

Compositional Learning of Dynamical System Models Using Port-Hamiltonian Neural Networks

Cyrus Neary

Ufuk Topcu

The University of Texas at Austin, United States

CNEARY@UTEXAS.EDU

UTOPCU@UTEXAS.EDU

Editors: N. Matni, M. Morari, G. J. Pappas

Abstract

Many dynamical systems—from robots interacting with their surroundings to large-scale multi-physics systems—involve a number of interacting subsystems. Toward the objective of learning composite models of such systems from data, we present i) a framework for compositional neural networks, ii) algorithms to train these models, iii) a method to compose the learned models, iv) theoretical results that bound the error of the resulting composite models, and v) a method to learn the composition itself, when it is not known a priori. The end result is a modular approach to learning: neural network submodels are trained on trajectory data generated by relatively simple subsystems, and the dynamics of more complex composite systems are then predicted without requiring additional data generated by the composite systems themselves. We achieve this compositionality by representing the system of interest, as well as each of its subsystems, as a *port-Hamiltonian neural network* (PHNN)—a class of neural ordinary differential equations that uses the port-Hamiltonian systems formulation as inductive bias. We compose collections of PHNNs by using the system’s physics-informed *interconnection structure*, which may be known a priori, or may itself be learned from data. We demonstrate the novel capabilities of the proposed framework through numerical examples involving interacting spring-mass-damper systems. Models of these systems, which include nonlinear energy dissipation and control inputs, are learned independently. Accurate compositions are learned using an amount of training data that is negligible in comparison with that required to train a new model from scratch. Finally, we observe that the composite PHNNs enjoy properties of port-Hamiltonian systems, such as cyclo-passivity—a property that is useful for control purposes.

Keywords: Physics-informed machine learning, port-Hamiltonian neural networks, neural ordinary differential equation, compositional deep learning

1. Introduction

Deep learning methods that use physics-based knowledge as inductive bias have recently shown promise in learning dynamical system models that respect physical laws and that generalize beyond the training dataset (Djeumou et al., 2022a; Menda et al., 2019; Gupta et al., 2020; Cranmer et al., 2020; Greydanus et al., 2019; Finzi et al., 2020; Zhong et al., 2021a). These methods, which often use neural networks to parametrize select terms in differential operators, are able to learn complex relationships from data while also yielding models that are compact and interpretable.

However, there remain barriers to the deployment of such algorithms in engineering applications. Many systems, from robots interacting with their surroundings to large-scale multiphysics systems, involve large numbers of interacting components. These interactions between subsystems

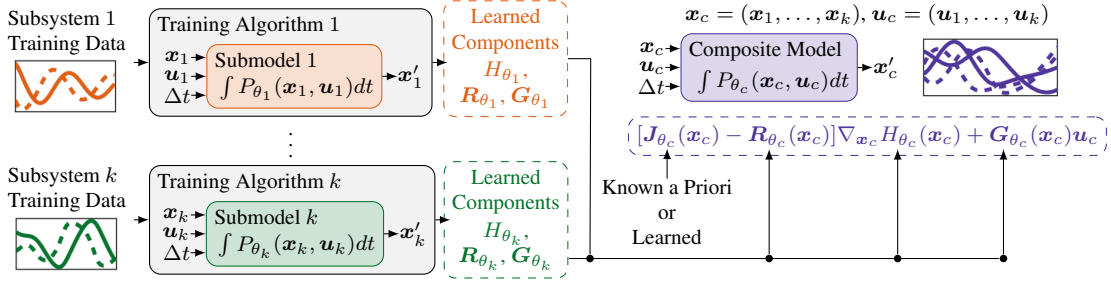


Figure 1: An illustration of the compositional learning framework. We train separate port-Hamiltonian neural networks (PHNNs) on data generated by individual subsystems, presented in §4. We then compose these submodels to construct another PHNN that models the composite system, presented in §5.

can increase the complexity of the overall system’s dynamics, rendering monolithic approaches to learning—which capture the entire system using a single model learned from data—challenging.

We present a framework and algorithms for learning and composing neural network models of dynamical systems. The framework models individual subsystems independently, and uses physics-informed interfaces between these submodels to capture their interactions. This compositional approach to learning provides a number benefits and novel capabilities that would not otherwise be possible. Firstly, it simplifies the learning problems to be solved. Submodels are trained on trajectory data generated by relatively simple subsystems. The dynamics of more complex composite systems are then predicted without requiring additional training. Secondly, it provides a modular framework for data-driven modeling. Previously learned component models can be composed in new ways to simulate different composite systems. Finally, it provides a natural way to compose data-driven models with models derived from first principles.

We achieve this compositionality by representing the system of interest, as well as each of its subsystems, as a *port-Hamiltonian neural network* (PHNN)—a class of deep learning models that use the port-Hamiltonian systems formalism (Duindam et al., 2009; Van Der Schaft et al., 2014) to inform the model’s structure. More specifically, PHNNs parametrize each subsystem’s Hamiltonian function, as well as how it dissipates energy, interacts with other subsystems, and how it responds to control inputs. We enforce known properties of the dissipation and interaction terms through the model’s construction. The PHNN’s output is obtained by numerically integrating a differential equation involving all of these terms to predict the system’s state at a future time.

Using the physics-informed structure provided by the PHNN, we present a method to compose collections of PHNNs in order to obtain models of the corresponding composite systems, and we provide upper bounds on this composite model’s prediction errors. Figure 1 illustrates the approach. The composite system’s Hamiltonian, dissipation term, and control input term are all obtained by combining the corresponding terms from the learned submodels. Interactions between the subsystems are captured by the *interconnection structure* of the composite system, which may be known a priori, or may itself be learned from data. In many cases this interconnection structure is given by a constant linear operator; we accordingly present a method to learn it via linear regression. In the general case, it may instead be parameterized using a neural network.

We demonstrate the novel capabilities of the proposed framework through numerical examples involving interacting spring-mass-damper systems. Models of these systems, which include nonlin-

ear energy dissipation and control inputs, are learned independently. The dynamics of the composite system are accurately predicted without additional training. If the system’s interconnection structure is unknown, we demonstrate that an accurate composition may be learned using an amount of training data that is negligible in comparison with that required to train a new model from scratch. Finally, we empirically observe that the proposed compositions of PHNNs exhibit the property of passivity—a property of port-Hamiltonian systems that is useful for control purposes.

2. Related Work

The port-Hamiltonian formulation of dynamical systems provides a rich mathematical framework that enables compositional modeling (Duindam et al., 2009; Van Der Schaft et al., 2014). This framework can be applied quite generally and has been used to model fluid-structure interactions (Cardoso-Ribeiro et al., 2015), aerial vehicles in contact scenarios (Rashad et al., 2022), and coupled gas and electricity distribution networks (Strehle et al., 2018). Furthermore, a wealth of existing methods and analysis for the nonlinear control of port-Hamiltonian systems already exists (Van Der Schaft, 2020). However, deriving a precise system model in port-Hamiltonian form can be challenging in many practical applications (Nagesh Rao et al., 2016; Cherifi et al., 2019). Furthermore, methods for data-driven identification and control of port-Hamiltonian systems have not yet been extensively explored (Nagesh Rao et al., 2016; Cherifi et al., 2022).

The inclusion of physics-based knowledge into neural network models of dynamical systems has, however, been studied extensively over the past several years. In particular, neural ordinary differential equations (NODEs) (Chen et al., 2018) provide a natural approach to incorporate physics-based knowledge as inductive bias in deep learning (Zhong et al., 2021a; Rackauckas et al., 2020). By using neural networks to parametrize differential equations, as opposed to directly fitting the available trajectory data, NODEs allow the user to harness an existing wealth of knowledge from applied mathematics, physics, and engineering (Djeumou et al., 2022a; Cranmer et al., 2020; Lutter et al., 2019; Gupta et al., 2020; Roehrl et al., 2020; Zhong et al., 2021b; Shi et al., 2019).

Of particular relevance to our work, *Hamiltonian neural networks* use the Hamiltonian formulation of dynamics to inform the structure of a neural ODE (Greydanus et al., 2019; Matsubara et al., 2020; Toth et al., 2020; Finzi et al., 2020). However, Hamiltonian-based models necessarily represent closed systems. By contrast, we study systems involving energy exchange, energy dissipation, and control inputs. Meanwhile, Xu et al. (2021); Furieri et al. (2022); Plaza et al. (2022) use neural networks to parametrize controllers for port-Hamiltonian systems. However, these works do not learn dynamics models—they assume the dynamics to be known and focus on control.

More closely related to our work, several recent papers also study neural ODEs that have a port-Hamiltonian structure (Zhong et al., 2020; Desai et al., 2021; Eidnes et al., 2023; Duong and Atanasov, 2021). However, none of these works study how physics-based knowledge may be used to compose deep learning models. By contrast, the primary focus of this work is to develop a framework, theoretical results, and methods that enable such compositional learning algorithms.

3. Port-Hamiltonian Systems

Port-Hamiltonian (PH) systems provide a versatile framework that enables the modeling and analysis of complex networks of interacting subsystems. Conceptually, PH systems are represented by their Hamiltonian functions, by their energy dissipation terms, and by a mathematical description




Figure 2: An illustrative coupled spring-mass-damper example. Two subsystems (left), with different spring and damping constants, are connected to obtain the composite system (right).

of the power-conserving interactions of their subsystems, called a *Dirac structure*. We refer to [Van Der Schaft et al. \(2014\)](#); [Duindam et al. \(2009\)](#) for further details.

In this work, we consider lumped parameter PH systems expressed in *explicit state-input-output* form; the system's state may be represented as a finite-dimensional vector, and its dynamics are given by equation 1. PH systems may be expressed in this form whenever there are no algebraic constraints on the system's state variables ([Donaire and Junco, 2009](#); [Dai et al., 2019](#)).

$$\dot{\mathbf{x}} = [\mathbf{J}(\mathbf{x}) - \mathbf{R}(\mathbf{x})] \nabla_{\mathbf{x}} H(\mathbf{x}) + \mathbf{G}(\mathbf{x}) \mathbf{u}, \quad \mathbf{y} = \mathbf{G}(\mathbf{x})^T \nabla_{\mathbf{x}} H(\mathbf{x}) \quad (1)$$

Here, $\mathbf{x} \in \mathbb{R}^n$ is the n -dimensional vector representing the system's state, $H(\mathbf{x})$ is the system's Hamiltonian function, $\mathbf{J}(\mathbf{x}) \in \mathbb{R}^{n \times n}$ is the skew symmetric interconnection matrix (i.e. $\mathbf{J}(\mathbf{x}) = -\mathbf{J}(\mathbf{x})^T$ for every $\mathbf{x} \in \mathbb{R}^n$), $\mathbf{R}(\mathbf{x}) \in \mathbb{R}^{n \times n}$ is the symmetric positive semi-definite dissipation matrix (i.e. $\mathbf{R}(\mathbf{x}) = \mathbf{R}(\mathbf{x})^T$ and $\tilde{\mathbf{x}}^T \mathbf{R}(\mathbf{x}) \tilde{\mathbf{x}} \geq 0$ for every $\mathbf{x}, \tilde{\mathbf{x}} \in \mathbb{R}^n$), $\mathbf{G}(\mathbf{x}) \in \mathbb{R}^{n \times m}$ is the control input matrix, $\mathbf{u} \in \mathbb{R}^m$ is the m -dimensional control input vector, and $\mathbf{y} \in \mathbb{R}^m$ is the corresponding output vector. Intuitively, while $H(\mathbf{x})$ describes the system's energy in terms of its state \mathbf{x} , $\mathbf{J}(\mathbf{x})$ encodes the energy-conserving interactions between the various components of the system and $\mathbf{R}(\mathbf{x})$ encodes how these components dissipate energy.

3.1. An Illustrative Running Example: The Coupled Mass-Spring-Damper

Let PHS_1 and PHS_2 represent the two subsystems on the left of Figure 2, and let PHS_c represent the system resulting from their composition, illustrated on the right. Similarly, let \mathbf{x}_1 , \mathbf{x}_2 , \mathbf{u}_1 , and \mathbf{u}_2 denote the subsystem states and control inputs, and let $\mathbf{x}_c := (\mathbf{x}_1, \mathbf{x}_2)$, $\mathbf{u}_c := (\mathbf{u}_1, \mathbf{u}_2)$ denote the state and control input of the composite system.

We assume that the damping forces are nonlinear $F_{damp} = b\dot{q}^3$, similar to the example presented by [Lopes et al. \(2015\)](#). We define the subsystem states as $\mathbf{x}_i = (\Delta q_i, p_i)$ for $i = 1, 2$, where Δq_i is the elongation of the spring and p_i is the momentum of the mass. The dynamics of the subsystems may then be represented in the form of equation 1 with $H_i(\mathbf{x}_i) = \frac{p_i^2}{2m_i} + \frac{k\Delta q_i^2}{2}$ and with

$$\mathbf{J}_i = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}, \quad \mathbf{R}_i(\mathbf{x}_i) = \begin{bmatrix} 0 & 0 \\ 0 & b_i \frac{p_i^2}{m_i^2} \end{bmatrix}, \quad \mathbf{G}_i = \begin{bmatrix} 0 \\ 1 \end{bmatrix}. \quad (2)$$

The port-Hamiltonian system PHS_c representing the system on the right of Figure 2 may be obtained by *composing* PHS_1 and PHS_2 , the port-Hamiltonian representations of the subsystems. That is, the dynamics of PHS_c may also be written in the form of equation 1, where the composite Hamiltonian is given by $H_c(\mathbf{x}_c) = H_1(\mathbf{x}_1) + H_2(\mathbf{x}_2)$, the dissipation $\mathbf{R}_c(\mathbf{x}_c)$ and control input

\mathbf{G}_c matrices are obtained by stacking the matrices $\mathbf{R}_i(\mathbf{x}_i)$ and \mathbf{G}_i diagonally, and the composite interconnection term \mathbf{J}_c is obtained by stacking \mathbf{J}_i diagonally and by including an additional pair of off-diagonal interaction terms. These additional entries in \mathbf{J}_c encode the coupling between the subsystems; spring 2 exerts a force of $\frac{\partial H_2}{\partial \Delta q_2}$ on mass 1, and mass 1’s contribution to the rate of change in the elongation of spring 2 is given by $-\frac{\partial H_1}{\partial p_1}$. The composite system’s dynamics are given explicitly in the extended version of the paper (Neary and Topcu, 2022).

4. Port-Hamiltonian Neural Networks

Our objective is to train port-Hamiltonian neural networks (PHNN) to predict the dynamics of relatively simple subsystems (e.g., the individual spring-mass-damper systems from the left of Figure 2) and then to compose these learned models in order to simulate more complex systems (e.g., the coupled spring-mass-damper system on the right of Figure 2). In this section we present how to construct and train PHNN submodels. We present a method to compose the learned models in §5.

4.1. Constructing Port-Hamiltonian Neural Networks

A PHNN parametrizes the unknown terms in equation 1, and solves the resulting differential equation in order to predict the system’s future states. Let $P(\mathbf{x}, \mathbf{u})$ denote the right-hand side of the state equation in 1. We use $H_\theta(\mathbf{x})$, $\mathbf{R}_\theta(\mathbf{x})$, and $\mathbf{G}_\theta(\mathbf{x})$ to denote the parametrizations of the potentially unknown terms in $P(\mathbf{x}, \mathbf{u})$. We use $P_\Theta(\mathbf{x}, \mathbf{u})$ to denote the expression that results when each of the unknown terms is replaced with its parametrization, where Θ denotes the collection of all the individual parameter vectors.

Figure 3 illustrates the process of constructing and evaluating a PHNN. The Hamiltonian $H_\theta(\mathbf{x})$ is parametrized as a multi-layer perceptron (MLP) and its gradient $\nabla_{\mathbf{x}} H_\theta(\mathbf{x})$ with respect to \mathbf{x} is computed using automatic differentiation. The entries in the input matrix $\mathbf{G}_\theta(\mathbf{x}) \in \mathbb{R}^{n \times m}$ are either parametrized directly (if \mathbf{G} is known to be a constant matrix), or as the output of an MLP (if $\mathbf{G}(\mathbf{x})$ varies with \mathbf{x}). The dissipation term $\mathbf{R}_\theta(\mathbf{x})$ is similarly either parametrized as a constant matrix or as the output of an MLP. However, we additionally enforce the positive semi-definiteness of $\mathbf{R}_\theta(\mathbf{x})$ by parametrizing its Cholesky decomposition, instead of parametrizing its entries directly. That is, similarly to as proposed by Zhong et al. (2020), we define $\mathbf{R}_\theta(\mathbf{x}) := \mathbf{L}_\theta(\mathbf{x})\mathbf{L}_\theta(\mathbf{x})^T$ for some parametrized lower-triangular matrix $\mathbf{L}_\theta(\mathbf{x}) \in \mathbb{R}^{n \times n}$ with non-negative diagonal entries.

Note that in this section, similar to all existing works involving port-Hamiltonian neural networks (Greydanus et al., 2019; Zhong et al., 2020; Desai et al., 2021; Eidnes et al., 2023), we assume that the interaction term $\mathbf{J}(\mathbf{x})$ is known a priori. It is possible to learn a skew-symmetric parametrization of $\mathbf{J}(\mathbf{x})$ along with all of the other terms in the PHNN. However, doing so necessitates additional considerations that are beyond the scope of the current work, when composing the resulting models.

PHNNs enforce the cyclo-passivity property by construction. The *cyclo-passivity* property enjoyed by port-Hamiltonian systems ensures that $dH/dt \leq \mathbf{y}^T \mathbf{u}$ —the rate of change of the system’s energy cannot exceed the externally-provided power (Van Der Schaft, 2000). Regardless of the output of the parametrized Hamiltonian $H_\theta(\mathbf{x})$, this property is guaranteed by the skew-symmetry of $\mathbf{J}(\mathbf{x})$ together with the positive semi-definiteness of the dissipation term $\mathbf{R}_\theta(\mathbf{x})$, which we enforce through the PHNN’s construction. When dissipation and control inputs are both present, $dH/dt = h^T \mathbf{J}(\mathbf{x})h - h^T \mathbf{R}_\theta(\mathbf{x})h + h^T \mathbf{G}_\theta(\mathbf{x})\mathbf{u} \leq \mathbf{y}^T \mathbf{u}$. Here, we use h to denote $\nabla_{\mathbf{x}} H_\theta(\mathbf{x})$.




Figure 3: An illustration of the construction and evaluation of a port-Hamiltonian Neural Network.

4.2. Evaluating Port-Hamiltonian Neural Networks

The input to a PHNN is a tuple $(\mathbf{x}, \mathbf{u}, t, T)$ consisting of the current state \mathbf{x} , a control input \mathbf{u} , the current time t , and a prediction time T with $t < T$. The output $\hat{\mathbf{x}}_T := PHNN_{\Theta}(\mathbf{x}, \mathbf{u}, t, T)$ of the PHNN parametrized by Θ is then given by $ODESolve(P_{\Theta}(\cdot, \mathbf{u}), \mathbf{x}, t, T) \approx \mathbf{x} + \int_t^T P_{\Theta}(\mathbf{x}_s, \mathbf{u}) ds$, where we use the subscript notation \mathbf{x}_t to denote $\mathbf{x}(t)$. Here, $ODESolve(P_{\Theta}(\cdot, \mathbf{u}), \mathbf{x}, t, T)$ is a numerical solution to the ordinary differential equation specified by $P_{\Theta}(\mathbf{x}, \mathbf{u})$ over the window of time $[t, T]$. We note that the particular algorithm used to evaluate $ODESolve(\cdot)$ influences the model’s accuracy and the computational cost of forward evaluations of the model (Djeumou et al., 2022b). In this work, we use fixed-timestep RK4 to evaluate $ODESolve(\cdot)$ in all experiments.

4.3. Training Port-Hamiltonian Neural Networks

We assume that a finite dataset \mathcal{D} of system trajectories—time-series data of states and control inputs—is available in lieu of the system’s model. That is, we are given a set $\mathcal{D} = \{\tau_1, \dots, \tau_{|\mathcal{D}|}\}$ of trajectories $\tau = \{(\mathbf{x}_0, \mathbf{u}_0, t_0), \dots, (\mathbf{x}_{|\tau|}, \mathbf{u}_{|\tau|}, t_{|\tau|})\}$, where \mathbf{x}_i is the state at time t_i , \mathbf{u}_i is the control input applied from time t_i until t_{i+1} , and $t_0 < t_1 < \dots < t_{|\tau|}$ is an increasing sequence of points in time. Note that we are using $P_{\Theta}(\cdot)$ to parametrize the time derivative $\dot{\mathbf{x}}$ of the system’s state, which is not explicitly included in the dataset.

Given a dataset \mathcal{D} , we construct the loss function $\mathcal{L}(\Theta, \mathcal{D})$ of the PHNN to capture the error in the model’s predictions of the future states, for any given norm, as

$$\mathcal{L}(\Theta, \mathcal{D}) = \sum_{\tau_i \in \mathcal{D}} \sum_{(\mathbf{x}_i, \mathbf{u}_i, t_i) \in \tau_i} \|PHNN_{\Theta}(\mathbf{x}_i, \mathbf{u}_i, t_i, t_{i+1}) - \mathbf{x}_{i+1}\|^2. \quad (3)$$

Finally, we search for local minima of $\mathcal{L}(\Theta, \mathcal{D})$ using gradient-based techniques, where $\nabla_{\Theta} \mathcal{L}(\Theta, \mathcal{D})$ may be computed using either direct automatic differentiation through $ODESolve(\cdot)$, or using the adjoint sensitivity method (Pontryagin, 1987; Chen et al., 2018).

5. Composing Port-Hamiltonian Neural Networks

While in §4 we introduced PHNNs and methods to train them, in this section we present a method to compose previously learned PHNNs in order to predict the dynamics of larger composite systems. In the context of the example from §3.1, we use the methods from §4 to learn PHNN models of the individual spring-mass-dampers, and we use the methods in this section to compose these models in order to predict the dynamics of the coupled system.

5.1. Composing Models Using a Known Interconnection Structure

We construct a model of the composite system $PHNN_{c,\Theta_c}^C(\cdot)$ by combining the terms of the learned submodels $PHNN_{i,\Theta}(\cdot)$, for $i = 1, \dots, k$, as illustrated in Figure 1. Let $\mathbf{x}_i \in \mathbb{R}^{n_i}$ and $\mathbf{u}_i \in \mathbb{R}^{m_i}$ represent the state and control input vectors of subsystem i . We define the state and control input vectors of the composite system to be $\mathbf{x}_c := [\mathbf{x}_1^T, \dots, \mathbf{x}_k^T]^T \in \mathbb{R}^{n_c}$ and $\mathbf{u}_c := [\mathbf{u}_1^T, \dots, \mathbf{u}_k^T]^T \in \mathbb{R}^{m_c}$, where $n_c = n_1 + \dots + n_k$ and $m_c = m_1 + \dots + m_k$.

The Hamiltonian of the composite system is defined as the sum of the subsystem Hamiltonians, $H_{c,\theta}(\mathbf{x}_c) := H_{1,\theta}(\mathbf{x}_1) + \dots + H_{k,\theta}(\mathbf{x}_k)$. Note that we assume that each subsystem’s Hamiltonian function $H_i(\mathbf{x}_i)$ depends only on the corresponding state vector \mathbf{x}_i . The dissipation and control input terms are defined as $\mathbf{R}_{c,\theta}(\mathbf{x}_c) := \text{Diag}(\mathbf{R}_{1,\theta}(\mathbf{x}_1), \dots, \mathbf{R}_{k,\theta}(\mathbf{x}_k))$ and $\mathbf{G}_{c,\theta}(\mathbf{x}_c) := \text{Diag}(\mathbf{G}_{1,\theta}(\mathbf{x}_1), \dots, \mathbf{G}_{k,\theta}(\mathbf{x}_k))$, respectively. Here, we use $\text{Diag}(\mathbf{A}_1, \dots, \mathbf{A}_k)$ to represent the matrix that results from using the submatrices $\mathbf{A}_1, \dots, \mathbf{A}_k$ to define the blocks of entries along the matrix diagonal, similarly to as in the example from §3.1. Finally, the interconnection term of the composition is given by $\mathbf{J}_c(\mathbf{x}_c) := \text{Diag}(\mathbf{J}_1(\mathbf{x}_1), \dots, \mathbf{J}_k(\mathbf{x}_k)) + \mathbf{C}(\mathbf{x}_c)$, where $\mathbf{C}(\mathbf{x}_c) \in \mathbb{R}^{n_c \times n_c}$ is a skew symmetric *composition* matrix encoding energy-conserving interactions between the state variables of the various subsystems. Furthermore, we define the blocks of entries along the diagonal of $\mathbf{C}(\mathbf{x}_c)$ to be zero—the internal subsystem interactions defined by $\mathbf{J}_i(\mathbf{x}_i)$ should not be altered by $\mathbf{C}(\mathbf{x}_c)$. We note that the interconnection and dissipation terms of the composite model retain their properties of skew-symmetry and positive semidefiniteness, respectively.

5.2. Learning Compositions of Port-Hamiltonian Neural Network Submodels

If the composition term $\mathbf{C}(\mathbf{x}_c)$ is unknown, we propose to learn it using additional data gathered from the composite system. That is, in addition to the datasets $\mathcal{D}_1, \dots, \mathcal{D}_k$ used to train the subsystem models, we assume access to a dataset $\mathcal{D}_c = \{\tau_1^c, \dots, \tau_{|\mathcal{D}_c|}^c\}$ of trajectories $\tau^c = \{(\mathbf{x}_{c,0}, \mathbf{u}_{c,0}, t_0), \dots, (\mathbf{x}_{c,|\tau|}, \mathbf{u}_{c,|\tau|}, t_{|\tau|})\}$, where $\mathbf{x}_{c,i}$ and $\mathbf{u}_{c,i}$ represent the composite state and control input at time t_i , respectively.

Let $\mathbf{C}_\phi(\mathbf{x}_c)$ denote a parametric model of the unknown composition matrix, defined in terms of the parameter vector ϕ . Given a collection of pre-trained submodels $PHNN_{i,\Theta}(\cdot)$ for $i = 1, \dots, k$, the composite model $PHNN_{c,\Theta_c}^{C_\phi}(\cdot)$ is constructed as defined in §5.1, with the exception of the parametrized composition term $\mathbf{C}_\phi(\mathbf{x}_c)$ being used in place of its ground-truth counterpart.

To learn the parameters ϕ using the dataset \mathcal{D}_c , we define a loss function $\mathcal{L}^{comp}(\phi, \Theta_c, \mathcal{D}_c)$ similarly to as in equation 3, where Θ_c denotes the collection of all of the parameter vectors Θ_i from the subsystem PHNNs. In general, $\mathbf{C}_\phi(\mathbf{x}_c)$ will be a function of \mathbf{x}_c and we parametrize its entries as the output of a neural network. $\mathcal{L}^{comp}(\phi, \Theta_c, \mathcal{D}_c)$ can then be minimized by fixing the pre-trained values of Θ_c and performing gradient descent over ϕ .

However, when the interconnection term $\mathbf{J}_c(\mathbf{x}_c)$ is a constant matrix, \mathbf{C}_ϕ may be parameterized as constant skew symmetric matrix. In such scenarios, so long as the timestep $t_{i+1} - t_i$ between the recorded datapoints is sufficiently small, each datapoint yields the following collection of linear equations in the unknown entries of \mathbf{C}_ϕ .

$$\frac{\mathbf{x}_{i+1} - \mathbf{x}_i}{t_{i+1} - t_i} - \left[\tilde{\mathbf{J}}_c - \mathbf{R}_{c,\theta}(\mathbf{x}_{c,i}) \right] \nabla_{\mathbf{x}_c} H_{c,\theta}(\mathbf{x}_{c,i}) + \mathbf{G}_{c,\theta}(\mathbf{x}_{c,i}) \mathbf{u}_{c,i} \approx \mathbf{C}_\phi \nabla_{\mathbf{x}_c} H_{c,\theta}(\mathbf{x}_{c,i}), \quad (4)$$

where we use $\tilde{\mathbf{J}}_c$ as shorthand for $\text{Diag}(\mathbf{J}_1, \dots, \mathbf{J}_k)$. Dataset \mathcal{D}_c thus yields a least squares problem that we solve to obtain the entries of \mathbf{C}_ϕ .

5.3. Bounding the Errors of Compositions of Port-Hamiltonian Neural Networks

Theorem 1 *Suppose the true dynamics of each subsystem may be written in port-Hamiltonian form as $P_i(\mathbf{x}_i, \mathbf{u}_i)$ for $i = 1, 2, \dots, k$. Furthermore, suppose the composite system of interest may be represented as a composition $P_c^C(\mathbf{x}_c, \mathbf{u}_c)$ of the port-Hamiltonian subsystems defined by the composition term $C(\mathbf{x}_c)$. Let $C_\phi(\mathbf{x}_c)$ denote the learned composition term and let $C_\phi(\mathbf{x}_c)_{i,j}$ denote the submatrix that defines the interactions between subsystems i and j . Suppose that for every $i, j = 1, \dots, k$ with $i \neq j$, we have $\|P_i(\mathbf{x}_i, \mathbf{u}_i) - P_{i,\Theta}(\mathbf{x}_i, \mathbf{u}_i)\| \leq \varepsilon_i$ and $\|\nabla_{\mathbf{x}_i} H_i(\mathbf{x}_i) - \nabla_{\mathbf{x}_i} H_{i,\theta}(\mathbf{x}_i)\| \leq \eta_i$ for all $\mathbf{x}_i \in \Omega_{\mathbf{x}_i}$, $\mathbf{u}_i \in \Omega_{\mathbf{u}_i}$. Also suppose that $\|(C(\mathbf{x}_c)_{i,j} - C_\phi(\mathbf{x}_c)_{i,j}) \nabla_{\mathbf{x}_j} H_j(\mathbf{x}_j)\| \leq \gamma_{i,j}$ for all $\mathbf{x}_j \in \Omega_{\mathbf{x}_j}$ that are consistent with some composite state $\mathbf{x}_c \in \Omega_{\mathbf{x}_1} \times \dots \times \Omega_{\mathbf{x}_k}$. Then,*

$$\left\| P_c^C(\mathbf{x}_c, \mathbf{u}_c) - P_{c,\Theta_c}^{C_\phi}(\mathbf{x}_c, \mathbf{u}_c) \right\| \leq \sum_{i=1}^k [\varepsilon_i + 2 \sum_{j>i}^k (\gamma_{i,j} + \sigma_{i,j} \eta_j)], \quad (5)$$

for every $\mathbf{x}_c \in \Omega_{\mathbf{x}_1} \times \dots \times \Omega_{\mathbf{x}_k}$ and $\mathbf{u}_c \in \Omega_{\mathbf{u}_1} \times \dots \times \Omega_{\mathbf{u}_k}$. Here, $\sigma_{i,j} := \max_{\mathbf{x}_c \in \Omega_{\mathbf{x}_c}} \|C_\phi(\mathbf{x}_c)_{i,j}\|$ where $\|C_\phi(\mathbf{x}_c)_{i,j}\| := \sup\{\|C_\phi(\mathbf{x}_c)_{i,j} y\| \text{ s.t. } y \in \mathbb{R}^{n_j}, \|y\| = 1\}$ is the matrix norm of $C_\phi(\mathbf{x}_c)_{i,j}$.

In words, Theorem 1 tells us that the prediction error of the composite PHNN is bounded by the error introduced by its component models and the error introduced by the influence that its component models have on each other. The latter is a function of the error in our estimate of the composition term $C_\phi(\mathbf{x})$ and in the gradients of the subsystem Hamiltonian functions $\nabla_{\mathbf{x}_i} H_{i,\theta}(\mathbf{x}_i)$. We note that when $C(\mathbf{x})$ is known a priori this result still holds with all $\gamma_{i,j} = 0$. This result ensures that any improvements to the accuracy of the submodels during training will improve the prediction accuracy of the composite model as well. The proof is provided in (Neary and Topcu, 2022).

6. Numerical Experiments

For detailed descriptions of the dataset generation, the employed neural network architectures, and the training algorithms used to generate the results, we refer the reader to (Neary and Topcu, 2022). Project code is available at github.com/cyrusneary/compositional_port_hamiltonian_NNs.

6.1. Composing Port-Hamiltonian Neural Networks Using a Known Composition Term

We begin by considering a numerical simulation of the coupled spring-mass-dampers illustrated in Figure 2. Recall that this example includes external control forces and nonlinear dissipation.

We first consider the scenario in which the composition term C is known a priori. We learn the subsystem models $PHNN_{1,\Theta}(\cdot)$ and $PHNN_{2,\Theta}(\cdot)$ using separate datasets \mathcal{D}_1 and \mathcal{D}_2 , each of which contains 100 trajectories with randomly sampled initial states. Throughout training we occasionally freeze the network parameters and compose the resulting submodels to obtain a composite model $PHNN_{c,\Theta}^C(\cdot)$.

Figure 4 illustrates the result of evaluating the prediction error of both submodels, and of the composite model, on separate testing datasets. The prediction loss measures the average Euclidean distance between the model-predicted future states and the true next states. We plot the median loss values over 10 independent experimental runs, while the shaded regions enclose the 25th and 75th percentiles. Each testing dataset contains 20 system trajectories beginning from randomly sampled initial states, and with different control inputs than the training datasets.

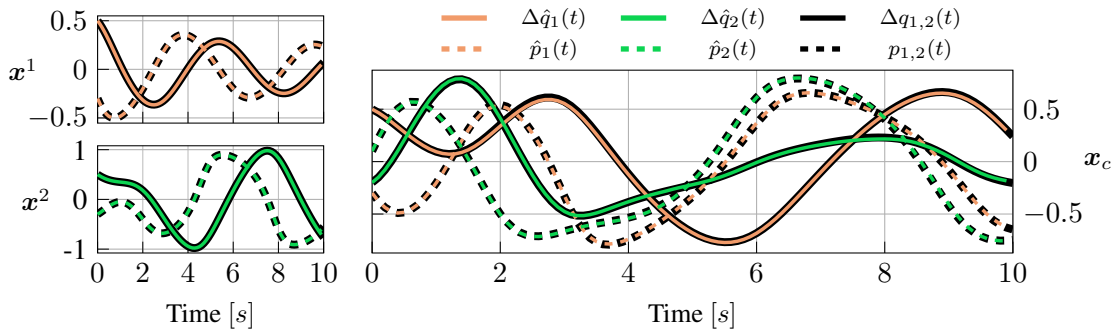


Figure 5: The composite PHNN accurately predicts state trajectories, even when the unknown composition term \mathcal{C}_ϕ is inferred using only four transition datapoints. Left: Predicted subsystem dynamics. Right: Predicted composite system dynamics.

The composite model accurately predicts dynamics without data from the composite system. We emphasize that the composite model’s loss values (blue) in Figure 4 are not the result of training a separate model using composite system data. Instead, the figure shows the result of learning the subsystem models independently (orange and green), and using a physics-informed composition of the submodels to accurately predict the dynamics of a more complex composite system for which we have no data. We additionally observe from Figure 4, that as the prediction losses of the submodels decrease during training, the loss of the composite model decreases as well. This observation empirically demonstrates Theorem 1—the composite model’s error should decrease with the error of the subsystem PHNNs.




Figure 4: Prediction loss values throughout training. The composite model makes accurate predictions without any training data from the composite system itself.

Finally, we note that training a baseline model for comparison in this example is not possible; the idea of composing deep learning submodels of dynamical systems is novel, and no monolithic baseline model could possibly learn to predict the composite dynamics without training data.

6.2. Learning the Interactions Between Subsystem Port-Hamiltonian Neural Networks

We now proceed to the setting in which the composition term \mathcal{C} is unknown—we do not have a priori knowledge of how the subsystems are influencing each other. Instead, as described in §5.2, we assume access to some small dataset of \mathcal{D}_c of observations of the dynamics of the composite system. Specifically, \mathcal{D}_c contains only four datapoints, each of which corresponds to a single-timestep transition beginning from a randomly sampled composite state. Using these datapoints, along with the pretrained submodels $PHNN_{1,\Theta}(\cdot)$ and $PHNN_{2,\Theta}(\cdot)$, we learn \mathcal{C}_ϕ by solving the least squares problem described in §5.2.

We accurately learn the unknown composition term using a negligible amount of data.

Figure 5 illustrates state trajectories predicted by the subsystem, and composite, PHNNs. In all of the subplots, the true dynamics (black) are matched very closely by the PHNN-predicted dynamics (orange and green); we accurately learn the composition term C_ϕ using only four datapoints from the composite system. Monolithic approaches cannot learn effective models from such a limited dataset.




Figure 6: Time derivative of the system’s Hamiltonian with no external control inputs.

The learned composite model enjoys properties of port-Hamiltonian systems. Figure 6 illustrates the time derivative of the learned composite Hamiltonian along a predicted trajectory without external control inputs. We note that for the entire trajectory, $dH_{c,\theta}/dt \leq 0$, which provides an empirical demonstration of the cyclo-passivity of the composite PHNN discussed in §4.

6.3. Predicting the Dynamics of Ten Interacting Subsystems Without Additional Training

We now consider ten interacting spring-mass-damper systems, most of which are identical to subsystem 1 from Figure 2, while the rest match subsystem 2. A sinusoidal force acts on one end of the chain. Figure 7 illustrates the results of composing the previously trained PHNN submodels to predict this system’s dynamics. The predicted state trajectory of each subsystem is plotted as a separate colored curve.




Figure 7: The composite PHNN accurately predicts the dynamics of 10 interacting subsystems with no additional training.

The composite PHNN accurately predicts the wave-like propagation of energy between the subsystems (the colored PHNN predictions overlay the true dynamics in black). We emphasize that this behavior is captured without any additional training of the subsystem models, and without access to any data from the ten-component system; PHNNs can be composed in a modular fashion to simulate entirely new complex systems.

7. Conclusions

In this work we present a framework, algorithms, and theoretical results for the compositional learning of dynamical system models via port-Hamiltonian neural networks (PHNNs). This work presents a first step towards learning modular neural network parametrizations of control systems that can be trained and tested independently, and that can be re-used in new contexts. We demonstrate that by using the structure of port-Hamiltonian systems as inductive bias, we may independently learn submodels on data generated by relatively simple subsystems, and then accurately predict the dynamics of more complex composite systems while using little to no data from the composite system itself. Future work will aim to learn compositional PHNNs from video observations, and also to use the compositional models for control.

Acknowledgments

This work was supported in part by AFOSR FA9550-19-1-0005, ARO W911NF-20-1-0140, and NSF 2214939.

References

- Flávio Luiz Cardoso-Ribeiro, Denis Matignon, and Valérie Pommier-Budinger. Modeling of a fluid-structure coupled system using port-hamiltonian formulation. *IFAC-PapersOnLine*, 48(13): 217–222, 2015.
- Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. In *Advances in Neural Information Processing Systems*, volume 31, 2018.
- Karim Cherifi, Volker Mehrmann, and Kamel Hariche. Numerical methods to compute a minimal realization of a port-hamiltonian system. *arXiv:1903.07042*, 2019.
- Karim Cherifi, Pawan Kumar Goyal, and Peter Benner. A non-intrusive method to inferring linear port-hamiltonian realizations using time-domain data. *Electronic Transactions on Numerical Analysis: Special Issue SciML*, 56:102–116, 2022.
- Miles Cranmer, Sam Greydanus, Stephan Hoyer, Peter Battaglia, David Spergel, and Shirley Ho. Lagrangian neural networks. In *Workshop on Integration of Deep Neural Models and Differential Equations, International Conference on Learning Representations*, 2020.
- Siyuan Dai, Zhenkai Zhang, and Xenofon Koutsoukos. A model-based design approach for simulation and virtual prototyping of automotive control systems using port-hamiltonian systems. *Software & Systems Modeling*, 18(3):1637–1653, 2019.
- Shaan A Desai, Marios Mattheakis, David Sondak, Pavlos Protopapas, and Stephen J Roberts. Port-hamiltonian neural networks for learning explicit time-dependent dynamical systems. *Physical Review E*, 104(3):034312, 2021.
- Franck Djeumou, Cyrus Neary, Eric Goubault, Sylvie Putot, and Ufuk Topcu. Neural networks with physics-informed architectures and constraints for dynamical systems modeling. In *Proceedings of The 4th Annual Learning for Dynamics and Control Conference*, volume 168 of *Proceedings of Machine Learning Research*, 2022a.
- Franck Djeumou, Cyrus Neary, Eric Goubault, Sylvie Putot, and Ufuk Topcu. Taylor-lagrange neural ordinary differential equations: Toward fast training and evaluation of neural odes. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence*, 2022b.
- Alejandro Donaire and Sergio Junco. Derivation of input-state-output port-hamiltonian systems from bond graphs. *Simulation Modelling Practice and Theory*, 17(1):137–151, 2009.
- Vincent Duindam, Alessandro Macchelli, Stefano Stramigioli, and Herman Bruyninckx. *Modeling and control of complex physical systems: the port-Hamiltonian approach*. Springer Science & Business Media, 2009.

- Thai Duong and Nikolay Atanasov. Hamiltonian-based neural ode networks on the se (3) manifold for dynamics learning and control. In *Proceedings of Robotics: Science and Systems*, 2021.
- Sølve Eidnes, Alexander J. Stasik, Camilla Sterud, Eivind Bøhn, and Signe Riemer-Sørensen. Pseudo-hamiltonian neural networks with state-dependent external forces. *Physica D: Nonlinear Phenomena*, 446:133673, 2023.
- Marc Finzi, Ke Alexander Wang, and Andrew G Wilson. Simplifying hamiltonian and lagrangian neural networks via explicit constraints. In *Advances in Neural Information Processing Systems*, volume 33, 2020.
- Luca Furieri, Clara Lucía Galimberti, Muhammad Zakwan, and Giancarlo Ferrari-Trecate. Distributed neural network control with dependability guarantees: a compositional port-hamiltonian approach. In *Proceedings of The 4th Annual Learning for Dynamics and Control Conference*, volume 168 of *Proceedings of Machine Learning Research*, 2022.
- Samuel Greydanus, Misko Dzamba, and Jason Yosinski. Hamiltonian neural networks. In *Advances in Neural Information Processing Systems*, volume 32, 2019.
- Jayesh K. Gupta, Kunal Menda, Zachary Manchester, and Mykel Kochenderfer. Structured mechanical models for robot learning and control. In *Proceedings of the 2nd Conference on Learning for Dynamics and Control*, volume 120 of *Proceedings of Machine Learning Research*, 2020.
- Nicolas Lopes, Thomas Hélie, and Antoine Falaize. Explicit second-order accurate method for the passive guaranteed simulation of port-hamiltonian systems. *IFAC-PapersOnLine*, 48(13):223–228, 2015.
- Michael Lutter, Christian Ritter, and Jan Peters. Deep lagrangian networks: Using physics as model prior for deep learning. In *International Conference on Learning Representations*, 2019.
- Takashi Matsubara, Ai Ishikawa, and Takaharu Yaguchi. Deep energy-based modeling of discrete-time physics. In *Advances in Neural Information Processing Systems*, volume 33, 2020.
- Kunal Menda, Jayesh K Gupta, Zachary Manchester, and Mykel J Kochenderfer. Structured mechanical models for efficient reinforcement learning. In *Workshop on Structure and Priors in Reinforcement Learning, International Conference on Learning Representations*, 2019.
- Subramanya P. Nagesh Rao, Gabriel A. D. Lopes, Dimitri Jeltsema, and Robert Babuška. Port-hamiltonian systems in adaptive and learning control: A survey. *IEEE Transactions on Automatic Control*, 61(5):1223–1238, 2016.
- Cyrus Neary and Ufuk Topcu. Compositional learning of dynamical system models using port-hamiltonian neural networks. *arXiv:2212.00893*, 2022.
- Santiago Sanchez-Escalonilla Plaza, Rodolfo Reyes-Báez, and Bayu Jayawardhana. Total energy shaping with neural interconnection and damping assignment - passivity based control. In *Proceedings of The 4th Annual Learning for Dynamics and Control Conference*, volume 168 of *Proceedings of Machine Learning Research*, 2022.
- Lev Semenovich Pontryagin. *Mathematical theory of optimal processes*. CRC Press, 1987.

- Christopher Rackauckas, Yingbo Ma, Julius Martensen, Collin Warner, Kirill Zubov, Rohit Sudeepkar, Dominic Skinner, Ali Ramadhan, and Alan Edelman. Universal differential equations for scientific machine learning. *arXiv:2001.04385*, 2020.
- Ramy Rashad, Davide Bicego, Jelle Zult, Santiago Sanchez-Escalonilla, Ran Jiao, Antonio Franchi, and Stefano Stramigioli. Energy aware impedance control of a flying end-effector in the port-hamiltonian framework. *IEEE Transactions on Robotics*, 38(6):3936–3955, 2022.
- Manuel A Roehrl, Thomas A Runkler, Veronika Brandtstetter, Michel Tokic, and Stefan Obermayer. Modeling system dynamics with physics-informed neural networks based on lagrangian mechanics. *IFAC-PapersOnLine*, 53(2):9195–9200, 2020.
- Guanya Shi, Xichen Shi, Michael O’Connell, Rose Yu, Kamyar Azizzadenesheli, Animashree Anandkumar, Yisong Yue, and Soon-Jo Chung. Neural lander: Stable drone landing control using learned dynamics. In *International Conference on Robotics and Automation*, 2019.
- Felix Strehle, Martin Pfeifer, Lukas Kölsch, Charlotte Degünther, Johannes Ruf, Lisa Andresen, and Sören Hohmann. Towards port-hamiltonian modeling of multi-carrier energy systems: A case study for a coupled electricity and gas distribution system. *IFAC-PapersOnLine*, 51(2):463–468, 2018.
- Peter Toth, Danilo J. Rezende, Andrew Jaegle, Sébastien Racanière, Aleksandar Botev, and Irina Higgins. Hamiltonian generative networks. In *International Conference on Learning Representations*, 2020.
- Arjan Van Der Schaft. *L2-gain and passivity techniques in nonlinear control*. Springer, 2000.
- Arjan Van Der Schaft. Port-hamiltonian modeling for control. *Annual Review of Control, Robotics, and Autonomous Systems*, 3(1):393–416, 2020.
- Arjan Van Der Schaft, Dimitri Jeltsema, et al. Port-hamiltonian systems theory: An introductory overview. *Foundations and Trends in Systems and Control*, 1(2-3):173–378, 2014.
- Liang Xu, Muhammad Zakwan, and Giancarlo Ferrari-Trecate. Neural energy casimir control. *arXiv:2112.03339*, 2021.
- Yaofeng Desmond Zhong, Biswadip Dey, and Amit Chakraborty. Dissipative symODEN: Encoding hamiltonian dynamics with dissipation and control into deep learning. In *Workshop on Integration of Deep Neural Models and Differential Equations, International Conference on Learning Representations*, 2020.
- Yaofeng Desmond Zhong, Biswadip Dey, and Amit Chakraborty. Benchmarking energy-conserving neural networks for learning dynamics from data. In *Proceedings of the 3rd Conference on Learning for Dynamics and Control*, volume 144 of *Proceedings of Machine Learning Research*, 2021a.
- Yaofeng Desmond Zhong, Biswadip Dey, and Amit Chakraborty. Extending lagrangian and hamiltonian neural networks with differentiable contact models. In *Advances in Neural Information Processing Systems*, volume 34, 2021b.