
Zero-Shot Transfer of Neural ODEs

Tyler Ingebrand, Adam J. Thorpe, Ufuk Topcu

University of Texas at Austin
Austin, TX 78712

Abstract

Autonomous systems often encounter environments and scenarios beyond the scope of their training data, which underscores a critical challenge: the need to generalize and adapt to unseen scenarios in real time. This challenge necessitates new mathematical and algorithmic tools that enable adaptation and zero-shot transfer. To this end, we leverage the theory of function encoders, which enables zero-shot transfer by combining the flexibility of neural networks with the mathematical principles of Hilbert spaces. Using this theory, we first present a method for learning a space of dynamics spanned by a set of neural ODE basis functions. After training, the proposed approach can rapidly identify dynamics in the learned space using an efficient inner product calculation. Critically, this calculation requires no gradient calculations or retraining during the online phase. This method enables zero-shot transfer for autonomous systems at runtime and opens the door for a new class of adaptable control algorithms. We demonstrate state-of-the-art system modeling accuracy for two MuJoCo robot environments and show that the learned models can be used for more efficient MPC control of a quadrotor.

1 Introduction

Models that are adaptable, generalizable, and capable of learning online from minimal data are essential for autonomy. These models must adapt to unseen tasks and environments at runtime without relying upon *a priori* parameterizations. For example, consider an autonomous UAV delivery robot navigating in a dense, urban environment through varying wind patterns and carrying uncertain payloads. This scenario requires rapid adaptation to ensure safe and correct operation because conditions can change unpredictably and online model updates are impractical. While prior works can control autonomous systems in a single setting, they fail to adapt to the continuum of real-life scenarios. The key challenge is enabling *zero-shot transfer* of learned models, where models quickly adapt to new data provided at runtime without retraining.

We present a method for modeling differential equations by learning a set of basis functions parameterized by neural ODEs. Our key insight is to learn a *space* of functions that captures feasible behaviors of the system. By focusing on learning the structure of the space of differential equations, our approach implicitly learns how the dynamics change due to changes in the environment. Our approach is based on the theory of function encoders [14], a framework for zero-shot transfer that has been applied to task transfer in reinforcement learning contexts.

We formulate the space of learned functions as a linear space equipped with an inner product (e.g. a Hilbert space), and learn a set of basis functions over this space, where each basis function is represented by a neural ODE. This structure offers an efficient way to approximate online dynamics via a linear combination of the basis functions.

By representing functions in a Hilbert space and pre-training on a suite of functions, we can quickly identify the basis functions' coefficients for a new dynamical system at runtime using minimal data. This is useful, for instance, in scenarios where we can pre-train offline in simulation, but need to

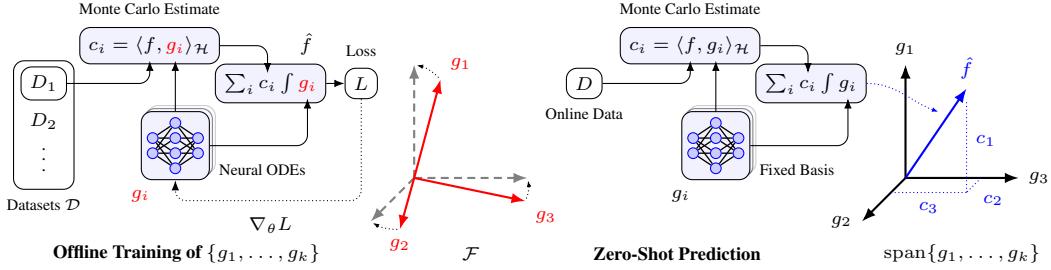


Figure 1: An illustration of our approach. The training phase uses a set of datasets \mathcal{D} to train basis functions $\{g_1, \dots, g_k\}$ to span \mathcal{F} . The zero-shot phase uses online data to identify the coefficients for a new function, which can be estimated as a linear combination of the basis functions.

quickly identify dynamics online from a single trajectory in a zero-shot manner. This strategy greatly reduces the computational overhead associated with adapting or re-training a neural ODE to new tasks at runtime. The ability to efficiently encode the behavior of a system at runtime without retraining is a key component of our approach. Our approach is outlined in Figure 1.

Our approach yields models which generalize to a large set of possible system behaviors and achieves better long-term prediction accuracy than neural ODEs alone. We showcase our approach on a problem of predicting the behavior of a first-order ODE system with Van der Pol dynamics. We demonstrate the scalability of our approach on two MuJoCo robotics environments [28]. Finally, we test the feasibility of using the learned model for downstream tasks such as model-predictive control (MPC). Our results show that our model achieves significantly better long-horizon prediction accuracy compared to the nearest baseline. Additionally, the MPC controller using our model has a lower slew rate, indicating that the improved model accuracy leads to more efficient control decisions.

1.1 Contributions

Representing Spaces of Dynamical Systems: We propose a novel framework for representing spaces of dynamical systems, i.e. induced by hidden system parameters, variations in the underlying physics, or changing environmental features. Using a large-scale set of data collected offline, we learn a collection of neural networks that act as a functional basis. This approach is based in the theory of function encoders [14]. Yet the extension to using neural ODEs as basis functions is non-trivial and poses several challenges, mainly from the need to integrate the learned model.

A Method for Online Adaptation: We construct a method for adapting neural ODE estimations based on online data without gradient updates, i.e. zero-shot transfer of system models. Our approach overcomes a significant challenge in learning behaviors of differential equations and dynamical systems. By offloading the computational effort to the training phase, we enable rapid online identification, adaptation, and prediction without retraining.

Empirical Results: We demonstrate accurate long-horizon predictions in challenging robotics tasks and show these models can be used for online control of quadrotor systems. We assess the quality of the approach in three areas and answer the following questions: 1) How well does our approach adapt to new dynamics online? 2) How does our approach compare to existing approaches for long-horizon prediction tasks? and 3) Does our approach work for downstream tasks such as control? We show that function encoders using neural ODEs as basis functions consistently outperform existing approaches.

2 Background

2.1 Neural ODEs

Consider an ordinary differential equation $\dot{x}(t) = f(x(t), t)$, where f is Lipschitz continuous and $x(t) \in \mathbb{R}^n$ is the state at time t . Given an initial condition $x(t_0)$, the ODE solution can be written as

$$x(t_f) = x(t_0) + \int_{t_0}^{t_f} f(x(\tau), \tau) d\tau. \quad (1)$$

Note that in general, the explicit dependence of $f(x(t), t)$ on t can be removed by augmenting the state x to include t . As such, we omit t throughout.

Neural ODEs [5] parameterize the function f as a neural network. In particular, neural ODEs solve the ODE using an off-the-shelf integrator and optimize the neural network with respect to a prediction loss. The training procedure requires a dataset $D = \{(t_i, x(t_i))\}_{i=1}^d$ which is used to train the model via a supervised objective, such as mean squared error, back-propagated through the integrator.

Neural ODEs have demonstrated impressive accuracy for long-horizon predictions of continuous-time systems. Furthermore, they can be trained on trajectories with irregular time intervals between samples [18], and generalize better than multi-layer perceptron models. However, they lack adaptability and need to be retrained for every scenario.

2.2 Function Encoders

To achieve zero-shot transfer, we employ the theory of *function encoders* [14]. While typical machine learning approaches learn a single function, function encoders learn basis functions to span a *space* of functions. This allows the function encoder to achieve zero-shot transfer within this space by identifying the coefficients of the basis functions for any function in the space at runtime. Once the coefficients have been identified, the function can be reproduced as a linear combination of basis functions. Despite the broad applicability of function encoders, the extension to capture solutions to differential equations via neural ODEs is non-trivial.

Formally, consider a function space $\mathcal{F} = \{f \mid f : \mathcal{X} \rightarrow \mathbb{R}^m\}$ where $\mathcal{X} \subset \mathbb{R}^n$. Instead of learning a single function, function encoders learn k basis functions g_1, g_2, \dots, g_k that are parameterized by neural networks in order to span \mathcal{F} [14]. Define the inner product of \mathcal{F} as $\langle f, g \rangle_{\mathcal{F}} = \int \langle f(x), g(x) \rangle dx$. Then, the functions $f \in \mathcal{F}$ can be represented as a linear combination of basis functions,

$$f(x) = \sum_{i=1}^k c_i g_i(x \mid \theta_i), \quad (2)$$

where $c \in \mathbb{R}^k$ are real coefficients and θ_i are the network parameters for g_i . Let V be the volume of \mathcal{X} . For any function $f \in \mathcal{F}$, an empirical estimate of the coefficients c can be calculated using data $\{(x_j, f(x_j))\}_{j=1}^m$ via a Monte-Carlo estimate of the inner product,

$$c_i = \langle f, g_i \rangle \approx \frac{V}{m} \sum_{j=1}^m \langle f(x_j), g_i(x_j \mid \theta_i) \rangle. \quad (3)$$

The basis functions are trained using a set of datasets, $\mathcal{D} = \{D_1, D_2, \dots\}$, where each dataset $D_i = \{(x_j, f_i(x_j))\}_{j=1}^m$ consists of input-output pairs corresponding to a single function $f_i \in \mathcal{F}$. For each function f_i and corresponding dataset D_i , first compute the coefficients $\{c_1, c_2, \dots, c_k\}$ via (3) and then obtain an empirical estimate of f_i via (2). Then compute the error of the estimate of f_i though the norm induced by the inner product of \mathcal{F} . The loss function is simply the sum of the losses for all f_i , which is minimized via gradient descent. For more details, see [14].

After training, the basis functions are fixed, and the coefficients c of a new function $f \in \text{span}\{g_1, \dots, g_k\}$ are computed via (3) or via least-squares, using data collected online. This is key for efficient, online calculations, since the approximation in (3) is effectively a sample mean and requires no gradient calculations.

Orthogonality of the Basis Functions: Note that function encoders do not enforce orthogonality of the basis functions explicitly [14]. Using Gram-Schmidt to orthonormalize the basis functions during training can significantly increase the training time and is computationally intensive. Instead, during training the coefficients are computed using (3) presuming that the basis functions are orthogonal. This is key. This causes the basis vectors to naturally become more orthogonal as training progresses since the loss implicitly penalizes the basis functions if they are not orthonormal. See Appendix I.

3 Function Encoders With Neural ODEs as Basis Functions

Consider a space \mathcal{F} of Lipschitz continuous dynamical systems $f : \mathcal{X} \rightarrow \mathcal{X}$. The space of dynamical systems can arise, for instance, due to uncertain parameters, minor variations in a first-order physics

model, or changing environmental features. Given an initial condition $x(t_0)$, our goal is to estimate the state $x(t_f)$ at a future time $t_f > t_0$. The integral form of the initial value problem is given by (1).

Our approach can be separated into two distinct phases: *offline training* and *zero-shot prediction*. During *offline training*, we presume that we have access to an offline dataset $\mathcal{D} = \{D_1, D_2, \dots\}$, where D_i is a realization of a trajectory from a function $f_i \in \mathcal{F}$. During *zero-shot prediction*, we seek to predict a previously unseen function f and have access to a minimal trajectory D taken from f . We seek to learn a set of basis functions g_1, \dots, g_k that span \mathcal{F} , where k is a user-specified hyper-parameter. By learning a set of basis functions that span the space \mathcal{F} , we obtain a means to represent the behavior of any dynamical system in the space. However, we do not observe measurements of f directly since we cannot typically measure the instantaneous derivative \dot{x} of a dynamical system. Instead, we will equivalently learn a set of neural ODE basis functions such that the underlying neural networks which are being integrated correspond to g_1, \dots, g_k . We then seek to compute a representation of a new function using data collected online.

3.1 Computing a Set of Neural ODE Basis Functions

For every $f \in \mathcal{F}$, we define the integral term in (1) as a function $H : \mathcal{X} \times \mathcal{T} \rightarrow \mathcal{X}$, given by,

$$H(x(t_0), t_f) := \int_{t_0}^{t_f} f(x(\tau)) d\tau. \quad (4)$$

We model the dynamical system f using a function encoder as in (2). Using (2) in (4), and by the linearity of the definite integral, we have that,

$$H(x(t_0), t_f) = \int_{t_0}^{t_f} \left[\sum_{i=1}^k c_i g_i(x(\tau) \mid \theta_i) \right] d\tau = \sum_{i=1}^k c_i \int_{t_0}^{t_f} g_i(x(\tau) \mid \theta_i) d\tau = \sum_{i=1}^k c_i G_i(x(t_0), t_f), \quad (5)$$

where g_i is neural network parameterized by θ_i and $G_i(x(t_0), t_f) := \int_{t_0}^{t_f} g_i(x(\tau) \mid \theta_i) d\tau$. One interpretation of the above equation is that we can represent H as a weighted combination of basis functions G_i , and the problem of learning a set of basis functions g_1, \dots, g_k can equivalently be viewed as learning a set of neural ODEs G_1, \dots, G_k . Thus, we define the Hilbert space \mathcal{H} of functions H as in (4) and equip it with the following inner product,

$$\langle H, G \rangle_{\mathcal{H}} := \int \langle H(z, t), G(z, t) \rangle_{\mathcal{X}} d(z, t). \quad (6)$$

We then learn basis functions G_1, \dots, G_k spanning \mathcal{H} where each basis function is a neural ODE.

From (6), the coefficients of a function $H \in \mathcal{H}$ are given by $c_i = \langle H, G_i \rangle_{\mathcal{H}}$. However, from [14], computing the inner product exactly is generally intractable in high-dimensional spaces. We can empirically estimate the coefficients c_i using a trajectory of (potentially irregularly) sampled states $\{x(t) \mid t = t_0, \dots, t_m\}$ from a dynamics function $f \in \mathcal{F}$. Using the trajectory, we form the dataset $D = \{(x(t_j), x(t_{j+1}))\}_{j=0}^{m-1}$. We can compute c using D via a Monte-Carlo estimate of the inner product in (6),

$$c_i = \langle H, G_i \rangle_{\mathcal{H}} \approx \frac{V}{m} \sum_{j=0}^{m-1} \left\langle x(t_{j+1}) - x(t_j), G_i(x(t_j), t_{j+1}) \right\rangle_{\mathcal{X}}, \quad (7)$$

where V is the volume of the region of integration, and following from (1),

$$x(t_{j+1}) - x(t_j) = \int_{t_j}^{t_{j+1}} f(\tau) d\tau = H(x(t_j), t_{j+1}). \quad (8)$$

In other words, we substitute the difference between states $x(t_{j+1}) - x(t_j)$ for $H(x(t_j), t_{j+1})$ in (7).

Let $\mathcal{D} = \{D_1, D_2, \dots\}$ be a set of datasets, where each $D_\ell = \{(x(t_j), x(t_{j+1}))\}_{j=0}^{m-1}$ is collected from a trajectory from a function $f_\ell \in \mathcal{F}$. For each dataset D_ℓ , we compute the coefficients c_1, \dots, c_k according to (7). The coefficients can be used to approximate the corresponding H_ℓ via (5). We then evaluate the error of H_ℓ using the dataset D_ℓ and minimize its loss via gradient descent. This is done

Algorithm 1 Training Function Encoders with Neural ODE Basis Functions

```

1: Input: Set of datasets  $\mathcal{D}$ , number of basis functions  $k$ , learning rate  $\alpha$ 
2: Output: Neural ODE basis functions  $G_1, G_2, \dots, G_k$ 
3: Initialize  $g_1, g_2, \dots, g_k$  as neural networks with parameters  $\theta = \{\theta_1, \theta_2, \dots, \theta_k\}$ 
4: while not converged do
5:   loss  $L = 0$ 
6:   for all  $D_\ell \in \mathcal{D}$  do
7:     for  $i \in 1, \dots, k$  do
8:        $c_i \approx \frac{V}{m} \sum_{j=0}^{m-1} \langle x(t_{j+1}) - x(t_j), G_i(x(t_j), t_{j+1} - t_j) \rangle x$ 
9:     end for
10:     $L = L + \sum_{j=0}^{m-1} \|(x(t_{j+1}) - x(t_j)) - \sum_{i=1}^k c_i G_i(x(t_j), t_{j+1} - t_j)\|^2$ 
11:   end for
12:    $\theta = \theta - \alpha \nabla_\theta L$ 
13: end while

```

for multiple functions $f \in \mathcal{F}$ at each gradient update to ensure the basis learns the space rather than a single function. We present this as Algorithm 1.

Applying Algorithm 1 yields basis functions which span the space of dynamical systems, where each basis function is a neural ODE. This space describes possible behaviors of the system, where variations in environmental parameters, physics, etc. correspond to a particular dynamics function within this space. Therefore, this algorithm represents complicated system behaviors simply as a vector within a Hilbert space. Section 3.2 shows how to use these basis functions for zero-shot dynamics prediction from small amounts of online data.

3.2 Efficient Online Transfer Without Retraining

After training, we fix the parameters of the basis functions g_1, \dots, g_k , and can compute the coefficient representation $c \in \mathbb{R}^k$ for any function $f \in \text{span}\{g_1, \dots, g_k\}$ via (7). If \mathcal{D} is rich enough to capture the various behaviors of the systems in \mathcal{F} , then we can estimate the behavior of any dynamics $f \in \mathcal{F}$.

Given data collected online from a single trajectory, we can compute the coefficients using the Monte-Carlo estimate of the inner product as in (7). This approximation is a crucial component of the approach. It allows the inner product to be computed from data through an operation that is effectively a sample mean. Therefore, this approach can be computed online quickly even for large amounts of data. Then, given the coefficients, the future states of the system can be predicted using (5). These properties allow the neural ODE to achieve zero-shot transfer. Identifying the coefficients only requires inner product calculations, vector addition, and scalar multiplication, and so it can be computed online without any gradient updates.

The Residuals Method: The zero vector of the coefficients space corresponds to the zero function. Since the feasible dynamics are differentiated by their coefficients, it is numerically convenient if the coefficients corresponding to all feasible systems are centered around zero.

Thus, we can re-center the space of coefficients around the center of the *cluster* of feasible dynamics. This is done by first modeling the average dynamics F_{avg} in the set of datasets \mathcal{D} , and then learning the residuals between each function and F_{avg} . In other words, the basis functions are trained to span the function space corresponding to $R(x(t_0), t_f) = x(t_f) - x(t_0) - F_{avg}(x(t_0), t_f)$. This method can achieve better accuracy, but requires learning one additional neural ODE, F_{avg} . Alternatively, an approximate dynamics model based on prior knowledge can be used as F_{avg} . We describe the training procedure for this approach in Algorithm 2.

3.3 Incorporating Zero-Order Hold Control Inputs

We can account for a zero-order hold (ZOH) control input $u \in \mathcal{U} \subset \mathbb{R}^p$ with minimal modifications. A ZOH control input is given by a piecewise constant function, meaning it is held constant over the period of integration. Given controlled dynamics, $f : \mathcal{X} \times \mathcal{U} \rightarrow \mathcal{X}$, we modify the corresponding functions H to incorporate a constant input, $H(x(t_0), u, t_f) = \sum_{i=1}^k c_i \int_{t_0}^{t_f} g_i(x(t_0), u \mid \theta_i) d\tau$. Then, using trajectory data that also includes the controls applied at each time interval, we can estimate

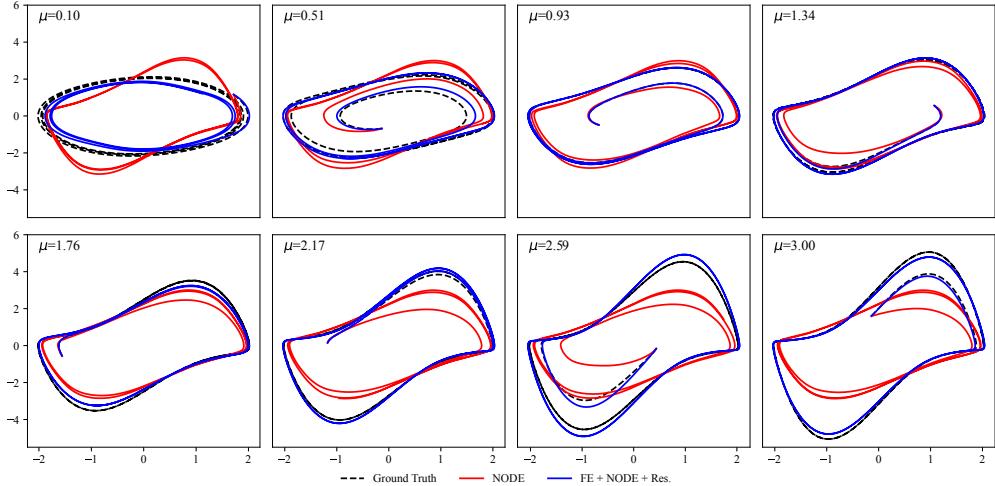


Figure 2: The approximated dynamics for different Van der Pol systems, where the parameter μ is varied. This plot shows that a NODE can only fit a single Van der Pol system, whereas FE + NODE + Res can fit a space of Van der Pol systems from 5000 example data points.

the coefficients using datasets $D = \{(x(t_j), u_j, x(t_{j+1}))\}_{j=0}^{m-1}$, substituting $g_i(x(\tau), u, | \theta_i)$ in (5) and (7). The remainder of the training procedure is unchanged.

4 Numerical Experiments

We demonstrate the effectiveness of our approach for predicting and controlling dynamical systems through several numerical experiments. We first demonstrate that the approach can adapt to different dynamics using a Van Der Pol oscillator system. We then show long-horizon prediction accuracy on challenging MuJoCo robotics experiments and compare to neural ODEs (NODE) [5] and function encoders as in [14] using the residuals method (FE + Res). Lastly, we show the learned models are sufficiently accurate for downstream tasks on a difficult control task using a quadrotor system. The source code is available at <https://github.com/tyler-ingebbrand/NeuralODEFunctionEncoder>.

Current off-the-shelf integrators with adaptive step sizes do not support efficient batch calculations. Because this algorithm involves training numerous neural ODEs on a large amount of data, the ability to train on data in batches is required. Therefore, we implement an RK4 integrator since it can be efficiently computed for multiple data points in parallel. The Van der Pol visualization uses 11 basis functions while the MuJoCo and Drone experiments use 100. For ablations on how the hyper-parameters affect results, see Appendix G.

4.1 Visualization on a Van der Pol Oscillator

We first demonstrate that our approach can adapt to a space of dynamics that vary according to a nonlinear parameter. The Van der Pol dynamics are defined as, $\dot{x} = y$, $\dot{y} = \mu(1 - x^2)y - x$, where $[x, y]^\top \in \mathbb{R}^2$ is the state, and μ is a hidden parameter. We collect multiple datasets D_ℓ where μ is fixed for the duration of the trajectory, but varies between trajectories. We train basis functions using Algorithm 1. We then compute the coefficients via (7) and approximate the dynamics via (5).

We plot the results in Figure 2. As expected, we observe that our proposed approach can predict the dynamics of a space of Van der Pol systems without retraining. We can also see that a single neural ODE trained on the same data can only fit a single function. Therefore, its prediction corresponds most closely with the behavior of a single Van der Pol system that has the mean μ value. This illustrates that our approach is capable of adapting to different dynamics at runtime.

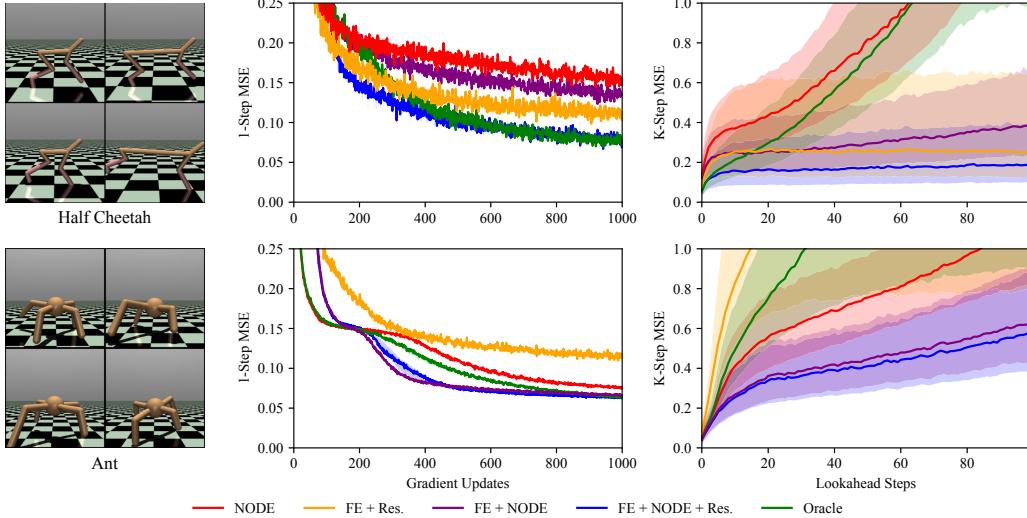


Figure 3: Model performance on predicting the dynamics of MuJoCo robotics environments with hidden parameters. 200 example data points are given to identify dynamics. The results show that $\text{FE} + \text{NODE} + \text{Res.}$ makes accurate, long-horizon predictions even in the presence of hidden parameters. Evaluation is over 5 seeds, shaded regions show the first and third quartiles around the median.

4.2 Long-Horizon Prediction on MuJoCo Environments

We evaluate the performance of our proposed approach on the Half-Cheetah and Ant environments [28], shown in Figure 3. The hidden environmental parameters are the length of the limbs, the friction coefficient, and the control authority. Of the two environments, Ant is more difficult due to its higher degrees of freedom. For training, we collect a dataset of trajectories where the hidden parameters are unobserved, but held constant throughout the duration of a given trajectory. After training, we use 200 datapoints, equivalent to about seven seconds of data for a system running at 30 Hertz, and use only this data to identify the dynamics. Note this online phase is computationally simple, and can be done in only milliseconds on a GPU.

Neural ODEs (NODE) perform poorly because they have no mechanism to condition the prediction on the hidden-parameters. Effectively, NODE learns the mean dynamics over all dynamics functions in the training set. Function encoders using the residuals method (FE + Res) can implicitly condition their predictions on the hidden parameters through the coefficient calculation, though they are unable to achieve accurate long horizon predictions on the more challenging Ant problem. This is because it lacks the inductive bias of neural ODEs. Our approach (FE + NODE) can both implicitly condition the predictions on the hidden parameters through data, but also benefits from the inductive bias of neural ODEs. We see that the residuals method performs best out of all approaches in both environments. This is because the average model significantly reduces the epistemic uncertainty and provides a meaningful baseline from which to center the training. The average model acts as a good inductive bias and makes it easier to distinguish between the learned functions during training. We additionally compare against an oracle prediction approach (Oracle), which has access to the hidden parameters as an additional input with a neural ODE as the underlying architecture. While its 1-step prediction accuracy demonstrates good empirical performance, the long-horizon predictions are unstable. This is because Oracle is required to generalize to an entire space of dynamics with one NODE, which is a complex and difficult function to learn.

4.3 Realistic Robotics Experiments and Control of a Quadrotor System

Lastly, we seek to test the accuracy of our approach for use on downstream tasks such as control on a realistic example using a robotic system. We seek to determine if the learned models are sufficiently accurate for model-based control in the presence of hidden parameters. We use a simulated quadrotor system using PyBullet [31], which is a highly nonlinear control system. We use the quadrotor’s mass as a hidden parameter. The goal is to predict the future state of the quadrotor system under any hidden

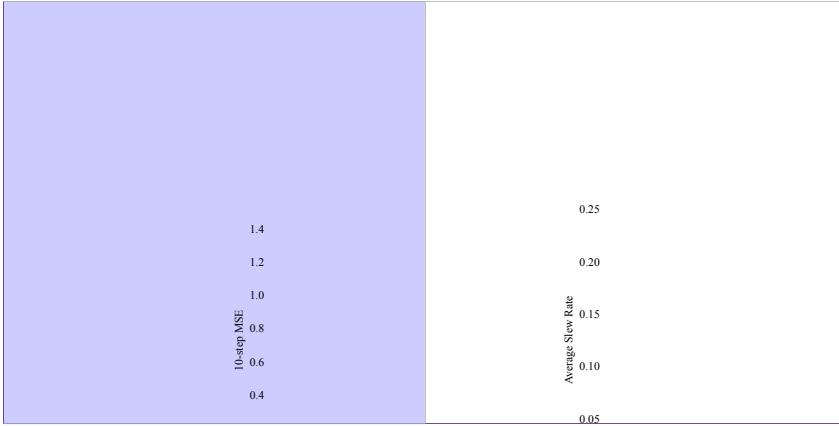


Figure 4: Model performance on the PyBullet quadrotor environment with varying mass. Function encoders improve model performance across varying masses. Shaded region is 1st and 3rd quartiles over 200 trajectories (left) and over 5 trajectories (middle, right).

Figure 5: Qualitative analysis of the difference in control between NODEs and our approach. Two trajectories with the same initial position but different masses are shown. NODE is unaware of the mass, and so its z position requires constant correction. In contrast, FE + NODE (+Res) accounts for the mass through the coefficients, meaning it is more accurate and requires fewer corrections.

parameters, using 2000 data points collected online to identify the dynamics. Then, given the learned model of system behavior, we seek to control the quadrotor using gradient-based model predictive control (MPC) to reach a pre-specified hover point. The results are plotted in Figure 4.

The results show that FE + NODE + Res outperforms competing approaches at long-horizon predictions. Furthermore, we plot the 10-step MSE as a function of mass in Figure 4, and we observe that FE + NODE + Res accurately predicts the system behavior across varying masses. We observe a slight decay in performance for low masses which are more sensitive to control inputs during simulation, which causes the simulated trajectories to diverge more from the rest of the observed data. The neural ODE (NODE) performs poorly for different masses, and its performance decays quickly as the dynamics deviate from the mean behavior.

Lastly, we see that this prediction accuracy translates to the downstream performance of an MPC controller. While all approaches are sufficiently accurate for control due to the fact that MPC is partially robust to model inaccuracies, the prediction accuracy has a corresponding impact on the task performance. Neural ODEs (NODE) demonstrate a high slew rate, which reflects the need for repeated positional corrections necessitated by taking bad actions. In contrast, FE + NODE and FE + NODE + Res have lower slew rates as they make more accurate decisions. We plot two example trajectories that demonstrate this behavior in Figure 5.

5 Scope & Limitations

Overhead: Our approach incurs a cost of either increased inference time or memory, depending on if the basis functions are integrated sequentially or in parallel. We integrate them sequentially during

training to reduce memory overhead and allow for larger batch sizes, while integrating in parallel at execution to prioritize inference speed. See Appendix D.

Data dependency: In order to make efficient, online approximations of a new function $f \in \mathcal{F}$ without gradient calculations, the basis functions must be trained to span the space of possible dynamics. To do so, there must be sufficient example datasets of possible dynamics under fixed hidden parameters in \mathcal{D} . This implies a larger amount of data must be collected to learn a space of dynamics than would be needed to learn a single dynamics function.

Integration: The training procedure trains the basis functions on short time intervals $t_f - t_0$, in which $x(t) \in \mathcal{X}$. The basis functions have only been trained for inputs in the space of \mathcal{X} , where their behavior outside of \mathcal{X} is unpredictable. As a result, it is necessary to either integrate each basis function for a short time interval before calculating the state according to (5), or to integrate the basis functions as described in B. Integrating the basis functions over long horizons without calculating the state of the system during intermediate steps may lead the predicted state to leave \mathcal{X} , at which point the behavior of that basis function becomes unpredictable.

6 Related Work

Basis Functions: Function approximation techniques often employ a linear model over a predefined set of basis functions. Techniques such as Taylor series, Fourier series, and orthogonal polynomial systems utilize an infinite set of basis functions, theoretically allowing perfect function representation [6, 23, 7]. However, in high-dimensional spaces, these techniques become impractical due to the exponential growth in the number of basis functions. Additionally, many approaches depend on the choice of a feature map or kernel to define the function space [29, 25], which imposes structure by selecting the class of functions to learn from (e.g. using radial basis functions [3]). These design choices necessarily introduce approximation errors through the choice of function class, or may depend on prior domain knowledge, which may not always be available. Methods for identifying system dynamics that use a large library of pre-defined basis functions, such as Koopman operators [2] or nonlinear system identification through sparse regression such as SINDy [4, 26, 15], have received considerable attention. Yet these approaches typically employ a finely-crafted finite dictionary of basis functions, which requires careful choice to achieve good data-driven performance [20]. Neural network approaches such as functional-link and orthogonal networks [30, 7] omit hidden layers and use gradient descent to learn a linear combination of features, encoding the function class into the network architecture, but fail to generalize well, and are not amenable to zero-shot transfer. In contrast, we compute the coefficients of the model through a well-defined inner product, which scales well with data and can be computed quickly. Furthermore, our basis functions are entirely learned from data during the training phase, similar to representation learning [1], and thus require no prior assumptions or domain knowledge.

Neural ODEs: In existing work, neural ODEs have proven to be a powerful tool for modeling dynamical systems [10, 24, 22, 11, 17] and stochastic differential equations [21, 13], but generally require an extensive data collection and training phase. While the model training can be enhanced [8] and the models can incorporate prior knowledge [9, 10] to reduce the training time, they inherently focus on a single system at a time. This inherently limits their ability to generalize across different systems without retraining. Notably, parameterized neural ODEs [19] pass the model parameters as an additional input to the neural ODE. This approach has been shown to achieve a form of transfer within the set of allowable parameters, but requires extensive knowledge of the system parameters and the structure of the dynamics, both at training and test time. In principle, our approach can be combined with these existing approaches to incorporate their distinct advantages, making our approach highly generalizable to different systems and modeling frameworks.

Deep Learning Techniques: Few-shot meta learning aims to solve a similar problem, where a learned model is adapted given an online dataset [12]. However, meta learning requires gradient updates to the learned models, which may be too slow for real-time control. Transformers are another technique that can adapt given an online dataset by feeding that dataset as input to the encoder side of the transformer. However, transformers have long forward pass times and scale quadratically with the amount of data [16], and so they are not amenable for model-based control. Domain randomization in reinforcement learning is another technique to generate a policy which is robust to a large set

of dynamics [27]. In contrast, our dynamics model *adapts* to the current dynamics, and thus our controller is adaptive, rather than robust.

7 Conclusion & Future Work

We introduced zero-shot neural ODEs, which accomplish both long-horizon predictions and zero-shot transfer. We demonstrated the performance of this approach on two challenging MuJoCo tasks and on the control of a quadrotor system. Our approach makes a significant step towards online adaptability of model-based control and has implications for the safe control of autonomous systems in the presence of uncertainty. In future work, we plan to address safety during training, perhaps using the properties of the Hilbert space to characterize the epistemic uncertainty. We also plan to explore theoretical extensions to stochastic differential equations and Hilbert spaces of probability measures.

8 Broader Impact

This approach demonstrates several clear benefits for enabling same-day adaptation, which is a critical need for autonomous systems that will be deployed in new, unstructured environments. Nevertheless, this approach will require a more thorough theoretical analysis before it can be deployed on actual robotics systems, e.g. to determine confidence or sample bounds to guarantee safety. Notably, this work is a step toward bridging the sim-to-real gap, though it remains unclear how well real-world systems will be represented by a set of basis functions learned in simulation.

9 Acknowledgements

Thank you to Dr. Cyrus Neary for helpful discussions. This material is based upon work supported by the National Science Foundation under NSF Grant Number 2214939. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation. This material is based upon work supported by the Air Force Office of Scientific Research under award number AFOSR FA9550-19-1-0005. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the U.S. Department of Defense.

References

- [1] Y. Bengio, A. Courville, and P. Vincent. Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8):1798–1828, 2013.
- [2] P. Bevanda, S. Sosnowski, and S. Hirche. Koopman operator dynamical models: Learning, analysis and control. *Annual Reviews in Control*, 52:197–212, 2021.
- [3] D. Broomhead and D. Lowe. Radial basis functions, multi-variable functional interpolation and adaptive networks. *Royal Signals and Radar Establishment Malvern (United Kingdom)*, RSRE-MEMO-4148, 03 1988.
- [4] S. L. Brunton, J. L. Proctor, and J. N. Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 113(15):3932–3937, 2016.
- [5] R. T. Q. Chen, Y. Rubanova, J. Bettencourt, and D. K. Duvenaud. Neural ordinary differential equations. In *Advances in Neural Information Processing Systems*, volume 31, 2018.
- [6] C. Cheng, Z. Peng, W. Zhang, and G. Meng. Volterra-series-based nonlinear system modeling and its engineering applications: A state-of-the-art review. *Mechanical Systems and Signal Processing*, 87A:340–364, 2017.
- [7] S. Dehuri and S.-B. Cho. A comprehensive survey on functional link neural networks and an adaptive PSO–BP learning for CFLNN. *Neural Computing and Applications*, 19(2):187–205, 2010.

- [8] F. Djeumou, C. Neary, E. Goubault, S. Putot, and U. Topcu. Taylor-lagrange neural ordinary differential equations: Toward fast training and evaluation of neural odes. In *IJCAI*, 2022.
- [9] F. Djeumou, C. Neary, E. Goubault, S. Putot, and U. Topcu. Neural networks with physics-informed architectures and constraints for dynamical systems modeling. In *L4DC*, 2022.
- [10] F. Djeumou, C. Neary, and U. Topcu. How to learn and generalize from three minutes of data: Physics-constrained and uncertainty-aware neural stochastic differential equations. In *CoRL*, 2023.
- [11] E. Dupont, A. Doucet, and Y. W. Teh. Augmented neural odes. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [12] C. Finn, P. Abbeel, and S. Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning, ICML*, volume 70, pages 1126–1135. PMLR, 2017.
- [13] L. Hodgkinson, C. van der Heide, F. Roosta, and M. W. Mahoney. Stochastic continuous normalizing flows: training SDEs as ODEs. In *Proceedings of the Thirty-Seventh Conference on Uncertainty in Artificial Intelligence*, volume 161 of *Proceedings of Machine Learning Research*, pages 1130–1140. PMLR, 27–30 Jul 2021.
- [14] T. Ingebrand, A. Zhang, and U. Topcu. Zero-shot reinforcement learning via function encoders. In *ICML*. ICML, 2024.
- [15] K. Kaheman, J. N. Kutz, and S. L. Brunton. SINDy-PI: A robust algorithm for parallel implicit sparse identification of nonlinear dynamics. *CoRR*, 2020.
- [16] F. D. Keles, P. M. Wijewardena, and C. Hegde. On the computational complexity of self-attention. In *International Conference on Algorithmic Learning Theory*, volume 201 of *Proceedings of Machine Learning Research*, pages 597–619, 2023.
- [17] J. Kelly, J. Bettencourt, M. J. Johnson, and D. K. Duvenaud. Learning differential equations that are easy to solve. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 4370–4380. Curran Associates, Inc., 2020.
- [18] P. Kidger, J. Morrill, J. Foster, and T. Lyons. Neural controlled differential equations for irregular time series. In *Advances in Neural Information Processing Systems*, volume 33, pages 6696–6707, 2020.
- [19] K. Lee and E. J. Parish. Parameterized neural ordinary differential equations: Applications to computational physics problems. *Proceedings of the Royal Society A*, 477(2253), 2021.
- [20] Q. Li, F. Dietrich, E. M. Boltt, and I. 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.
- [21] X. Liu, T. Xiao, S. Si, Q. Cao, S. Kumar, and C. Hsieh. Neural SDE: stabilizing neural ODE networks with stochastic noise. *CoRR*, 2019.
- [22] Y. Lu, A. Zhong, Q. Li, and B. Dong. Beyond finite layer neural networks: Bridging deep architectures and numerical differential equations. In *ICML*, 2018.
- [23] J. Patra and A. Kot. Nonlinear dynamic system identification using Chebyshev functional link artificial neural networks. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 32(4):505–511, 2002.
- [24] M. Raissi, P. Perdikaris, and G. E. Karniadakis. Multistep neural networks for data-driven discovery of nonlinear dynamical systems, 2018.
- [25] B. Schölkopf and A. J. Smola. *Learning with Kernels: support vector machines, regularization, optimization, and beyond*. Adaptive computation and machine learning series. MIT Press, 2002.

- [26] J. N. K. Steven L. Brunton, Joshua L. Proctor. Sparse identification of nonlinear dynamics with control (SINDYC). *IFAC-PapersOnLine*, Volume 49, Issue 18:710–715, 2016.
- [27] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel. Domain randomization for transferring deep neural networks from simulation to the real world. In *IROS*, pages 23–30. IEEE, 2017.
- [28] E. Todorov, T. Erez, and Y. Tassa. MuJoCo: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5033, 2012.
- [29] C. K. Williams and C. E. Rasmussen. *Gaussian processes for machine learning*. MIT press Cambridge, MA, 2006.
- [30] S.-S. Yang and C.-S. Tseng. An orthogonal neural network for function approximation. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 26(5):779–785, 1996.
- [31] Z. Yuan, A. W. Hall, S. Zhou, L. Brunke, M. Greeff, J. Panerati, and A. P. Schoellig. Safe-control-gym: A unified benchmark suite for safe learning-based control and reinforcement learning in robotics. *IEEE Robotics and Automation Letters*, 7(4):11142–11149, 2022.

A Hardware

All experiments use an Intel 9th Generation i9 CPU and a Nvidia 2060 GPU with 6GB of memory.

B Faster Integration

When computing the approximate inner product between the true system and each basis function, it is necessary to compute $\tilde{G}_i(x(t_0), t_f)$ for every tuple in the dataset D . However, once the coefficients c have been computed, it is, in theory, no longer necessary to integrate the basis functions separately. From (5), we can approximate H as

$$H(x(t_0), t_f) = \int_{t_0}^{t_f} \left[\sum_{i=1}^k c_i g_i(x(\tau) \mid \phi_i) \right] d\tau, \quad (9)$$

by the linearity of the integral. In effect, we are summing the gradients of each basis function rather than the basis functions themselves. As a result, inference for a specific set of coefficients can be decreased from requiring k integrations to only a single integration. However, we find this method to make less accurate predictions in practice, which may be due to our choice of integrator. This trick may be better suited to variable step-size integrators, which benefit more from reduced calls to the integrator than RK4 does.

C Method of Integration

We leverage RK4 as the default integrator for this work, as it can run a forward pass in milliseconds. There are more accurate integrators available, such as `odeint`. However, there is inherently a trade off with respect to both training time and execution time. Integrators such as adaptive step size solvers can potentially make 20 or more calls to the neural ODE during a forward pass, while RK4 makes only 4. The increased number of neural ODE forward passes greatly increases memory usage and compute time. We experimented with more accurate integrators, but ultimately found this tradeoff to be unfavorable. Future work should investigate integrators that are fast, but achieve better accuracy than RK4.

D Overhead

While the function encoder algorithm alone has minimal overhead relative to a MLP, this is not the case for neural ODEs due to the need for integration. The k neural ODEs may either be integrated sequentially or in parallel. If they are integrated sequentially, the memory overhead is lower, especially with respect to back-propagation. Thus, we find this useful for *offline training*, where the sequential method allows us to compute gradients for a larger batch of data at the cost of training time. In contrast, *online execution* generally favors inference speed over memory overhead. Thus, we use the parallel method for online inference in the drone example, which requires much more memory but only a small overhead of inference time relative to neural ODEs alone.

E Residuals Method Algorithm

Training a function encoder with the residuals method requires two separate loss functions. The first loss function trains F_{avg} on data from all datasets D_ℓ , which means it effectively learns the expectation of F given the training set. This loss function can be skipped if F_{avg} is a fixed function based on prior knowledge.

The second loss function trains the basis functions. Unlike in Algorithm 1, the function being learned is $x(t_{j+1}) - x(t_j) - F_{avg}(x(t_j), t_{j+1} - t_j)$. In other words, the residual between the data and the average function. This loss is only used to train the basis functions, it is **not** used to train the average function. See Algorithm 2.

Algorithm 2 The Residuals Method

```
1: Input: Set of datasets  $\mathcal{D}$ , number of basis functions  $k$ 
2: Output: Average function  $F_{avg}$  and Neural ODE basis functions  $G_1, G_2, \dots, G_k$ 
3: Initialize  $f_{avg}$  and  $g_1, g_2, \dots, g_k$  as neural networks with parameters  $\bar{\theta}$  and  $\theta = \{\theta_1, \theta_2, \dots, \theta_k\}$ 
4: while not converged do
5:   // Train Average Function
6:   loss  $L_1 = 0$ 
7:   for all  $D_\ell \in \mathcal{D}$  do
8:      $L_1 = L_1 + \sum_{j=1}^{m-1} \|(x(t_{j+1}) - x(t_j)) - F_{avg}(x(t_j), t_{j+1} - t_j)\|^2$ 
9:   end for
10:   $\bar{\theta} = \bar{\theta} - \alpha \nabla_{\bar{\theta}} L_1$ 
11:  // Train Basis Functions
12:  loss  $L_2 = 0$ 
13:  for all  $D_\ell \in \mathcal{D}$  do
14:    for  $i \in 1, \dots, k$  do
15:       $c_i \approx \frac{V}{m-1} \sum_{j=1}^{m-1} \langle (x(t_{j+1}) - x(t_j) - F_{avg}(x(t_j), t_{j+1} - t_j), G_i(x(t_j), t_{j+1} - t_j) \rangle x$ 
16:    end for
17:     $L_2 = L_2 + \sum_{j=1}^{m-1} \|(x(t_{j+1}) - x(t_j) - F_{avg}(x(t_j), t_{j+1} - t_j) - \sum_{i=1}^k c_i G_i(x(t_j), t_{j+1} - t_j)\|^2$ 
18:  end for
19:   $\theta = \theta - \alpha \nabla_{\theta} L_2$ 
20: end while
```

F Implementation Details

All baselines use the same training scheme. We use an ADAM optimizer with a learning rate of $1e - 3$, and gradient clipping with a max norm of 1. NODE baselines uses 4 hidden layers of size 512, while FE + NODE baselines uses 4 hidden layers of size 51 for each basis function. Note this leads to approximately the same number of parameters for both approaches because the number of hidden parameters scales quadratically with the size of the hidden layers. All baselines train on 50 functions per gradient update via gradient accumulation. States are normalized to have 0 mean and unit variance.

A random policy is used to collect data for the MuJoCo environments. A PID-based exploratory policy, which moves to random nearby points, is used to collect data for the quadrotor since a random policy collides with the floor. Evaluations are done on a holdout set collected through the same means.

All quadrotor baselines use the same MPC controller. The controller optimizes the actions through a combined sampling, gradient descent over 100 iterations. The episode is 100 steps, while the planning horizon is 10 steps. The controller optimizes 100 sample trajectories in parallel, and ultimately chooses the best one. Warm starting is used for following MPC calls to improve performance. The cost function penalizes distance to the objective point, deviance from a stable horizontal position, velocity, and the difference between torques on each rotor.

G Hyper-Parameter Ablations

G.1 Number of Basis Functions

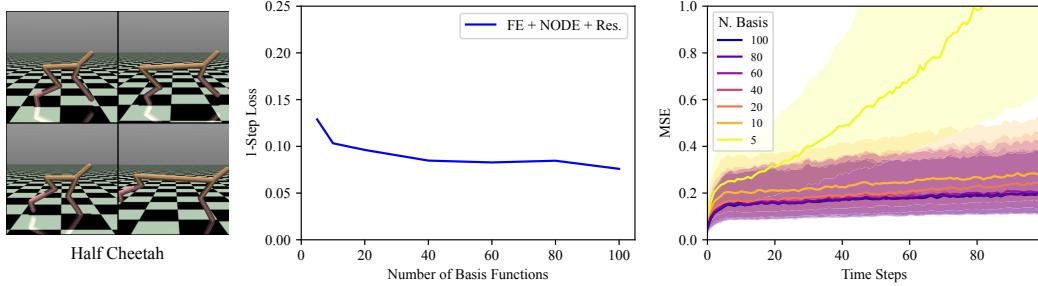


Figure 6: We ablate the effect of the number of basis functions (k) on the performance of the learned model. Results are shown for the FE + NODE + Res. algorithm applied to the Half Cheetah environment. The results indicate the the proposed approach is insensitive to the number of basis functions around $k = 100$, while performance eventually decays as k approaches 0.

G.2 Number of Example Data Points

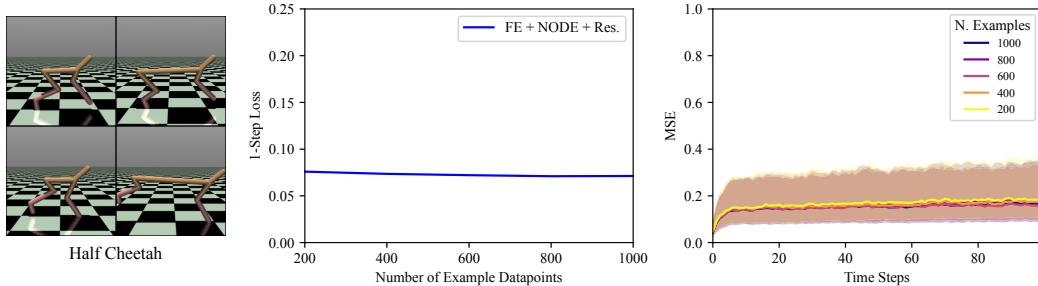
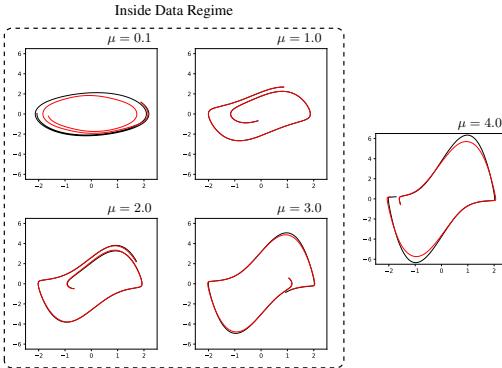


Figure 7: We ablate the effect of the number of example data points on the performance of the learned model. Results are shown for the FE + NODE + Res. algorithm applied to the Half Cheetah environment. The results indicate the the proposed approach is insensitive to increasing example dataset sizes, which suggests that 200 data points is sufficient for the coefficients to converge.

H Generalization

Figure 3: This figure shows the generalization capabilities of the proposed method. The black line indicates the ground truth Van Der Pol dynamics, and the red line shows an approximation. The model was trained on $\mu \in [0.1, 3.0]$. The left side of the figure shows Van Der Pol dynamics for values of μ that are within the distribution of training environments, though each environment is unseen. The right shows the approximation for $\mu = 4.0$, which lies outside of the training distribution. The figure shows that the function encoder is able to reasonably generalize outside of its training set in this example.



I Orthonormality

Consider a set of basis functions g_1, \dots, g_k . Suppose that g_1, \dots, g_k is not orthonormal. Now consider a function f , and suppose f happens to be in the span of g_1, \dots, g_k . Then f can be expressed as $f = b^\top g$, where b is a set of coefficients and g is the concatenation of g_1, \dots, g_k . The coefficients are calculated via the inner product,

$$c^\top = \begin{bmatrix} \langle f, g_1 \rangle \\ \vdots \\ \langle f, g_k \rangle \end{bmatrix} = \begin{bmatrix} \langle b^\top g, g_1 \rangle \\ \vdots \\ \langle b^\top g, g_k \rangle \end{bmatrix} = b^\top \begin{bmatrix} \langle g_1, g_1 \rangle & \dots & \langle g_1, g_k \rangle \\ \vdots & \ddots & \vdots \\ \langle g_k, g_1 \rangle & \dots & \langle g_k, g_k \rangle \end{bmatrix}$$

The loss function $L = |f - \hat{f}|^2 = |f - c^\top g|^2$. If $c = b$, then the loss will be 0. Observe that $c = b$ if and only if the Gram matrix is identity, and the Gram matrix is identity only for an orthonormal basis. In other words, the minimizer of the loss function is an orthonormal basis. Thus, in order for gradient descent to decrease loss, the basis functions converge towards orthonormality [14]. This intuition is empirically validated in [14], Appendix A.5.

As a final note, the coefficients can be computed via least squares after training. Least squares does not require an orthonormal basis as it uses the Gram matrix to account for the inner products between basis functions.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: We claim to learn a space of dynamics spanned by a set of neural ODEs. These claims are justified in the methods §3, where we show theoretical results, and empirically validated in §4.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: We discuss the key limitations of our approach in §5, including: computational overhead, dependence on training data, and the constraints on integration.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: Our method does not rely upon any new theorems or proofs. We reference existing results and theorems where appropriate in the text.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We provide a comprehensive description of the experimental results in §4, including: the architecture, integration scheme, the number of basis functions used, and descriptions of the data used for training. Further implementation details are provided in the appendix in §F.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We provide code as a zip file in the initial submission. A link to a github repository will be provided in the final version.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: Implementation details are provided in the experimental results section §4 and comprehensive details are provided in the appendix §F.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: We provide confidence intervals for the first and third quartile in all figures. The significance of the error bars with respect to parameter variation and random seeds is thoroughly described in the text.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer “Yes” if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [\[Yes\]](#)

Justification: We describe the computing resources used to compute all results in the appendix in §A. Computational overhead is discussed in §5.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [\[Yes\]](#)

Justification: Our work conforms with the NeurIPS guidelines and Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [\[Yes\]](#)

Justification: Our work has positive implications for the design and modeling of autonomous systems. We briefly discuss the broader impacts of our work in §8

Guidelines:

- The answer NA means that there is no societal impact of the work performed.

- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: We do not provide a model or dataset with a high risk for misuse.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [NA]

Justification: Our experiments do not rely upon any restricted licenses or code, and we provide citations to key libraries used.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.

- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: We do not release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: Our paper does not involve crowdsourced data or research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: Our paper does not involve crowdsourced data or research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.

- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.