way to extend these architectures for time-varying nonautonomous systems is to include time as an additional input. This gives rise to the baseline network (baseline NN) represented by Fig. 1(b) and a time-dependent HNN (TDHNN) shown by Fig. 1(d).

Although the baseline NN and TDHNN incorporate time, they do not provide information for the underlying dynamics of the system. On the other hand, the pHNN is able to extract and provide information about the stationary Hamiltonian, the driving force, and the damping term. Moreover, the pHNN consistently outperforms all the network architectures shown in Figs. 1(a)–1(d) across the applications investigated in this study.

B. Network optimization

The training of the pHNN consists of feeding the inputs $[\mathbf{q}, \mathbf{p}, t]$ into the model. The first component, \mathcal{H}_{θ_1} , consists of three hidden layers designed to predict a stationary $\hat{\mathcal{H}}_{\theta_1}$ from $[\mathbf{q}, \mathbf{p}]$ input data. The second component, \hat{N}_{θ_3} , consists of a single weight parameter (i.e., θ_3 is a single node) designed to learn the damping. The third neural-network unit \hat{F}_{θ_2} solely depends on t and consists of three hidden layers designed to predict a time-varying force. The output of each component is transformed through Eq. (4) to obtain predictions of the state time derivatives $[\hat{\mathbf{q}}, \hat{\mathbf{p}}]$. Using these predicted quantities we construct the loss function for optimizing

the pHNN:

$$\mathcal{L} = \|\hat{\mathbf{q}}_t - \dot{\mathbf{q}}_t\|_2^2 + \|\hat{\mathbf{p}}_t - \dot{\mathbf{p}}_t\|_2^2 + \lambda_F \|\hat{\mathbf{F}}_{\theta_2}\|_1 + \lambda_N \|\hat{N}_{\theta_3}\|_1, \quad (5)$$

where $[\dot{\mathbf{q}}, \dot{\mathbf{p}}]$ are known ground truth data. The first two components of the left-hand side of Eq. (5) minimize the difference between the predicted and ground truth state time derivatives with a squared error loss. The last two components in Eq. (5) are the forcing and damping terms that are added to the loss function with an L_1 penalty when using pHNN. Using an L_1 penalty on these terms encourages the network to learn simpler models. We empirically found that this technique prevents the pHNN from learning spurious force and damping terms in unforced and undamped systems compared to an L_2 penalty. The regularization parameters λ_F and λ_N were determined via grid search (see Supplemental Material, Appendix A [29]). We use 200 nodes per hidden layer and find that most activation functions, including $\tanh(\cdot)$, $\sin(\cdot)$ and $\cos(\cdot)$ yield comparable results.

We generate our data using an RK4 integrator given some initial conditions for each system. We use a small Δt to evaluate the integral and ensure ground truth data is generated with $\text{rtol} = 10^{-10}$. The gradients of the state at each integration step are computed using the underlying differential equation. Details about the sampling of the initial conditions are described independently for each system in Sec. IV.

We note that in some settings it might be hard to obtain the ground truth state time-derivatives $[\dot{\mathbf{q}}, \dot{\mathbf{p}}]$ for training. A natural way to address this problem is to embed an integrator into the training, similar to NeuralODE [11]. As such, we also run our study with an embedded RK-4 integrator. Our method is still the most performant when all the methods incorporate an embedded RK-4 integrator (see Appendix in Supplemental Material [29]) and the loss function is $([\hat{q}, \hat{p}]_{t+1} - [q, p]_{t+1})^2$. In other words, our system can learn from either state time derivative data or directly from state data given an embedded integrator. We were motivated to use gradient data since HNN is trained in this way and we wanted a fair comparison.

C. Testing

Once trained, each of the networks in Fig. 1 can approximate $[\dot{\mathbf{q}}, \dot{\mathbf{p}}]$. As such, these networks can be used in a time integrator to evolve initial conditions in the test set. We refer to the integration for $t \in (0, T_{\max})$ as the state roll out. We measure the performance of the network by comparing the predicted state roll out with the ground truth. Specifically, we assess the networks performance by computing the mean-squared error (MSE) of the predicted state variables and predicted energy (Hamiltonian) across the integration time,

$$\text{MSE}_{\text{state}} = \frac{1}{N} \sum_{i=1}^N (\mathbf{q}_i - \hat{\mathbf{q}}_i)^2 + \frac{1}{N} \sum_{i=1}^N (\mathbf{p}_i - \hat{\mathbf{p}}_i)^2 \quad (6)$$

$$\text{MSE}_{\text{energy}} = \frac{1}{N} \sum_{i=1}^N (\mathcal{H}(\mathbf{q}_i, \mathbf{p}_i) - \mathcal{H}(\hat{\mathbf{q}}_i, \hat{\mathbf{p}}_i))^2, \quad (7)$$

where $N = T_{\max}/\Delta t$, with Δt the time step size. These terms are computed for multiple initial conditions during inference

and averaged across them. A guide to the training process is outlined in Appendix B (see Supplemental Material [29]).

IV. RESULTS

We benchmark the performance of pHNNs against the other networks shown in Fig. 1. We evaluate the methods over datasets that cover simple time-independent systems to complex chaotic damped and driven dynamical systems. The results are presented in order of increasing complexity from a model perspective.

A. Simple mass-spring system

We begin our analysis with a simple mass-spring system (harmonic oscillator), obeying Hooke's law from classical physics with no force or damping. The Hamiltonian that describes such a system reads:

$$\mathcal{H} = \frac{1}{2}kq^2 + \frac{p^2}{2m}, \quad (8)$$

where k is the spring constant and m denotes the mass. In this one-dimensional system the position and momentum are scalar functions of time.

Training. Without loss of generalization, we set $k = m = 1$ for our experiments. We randomly sample 25 initial training conditions $[q_0, p_0]$ that satisfy $q_0^2 + p_0^2 = r_0^2$ where $1 \leq r_0 \leq 4.5$, which corresponds to sampling initial conditions with energies in the range $[1, 4.5]$. We evolve each initial state using a RK-4 integrator with $\Delta t = 0.05$ and $T_{\max} = 3.05$.

Testing. We evaluate the performance of the NNs by sampling 25 random initial conditions in the same way as training. We investigate the simple harmonic oscillator system and show that learning a separate, regularized forcing term results in better state and energy predictions in comparison to TDHNN and the baseline NN. Learning a separate forcing term and regularizing pHNN keeps the time component independent of the Hamiltonian and therefore allows us to closely match the performance of the standard time-implicit HNN.

In particular, in Fig. 2(a) we show the state and energy MSE for an initial state from the testing set. Figure 2(b) presents the predicted force (blue solid line) and damping over time $v \frac{\partial \mathcal{H}}{\partial p}$ (red solid curve), while black dots corresponds to ground truth observations. We observe that the error in the predictions is of the order of 10^{-5} for recovering the force function and of 10^{-8} for the damping term. In Fig. 2(c) we report the state and energy MSE averaged along all the testing initial states, where the black lines in the histograms represent the error bars of the statistics.

B. Damped mass-spring system

We extend the simple mass-spring system to include a damping term that reduces the initial energy of the system over time. The inclusion of this term violates energy conservation and therefore we cannot write a scalar Hamiltonian for such a system (see Supplemental Material, Appendix C [29] for details). Knowing this *a priori* already gives us an indication that the HNN will perform poorly on such a system.

Training. We have 20 initial training conditions, with position and momentum uniformly sampled in $[-1, 1]^2$. Each

trajectory is evolved until $T_{\max} = 30.1$ with a $\Delta t = 0.1$. We fix the damping coefficient $\nu = 0.3$ without loss of generality.

Testing. At inference, we compute the average roll-out MSE of 25 unseen initial conditions sampled in the same manner as the training data. Figure 3(a) outlines the predicted force and damping for an arbitrary initial state, while in Fig. 3(b) we report the average state and energy roll-out MSE.

We observe in Fig. 3 that both baseline NN and pHNN recover the dynamics well, whereas the HNN (as expected) and TDHNN struggle to learn the dynamics of the damped system. This failure happens because there is no direct way of writing a scalar Hamiltonian with damping and thus, both the HNN and TDHNN cannot learn the underlying dynamics. This observation indicates that a partial inductive bias is not enough to make a network robust at predicting the dynamics well. A second observation shows that the pHNN learns a nonzero oscillating force. That implies a possible leaking of the information from the predicted Hamiltonian into the predicted force and vice versa. This arises because there is an identifiability challenge inferring the damping and force exclusively from the state data we provide. In spite of this challenge, we find that the pHNN converges to forcing and damping terms consistent with the ground truth generating terms and sufficient to inform us about the underlying dynamics as well as to evolve initial states with small numerical error.

C. Forced mass-spring system

We complete the investigation of the simple mass-spring problem by including a driven time-dependent force that controls the system. To understand the effect of the force we consider an undamped driven oscillator system. Typically, while we cannot write the Hamiltonian for a damped system, we can write one for a forced system. We study two cases of forced mass-spring systems. The first has the following Hamiltonian form:

$$\mathcal{H} = \frac{1}{2}kq^2 + \frac{p^2}{2m} - qF_0 \sin(\omega t). \quad (9)$$

The second has a more complex force described by the Hamiltonian:

$$\mathcal{H} = \frac{1}{2}kq^2 + \frac{p^2}{2m} - qF_0 \sin(\omega t) \sin(2\omega t), \quad (10)$$

where F_0 and ω is the amplitude and frequency of the external force term. The forced mass-spring system is typically used to study resonance effects, e.g., in material science, and plays an important role in a wide range of applications including music, bridge design, and molecular excitation making it an important system to investigate.

Training. In both systems of Hamiltonian Eqs. (9) and (10), we use 20 initial conditions, where the initial state $[q_0, p_0]$ is sampled such that $q_0^2 + p_0^2 = r_0^2$ where $1 \leq r_0 \leq 4.5$. For the force term we set $F_0 = 1$ and $\omega = 3$, and without loss of generality we set $k = m = 1$. The states are rolled out to $T_{\max} = 10.01$ at a $\Delta t = 0.01$.

Testing. At inference, we compute the roll out of 25 unseen initial conditions in the same range as the training

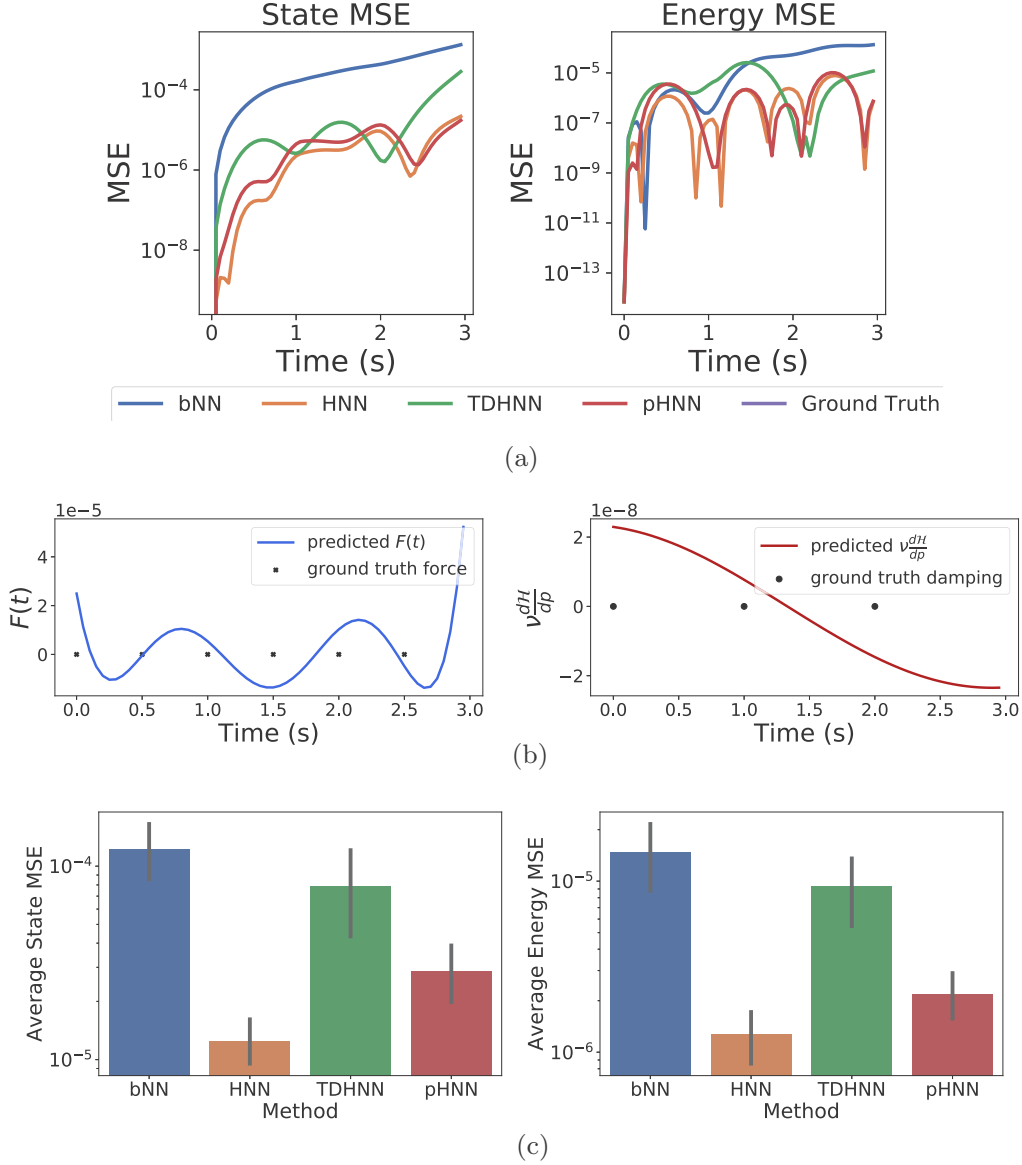


FIG. 2. The simple mass-spring system has no explicit time dependence. We see that the pHNN can almost recover the dynamics as well as the HNN. While the pHNN does learn a non-zero force and damping term, their contribution to $\frac{dp}{dt}$ is small. (a) State and energy MSE as a function of time of an initial condition in the test set, (b) Lrnt force and damping terms by pHNN, and (c) State and energy MSE averaged across 25 initial test states (error bars showing $\pm 1\sigma$).

data. Figures 4(a) and 5(a) demonstrate the predicted forces and dissipation terms for the systems of Eqs. (9) and (10), respectively. Accordingly, in Figs. 4(b) and 5(b) we report the average state and energy roll-out MSE for each system.

We study both systems to illustrate that while the baseline NN performs relatively well in comparison to the pHNN when a simple force is considered such as in the system of Eq. (9), a more complex force significantly hurts its performance in terms of state or energy MSE shown by Fig. 5(b). More importantly, we read in Figs. 4(a) and 5(a) that for both systems pHNN can recover the ground truth force quite precisely while it learns very small spurious damping terms that do not significantly contribute to the dynamics; the contribution to $\frac{dp}{dt}$ term is of the order 10^{-5} , which is practically negligible.

D. Duffing equation

Another problem that we investigate is given by the Duffing equation, a nonlinear dynamical system that includes both forcing and damping. The unforced and undamped stationary Hamiltonian $\mathcal{H}_{\text{stat}}$ of the Duffing system is given by:

$$\mathcal{H}_{\text{stat}} = \frac{p^2}{2m} + \alpha \frac{q^2}{2} + \beta \frac{q^4}{4}. \quad (11)$$

Unlike the simple mass-spring system, the Duffing equation has an additional quadratic function of q that makes the system nonharmonic. The shape of the potential function can be tuned to be a double well or a single well based on the coefficients α and β . The general Duffing equation includes a time variant force and a damping term proportional to $\partial \mathcal{H}_{\text{stat}} / \partial q$. Typically the Duffing nonlinear equation of motion is written

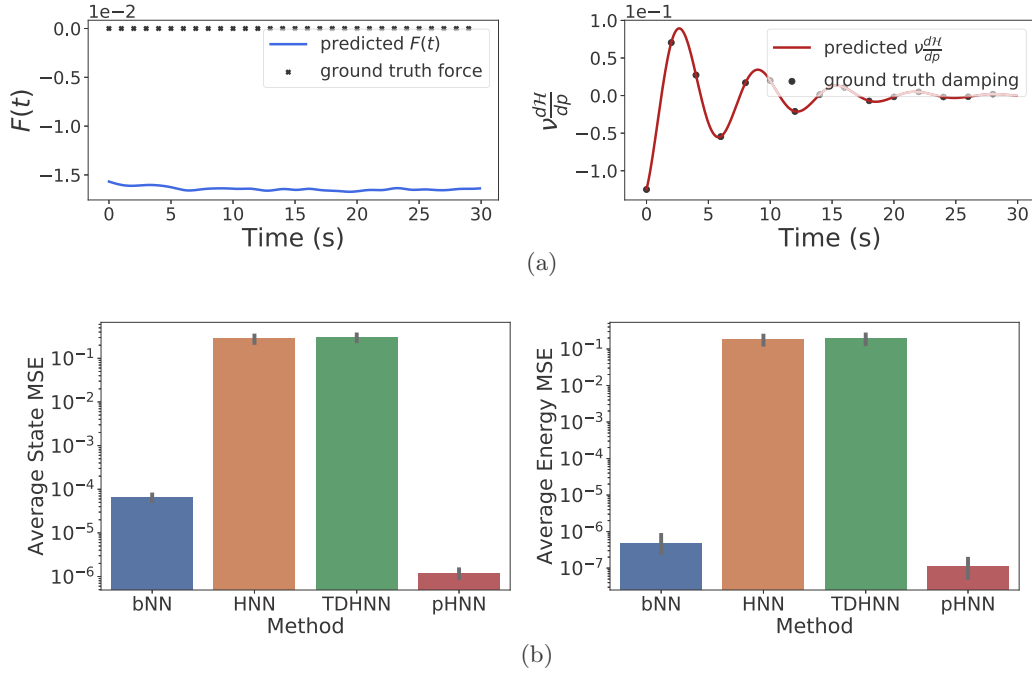


FIG. 3. Damped mass-spring setting: The baseline NN and pHNN recover the underlying dynamics well. pHNN is also able to accurately learn the damping coefficient since the predicted damping is indistinguishable from the ground truth. (a) Learned force and damping terms by pHNN and (b) State and energy MSE averaged across 25 initial test states (error bars showing $\pm 1\sigma$).

as:

$$\ddot{q} = -\delta\dot{q} - \alpha q - \beta q^3 + \gamma \sin(\omega t). \quad (12)$$

Different combinations of parameters $\alpha, \beta, \delta, \gamma, \omega$ make the Duffing system either chaotic or nonchaotic. We study both regimes. The Duffing equation reveals numerous phenomena of practical importance including frequency hysteresis (e.g., in magnets), elasticity, and chaos theory.

1. Nonchaotic regime

Given a set of initial parameters for the Duffing equation: $\alpha = -1, \beta = 1, \delta = 0.3, \gamma = 0.2, \omega = 1.2$ we can obtain training data in a nonchaotic regime of the Duffing system.

Training. We uniformly sample initial conditions in $[-1, 1]^2$ and use 25 initial conditions for training, rolled out to $T_{\max} = 10.01$ with $\Delta t = 0.01$.

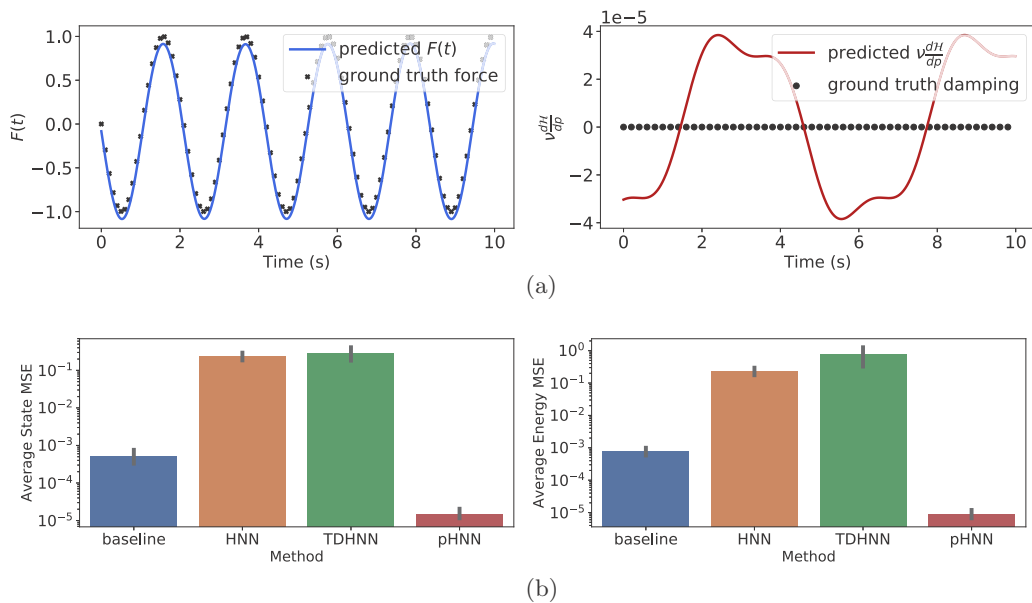


FIG. 4. Forced mass-spring of Eq. (9): Standard HNN cannot learn the underlying dynamics as it has no explicit-time dependence. pHNN shows the best performance as it explicitly learns a time-dependent force. (a) Learned force and damping terms by pHNN and (b) Roll-out state and energy MSE averaged across 25 initial test states (error bars showing $\pm 1\sigma$).

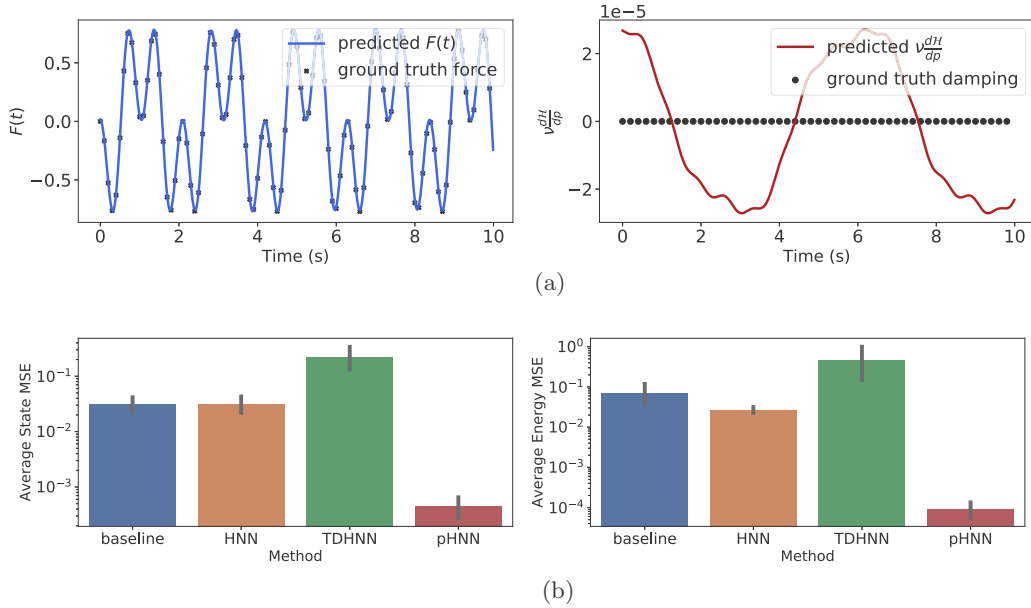


FIG. 5. Forced mass-spring system of Eq. (10): pHNN is able to recover a nonharmonic force and evolves testing initial states better than the other models. (a) pHNN recovers force and damping and (b) Roll-out state and energy MSE averaged across 25 initial test states (error bars showing $\pm 1\sigma$).

Testing. We integrate 25 unseen initial conditions at inference using the same T_{\max} and Δt used to generate the training set. We evaluate all the neural network models on this testing set and present the results in Fig. 6. We observe in Fig. 6(a) that pHNN accurately recovers the underlying force and damping. Moreover, Fig. 6(b) indicates the pHNN

outperforms the other models used in this study. We further assess the network performance by inspecting the predicted Hamiltonian. In Fig. 7 we outline the learned $\mathcal{H}_{\text{stat}}$ as a function of q and p that comprise the phase space. We observe that pHNN can learn the functional form of $\mathcal{H}_{\text{stat}}$, outperforming the other architectures.

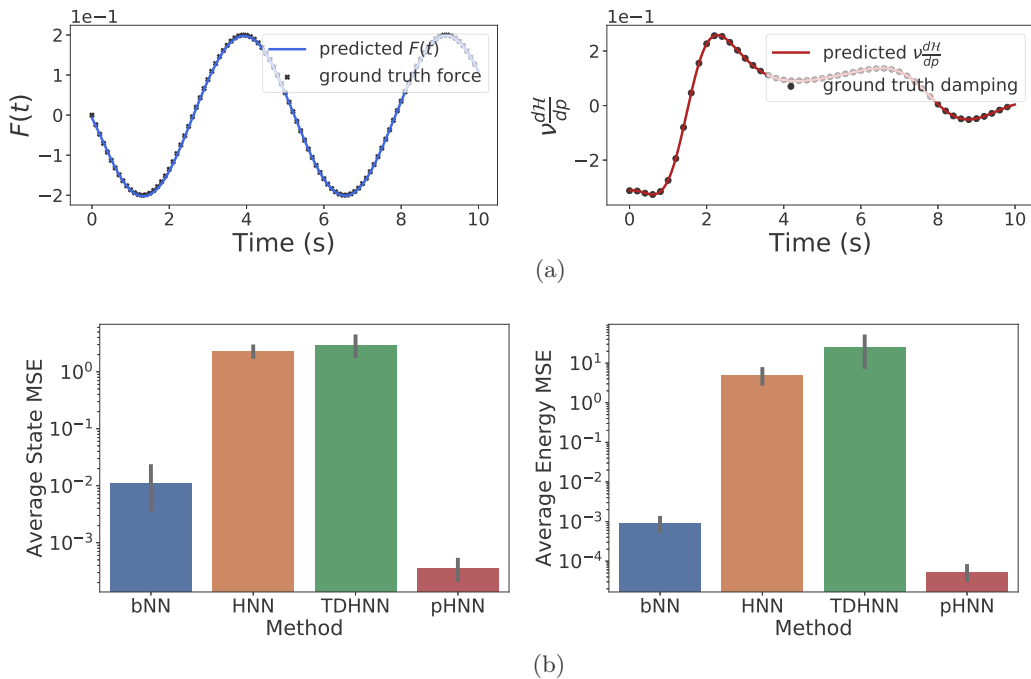


FIG. 6. Duffing system (nonchaotic): pHNN significantly outperforms the other methods and is able to extract the ground truth force and damping coefficient. (a) Learned force and damping terms of pHNN and (b) State and energy MSE averaged across 25 initial test states (error bars showing $\pm 1\sigma$).




FIG. 7. Nonchaotic Duffing setting: Learned $\mathcal{H}_{\text{stat}}$ in q - p plane (phase space) across different network architectures. HNN and TDHNN learn distorted Hamiltonians that strongly depend on the input time-variable, whereas pHNN is able to recover a nondistorted Hamiltonian.

2. Chaotic regime

The choice of the parameters: $\alpha = 1, \beta = 1, \delta = 0.1, \gamma = 0.39, \omega = 1.4$, yields chaotic behavior in the Duffing system. Chaotic trajectories are highly sensitive to initial conditions and thus, it is much more difficult to learn from a chaotic system than from a nonchaotic.

Training. 20 initial conditions, sampled uniformly in $[-1, 1]^2$ each rolled out for one period $T = 2\pi/\omega$ where $\Delta t = T/100$ resulting in 2000 training points.

Testing. We test our system by assessing whether it is visually able to recover the ground truth Poincaré section of an initial condition and we focus our attention on baseline NN and pHNN since they are the most performant in the non-chaotic regime. The Poincaré map (or section) of a trajectory is measured by plotting the position and momentum values at regular intervals governed by the period of the forcing term. For example, a simple mass-spring system will generate a single point in phase space when measured at regular intervals, while a chaotic system generates a more complex map. To visually assess the performance of our network through a Poincaré map, we test the system on a single initial condition, not used in the training set, rolled out to $T_{\text{max}} = 18000$ with the same Δt as in the training phase. In order to integrate our system to such a large T_{max} for this example we work under the assumption that we have explicit knowledge of the period of the force, and as such, we normalize the time variable with the period. We emphasize that while prior knowledge of ω assists the training, it is only used to extract the Poincaré map. This is necessary as the models are not explicitly trained on time steps beyond $2\pi/\omega$. The results are demonstrated in Fig. 8. In our study we find that the visual similarity between the section

generated by pHNN is much closer to the ground truth than baseline. A simple quantitative measure of this similarity can be computed using the MSE between the 2D histogram plots. The pHNN results in 1.61 and bNN in 4.05. The outcome suggests that the pHNN can indeed be used to model chaos even after being trained with only a few data points from a chaotic trajectory. We believe this is a valuable result, since it implies that the pHNN can be used to model chaotic behavior.

E. Relativistic system

In the final experiment presented in this study, we go beyond nonrelativistic classical mechanics and explore a different form of Hamiltonians. In particular, we investigate the motion of a driven relativistic particle in a nonlinear double well potential, which is mathematically represented by the Duffing equation in a relativistic framework. The Hamiltonian under consideration is

$$\mathcal{H} = c\sqrt{p^2 + m_0^2c^2} + \frac{\alpha}{2}q^2 + \frac{\beta}{4}q^4 - q\gamma \sin(\omega t), \quad (13)$$

where c is the speed of light that typically is set to 1. For simplicity, we also set the rest mass $m_0 = 1$, though our framework naturally accounts for other values.

Training. We train on 25 initial conditions, uniformly sampled in $[0, 2]^2$. We consider $T_{\text{max}} = 20.01, \Delta t = 0.01$, and the parameters $\alpha = 1, \beta = 1, \delta = 0, \gamma = 0.2, \omega = 1.2$.

Testing. Using the same parameters as training, we roll out 25 unseen initial conditions and present the results in Fig. 9. We observe that pHNN is able to recover the underlying force and outperforms the other architectures in generating the




FIG. 8. Poincaré sections of a chaotic Duffing oscillator. Baseline NN (left) and pHNN (middle) are trained for 20000 iterations with 2000 data points. The left and middle images indicate the predicted Poincaré map for a initial state not used in networks optimization. The pHNN significantly outperforms the baseline NN at recovering the ground truth Poincaré section (right).

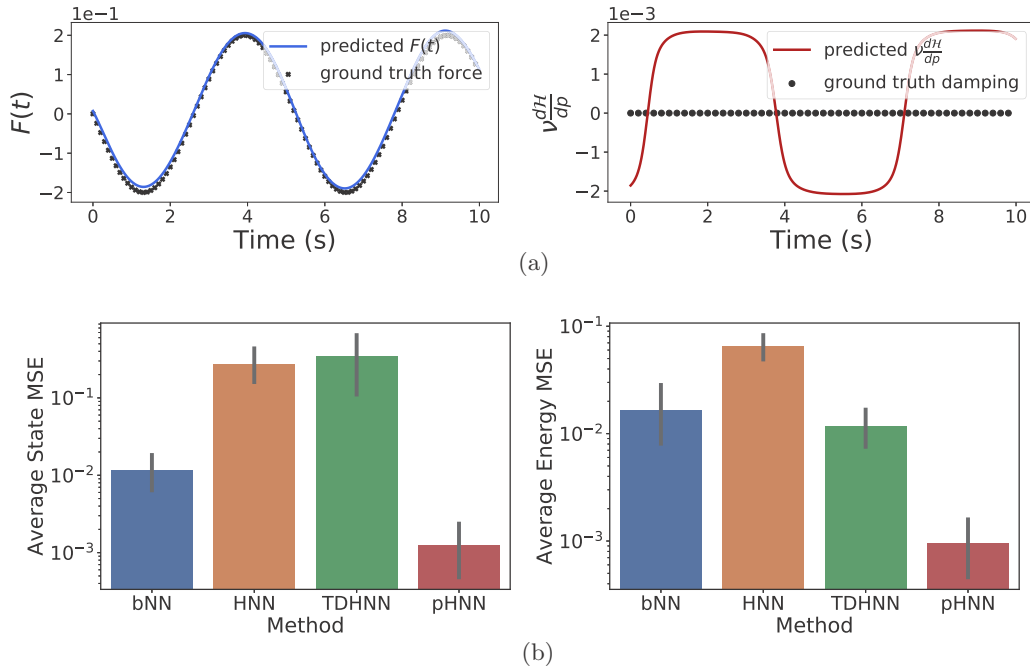


FIG. 9. Learned dynamics of a relativistic Duffing system. (a) Learned force and damping terms of pHNN and (b) Roll-out state and energy MSE averaged across 25 initial test states.

temporal state of previously unknown initial conditions. This experiment is evidence that pHNN can discover the dynamics of temporal systems independent of the form of the underlying Hamiltonian and consistently outperforms the other architectures investigated in this study.

V. DISCUSSION

We have shown that pHNN outperforms other approaches in learning complex physical systems (details shown in the Supplemental Material, Appendixes D and E [29]), as well as being able to recover the underlying stationary Hamiltonian, the external time varying force, and the damping term of non-autonomous systems. One challenge in achieving this result is fine tuning the λ_F and λ_N regularization coefficients for the force and damping. In the Duffing setting where we have both terms, it is possible to learn a shifted force, i.e., $F = F_0 + \epsilon$. This is possible because the state vectors \mathbf{q}, \mathbf{p} do not provide enough information to simultaneously identify both the Hamiltonian and the force resulting in a leak of information between \mathcal{H} and F . Nevertheless, we find that a reasonable force and damping term are generally learned, which are sufficient to reveal the underlying dynamics.

While it may be argued that pHNN is constrained to learn and predict within the training time horizon, we believe our method is still versatile at informing us of periodic forcing, since we can inspect the force over time and, essentially, learn the period of the underlying force. This, in turn, can readily be used to renormalize the time variable at periodic intervals to integrate the system beyond the training time as we showed in the Poincaré map of Fig. 8.

We also run the entire set of experiments on noisy input state vectors where the noise is sampled from $\mathcal{N}(0, \sigma)$ with $\sigma \in [0.01, 0.1, 0.5]$. The details of the results can be found in

Appendix F [29]. We find that even with the addition of noise to the input state vector, pHNN outperforms other methods.

We also carry out a simple study on a two-body coupled spring system where one of the masses is forced by a cosine varying signal. The results indicate that pHNN performs the best and has potential to scale to larger domains (see Appendix [29]).

VI. CONCLUSION

We have shown that learning the dynamics of time-dependent nonautonomous systems can be achieved with pHNN, a versatile neural network embedded with the Port-Hamiltonian formulation. Our experimental investigation demonstrates that pHNN outperforms extensions of existing methods in numerous settings. Specifically, we outlined that the proposed network not only learns the underlying dynamics of a simple mass-spring system, achieving comparable performance to the efficient HNN architecture, but it extends to more complex nonlinear forced and damped physical systems. Furthermore, using pHNN we were able to show, with minimal training data, the ability to recover the Poincaré section of a chaotic driven system. Unlike existing methods, pHNN is able to identify systems where minimal state data is available and reveal the functional form of the controlling force, damping, and underlying stationary Hamiltonian. Collectively, these results form a strong basis for further advances in learning complex systems including, but not limited to, chemical bond forces, robotic motion, and more general controlled dynamics without explicit knowledge of the force and damping.

ACKNOWLEDGMENT

We would like to thank the Rhodes Trust for supporting this research.

- [1] K. Hornik, M. Stinchcombe, and H. White, Multilayer feed-forward networks are universal approximators, *Neural Netw.* **2**, 359 (1989).
- [2] K. He, G. Gkioxari, P. Dollár, and R. Girshick, Mask R-CNN, in *Proceedings of the 2017 IEEE International Conference on Computer Vision (ICCV)* (IEEE, Piscataway, NJ, 2017), pp. 2980–2988.
- [3] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, BERT: Pre-training of deep bidirectional transformers for language understanding, in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Long and Short Papers)* (Association for Computational Linguistics, Minneapolis, Minnesota, 2019), Vol. 1, pp. 4171–4186.
- [4] M. Toussaint, K. Allen, K. Smith, and J. Tenenbaum, Differentiable physics and stable modes for tool-use and manipulation planning, in *Robotics: Science and Systems XIV* (Robotics: Science and Systems Foundation, Los Altos, 2018).
- [5] K. Yao, J. E. Herr, D. Toth, R. Mckintyre, and J. Parkhill, The TensorMol-0.1 model chemistry: A neural network augmented with long-range physics, *Chem. Sci.* **9**, 2261 (2018).
- [6] S. Greydanus, M. Dzamba, and J. Yosinski, Hamiltonian neural networks, in *Advances in Neural Information Processing Systems 32*, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. d. Alché-Buc, E. Fox, and R. Garnett (Curran Associates, Red Hook, 2019) pp. 15379–15389.
- [7] A. Pukrittayakamee, M. Malshe, M. Hagan, L. M. Raff, R. Narulkar, S. Bukkapatnum, and R. Komanduri, Simultaneous fitting of a potential-energy surface and its corresponding force fields using feedforward neural networks, *J. Chem. Phys.* **130**, 134101 (2009).
- [8] M. Mattheakis, D. Sondak, A. S. Dogra, and P. Protopapas, Hamiltonian neural networks for solving equations of motion, [arXiv:2001.11107](https://arxiv.org/abs/2001.11107).
- [9] M. Cranmer, S. Greydanus, S. Hoyer, P. Battaglia, D. Spergel, and S. Ho, Lagrangian neural networks, [arXiv:2003.04630](https://arxiv.org/abs/2003.04630).
- [10] M. Lutter, C. Ritter, and J. Peters, Deep Lagrangian networks: Using physics as model prior for deep learning, *International Conference on Learning Representations* (2019), <https://openreview.net/forum?id=BklHjpCqKm>.
- [11] R. T. Q. Chen, Y. Rubanova, J. Bettencourt, and D. K. Duvenaud, Neural ordinary differential equations, in *Advances in Neural Information Processing Systems 31*, edited by S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Curran Associates, Red Hook, 2018), pp. 6571–6583.
- [12] M. Raissi, P. Perdikaris, and G. E. Karniadakis, Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations, *J. Comput. Phys.* **378**, 686 (2019).
- [13] A. Choudhary, J. F. Lindner, E. G. Holliday, S. T. Miller, S. Sinha, and W. L. Ditto, Physics-enhanced neural networks learn order and chaos, *Phys. Rev. E* **101**, 062207 (2020).
- [14] P. Toth, D. J. Rezende, A. Jaegle, S. Racanière, A. Botev, and I. Higgins, Hamiltonian generative networks, in *International Conference on Learning Representations* (2020), <https://openreview.net/forum?id=HJenn6VFvB>.
- [15] P. Battaglia, R. Pascanu, M. Lai, D. J. Rezende, and K. Kavukcuoglu, Interaction networks for learning about objects, relations and physics, in *Proceedings of the 30th International Conference on Neural Information Processing Systems* (Curran Associates, Barcelona, Spain, 2016), pp. 4509–4517.
- [16] A. Sanchez-Gonzalez, V. Bapst, K. Cranmer, and P. Battaglia, Hamiltonian graph networks with ODE integrators, [arXiv:1909.12790](https://arxiv.org/abs/1909.12790).
- [17] A. van der Schaft, Port-Hamiltonian systems: An introductory survey, in *Proceedings of the International Congress of Mathematicians Madrid*, edited by M. Sanz-Solé, J. Soria, J. L. Varona, and J. Verdera (European Mathematical Society Publishing House, Zuerich, 2007), pp. 1339–1365.
- [18] R. Ortega, A. van der Schaft, B. Maschke, and G. Escobar, Interconnection and damping assignment passivity-based control of port-controlled hamiltonian systems, *Automatica* **38**, 585 (2002).
- [19] J. A. Acosta, R. Ortega, A. Astolfi, and A. D. Mahindrakar, Interconnection and damping assignment passivity-based control of mechanical systems with underactuation degree one, *IEEE Trans. Auto. Control* **50**, 1936 (2005).
- [20] M. Zheng, T. Yuan, and T. Huang, *Time-Varying Impedance Control of Port Hamiltonian System with a New Energy-Storing Tank* (Hindawi, London, 2018).
- [21] K. Cherifi, An overview on recent machine learning techniques for Port Hamiltonian systems, *Physica D* **411**, 132620 (2020).
- [22] Y. D. Zhong, B. Dey, and A. Chakraborty, Symplectic ODE-Net: Learning Hamiltonian dynamics with control, in *International Conference on Learning Representations* (2020), <https://openreview.net/forum?id=ryxmb1rKDS>.
- [23] D. Maclaurin, D. Duvenaud, and R. P. Adams, Autograd Eortless Gradients in Numpy3.
- [24] E. Dupont, A. Doucet, and Y. W. Teh, Augmented neural ODEs, in *Advances in Neural Information Processing Systems*, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett (Curran Associates, Red Hook, 2019), Vol. 32.
- [25] A. Zhu, P. Jin, and Y. Tang, Deep Hamiltonian networks based on symplectic integrators, [arXiv:2004.13830](https://arxiv.org/abs/2004.13830).
- [26] A. Sanchez-Gonzalez, N. Heess, J. T. Springenberg, J. Merel, M. A. Riedmiller, R. Hadsell, and P. W. Battaglia, Graph networks as learnable physics engines for inference and control, in *ICML*, edited by J. G. Dy and A. Krause, Proceedings of Machine Learning Research (PMLR, 2018), Vol. 80, pp. 4467–4476.
- [27] A. Sanchez-Gonzalez, J. Godwin, T. Pfaff, R. Ying, J. Leskovec, and P. W. Battaglia, Learning to simulate complex physics with graph networks, in *Proceedings of the 37th International Conference on Machine Learning* (PMLR, 2020), Vol. 119, pp. 8459–8468.
- [28] M. Finzi, K. A. Wang, and A. G. Wilson, Simplifying Hamiltonian and Lagrangian neural networks via explicit constraints, in *Advances in Neural Information Processing Systems*, edited by H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (Curran Associates, Red Hook, 2020), Vol. 33, pp. 13880–13889.
- [29] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevE.104.034312> for information on additional experiments with embedded integrators and hyperparameter exploration.