

Neural Networks with Physics-Informed Architectures and Constraints for Dynamical Systems Modeling

Franck Djeumou^{*1}

Cyrus Neary^{*1}

Eric Goubault²

Sylvie Putot²

Ufuk Topcu¹

FDJEUMOU@UTEXAS.EDU

CNEARY@UTEXAS.EDU

GOUBAULT@LIX.POLYTECHNIQUE.FR

PUTOT@LIX.POLYTECHNIQUE.FR

UTOPCU@UTEXAS.EDU

¹ *The University of Texas at Austin, United States*

² *LIX, CNRS, École Polytechnique, Institut Polytechnique de Paris, France*

Editors: R. Firoozi, N. Mehr, E. Yel, R. Antonova, J. Bohg, M. Schwager, M. Kochenderfer

Abstract

Effective inclusion of physics-based knowledge into deep neural network models of dynamical systems can greatly improve data efficiency and generalization. Such a priori knowledge might arise from physical principles (e.g., conservation laws) or from the system’s design (e.g., the Jacobian matrix of a robot), even if large portions of the system dynamics remain unknown. We develop a framework to learn dynamics models from trajectory data while incorporating a priori system knowledge as inductive bias. More specifically, the proposed framework uses physics-based side information to inform the structure of the neural network itself, and to place constraints on the values of the outputs and the internal states of the model. It represents the system’s vector field as a composition of known and unknown functions, the latter of which are parametrized by neural networks. The physics-informed constraints are enforced via the augmented Lagrangian method during the model’s training. We experimentally demonstrate the benefits of the proposed approach on a variety of dynamical systems – including a benchmark suite of robotics environments featuring large state spaces, non-linear dynamics, external forces, contact forces, and control inputs. By exploiting a priori system knowledge during training, the proposed approach learns to predict the system dynamics two orders of magnitude more accurately than a baseline approach that does not include prior knowledge, given the same training dataset.

Keywords: Physics-constrained learning; neural ordinary differential equations; nonlinear system identification; dynamical systems.

1. Introduction

Owing to their tremendous capability to learn complex relationships from data, neural networks offer promise in their ability to model unknown dynamical systems from trajectory observations. Such models of system dynamics can then be used to synthesize control strategies, to perform model-based reinforcement learning, or to predict the future values of quantities of interest.

However, purely data-driven approaches to learning can result in poor data efficiency and in model predictions that violate physical principles. These deficiencies become particularly emphasized when the training dataset is relatively small – the neural network must learn to approximate

* These authors contributed equally.

Project code is freely available at https://github.com/wuwushrek/physics_constrained_nn.

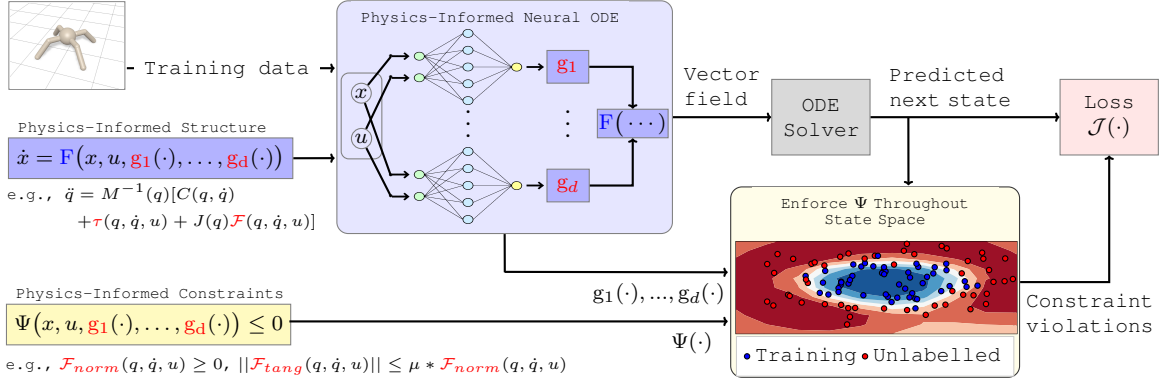


Figure 1: An illustration of the proposed framework. A priori physics knowledge is captured by a structured representation of the vector field (blue) – unknown components are represented by neural networks (red). Physics-based constraints are enforced on the outputs of the model not only on the labeled training datapoints, but also on any unlabeled points within the state space where the constraints are known to hold true (yellow).

a high-dimensional and non-linear map from a limited number of state trajectories. This limited reliability in the scarce data regime can render neural-network-based dynamics models impractical for the aforementioned applications.

On the other hand, useful a priori system knowledge is often available, even in circumstances when the exact dynamics remain unknown. Such knowledge might stem from a variety of sources – basic principles of physics, geometry constraints arising from the system’s design, or empirically validated invariant sets in the state space. The central thesis of this paper is that effective inclusion of a priori knowledge into the training of deep neural network models of dynamical systems can greatly improve data efficiency and model generalization to previously unseen regions of the state space, while also ensuring that the learned model respects physical principles.

We develop a framework to incorporate a wide variety of potential sources of a priori knowledge into data-driven models of dynamical systems. The framework uses neural networks to parametrize the dynamical system’s vector field and uses numerical integration schemes to predict future states, as opposed to parametrizing a model that predicts the next state directly. Physics-based knowledge is incorporated as inductive bias in the model of the vector field via two distinct mechanisms:

1. Physics-informed model structure. We represent the system’s vector field as compositions of unknown terms that are parametrized as neural networks and known terms that are derived from a priori knowledge. For example, Figure 1 illustrates a robotics environment in *Brax* (Freeman et al., 2021) whose equations of motion we assume to be partially known; the robot’s mass matrix is available, while the remaining terms in its equations of motion (i.e. its actuation and contact forces) must be learned. Neural networks representing the unknown terms are composed with the known mass matrix to obtain a model of the system’s vector field, as is illustrated in Figure 1.

2. Physics-informed constraints. We additionally enforce physics-based constraints on the values of the outputs and the internal states of the model. Such constraints could, for example, encode known system equilibria, invariants, or symmetries in the dynamics. We note that while only a limited number of datapoints may be available for supervised learning, constraints derived from a




Figure 2: Incorporation of a priori knowledge results in a significant improvement in data efficiency and out-of-sample prediction accuracy. *K1* corresponds to the incorporation of basic vector field knowledge: $\frac{dx}{dt} = v$. *K2* corresponds to knowledge of the mass matrix of the robot. The *Testing Loss* plot illustrates the model’s prediction errors during training. The *Avg. Rollout Error* plot illustrates a measure of the error over an entire predicted trajectory beginning from an out-of-sample initial state, after training has converged.

priori knowledge will hold over large sub-sets of the state space, and potentially over the state space in its entirety. Ensuring that the learned model satisfies the relevant constraints throughout the state space, while also fitting the available trajectory data, leads to a semi-supervised learning scheme that generalizes the available training data to the unlabeled portions of the state space.

We train the model on time-series data of state and control input observations. The physics-informed model of the vector field is integrated in time to obtain predictions of the future state. These predictions are in turn compared against the true future state values obtained from training data to define the model’s loss. During the model’s training, we enforce the physics-informed constraints on a collection of points in the state space, both labeled and unlabeled, via the augmented Lagrangian method (Lu et al., 2021; Hestenes, 1969).

We experimentally demonstrate the effectiveness of the proposed approach on a double pendulum system as well as on a suite of controlled rigid multi-body systems simulated in Brax (Freeman et al., 2021). This suite includes systems with nonlinear dynamics, control inputs, contact forces, and states with hundreds of dimensions. We are the first to consider such a suite of experiments in the context of physics-informed neural networks. We consider symmetries in the system vector fields, knowledge of the mass and Jacobian matrices of the robotic systems, and constraints that encode the laws of friction to help learn contact forces. We examine the extent to which these different sources of side information aid in improving the data efficiency and the out-of-sample prediction accuracy of the learned models. By exploiting a priori system knowledge during training, the proposed approach learns to predict system dynamics more than two orders of magnitude more accurately than a baseline approach that does not include any prior knowledge, as illustrated in Figure 2. The experiments demonstrate the viability of the proposed framework for the efficient learning of complex systems that are often used as challenging benchmark environments for deep reinforcement learning and model-based control (Duan et al., 2016; Freeman et al., 2021).

2. Physics-Informed Neural Architectures

We consider arbitrary non-linear dynamics that can be expressed in terms of an ordinary differential equation (ODE) of the form $\dot{x} = h(x, u)$. The state $x : \mathbb{R}_+ \mapsto \mathcal{X}$ is a continuous-time signal,

the control input $u : \mathbb{R}_+ \mapsto \mathcal{U}$ is a (possibly non-continuous) signal of time, and the vector field $h : \mathcal{X} \times \mathcal{U} \mapsto \mathbb{R}^n$ is assumed to be unknown and possibly non-linear. Here, n denotes the dimension of the state. We note that it is straightforward to explicitly include time dependency in the ODE representing the system dynamics, however, we omit it for clarity and notational simplicity.

Throughout the paper, 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 $\tau = \{(x_0, u_0), (x_1, u_1), \dots, (x_T, u_T)\}$, where $x_i = x(t_i)$ is the state at time t_i , $u_i = u(t_i)$ is the control input applied at time t_i , and $t_0 < t_1 < \dots < t_T$ is an increasing sequence of points in time. We seek to learn a function to predict future state values $x_{k+1}, \dots, x_{k+n_r+1}$, given the current state value x_k and a sequence of control inputs u_k, \dots, u_{k+n_r} .

We propose to parametrize and to learn the unknown vector field h , as opposed to learning a prediction function that maps directly from a history of states and control inputs to the predicted next state. This choice to parametrize a model of the vector field is primarily motivated by the two following points: (1) a priori knowledge derived from physical principles is typically most easily expressed in terms of the system’s vector field; (2) [Chen et al. \(2018\)](#) recently demonstrated that parametrizing the vector field results in models that are able to approximate time-series data with more accuracy than approaches that directly estimate the next state. We thus build a framework capable of learning the vector field h from the dataset \mathcal{D} of trajectories, while also taking advantage of prior knowledge surrounding the functional form of the vector field.

Compositional Representation of the Vector Field. We represent the vector field of the dynamical system as the composition of a known function – derived from a priori knowledge – and a collection of unknown functions that must be learned from data.

$$\dot{x} = h(x, u) = F(x, u, g_1(\cdot), \dots, g_d(\cdot)) \quad (1)$$

Here, F is a known differentiable function encoding available prior knowledge on the system’s model. The functions g_1, \dots, g_d encode the unknown terms within the underlying model. The inputs to these functions could themselves be arbitrary functions of the states and control inputs, or even of the outputs of the other unknown terms.

End-To-End Training of the Model. Using the available training data, we learn the unknown functions $g_{\theta_1}, g_{\theta_2}, \dots, g_{\theta_d}$ in an end-to-end fashion. We parametrize g_1, g_2, \dots, g_d by a collection of neural networks $g_{\theta_1}, g_{\theta_2}, \dots, g_{\theta_d}$, where $\theta_1, \theta_2, \dots, \theta_d$ are the parameter vectors of each of the individual networks. For notational simplicity, we define $\Theta = (\theta_1, \dots, \theta_d)$ and we denote the collection of neural networks $g_{\theta_1}, g_{\theta_2}, \dots, g_{\theta_d}$ by G_Θ .

For fixed values of the parameters Θ , we integrate the current estimate of the dynamics $F(\cdot, \cdot, G_\Theta)$ using a differentiable ODE solver to obtain predictions of the future state. More specifically, given an increasing sequence of n_r distinct points in time $t_i < t_{i+1} < \dots < t_{i+n_r}$, an initial state x_i , and a sequence of control input values u_i, \dots, u_{i+n_r} , we use (2) to solve for the model-predicted sequence of future states $x_{i+1}^\Theta, \dots, x_{i+n_r+1}^\Theta$. We assume the control input has a constant value u_i over the time interval $[t_i, t_{i+1})$.

$$x_{i+1}^\Theta, \dots, x_{i+n_r+1}^\Theta = \text{ODESolve}(x_k, u_i, \dots, u_{i+n_r}; F(\cdot, \cdot, G_\Theta)) \quad (2)$$

The loss function $\mathcal{J}(\Theta)$ can then be constructed to minimize the state prediction error over the fixed rollout horizon n_r , as in $\mathcal{J}(\Theta) = \sum_{\tau_l \in \mathcal{D}} \sum_{(x_i, u_i) \in \tau_l} \sum_{j=i}^{i+n_r} \|x_{j+1}^\Theta - x_{j+1}\|^2$, for a given norm.

Finally, assuming that F is differentiable, we can update the weights in Θ using either automatic differentiation, or using the adjoint sensitivity method (Chen et al., 2018). We note that for each datapoint (x_i, u_i) within each trajectory τ_l , we use the model to roll out predictions of the next n_r states, which are in turn used to define the loss function. This use of rollouts instead of one-step predictions while defining the loss function results in a reduction in error accumulation.

3. Physics-Informed Neural Network Constraints

In addition to using a priori physics-based knowledge to dictate the structure of the neural network, we also use such knowledge to derive constraints on the outputs and the internal states of the model. More formally, suppose we derive a particular physics-informed model $F(x, u, G_\Theta)$ of the system's vector field with unknown terms parametrized by the collection of neural networks G_Θ . Recall that $\Theta = (\theta_1, \theta_2, \dots, \theta_d)$ and that by G_Θ we denote $g_{\theta_1}, \dots, g_{\theta_d}$. Our objective is to solve for a set of parameter values minimizing the loss $\mathcal{J}(\Theta)$ over the training dataset while also satisfying all of the known physics-based constraints. That is, we aim to solve the optimization problem (3)-(4).

$$\min_{\Theta} \mathcal{J}(\Theta) \text{ s.t. } \Phi_i(x, u, G_\Theta) = 0, \forall (x, u) \in \mathcal{C}_{\Phi_i}, i \in [v], \quad (3)$$

$$\Psi_j(x, u, G_\Theta) \leq 0, \forall (x, u) \in \mathcal{C}_{\Psi_j}, j \in [l] \quad (4)$$

Here, $\Phi_i(x, u, G_\Theta)$ and $\Psi_j(x, u, G_\Theta)$ are differentiable functions capturing the physics-informed equality and inequality constraints respectively. We use $[v] = \{1, \dots, v\}$ and $[l] = \{1, \dots, l\}$ to denote the sets indexing these constraints. For each constraint, we additionally assume there is some sub-set $\mathcal{C}_{\Phi_i}, \mathcal{C}_{\Psi_j} \subseteq \mathcal{X} \times \mathcal{U}$, also derived from a priori knowledge, over which the constraint should hold true. As elements $(x, u) \in \mathcal{C}_{\Phi_i}, \mathcal{C}_{\Psi_j}$ are not necessarily included in the trajectory data used for training, the constraints thus provide useful information about potentially unlabeled points.

Enforcing Constraints Throughout the State Space. While (3)-(4) specifies that each constraint should hold over some known subset of the state space, this formulation leads to a possibly infinite number of constraints. To approximately solve the problem numerically, we instead select a finite set $\Omega = \{(x_1, u_1), (x_2, u_2), \dots, (x_{|\Omega|}, u_{|\Omega|})\} \subseteq \mathcal{X} \times \mathcal{U}$ of points at which we enforce the constraints. That is, we replace the possibly infinite sets of constraints given by (3)-(4) with the finite set of constraints given by (5)-(6).

$$\Phi_i(x, u, G_\Theta) = 0, \forall (x, u) \in \Omega \cap \mathcal{C}_{\Phi_i}, i \in [v] \quad (5)$$

$$\Psi_j(x, u, G_\Theta) \leq 0, \forall (x, u) \in \Omega \cap \mathcal{C}_{\Psi_j}, j \in [l] \quad (6)$$

Intuitively, by optimizing the loss function $\mathcal{J}(\Theta)$ in (3) subject to (5)-(6), we are finding the solution that fits the training data as well as possible, while also satisfying all of the relevant constraints at each point within a finite set Ω of representative states throughout $\mathcal{X} \times \mathcal{U}$.

The Augmented Lagrangian Method. In order to solve this optimization problem, we use a stochastic gradient descent (SGD) variant of the augmented Lagrangian method, as proposed by (Lu et al., 2021; Toussaint, 2014) for the numerical solution of constrained optimization problems. For each constraint in (5)-(6), we define a separate Lagrange variable. That is, we define an individual variable $\lambda_{i,k}^{\text{eq}}$ for each equality constraint $\Phi_i, i \in [v]$ evaluated at each point $(x_k, u_k) \in \Omega \cap \mathcal{C}_{\Phi_i}$, and similarly $\lambda_{j,k}^{\text{ineq}}$ at each point $(x_k, u_k) \in \Omega \cap \mathcal{C}_{\Psi_j}$ for all the inequality constraints Ψ_j . The augmented Lagrangian of the optimization problem is given by (7).

$$\begin{aligned} \mathcal{L}(\Theta, \mu, \lambda) = & \mathcal{J}(\Theta) + \sum_{\substack{i \in [v] \\ (x_k, u_k) \in \Omega \cap \mathcal{C}_{\Phi_i}}} \mu \Phi_i(x_k, u_k, G_\Theta)^2 + \sum_{\substack{i \in [v] \\ (x_k, u_k) \in \Omega \cap \mathcal{C}_{\Phi_i}}} \lambda_{i,k}^{\text{eq}} \Phi_i(x_i, u_i, G_\Theta) \quad (7) \\ & + \sum_{\substack{j \in [l] \\ (x_k, u_k) \in \Omega \cap \mathcal{C}_{\Psi_j}}} \mu [\lambda_{j,k}^{\text{ineq}} > 0 \vee \Psi_j > 0] \Psi_j(x_k, u_k, G_\Theta)^2 + \sum_{\substack{j \in [l] \\ (x_k, u_k) \in \Omega \cap \mathcal{C}_{\Psi_j}}} \lambda_{j,k}^{\text{ineq}} \Psi_j(x_k, u_k, G_\Theta) \end{aligned}$$

The Proposed Training Algorithm. Algorithm 1 outlines the proposed approach to optimizing the loss function $\mathcal{J}(\Theta)$ in (3) subject to (5)–(6) via the augmented Lagrangian method. In the algorithm, we use the notation $x_+ = \max\{0, x\}$. We initialize values for μ and λ and minimize $\mathcal{L}(\Theta, \mu, \lambda)$ via SGD over Θ while holding the values of μ and λ fixed. To prevent the gradient descent from getting stuck in local minima, Algorithm 1 randomly samples a subset of points from Ω at each gradient update, instead of including the entire set in the definition of $\mathcal{L}(\Theta, \mu, \lambda)$. This random selection of points at which to evaluate the constraint violations is akin to the random sampling of minibatches of training data during traditional SGD. Once this inner loop SGD has converged, we update the values of μ and λ according to the update rules outlined in Algorithm 1. This process is repeated until the constraints are all satisfied, to within an allowed tolerance value ϵ .

Algorithm 1 Training Algorithm

Input: $F(\cdot)$, $\{\Phi_i(\cdot)\}_{i \in [v]}$, $\{\Psi_j(\cdot)\}_{j \in [l]}$, \mathcal{D} , Ω

Parameter: ϵ , μ_0 , μ_{mult} , $N_{\text{trainBatch}}$, $N_{\text{constrBatch}}$

Output: Model parameters Θ .

- 1: Initialize parameters Θ ; $\lambda_{i,k}^{\text{eq}} \leftarrow 0$; $\lambda_{j,k}^{\text{ineq}} \leftarrow 0$; $\mu \leftarrow \mu_0$.
 - 2: **while** $\sum_i \sum_k |\Phi_i(x_k, u_k, G_\Theta)| + \sum_j \sum_k (\Psi_j(x_k, u_k, G_\Theta))_+ \geq \epsilon$ **do**
 - 3: **while not** `SGDStoppingCriterion()` **do**
 - 4: $\mathcal{D}_{\text{batch}} \leftarrow \text{Sample}(\mathcal{D}, N_{\text{trainBatch}})$; $\Omega_{\text{batch}} \leftarrow \text{Sample}(\Omega, N_{\text{constrBatch}})$
 - 5: $\Theta \leftarrow \text{optimUpdate}(\mathcal{L}, \Theta, \mathcal{D}_{\text{batch}}, \Omega_{\text{batch}})$
 - 6: **end while**
 - 7: $\lambda_{i,k}^{\text{eq}} \leftarrow \lambda_{i,k}^{\text{eq}} + 2 * \mu * \Phi_i(x_k, u_k, G_\Theta)$; $\lambda_{j,k}^{\text{ineq}} \leftarrow (\lambda_{j,k}^{\text{ineq}} + 2 * \mu * \Psi_j(x_k, u_k, G_\Theta))_+$
 - 8: $\mu \leftarrow \mu * \mu_{\text{mult}}$
 - 9: **end while**
-

4. Experimental Results

4.1. Learning the Dynamics of a Double Pendulum

The dynamics of the double pendulum are described by (8).

$$\ddot{\phi}_1 = (g_1 - \alpha_1 g_2)/(1 - \alpha_1 \alpha_2); \quad \ddot{\phi}_2 = (-\alpha_2 g_1 + g_2)/(1 - \alpha_1 \alpha_2) \quad (8)$$

Here, ϕ_1 and ϕ_2 specify the angular position of the first and second links of the pendulum, $\alpha_1(\phi_1, \phi_2) \propto \cos(\phi_1 - \phi_2)$, $\alpha_2(\phi_1, \phi_2) \propto \cos(\phi_2 - \phi_1)$, and $g_1(\phi_1, \phi_2, \dot{\phi}_1, \dot{\phi}_2)$, $g_2(\phi_1, \phi_2, \dot{\phi}_1, \dot{\phi}_2)$ are both complicated trigonometric functions of the state variables. By setting $x_1 = \phi_1$, $x_2 = \phi_2$, $x_3 = \dot{\phi}_1$, and $x_4 = \dot{\phi}_2$ we can thus express the dynamics as $\dot{x} = h(x)$, where $h_1(x) = x_3$, $h_2(x) = x_4$, and $h_3(x)$, $h_4(x)$ are both given by (8). We refer to the extended version of the paper for more details surrounding the dynamics (Djeumou et al., 2021).




Figure 3: Double pendulum numerical results. A single trajectory that is 3 seconds in length is used as training data. The shaded areas illustrate the variance of the results over five runs.

Defining a Baseline for Comparison. As a baseline, we parametrize the vector field h by a single neural network g_{θ_1} . In terms of the framework outlined in the previous sections, we thus have $\Theta = \theta_1$ and $F(x, G_{\Theta}) = g_{\theta_1}(x)$.

(K1) Incorporating Knowledge of Equation (8). In order to demonstrate how arbitrary knowledge surrounding the system’s vector field can be incorporated into the dynamics model, we assume that equation (8) is known, along with $\alpha_1(\cdot)$ and $\alpha_2(\cdot)$. However, we assume $g_1(\cdot)$ and $g_2(\cdot)$ are unknown. We parametrize g_1 and g_2 by two neural networks g_{θ_1} and g_{θ_2} , respectively. The proposed model for the vector field may thus be expressed as $F_1(x, G_{\Theta}) = x_3$, $F_2(x, G_{\Theta}) = x_4$, and with $F_3(x, G_{\Theta})$, $F_4(x, G_{\Theta})$ both being functions of α_1 , α_2 , g_{θ_1} , g_{θ_2} , as in equation (8).

(K2) Incorporating Symmetry Constraints on g_1 and g_2 . To demonstrate that the proposed training algorithm properly enforces constraints on the system’s dynamics, we impose equality constraints on g_{θ_1} , g_{θ_2} derived from symmetries in the vector field. In particular, we impose four separate equality constraints, an example of which is as follows: $g_1(x_{1:2}, x_{3:4}) = -g_1(-x_{1:2}, x_{3:4})$. These equality constraints are enforced at a collection of points Ω that are sampled uniformly throughout the state space.

Experimental Setup. Each neural network $g_{\theta_i}(x)$ is a multilayer perceptron (MLP) with ReLU activation functions and two hidden layers with 128 nodes each. We use a rollout horizon of $n_r = 5$ when defining the loss function and we randomly sample $|\Omega| = 10000$ points throughout the state space \mathcal{X} at which we enforce the above equality constraints. For further details on the setup of the experiment, we refer to the extended version of the paper (Djeumou et al., 2021).

Experimental Results. Figure 3 plots the results of training the various physics-informed models using only a single trajectory of data, which equivalent to 3 seconds worth of observations. We observe that in terms of prediction losses, (K1) performs better than the baseline, and (K2) performs better than both the baseline and (K1); its testing loss at convergence is roughly an order of magnitude lower than that of (K1). Furthermore, we observe that the constraint loss of (K2) – which measures the average value of the model’s violation of the equality constraints – is roughly two orders of magnitude lower than that of (K1). These results demonstrate the proposed framework’s effectiveness at incorporating a priori system knowledge into the dynamics model. As more detailed knowledge is included into the model, the testing loss at convergence is significantly reduced. Furthermore, the gap in constraint loss between (K1) and (K2) demonstrates the effectiveness of the proposed framework in learning to fit the training data while simultaneously enforcing constraints.



Figure 4: Illustration of the suite of simulated robotic systems used during testing.

4.2. Learning the Dynamics of Controlled, Rigid, Multi-Body Systems

In this section we apply the proposed framework to a suite of robotic systems simulated in Brax (Freeman et al., 2021). These systems feature control inputs, non-linear dynamics, and external forces. Figure 4 illustrates the specific robotic systems we use for testing.

The Governing Equations of Motion. The dynamics of all of the environments listed in Figure 4 can be described by the general equations of motion given in (9).

$$\ddot{q} = M(q)^{-1}[C(q, \dot{q}) + \tau(q, \dot{q}, u) + J(q)\mathcal{F}(q, \dot{q}, u)] \quad (9)$$

Here, q and \dot{q} represent the system state and its time derivative respectively, $M(q)$ represents the system’s mass matrix, $C(q, \dot{q})$ is a vector representing the Coriolis forces, $\tau(q, \dot{q}, u)$ is a vector representing the actuation forces (given the control input u), $J(q)$ is the system’s Jacobian matrix, and $\mathcal{F}(q, \dot{q}, u)$ is a vector representing the contact forces. By setting $x_1 = q$ and $x_2 = \dot{q}$, we may re-write the system in the form $\dot{x} = h(x, u)$.

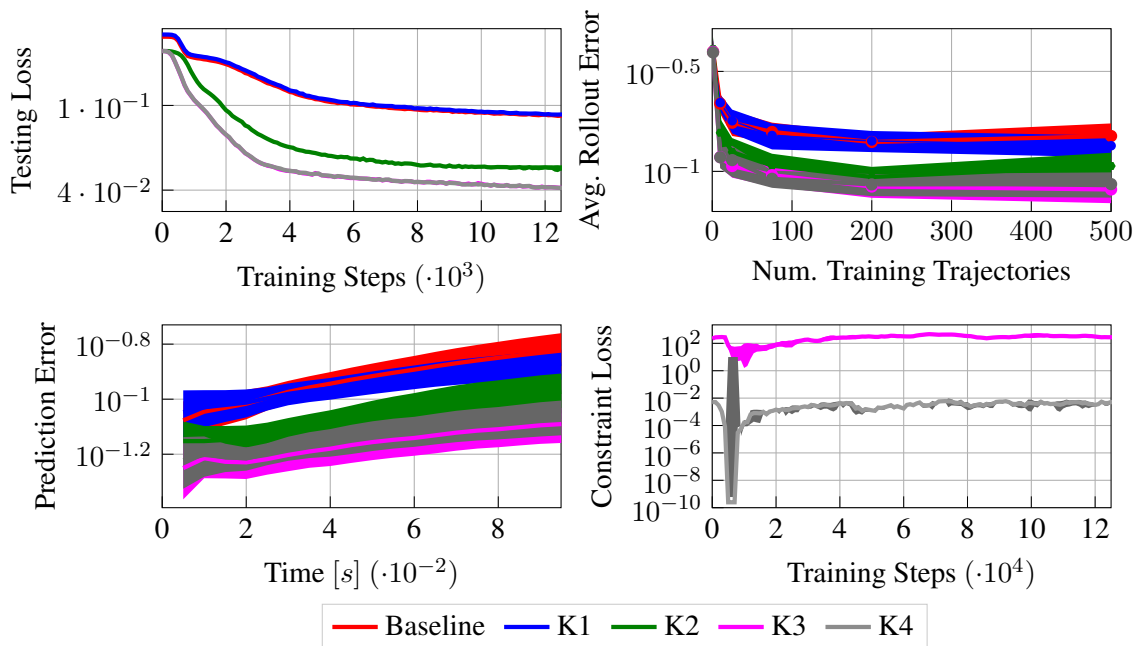
Defining a Baseline for Comparison. Similarly to the double pendulum baseline, we use a single neural network to parametrize the vector field, i.e. $F(x, u, G_\Theta) = g_{\theta_1}(x, u)$.

(K1) Incorporating Basic Knowledge on the Vector Field. A relatively simple piece of a priori knowledge is that $\dot{x}_1 = x_2$. This follows directly from our formulation of the system’s dynamics, and it provides a useful piece of inductive bias surrounding the structure of the vector field. In our proposed framework, we again parametrize the vector field using a single neural network g_{θ_1} , however, we now have $F_1(x, u, G_\Theta) = x_2$ and $F_2(x, u, G_\Theta) = g_{\theta_1}(x, u)$.

(K2) Incorporating Knowledge of the Mass Matrix. We now assume the mass matrix $M(x_1)$ and the Coriolis force $C(x_1, x_2)$ terms are known a priori. To use this a priori knowledge, we parametrize the actuation force term and the contact force term using individual neural networks g_{θ_1} and g_{θ_2} . Our structured neural network model of the system dynamics may be written as follows: $F_1(x, u, G_\Theta) = x_2$ and $F_2(x, u, G_\Theta) = M^{-1}(x_1)C(x_1, x_2) + g_{\theta_1}(x, u) + g_{\theta_2}(x, u)$.

(K3) Incorporating Knowledge of the Jacobian Matrix. We assume, in addition to the a priori knowledge assumed in (K3), that the system’s Jacobian matrix $J(x_1)$ is also known: $F_1(x, u, G_\Theta) = x_2$, $F_2(x, u, G_\Theta) = M^{-1}(x_1)C(x, x_2) + g_{\theta_1} + J(x_1)g_{\theta_2}$.

(K4) Incorporating Contact Force Constraints. We use the laws of friction to derive constraints on the contact forces between the robotic systems and the ground. In particular, we have $\mathcal{F}_{norm}(x, u) \leq 0$ and $\|\mathcal{F}_{tang}(x, u)\| \geq \mu\mathcal{F}_{norm}(x, u)$, where \mathcal{F}_{norm} represents the normal component of the contact forces and \mathcal{F}_{tang} is the tangential force vector. We assume (K4) has access to the same amount of a priori knowledge as (K3), and that it additionally imposes the above inequality constraints on the neural network g_{θ_2} , which represents the contact force term.

Figure 5: Numerical results for the *Fetch* robotic system.

Experimental Setup. We provide the experimental results for the *Fetch* robotic system. This robot has a 143-dimensional state space and a 10-dimensional control input space. To parametrize each unknown term g_{θ_i} we use a multilayer perceptron (MLP) with ReLU activation functions and two hidden layers of 256 nodes each. We use a rollout horizon of $n_r = 5$ to define the loss function. Further details, as well as the experimental results for the other robots, are provided in the extended version of the paper (Djeumou et al., 2021).

Experimental Results. Figure 5 illustrates the results. Similarly to the double pendulum study, we observe from Figure 5 that as we increase the amount of a priori knowledge incorporated into the model, the average prediction error correspondingly decreases. In particular, the *Avg. Rollout Error* plot shows that by incorporating a priori knowledge of the mass matrix, (*K2*) consistently makes more accurate predictions than either the baseline or (*K1*), regardless of the number of trajectories included in the training dataset. From the *Constraint Loss* plot we observe a gap of four orders of magnitude between the constraint loss of (*K3*) and (*K4*): the proposed algorithm learns a model that effectively enforces the contact force constraints incorporated into (*K4*).

5. Related Work

Techniques for *nonlinear system identification* – the construction of mathematical models of dynamical systems from empirical observations – often rely on a priori knowledge of a suitable collection of nonlinear basis functions (Brunton et al., 2016a,b; Williams et al., 2014, 2015). However, choosing such sets of basis functions can be challenging, particularly for systems with complex dynamics and high-dimensional state spaces. Li et al. (2017); Takeishi et al. (2017); Champion et al. (2019); Yeung et al. (2019) thus use neural networks to learn appropriate collections of nonlinear functions to be used in conjunction with these techniques. More closely related to our work, Raissi et al.

(2018); Lu et al. (2018); Dupont et al. (2019); Kelly et al. (2020); Massaroli et al. (2020) use neural networks to represent differential equations. Chen et al. (2018) introduce the neural ODE; neural networks parametrizing ordinary differential equations. However, the above works do not consider how physics-based knowledge might be incorporated into neural networks, nor how physics-based constraints can be enforced during training. These are the primary focuses of this paper.

The inclusion of physics-based knowledge in the training of neural networks has, however, been studied extensively over the past several years. In particular, Lu et al. (2021); Raissi et al. (2019); Raissi (2018); Han et al. (2018); Sirignano and Spiliopoulos (2018); Long et al. (2018, 2019) use neural networks for the solution and discovery of partial differential equations (PDE). Our work instead focuses on using neural networks to parametrize the vector field of dynamical systems.

More closely related to our work, a collection of recent papers study how physics-based knowledge can be incorporated into neural networks for the learning of dynamics models from trajectory data (Zhong et al., 2021a; Hernández et al., 2021; Rackauckas et al., 2020). Many of these works either use the Lagrangian or the Hamiltonian formulation of dynamics to inform the structure of a neural ODE, as in Cranmer et al. (2020); Lutter et al. (2019); Roehrl et al. (2020); Zhong and Leonard (2020); Allen-Blanchette et al. (2020) vs. Greydanus et al. (2019); Matsubara et al. (2020); Toth et al. (2020). Finzi et al. (2020) use the method of Lagrange multipliers to explicitly enforce holonomic constraints on the state variables in Lagrangian and Hamiltonian neural networks. A number of works have also studied physics-informed neural ODEs in the context of learning control-oriented dynamics models (Zhong et al., 2020a,b; Roehrl et al., 2020; Duong and Atanasov, 2021; Gupta et al., 2020; Menda et al., 2019; Zhong et al., 2021b; Shi et al., 2019). We note however, that the majority of the above works employ a specific neural ODE structure, rendering them less flexible than the general framework that we propose. Furthermore, the examples explored in these works typically feature low-dimensional state spaces that often do not include energy dissipation or external forces. By contrast, we apply the proposed approach to a variety of high-dimensional and complex systems. Finally, the above works do not enforce general physics-based constraints on the dynamics model.

A number of other recent works have also studied imposing constraints on the outputs of deep neural networks (Márquez-Neila et al., 2017; Fioretto et al., 2021; Nandwani et al., 2019; Kervadec et al., 2019). Similarly to our work, Lu et al. (2021); Dener et al. (2020) consider an augmented Lagrangian approach to the enforcement of physics-based constraints. However, they do so in different contexts than ours; the former focuses on PDE-constrained inverse design problems while the latter uses neural networks to approximate the Fokker-Planck operator used in fusion simulations.

6. Conclusions

In this work we present a framework for the incorporation of a priori physics-based knowledge into neural network models of dynamical systems. This physics based knowledge is incorporated either by influencing the structure of the neural network, or as constraints on the model’s outputs and internal states. We present a suite of numerical experiments that exemplify the effectiveness of the proposed approach: the inclusion of increasingly detailed forms of side-information leads to increasingly accurate model predictions. We also demonstrate that the framework effectively enforces physics-based constraints on the outputs and internal states of the models. Future work will aim to use these physics-informed dynamics models for control.

Acknowledgments

This work was supported in part by ONR N00014-22-1-2254, AFOSR FA9550-19-1-0005, and NSF 1646522.

References

- Christine Allen-Blanchette, Sushant Veer, Anirudha Majumdar, and Naomi Ehrich Leonard. Lagnetvip: A lagrangian neural network for video prediction, 2020. URL <https://arxiv.org/abs/2010.12932>.
- Steven L Brunton, Joshua L Proctor, and J Nathan Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. Proceedings of the national academy of sciences, 113(15):3932–3937, 2016a.
- Steven L Brunton, Joshua L Proctor, and J Nathan Kutz. Sparse identification of nonlinear dynamics with control (sindyc). IFAC-PapersOnLine, 49(18):710–715, 2016b.
- Kathleen Champion, Bethany Lusch, J Nathan Kutz, and Steven L Brunton. Data-driven discovery of coordinates and governing equations. Proceedings of the National Academy of Sciences, 116(45):22445–22451, 2019.
- 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. Curran Associates, Inc., 2018. URL <https://proceedings.neurips.cc/paper/2018/file/69386f6bb1dfed68692a24c8686939b9-Paper.pdf>.
- Miles Cranmer, Sam Greydanus, Stephan Hoyer, Peter Battaglia, David Spergel, and Shirley Ho. Lagrangian neural networks, 2020. URL <https://arxiv.org/abs/2003.04630>.
- Alp Dener, Marco Andres Miller, Randy Michael Churchill, Todd Munson, and Choong-Seock Chang. Training neural networks under physical constraints using a stochastic augmented lagrangian approach, 2020. URL <https://arxiv.org/abs/2009.07330>.
- Franck Djeumou, Cyrus Neary, Eric Goubault, Sylvie Putot, and Ufuk Topcu. Neural networks with physics-informed architectures and constraints for dynamical systems modeling, 2021. URL <https://arxiv.org/abs/2109.06407>.
- Yan Duan, Xi Chen, Rein Houthoofd, John Schulman, and Pieter Abbeel. Benchmarking deep reinforcement learning for continuous control. In International conference on machine learning, pages 1329–1338. PMLR, 2016.
- Thai Duong and Nikolay Atanasov. Hamiltonian-based neural ode networks on the se (3) manifold for dynamics learning and control. In Robotics: Science and Systems (RSS), 2021.
- Emilien Dupont, Arnaud Doucet, and Yee Whye Teh. Augmented neural odes. In Advances in Neural Information Processing Systems, volume 32. Curran Associates, Inc., 2019. URL <https://proceedings.neurips.cc/paper/2019/file/21be9a4bd4f81549a9d1d241981cec3c-Paper.pdf>.

- 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, pages 13880–13889. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/9f655cc8884fda7ad6d8a6fb15cc001e-Paper.pdf>.
- Ferdinando Fioretto, Pascal Van Hentenryck, Terrence W. K. Mak, Cuong Tran, Federico Baldo, and Michele Lombardi. Lagrangian duality for constrained deep learning. In Yuxiao Dong, Georgiana Ifrim, Dunja Mladenić, Craig Saunders, and Sofie Van Hoecke, editors, Machine Learning and Knowledge Discovery in Databases. Applied Data Science and Demo Track, pages 118–135, Cham, 2021. Springer International Publishing. ISBN 978-3-030-67670-4.
- C. Daniel Freeman, Erik Frey, Anton Raichuk, Sertan Girgin, Igor Mordatch, and Olivier Bachem. Brax – a differentiable physics engine for large scale rigid body simulation, 2021. URL <https://arxiv.org/abs/2106.13281>.
- Samuel Greydanus, Misko Dzamba, and Jason Yosinski. Hamiltonian neural networks. In Advances in Neural Information Processing Systems, volume 32. Curran Associates, Inc., 2019. URL <https://proceedings.neurips.cc/paper/2019/file/26cd8ecadce0d4efd6cc8a8725cbd1f8-Paper.pdf>.
- Jayesh K Gupta, Kunal Menda, Zachary Manchester, and Mykel Kochenderfer. Structured mechanical models for robot learning and control. In Learning for Dynamics and Control, pages 328–337. PMLR, 2020.
- Jiequn Han, Arnulf Jentzen, and E Weinan. Solving high-dimensional partial differential equations using deep learning. Proceedings of the National Academy of Sciences, 115(34):8505–8510, 2018.
- Quercus Hernández, Alberto Badías, David González, Francisco Chinesta, and Elías Cueto. Structure-preserving neural networks. Journal of Computational Physics, 426:109950, 2021.
- Magnus R Hestenes. Multiplier and gradient methods. Journal of optimization theory and applications, 4(5):303–320, 1969.
- Jacob Kelly, Jesse Bettencourt, Matthew J Johnson, and David K Duvenaud. Learning differential equations that are easy to solve. In Advances in Neural Information Processing Systems, volume 33, pages 4370–4380. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/2e255d2d6bf9bb33030246d31f1a79ca-Paper.pdf>.
- Hoel Kervadec, Jose Dolz, Jing Yuan, Christian Desrosiers, Eric Granger, and Ismail Ben Ayed. Constrained deep networks: Lagrangian optimization via log-barrier extensions, 2019. URL <https://arxiv.org/abs/1904.04205>.
- Qianxiao Li, Felix Dietrich, Erik M Bollt, and Ioannis G Kevrekidis. Extended dynamic mode decomposition with dictionary learning: A data-driven adaptive spectral decomposition of the koopman operator. Chaos: An Interdisciplinary Journal of Nonlinear Science, 27(10):103111, 2017.

- Zichao Long, Yiping Lu, Xianzhong Ma, and Bin Dong. Pde-net: Learning pdes from data. In International Conference on Machine Learning, pages 3208–3216. PMLR, 2018.
- Zichao Long, Yiping Lu, and Bin Dong. Pde-net 2.0: Learning pdes from data with a numeric-symbolic hybrid deep network. Journal of Computational Physics, 399:108925, 2019.
- Lu Lu, Raphael Pestourie, Wenjie Yao, Zhicheng Wang, Francesc Verdugo, and Steven G Johnson. Physics-informed neural networks with hard constraints for inverse design. SIAM Journal on Scientific Computing, 43(6):B1105–B1132, 2021.
- Yiping Lu, Aoxiao Zhong, Quanzheng Li, and Bin Dong. Beyond finite layer neural networks: Bridging deep architectures and numerical differential equations. In International Conference on Machine Learning, pages 3276–3285. PMLR, 2018.
- Michael Lutter, Christian Ritter, and Jan Peters. Deep lagrangian networks: Using physics as model prior for deep learning. In International Conference on Learning Representations. OpenReview.net, 2019. URL <https://openreview.net/forum?id=BklHpjCqKm>.
- Stefano Massaroli, Michael Poli, Jinkyoo Park, Atsushi Yamashita, and Hajime Asama. Dissecting neural odes. In Advances in Neural Information Processing Systems, volume 33, pages 3952–3963. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/293835c2cc75b585649498ee74b395f5-Paper.pdf>.
- Takashi Matsubara, Ai Ishikawa, and Takaharu Yaguchi. Deep energy-based modeling of discrete-time physics. In Advances in Neural Information Processing Systems, volume 33, pages 13100–13111. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/98b418276d571e623651fc1d471c7811-Paper.pdf>.
- 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, pages 138–171, 2019.
- Pablo Márquez-Neila, Mathieu Salzmann, and Pascal Fua. Imposing hard constraints on deep networks: Promises and limitations, 2017. URL <https://arxiv.org/abs/1706.02025>.
- Yatin Nandwani, Abhishek Pathak, Mausam, and Parag Singla. A primal dual formulation for deep learning with constraints. In Advances in Neural Information Processing Systems, volume 32. Curran Associates, Inc., 2019. URL <https://proceedings.neurips.cc/paper/2019/file/cf708fc1decf0337aded484f8f4519ae-Paper.pdf>.
- 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, 2020. URL <https://arxiv.org/abs/2001.04385>.
- Maziar Raissi. Deep hidden physics models: Deep learning of nonlinear partial differential equations. The Journal of Machine Learning Research, 19(1):932–955, 2018.

- Maziar Raissi, Paris Perdikaris, and George Em Karniadakis. Multistep neural networks for data-driven discovery of nonlinear dynamical systems, 2018. URL <https://arxiv.org/abs/1801.01236>.
- Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. Journal of Computational Physics, 378:686–707, 2019.
- 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, pages 9784–9790, 2019.
- Justin Sirignano and Konstantinos Spiliopoulos. Dgm: A deep learning algorithm for solving partial differential equations. Journal of computational physics, 375:1339–1364, 2018.
- Naoya Takeishi, Yoshinobu Kawahara, and Takehisa Yairi. Learning koopman invariant subspaces for dynamic mode decomposition. In Advances in Neural Information Processing Systems, volume 30. Curran Associates, Inc., 2017. URL <https://proceedings.neurips.cc/paper/2017/file/3a835d3215755c435ef4fe9965a3f2a0-Paper.pdf>.
- Peter Toth, Danilo J. Rezende, Andrew Jaegle, Sébastien Racanière, Aleksandar Botev, and Irina Higgins. Hamiltonian generative networks. In International Conference on Learning Representations. OpenReview.net, 2020. URL <https://openreview.net/forum?id=HJenn6VFvB>.
- Marc Toussaint. A novel augmented lagrangian approach for inequalities and convergent any-time non-central updates, 2014. URL <https://arxiv.org/abs/1412.4329>.
- Matthew O. Williams, Clarence W. Rowley, and Ioannis G. Kevrekidis. A kernel-based approach to data-driven koopman spectral analysis, 2014. URL <https://arxiv.org/abs/1411.2260>.
- Matthew O Williams, Ioannis G Kevrekidis, and Clarence W Rowley. A data-driven approximation of the koopman operator: Extending dynamic mode decomposition. Journal of Nonlinear Science, 25(6):1307–1346, 2015.
- Enoch Yeung, Soumya Kundu, and Nathan Hodas. Learning deep neural network representations for koopman operators of nonlinear dynamical systems. In American Control Conference, pages 4832–4839. IEEE, 2019.
- Yaofeng Desmond Zhong and Naomi Leonard. Unsupervised learning of lagrangian dynamics from images for prediction and control. In Advances in Neural Information Processing Systems, volume 33, pages 10741–10752. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/79f56e5e3e0e999b3c139f225838d41f-Paper.pdf>.

Yaofeng Desmond Zhong, Biswadip Dey, and Amit Chakraborty. Symplectic ode-net: Learning hamiltonian dynamics with control. In International Conference on Learning Representations. OpenReview.net, 2020a. URL <https://openreview.net/forum?id=ryxmb1rKDS>.

Yaofeng Desmond Zhong, Biswadip Dey, and Amit Chakraborty. Dissipative symoden: Encoding hamiltonian dynamics with dissipation and control into deep learning. In ICLR 2020 Workshop on Integration of Deep Neural Models and Differential Equations, 2020b.

Yaofeng Desmond Zhong, Biswadip Dey, and Amit Chakraborty. Benchmarking energy-conserving neural networks for learning dynamics from data. In Learning for Dynamics and Control, pages 1218–1229. PMLR, 2021a.

Yaofeng Desmond Zhong, Biswadip Dey, and Amit Chakraborty. Extending lagrangian and hamiltonian neural networks with differentiable contact models, 2021b. URL <https://arxiv.org/abs/2102.06794>.