

TOWARDS CROSS DOMAIN GENERALIZATION OF HAMILTONIAN REPRESENTATION VIA META LEARNING

Yeongwoo Song¹ Hawoong Jeong^{1,2*}

¹Department of Physics, KAIST ²Center for Complex Systems, KAIST
ywsong1025@kaist.ac.kr, hjeong@kaist.edu

ABSTRACT

Recent advances in deep learning for physics have focused on discovering shared representations of target systems by incorporating physics priors or inductive biases into neural networks. While effective, these methods are limited to the system domain, where the type of system remains consistent and thus cannot ensure the adaptation to new, or unseen physical systems governed by different laws. For instance, a neural network trained on a mass-spring system cannot guarantee accurate predictions for the behavior of a two-body system or any other system with different physical laws. In this work, we take a significant leap forward by targeting cross domain generalization within the field of Hamiltonian dynamics. We model our system with a graph neural network (GNN) and employ a meta learning algorithm to enable the model to gain experience over a distribution of systems and make it adapt to new physics. Our approach aims to learn a unified Hamiltonian representation that is generalizable across multiple system domains, thereby overcoming the limitations of system-specific models. We demonstrate that the meta-trained model captures the generalized Hamiltonian representation that is consistent across different physical domains. Overall, through the use of meta learning, we offer a framework that achieves cross domain generalization, providing a step towards a unified model for understanding a wide array of dynamical systems via deep learning.

1 INTRODUCTION

Deep learning has succeeded in many application areas such as image classification, image generation, natural language processing, and so on (Reed et al., 2016; Tan & Le, 2019; Raffel et al., 2020; Gu et al., 2022). One of the major roles of such accomplishment was capable of parameterizing useful representations from data with neural networks (Chen et al., 2020; Van Den Oord et al., 2017; Hamilton et al., 2017). However, grafting deep learning to physics is yet another problem. They struggle to learn conservation laws or implicit physical geometries or symmetries. Although various studies make the model learn the conservative quantities or symmetries inside the system (Greydanus et al., 2019; Sanchez-Gonzalez et al., 2019; Liu & Tegmark, 2021), these approaches are system-specific, which means that if a model is trained for one system, it cannot easily adapt to another system with different physical laws. Since systems whose physics is unknown have more sparse data, these flaws make the standard supervised learning methods harder to learn the physics of the system.

One of the current extents of this limitation is crafting the generalization of a fixed type of system across various parameter environments (Lee et al., 2021; Yin et al., 2021; Kirchmeyer et al., 2022; Li et al., 2023). Here, one can think of an extension of building a unified model that is not limited to adapting to various parameter environments within a single system type, but instead generalizes across the vast landscape of different system types. In this paper, we focus our objectives on systems that Hamiltonian mechanics can formulate. Systems governed by Hamiltonian mechanics inhere symmetries, and their state trajectories are laid on the symplectic manifold (Liebermann & Marle, 2012; Arnol’d, 2013). Moreover, almost all of the physics in our nature has its own conservation laws, and Hamiltonian mechanics relates the state of the system to its corresponding conservative quantities (usually energy).

*Correspondence to Hawoong Jeong.

In this work, we search for the generalized representation that shares the essence of Hamiltonian mechanics that a neural network can learn. In doing so, we aim to make the model easily adaptable to unseen physics.

Our contributions are summarized as follows.

- Utilizing the meta learning algorithm, we extend its application to the data distribution across the Hamiltonian manifold, encompassing various types of physical systems. This enables our model to generalize effectively to system data with unknown physics, enhancing its versatility.
- We begin by evaluating the model’s performance on simpler examples, such as the mass-spring and pendulum systems. Subsequently, we expand our evaluation to more complex scenarios characterized by chaotic behavior, such as the Hénon-Heiles and magnetic-mirror systems. This demonstrates the robustness and applicability of our approach across a wide range of dynamical systems.
- To gain insights into the superior performance of the meta-trained model compared to other baselines, we conduct an in-depth analysis of the learned representations within the neural network. This investigation provides a deeper understanding of the underlying factors contributing to the model’s effectiveness due to meta learning.
- We use the term *domain* to refer to different types of physical systems governed by distinct laws. To the best of our knowledge, we are the first to explore cross-domain generalization across diverse dynamical system types enabling easy adaptation to new physics, surpassing prior studies limited to fixed system types.

2 PRELIMINARIES

2.1 HAMILTONIAN MECHANICS

Hamiltonian mechanics is one of the concrete frameworks to describe the state of a dynamical system. It begins by defining a vector composed of canonical coordinates $\mathbf{x} = (\mathbf{q}, \mathbf{p})$. Here \mathbf{q} and \mathbf{p} denotes the generalized coordinate and its canonical momentum respectively; $\mathbf{q} = (q_1, q_2, \dots, q_N)$, $\mathbf{p} = (p_1, p_2, \dots, p_N)$, where N is the degree of freedom of the given system. Then, Hamiltonian (\mathcal{H}) is defined such that Hamilton’s equation (Equation 1) is held. Usually, the Hamiltonian of the system corresponds to its energy where it is to be conserved.

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

Here, the Hamiltonian vector field $\mathbf{X}_{\mathcal{H}} = \left(\frac{\partial \mathcal{H}}{\partial \mathbf{p}}, -\frac{\partial \mathcal{H}}{\partial \mathbf{q}} \right) = J \nabla_{\mathbf{x}} \mathcal{H}_{\theta}$, where $J = \begin{bmatrix} 0 & I_n \\ -I_n & 0 \end{bmatrix}$, is known as symplectic, which allows the Hamiltonian to lie on a symplectic manifold. Symplectic manifold is one of the essences of differential geometry, which naturally arises from the formulations of Hamiltonian mechanics (Liebermann & Marle, 2012; Arnol’d, 2013). It provides the generalization of the phase space of a closed system, allowing one to probe the time evolution of the system.

2.2 LEARNING PHYSICAL DYNAMICS FROM NEURAL NETWORKS

Various works use neural networks to analyze the time evolution, or the conservation law of the physical system (Greydanus et al., 2019; Sanchez-Gonzalez et al., 2019; Cranmer et al., 2020; Kasim & Lim, 2022; Greydanus et al., 2023). For example, Hamiltonian Neural Network (HNN) (Greydanus et al., 2019) parameterizes the Hamiltonian of the system with a neural network. HNN learns the dynamics of the given system by inducing a physical bias to the loss function. Making use of Equation 1, HNN predicts the dynamics of the system incorporating the symplectic gradient inside the loss function as in Equation 2.

$$\mathcal{L}_{\text{HNN}} = \left\| \frac{\partial \mathcal{H}_{\theta}}{\partial \mathbf{p}} - \frac{\partial \mathbf{q}}{\partial t} \right\|_2 + \left\| \frac{\partial \mathcal{H}_{\theta}}{\partial \mathbf{q}} + \frac{\partial \mathbf{p}}{\partial t} \right\|_2 \quad (2)$$

As an alternative way to encode the physical bias into the objective function, the Variational Integrator Network (VIN) (Saemundsson et al., 2020) has been introduced, which preserves the geometric structure of physical systems. The architecture of VIN is designed to fit the discrete-time equations of motion of the given dynamical system. Thus, it provides interpretability and more efficient learning by preserving the manifold geometry inherent in physical systems directly in the model architecture.

While the theoretical understanding of inductive biases that underlie the advantages of neural networks remains elusive, recent studies attempt to shed light on the influence of such specific biases. For example, one demonstrates that contrary to conventional wisdom, the performance improvement in HNNs arises from the second-order structure, rather than the symplectic bias (Gruver et al., 2022). Nevertheless, a number of works focus on leveraging the inherent nature of physical systems, such as their symplectic structure and the principle of least action, to enhance the predictive capabilities of neural networks for system dynamics (Chen et al., 2021; Greydanus et al., 2023).

In opposition to implying relevant biases in neural networks, recent studies suggest a data-driven approach to learning the dynamics of one type of system generalized across different environments (Yin et al., 2021; Kirchmeyer et al., 2022). These focus on dynamical systems that are governed by the form of the following differential equation; $\frac{d\mathbf{x}_e(t)}{dt} = f_e(\mathbf{x}_e(t))$, where $\mathbf{x}(t)$ is a time-dependent state in the phase space \mathcal{X} , and $f_e : \mathcal{X} \rightarrow T\mathcal{X}$ maps $\mathbf{x}_e(t)$ to its time-derivative state $\frac{d\mathbf{x}_e(t)}{dt}$ which lies in the tangent space $T\mathcal{X}$. Here, the subscript e denotes an environment for a given dynamics (e.g. different $\alpha, \beta, \gamma, \delta$ parameter configuration of Lotka-Volterra system; $du/dt = \alpha u - \beta uv, dv/dt = \delta uv - \gamma v$).

2.3 META LEARNING FOR GENERALIZING DYNAMICAL SYSTEM

Meta learning, also referred to as "learning to learn," focuses on training a model that exhibits strong generalization across diverse data distributions, enabling it to effectively adapt to new tasks even with a small amount of previously unseen samples. As a result, meta learning can serve as a promising data-driven approach for generalizing the dynamics of physical systems. Among various approaches of meta learning, optimization-based methods are particularly versatile and compatible with differentiable models (Hospedales et al., 2020). For example, CoDA (Kirchmeyer et al., 2022) is a representative optimization-based meta learning method for multi-environment formulation of the generalization problem for dynamical systems. Particularly, previous studies (Lee et al., 2021; Li et al., 2023) have demonstrated the ability of a broadly applicable optimization-based meta-learning algorithm, Model-Agnostic Meta-Learning (MAML) (Finn et al., 2017), in uncovering the physical laws of a given system across different environments. There also exists non-MAML based methods for generalizing dynamics. Sequential Neural Processes (Singh et al., 2019) aims to build a meta-transfer learning framework for stochastic processes, DyAd (Wang et al., 2022) is a model-based meta learning method that learns the shared dynamics of a given domain, and meta-SLVM (Jiang et al., 2022) utilizes Bayesian meta learning for learning to adapt latent dynamic functions.

The approaches mentioned above are all limited to scenarios where the systems share the same functional form of Hamiltonians, such as a set of pendulum systems with varying physical parameters (e.g. mass and length). Expanding upon these limitations, we aim to investigate whether a shared neural representation can be discovered among diverse functional forms of Hamiltonians. In this study, we explore the portrayal of a symplectic manifold governed by Hamilton’s equation within neural networks by incorporating MAML with GNNs. Further details regarding the MAML algorithm can be found in Section 3.2 and 3.3.

3 METHODS

3.1 PREPARING THE DATASET FOR META LEARNING

We constructed a task distribution comprising six distinct types of physical systems: mass-spring, pendulum, Hénon-Heiles, magnetic-mirror, two-body, and three-body. Given that these systems feature separable Hamiltonians, which inherently limit the diversity in the kinetic energy terms, we intentionally selected systems with varying complexities in their potential energy terms. Starting from the simplest mass-spring system, we extend the type of system to more difficult systems such as Hénon-Heiles or magnetic-mirror system. The details for each system can be found in the Supplementary Materials (SM) A.1.

This variety of systems encompasses a range of complexities, featuring simple physical systems like the mass-spring and the pendulum, as well as more intricate systems exhibiting chaotic behavior such as the Hénon-Heiles and the magnetic-mirror. Additionally, the task distribution extends to systems involving more than two particles, specifically the two-body, and the three-body, thereby providing a comprehensive set of challenges. For all of our systems, we generated a dataset with $N = 10,000$ trajectories confined to the two-dimensional space, where the canonical coordinates $\mathbf{x} = (\mathbf{q}, \mathbf{p})$ as the input, and their corresponding time derivatives $\dot{\mathbf{x}} = (\dot{\mathbf{q}}, \dot{\mathbf{p}})$ as the output of a neural network. We provide the detailed procedure for generating the dataset in the SM A.2.

In the language of meta learning, several different types of systems will be used for meta training, while one of the remaining unselected systems will be used for meta testing (i.e. evaluation). We will denote each trajectory as $\mathcal{T}_{\mathbb{S},i}$, which is sampled from the task distribution $p(\mathcal{T}_{\mathbb{S}})$ (where $S \in \{\text{mass-spring, pendulum, Hénon-Heiles, magnetic-mirror, two-body, three-body}\}$, and $i \in \llbracket 1..N \rrbracket$). Then tasks from $\mathcal{T}_{\mathbb{S}^c}$ will compose the task distribution for meta training, and tasks from $\mathcal{T}_{\mathbb{S}}$ will form the data distribution for meta testing. For example, we will use the data distribution $\mathcal{T}_{\mathbb{S}^c}$; $\mathbb{S}^c = \{\text{mass-spring, pendulum, Hénon-Heiles}\}$ for meta training, $\mathcal{T}_{\mathbb{S}}$; $S = \{\text{magnetic-mirror}\}$ for evaluating the magnetic-mirror system. The exact dataset configuration for meta training scenario is described in SM A.3.

3.2 META TRAINING THE NEURAL NETWORK

We use a simple graph convolutional network (GCN) (Kipf & Welling, 2016) to parametrize the given physical system in which we implement it using PyTorch (Paszke et al., 2019). Using GNNs to describe a system of several particles mediates the model to handle various degree-of-freedom inputs such that the model can be fed with data from a diverse type of system, unlike from the previous works listed in Section 2.2, and 2.3 that dealt with fixed system scenarios. In line with the previous methods (Sanchez-Gonzalez et al., 2019; Bishnoi et al., 2023) that represented the physical systems as graphs, we represent the state variable (\mathbf{q}, \mathbf{p}) as the node feature, while currently, edge features are not utilized. Our model is composed of three GCN layers to extract features of the input states, preceded by three fully connected linear layers for the regression of Hamiltonian value. We choose the *mish* function as the non-linear activation function. Prior works (Greydanus et al., 2019; Sanchez-Gonzalez et al., 2019; Lee et al., 2021) have utilized *tanh* or *softplus* activations, hypothesizing that the *relu* activation may hinder parameter optimization due to its piecewise linear nature, as the HNN loss defined in Equation 2 involves calculating derivatives of the output with respect to the inputs. After conducting several trials, we observed that the *mish* activation produced more stable and improved results compared to *tanh* or *softplus* activations in our case. More details regarding the network architecture are further described in the SM B.1.

Upon the framework that we described so far, we represent our model as f_{θ} ; parameterized by neural network parameters θ . To account for the varying input (\mathbf{x}^i) and output ($\dot{\mathbf{x}}^i$) scales across different systems, we employ the log-cosh function as our loss function. This choice enhances the robustness of our model to data with diverse system dynamics, as it mitigates the impact of differing data scales.

$$\mathcal{L}_{\mathcal{T}_i}(f_{\theta}) = \sum_{(\mathbf{x}^i, \dot{\mathbf{x}}^i) \sim \mathcal{T}_i} \log \cosh(\dot{\mathbf{x}}^i - J \nabla_{\mathbf{x}} f_{\theta}(\mathbf{x}^i)) \quad (3)$$

In the inner loop of our variation of MAML, the model adapts to a new task $\mathcal{T}_{\mathbb{S}^c,i}$ by updating the parameters with one gradient step using $K = 50$ phase space points for each task.

$$\theta'_i = \theta - \alpha \nabla_{\theta} \mathcal{L}_{\mathcal{T}_{\mathbb{S}^c,i}}(f_{\theta}) \quad (4)$$

In the outer loop using the Adam optimizer (Kingma & Ba, 2014), meta optimization across systems is done by updating the parameters as below.

$$\theta \leftarrow \theta - \beta \nabla_{\theta} \sum_i \mathcal{L}_{\mathcal{T}_{\mathbb{S}^c,i}}(f_{\theta'_i}) \quad (5)$$

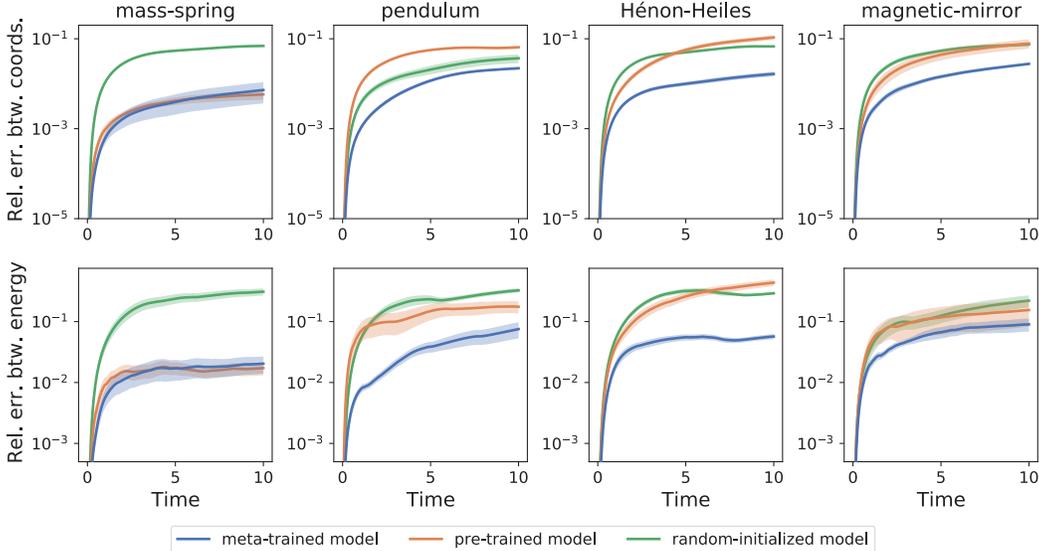


Figure 1: The relative error throughout the time rollout for the meta-trained, pre-trained, and random-initialized model at adaptation step 50 for the predicted coordinates (top) and energy (bottom). The solid line and the shaded area each represent the average and the standard error of 10 runs.

3.3 EVALUATING THE META-TRAINED MODEL

We evaluate the performance by making a few-step gradient descent on the unseen type of system during the meta training phase with $K = 50$ data points. The model adaptation to the unseen system is done by updating the parameters as follows.

$$\theta \leftarrow \theta - \alpha \nabla_{\theta} \mathcal{L}_{\mathcal{T}_{s,i}}(f_{\theta}) \quad (6)$$

Using the adapted model at each step, we integrate our model output using the LSODA integrator (Petzold, 1983) to achieve the system dynamics. Then, the performance was achieved by calculating the relative error between the predicted and ground truth trajectory points. We perform our experiments on NVIDIA RTX A6000 by performing 10 independent runs as default to ensure the stability of our results and report the average and standard error values. We again further describe the related details for the training and adaptation process in the SM B.2.

4 RESULTS

4.1 QUANTITATIVE ANALYSIS OF THE MODEL PERFORMANCE

For each system, we evaluated the adaptation performance on unseen systems for the meta-trained model and compared it with the pre-trained model that doesn’t utilize meta learning, and the scratch model which was adapted from random-initialized weights. We conducted up to 50 adaptation steps and computed the relative error between the predicted and the ground truth coordinates and energies at each time rollout and adaptation step. The measured relative error is defined as $\text{Err}(t) = \|\hat{z}(t) - z(t)\|_2 / (\|\hat{z}(t)\|_2 + \|z(t)\|_2)$, as described in (Finzi et al., 2020). This error metric enables us to quantify the error independent of the data scale and tends towards unity as the predictions become orthogonal to the ground truth or $\|\hat{z}\| \gg \|z\|$. Compared to metrics like mean squared error, the above notion of bounded relative error provides a more accurate representation of the prediction performance. Finally, we calculated the geometric moving average of the relative error to account for the multiplicative accumulation of the error from the numerical integration involved.

In Figure 1, the relative error in the predicted trajectories serves as an indicator of the overall accuracy of the adapted models, while the energy metric demonstrates that the physics of the system remains well-behaved during the prediction. Considering both predicted coordinates and energy, our results demonstrate that the meta-trained model robustly outperforms the other two baselines.

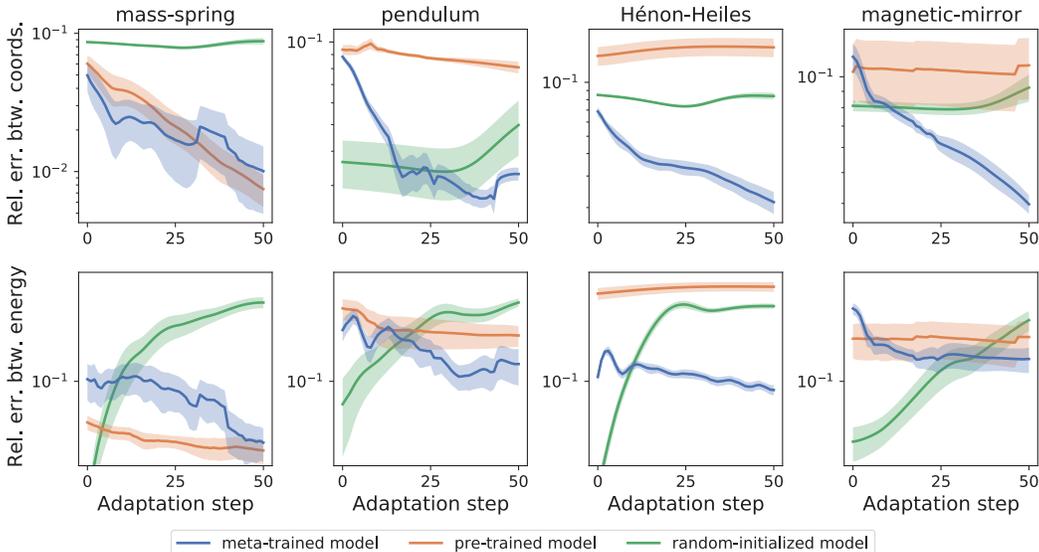


Figure 2: The relative error across the adaptation step for the meta-trained, pre-trained, and random-initialized model at adaptation step 50 for the predicted coordinates (top) and energy (bottom). The solid line and the shaded area each represent the average and the standard error of 10 runs.

We further examine the capability of fast adaptation of each model throughout the adaptation steps. Figure 2 illustrates the decreasing trend of the relative error as the number of adaptation steps increases from 0 to 50. Note that the meta-trained model requires significantly fewer adaptation steps compared to the randomly initialized model, which corresponds to the vanilla HNN model. Also, the relative error between the predicted coordinates shows that the fast adaptation shows a significant difference as the system becomes more complex. Within a gradient step of less than 100, the meta-trained model demonstrates faster and superior adaptation compared to both the pre-trained model without meta learning and the HNN baseline.

4.2 QUALITATIVE ANALYSIS OF THE PREDICTED DYNAMICS

In Section 4.1, we observed that the meta-trained model consistently achieved a lower relative error between the predicted and ground truth dynamics across both time rollouts and adaptation steps. However, a small error in the predicted trajectory does not always guarantee accurate system dynamics prediction. To provide more conclusive results, we examined the exact predicted dynamics of the three adapted models after 50 steps. We conducted this analysis for the mass-spring, pendulum, Hénon-Heiles, and magnetic-mirror systems, and the corresponding results are shown in Figure 3.

For the mass-spring system, both the meta-trained model and pre-trained model accurately predicted the coordinates (x, p_x) within 50 adaptation steps, while the random-initialized model required more steps to adapt. Here, it is worth noting the behavior of the prediction of the surplus coordinates (y, p_y) . In the real world, we cannot readily distinguish what coordinates remain physically meaningful among different phase variables. Thus, it would be beneficial if our trained neural network could also distinguish which one is necessary for describing the system. In the experiments, we confined the dynamics of the system to a two-dimensional space, thus y, p_y should remain constant for the mass-spring system. Comparing the predicted trajectories of the redundant coordinates between the meta-trained model and the pre-trained model, we observed that the meta-trained model had a much smaller scale in the prediction of unnecessary coordinates (y, p_y) .

Moving on to the pendulum system, the meta-trained model also provides better predictions than the other two baselines. While the pre-trained model performed better than the random-initialized model, the dynamics of θ and p_θ were still out of phase compared to the meta-trained model. We also observed a similar behavior in the prediction of the surplus coordinates (r, p_r) , as observed in the mass-spring system.

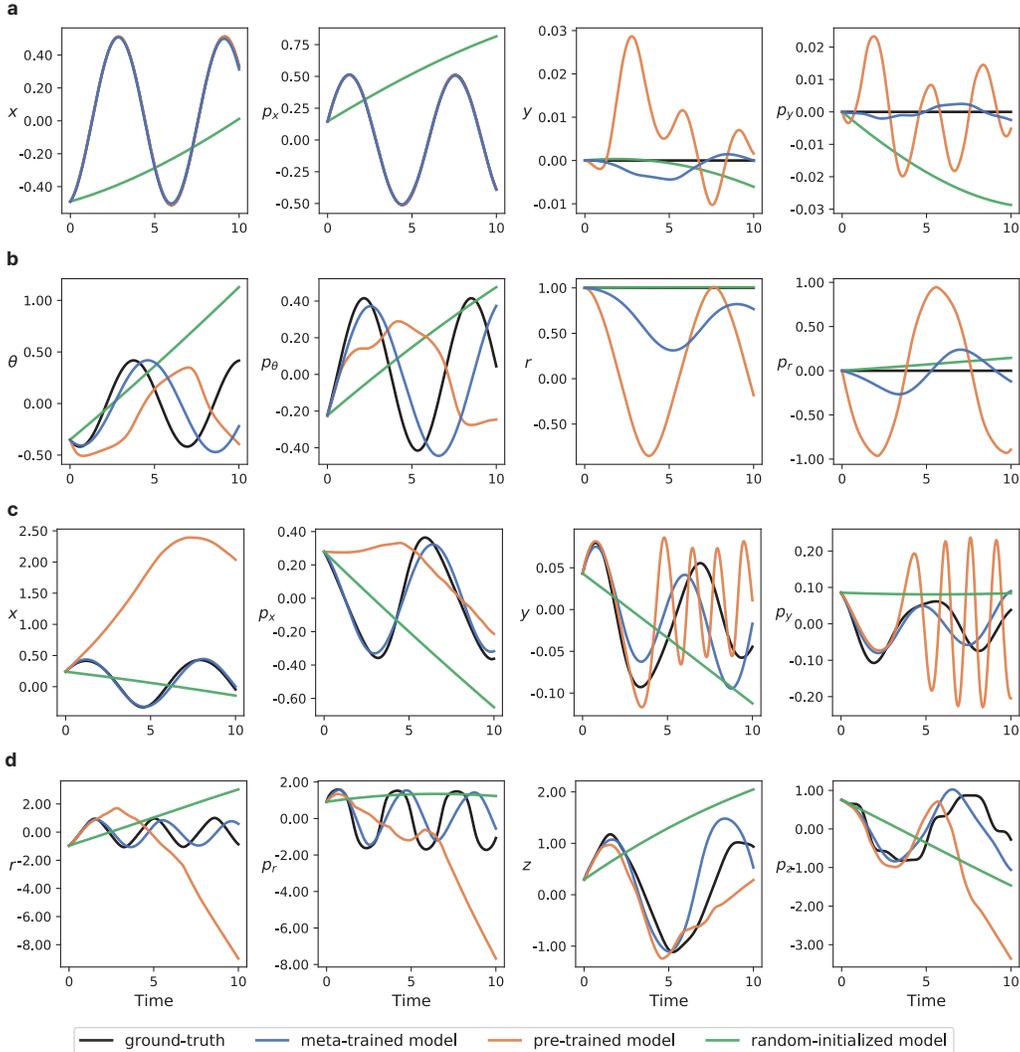


Figure 3: System dynamics through time rollout 0 to 10, generated from the meta-trained, pre-trained, and random-initialized model each after 50 adaptation steps. Each row (a) to (d) corresponds to the predictions with mass-spring, pendulum, Hénon-Heiles, and the magnetic mirror system respectively.

For the more complex systems exhibiting chaotic behavior, such as the Hénon-Heiles and magnetic-mirror system, all four coordinates are necessary to describe the dynamics. In these scenarios, the meta-trained model significantly outperformed the pre-trained model and random-initialized model. Although the meta-trained model did not perfectly match the trajectory, it captured the overall behavior of each coordinate dynamics more reasonably compared to the other baselines.

4.3 ANALYSIS ON THE LEARNED REPRESENTATION

Our objective in this section is to understand why the meta-trained model outperforms the other baselines, particularly the pre-trained model. The existing difference between the meta-trained model and the pre-trained model lies in the training strategy during the training process, while factors such as the total iterations, along with other training conditions remained the same. To investigate the impact of the training algorithm, we conducted the centered kernel alignment (CKA) (Kornblith et al., 2019) analysis on the three models. CKA offers a means to compare the representations of two layers within a neural network, providing a similarity score ranging from 0 to 1. Furthermore, it is widely recognized that early layers tend to learn general representations, while later layers tend to specialize in specific features (Yosinski et al., 2014; Raghu et al., 2019; 2020; Alabdulmohsin et al., 2021).

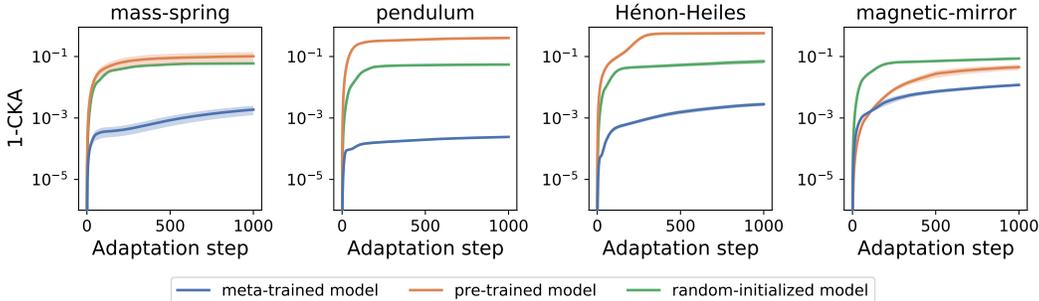


Figure 4: The measured CKAs in the last layer between the model before adapting (i.e. right after finishing meta-training or pre-training), and during the adaptation up to 1000 steps. The CKA values are transformed to 1-CKA to make a clearer visualization. The solid line and the shaded area each represent the average and the standard error of 10 runs.

Upon this observation, we examine the CKA between the model’s last layer before adaptation and its state at each adaptation step, ranging from 1 to 1000. Figure 4 shows these 1000 individual 1-CKA measurements, each comparing the model before adaptation to its state at a specific step during the adaptation process. In all four cases, the meta-trained model consistently exhibited substantially lower 1-CKA values compared to the other models. This finding suggests that the representations learned during the adaptation process closely resemble those of the meta-trained model when it serves as the starting point. Drawing from insights gained in the prior works utilized meta learning (Lee et al., 2021; Li et al., 2023), and the manifold hypothesis (Fefferman et al., 2016), we interpret the meta-trained model gives the representation of the fundamental structure of underlying Hamiltonian dynamics. Thus making the meta-trained model acquire a generalized understanding of the diverse domain of distinct physical law. During the evaluation, the model adeptly adapts to unknown dynamics by leveraging the generalized physics learned during the meta training phase. Consequently, this observation provides further support for our initial assumption that meta learning can capture general representations across a broad domain of Hamiltonian manifolds.

4.4 ABLATION ANALYSIS AND COMPARISON WITH EXISTING METHODS

First, we present the ablation studies to scrutinize the impact of both the choice of integrator and the number of data points K in the adaptation task. We experimented with symplectic integrators ranging from 1st to 4th order to ensure energy conservation compared to the LSODA integrator employed in our primary experiments. And, since we are considering scenarios where abundant data is available for known systems but sparse data for the target unknown system, we varied the test data size as $K \in (5, 10, 20, 50)$ to investigate the data efficiency of the meta-trained model relative to other baselines. The ablation results for varying integrator and K each detailed in SM C.1 and C.2, reveal several key findings. 1) LSODA integrator performs adequately for dynamics that do not involve large time scales, making it a reliable choice for most scenarios considered in this study. 2) Meta-trained model demonstrates the data efficiency when adapting to unseen tasks. This suggests that the meta-trained model is not only data-efficient but also robust, capable of adapting well when data for the target system is sparse.

Second, we demonstrate the generalization ability of the meta-trained model by not just testing the adaptation to a single-held-out system scenario as in the previous sections but also checking the adaptation performance on various types of unseen systems (i.e. performing the adaptation task to multi-held-out-system scenario). This task involved four models, each of which was initially trained for adaptation to one of the following systems; mass-spring, pendulum, Hénon-Heiles, and magnetic-mirror. We conducted an additional adaptation task using namely the 2D harmonic oscillator system and the Kepler system. We evaluated the performance of these models in terms of relative error for both predicted coordinates and energy, with detailed results presented in SM C.3. The results indicate that the meta-trained model is capable of effectively adapting to more than just a single type of unseen system, further suggesting the potential of a general model for real-world applications. This underscores the robustness and versatility of the generalized representation learned by the meta-trained model.

Third, we compared the meta-trained model to the existing baselines in domain generalization for learning the physical dynamics. Although there is none that aims to generalize across the type of dynamical systems as we have already discussed in Section 2.3, it is important to recognize that the methods from previous works are not appropriate in our broader scope of generalization. We considered CoDA (Kirchmeyer et al., 2022), which showed the best performance across prior methods for generalizing dynamics across multi-environment within a fixed system setting, and DyAd (Wang et al., 2022), which is a representative example of a non-optimization-based (model-based) meta-learning method as our baselines. The results in SM C.4, and C.5 show that the meta-trained model outperforms the CoDA and DyAd baselines for the generalization problem in cross domain settings. Moreover, note that in some cases, the existing baselines struggled to adapt.

4.5 LIMITATIONS OF OUR APPROACH

In our experiments, we treated all systems as conservative. However, data from unknown physics may not satisfy these conditions, raising concerns about its applicability to real-world scenarios. Correspondingly, we performed the adaptation task to a damped mass-spring system with a damping term $-c\dot{x}$. From the results in SM C.6, we can see that the non-conservative feature disturbs the adaptation. This is due to the damping term $-c\dot{x}$, Equation 1 is violated, which hinders the validity of Equation 2 during adaptation. Thus to make use of our approach, both the data used in meta training and adaptation should not differ in their intrinsic nature.

Furthermore, we extend our analysis to two-body and three-body systems, going beyond the scenarios on single-particle tasks, thereby opening the door to real-world applications. This extension allows us to explore the dynamics of two-body systems as a direct continuation, and three-body systems as a more challenging problem due to their chaotic behavior. However, we acknowledge that further refinement is required for the meta learning configuration for non-single-particle tasks. Here, we present the preliminary results for the two-body and three-body systems in the SM C.7.

5 DISCUSSION

While previous studies have focused on generalizing system dynamics across different environments under consistent physical laws, our work takes a step further by aiming for generalization across diverse system domains. Specifically, we tackle the challenge of identifying the shared representation of physical systems across various functional forms of Hamiltonian. In this perspective, our work presents unique contributions compared to existing literature. Here, we leverage the power of meta learning algorithm to uncover the representation of the Hamiltonian manifold within neural networks. In our comparison between the meta-trained model and other baseline models, we provide compelling evidence that the neural network learns the Hamiltonian of an unseen system by identifying the general representation for Hamiltonian dynamics itself, rather than directly constructing the functional form of the Hamiltonian of the given system. This finding highlights the ability of meta learning to capture and exploit the underlying structure of Hamiltonian mechanics. To support our claim, we conduct a CKA analysis on the trained networks. This analysis provides valuable insights into why the meta-trained model outperforms the other baselines across different physics. By examining the similarity between representations before and during the adaptation process, we gain an understanding of how the meta-trained model effectively learns and adapts to diverse physical systems. Overall, our work contributes to the field by demonstrating the efficacy of meta learning in discovering shared representations of physical systems across distinct functional forms of Hamiltonian.

In nature, each physical system operates according to its unique governing law, which is rooted in fundamental principles. In the scope of this work, systems with different functional forms of Hamiltonian are expressed with one Hamilton’s equation. Building upon this observation, we propose an approach that harnesses the power of meta learning algorithms to embed these fundamental physical principles into neural networks. Our method allows us to capture and represent these general physical principles within the neural network architecture, thereby providing a means to interpret neural networks in a physical context. A key advantage of our approach is that we imply minimal predefined physical priors or inductive biases. Instead, it leverages the data-driven capabilities of deep learning to uncover the underlying physics. This marks a significant step forward in developing a more comprehensive and general model for discovering physics through the lens of neural networks.

ACKNOWLEDGMENTS

This research was supported by the Basic Science Research Program through the National Research Foundation of Korea NRF-2022R1A2B5B02001752.

REFERENCES

- Ibrahim Alabdulmohsin, Hartmut Maennel, and Daniel Keysers. The impact of reinitialization on generalization in convolutional neural networks. *arXiv preprint arXiv:2109.00267*, 2021.
- Vladimir Igorevich Arnol'd. *Mathematical methods of classical mechanics*, volume 60. Springer Science & Business Media, 2013.
- Suresh Bishnoi, Ravinder Bhattoo, Jayadeva Jayadeva, Sayan Ranu, and N M Anoop Krishnan. Learning the dynamics of physical systems with hamiltonian graph neural networks. In *ICLR 2023 Workshop on Physics for Machine Learning*, 2023.
- James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Nectala, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pp. 1597–1607. PMLR, 2020.
- Yuhan Chen, Takashi Matsubara, and Takaharu Yaguchi. Neural symplectic form: Learning hamiltonian equations on general coordinate systems. *Advances in Neural Information Processing Systems*, 34:16659–16670, 2021.
- Miles Cranmer, Sam Greydanus, Stephan Hoyer, Peter Battaglia, David Spergel, and Shirley Ho. Lagrangian neural networks. *arXiv preprint arXiv:2003.04630*, 2020.
- Ch Efthymiopoulos, M Harsoula, and G Contopoulos. Resonant normal form and asymptotic normal form behaviour in magnetic bottle hamiltonians. *Nonlinearity*, 28(4):851, 2015.
- Charles Fefferman, Sanjoy Mitter, and Hariharan Narayanan. Testing the manifold hypothesis. *Journal of the American Mathematical Society*, 29(4):983–1049, 2016.
- Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, pp. 1126–1135. PMLR, 2017.
- Marc Finzi, Ke Alexander Wang, and Andrew G Wilson. Simplifying hamiltonian and lagrangian neural networks via explicit constraints. *Advances in neural information processing systems*, 33:13880–13889, 2020.
- Samuel Greydanus, Misko Dzamba, and Jason Yosinski. Hamiltonian neural networks. In *Advances in Neural Information Processing Systems*, volume 32, 2019.
- Samuel James Greydanus, Timothy David Strang, and Isabella Caruso. Nature's cost function: Simulating physics by minimizing the action. In *ICLR 2023 Workshop on Physics for Machine Learning*, 2023.
- Nate Gruver, Marc Anton Finzi, Samuel Don Stanton, and Andrew Gordon Wilson. Deconstructing the inductive biases of hamiltonian neural networks. In *International Conference on Learning Representations*, 2022.
- Shuyang Gu, Dong Chen, Jianmin Bao, Fang Wen, Bo Zhang, Dongdong Chen, Lu Yuan, and Baining Guo. Vector quantized diffusion model for text-to-image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10696–10706, 2022.
- Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. *Advances in neural information processing systems*, 30, 2017.

- Michel Hénon. Numerical exploration of hamiltonian systems. *Chaotic Behavior of Deterministic Systems (Les Houches, 1981)*, pp. 53–170, 1983.
- Michel Hénon and Carl Heiles. The applicability of the third integral of motion: some numerical experiments. *Astronomical Journal*, Vol. 69, p. 73 (1964), 69:73, 1964.
- Timothy Hospedales, Antreas Antoniou, Paul Micaelli, and Amos Storkey. Meta-learning in neural networks: A survey, 2020.
- Xiajun Jiang, Ryan Missel, Zhiyuan Li, and Linwei Wang. Sequential latent variable models for few-shot high-dimensional time-series forecasting. In *The Eleventh International Conference on Learning Representations*, 2022.
- Muhammad Firmansyah Kasim and Yi Heng Lim. Constants of motion network. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.), *Advances in Neural Information Processing Systems*, 2022.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.
- Matthieu Kirchmeyer, Yuan Yin, Jérémie Donà, Nicolas Baskiotis, Alain Rakotomamonjy, and Patrick Gallinari. Generalizing to new physical systems via context-informed dynamics model. *arXiv preprint arXiv:2202.01889*, 2022.
- Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey Hinton. Similarity of neural network representations revisited. In *International Conference on Machine Learning*, pp. 3519–3529. PMLR, 2019.
- Seungjun Lee, Haesang Yang, and Woojae Seong. Identifying physical law of hamiltonian systems via meta-learning. In *International Conference on Learning Representations*, 2021.
- Qiaofeng Li, Tianyi Wang, Vwani Roychowdhury, and M. K. Jawed. Metalearning generalizable dynamics from trajectories. *Phys. Rev. Lett.*, 131:067301, Aug 2023.
- Paulette Libermann and Charles-Michel Marle. *Symplectic geometry and analytical mechanics*, volume 35. Springer Science & Business Media, 2012.
- Ziming Liu and Max Tegmark. Machine learning conservation laws from trajectories. *Physical Review Letters*, 126(18):180604, 2021.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- Linda Petzold. Automatic selection of methods for solving stiff and nonstiff systems of ordinary differential equations. *SIAM journal on scientific and statistical computing*, 4(1):136–148, 1983.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 21(1):5485–5551, 2020.
- Aniruddh Raghu, Maithra Raghu, Samy Bengio, and Oriol Vinyals. Rapid learning or feature reuse? towards understanding the effectiveness of maml. In *International Conference on Learning Representations*, 2020.

- Maithra Raghu, Chiyuan Zhang, Jon Kleinberg, and Samy Bengio. Transfusion: Understanding transfer learning for medical imaging. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- Scott Reed, Zeynep Akata, Xinchun Yan, Lajanugen Logeswaran, Bernt Schiele, and Honglak Lee. Generative adversarial text to image synthesis. In *International conference on machine learning*, pp. 1060–1069. PMLR, 2016.
- Steindor Saemundsson, Alexander Terenin, Katja Hofmann, and Marc Deisenroth. Variational integrator networks for physically structured embeddings. In *International Conference on Artificial Intelligence and Statistics*, pp. 3078–3087. PMLR, 2020.
- Alvaro Sanchez-Gonzalez, Victor Bapst, Kyle Cranmer, and Peter Battaglia. Hamiltonian graph networks with ode integrators. *arXiv preprint arXiv:1909.12790*, 2019.
- Gautam Singh, Jaesik Yoon, Youngsung Son, and Sungjin Ahn. Sequential neural processes. *Advances in Neural Information Processing Systems*, 32, 2019.
- Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International conference on machine learning*, pp. 6105–6114. PMLR, 2019.
- Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017.
- Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020.
- Rui Wang, Robin Walters, and Rose Yu. Meta-learning dynamics forecasting using task inference. *Advances in Neural Information Processing Systems*, 35:21640–21653, 2022.
- Yuan Yin, Ibrahim Ayed, Emmanuel de Bézenac, Nicolas Baskiotis, and Patrick Gallinari. Leads: Learning dynamical systems that generalize across environments. *Advances in Neural Information Processing Systems*, 34:7561–7573, 2021.
- Jason Yosinski, Jeff Clune, Yoshua Bengio, and Hod Lipson. How transferable are features in deep neural networks? In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K.Q. Weinberger (eds.), *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014.