Zero-Shot Transfer of Neural ODEs

TYLER INGEBRAND, ADAM J. THORPE, UFUK TOPCU

Function Encoders: A visualization

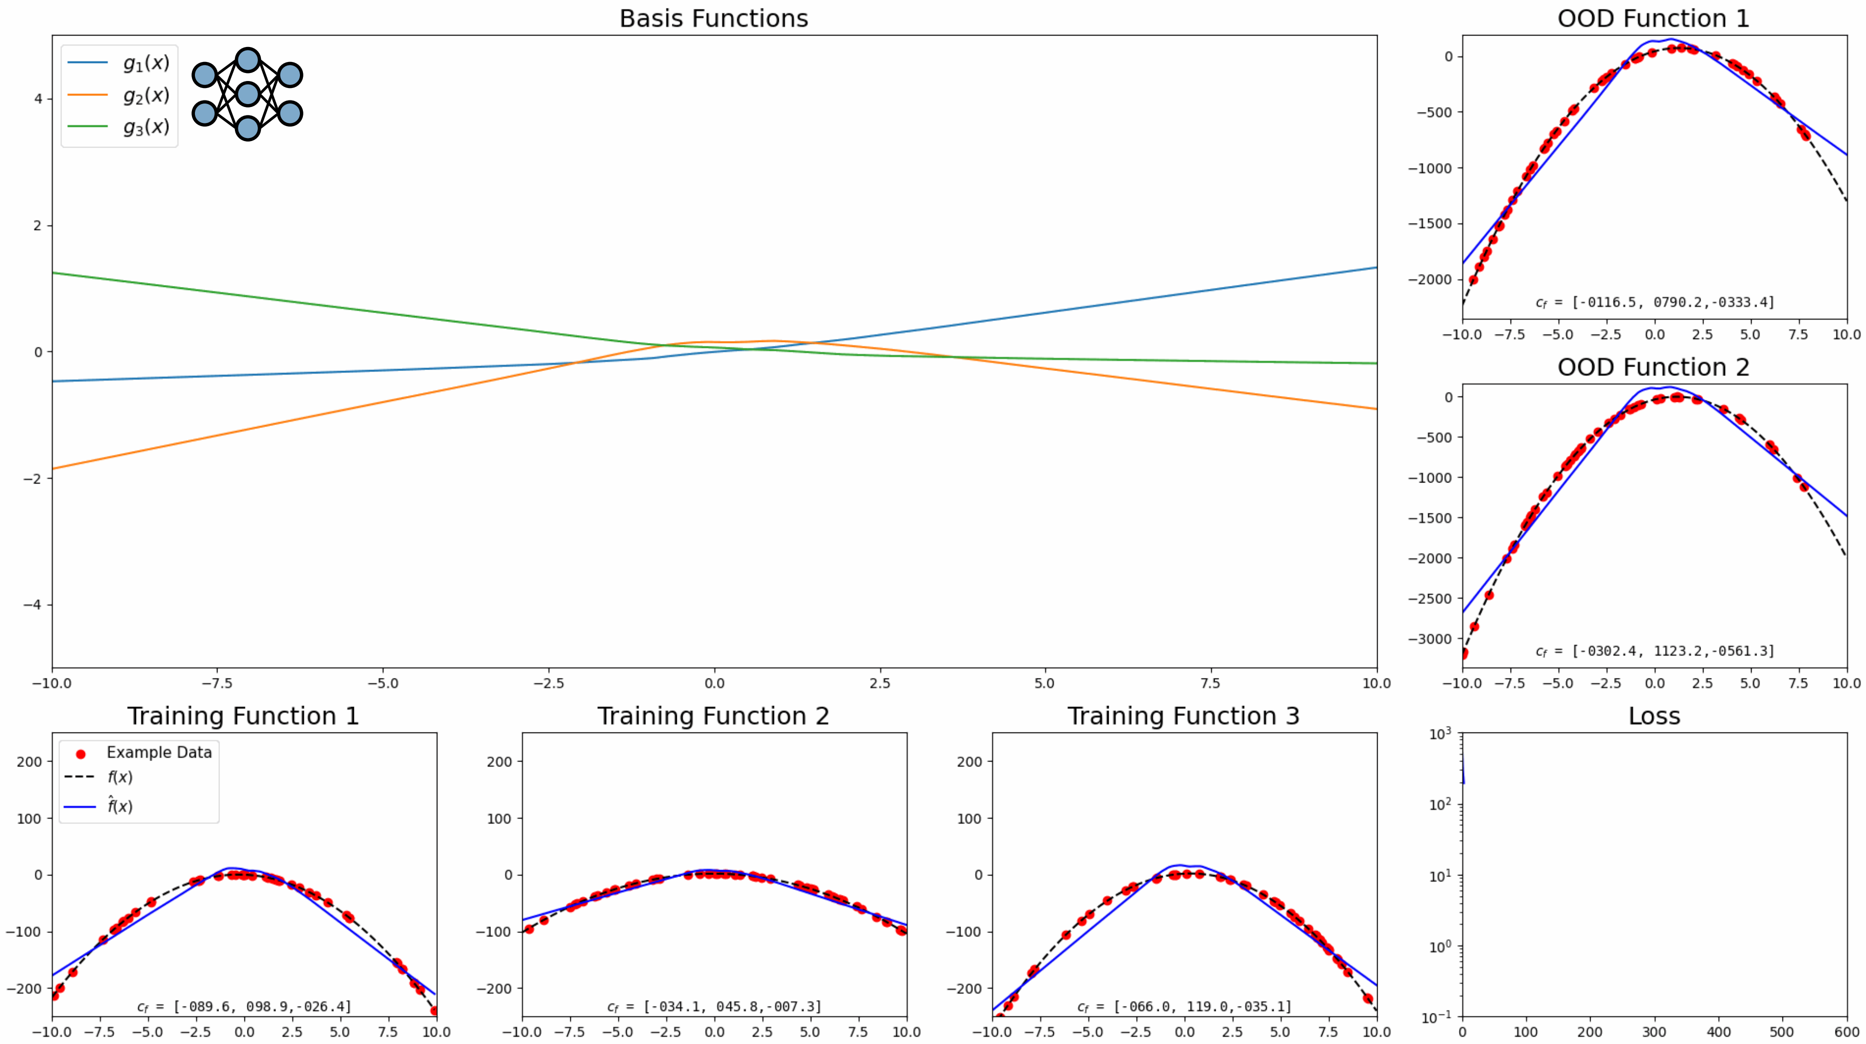

Improving Accuracy with Neural ODEs

$$x_{t+1} = f_\theta\,(x_t, u_t)$$

$$x(t_f) - x(t_0) = \int_{t_0}^{t_f} f_\theta(x(t), u(t))dt$$

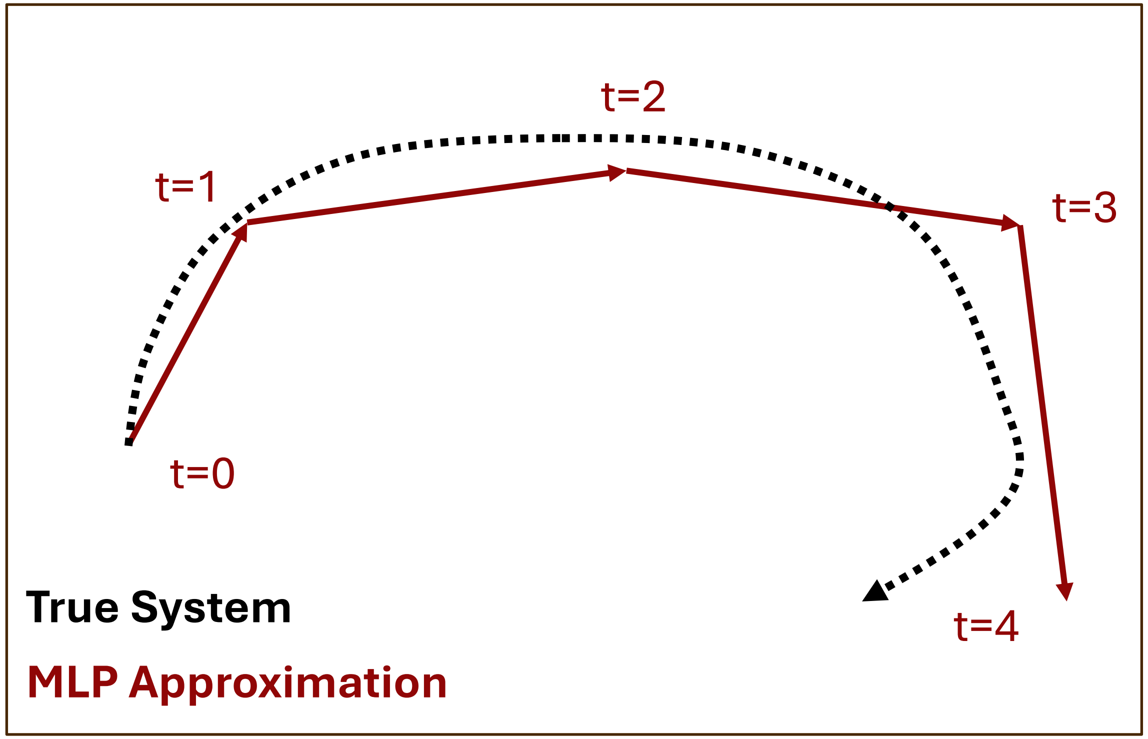

Modeling a dynamical system with a MLP predicts the state *only* at discrete time points.

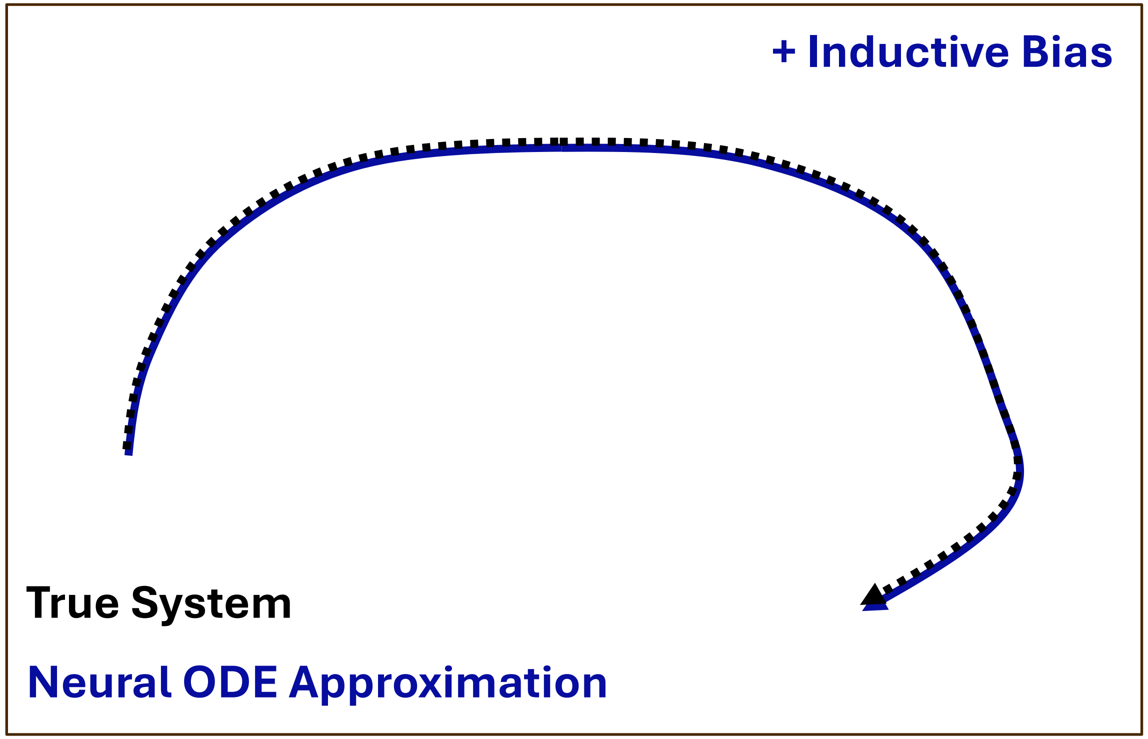

Modeling a dynamical system with a neural ODE predicts the state at *all* time points.

Hidden-Parameter Dynamics Prediction

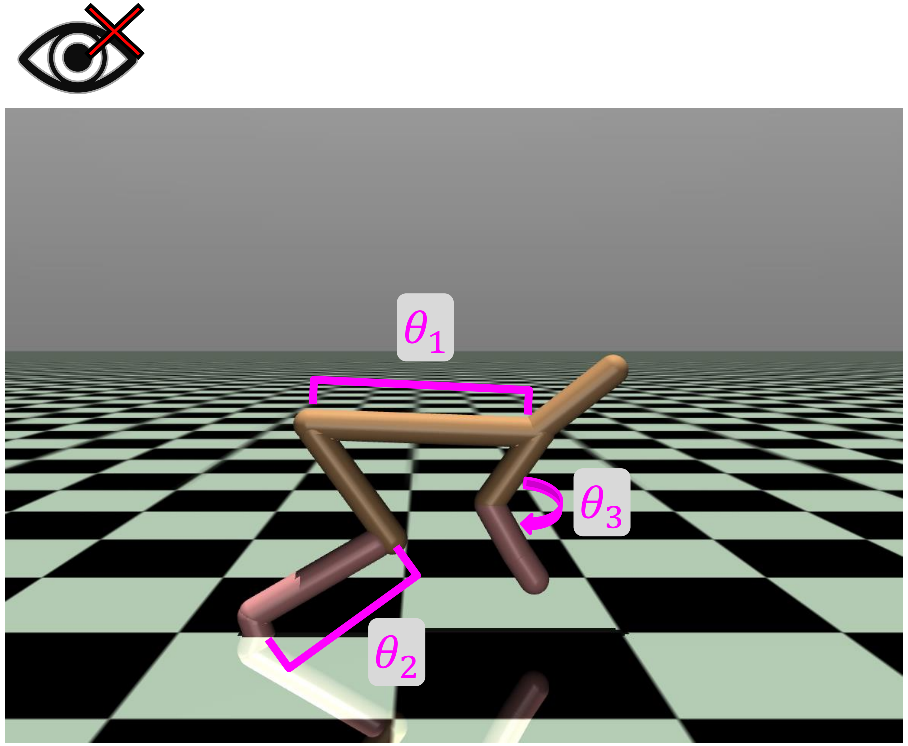

Segment lengths, control authority, and friction are varied every episode. Transition dynamics depend on θ, e.g. $s_{t+1} = T(s_t, a_t | \theta)$

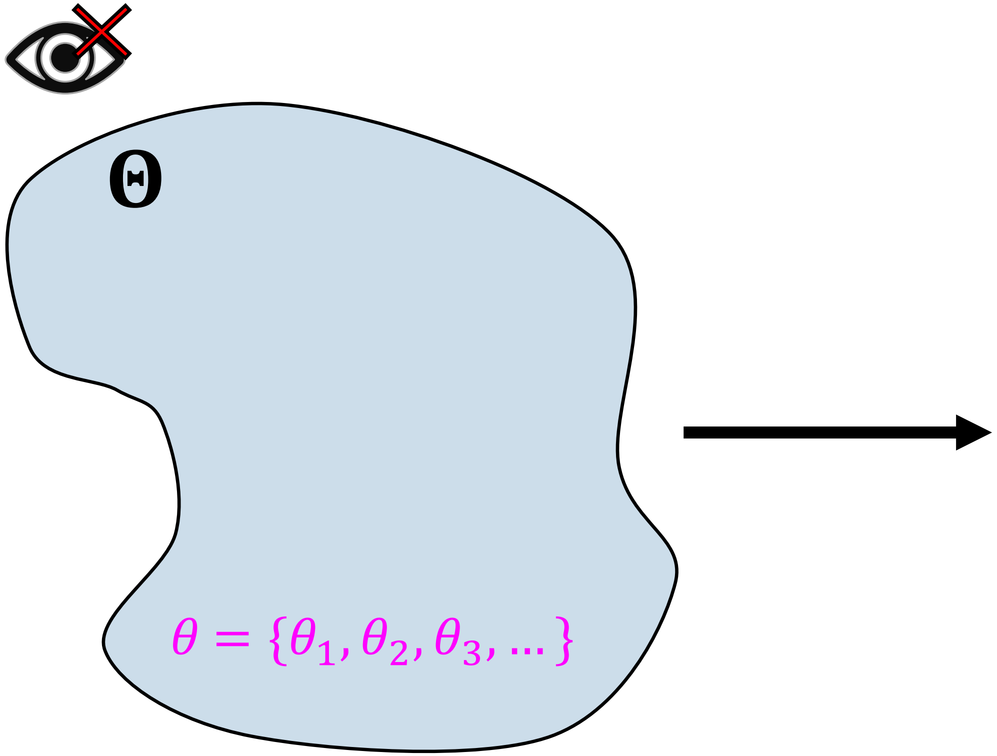

The set of possible hidden parameters is unknown

Hidden-Parameter Dynamics Prediction

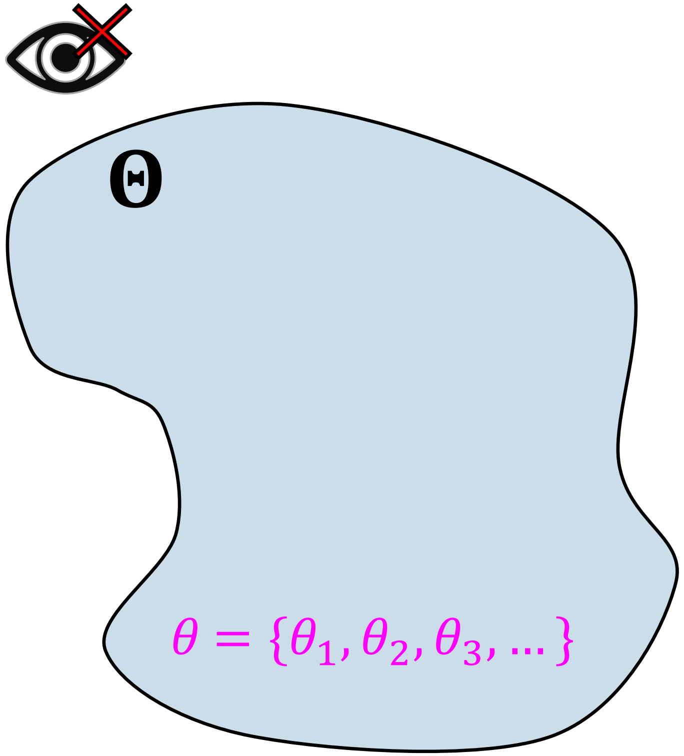

The set of possible hidden parameters is unknown

$T_\theta(s, a) = \mathrm{T}(\mathrm{s}, \mathrm{a}|\theta)$

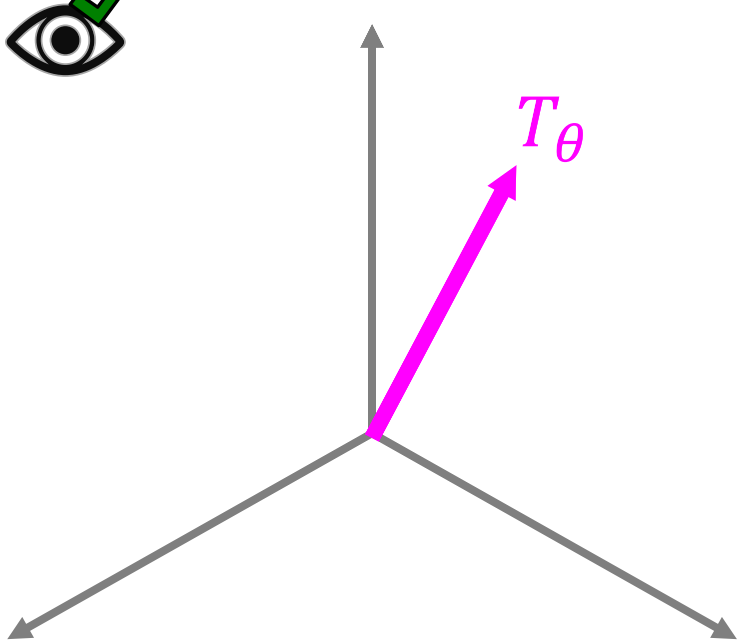

The induced dynamics function is observable each episode via data $\{s_t, a_t, T_\theta(s_t, a_t)\}_{t=1}^{H}$

Hidden-Parameter Dynamics Prediction

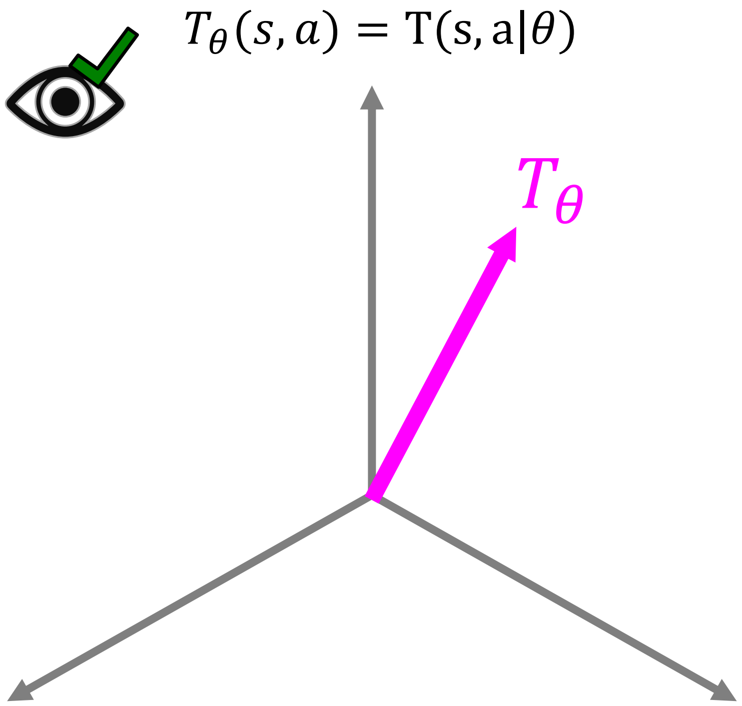

The induced dynamics function is observable each episode via data $\{s_t, a_t, T_\theta(s_t, a_t)\}_{t=1}^{H}$

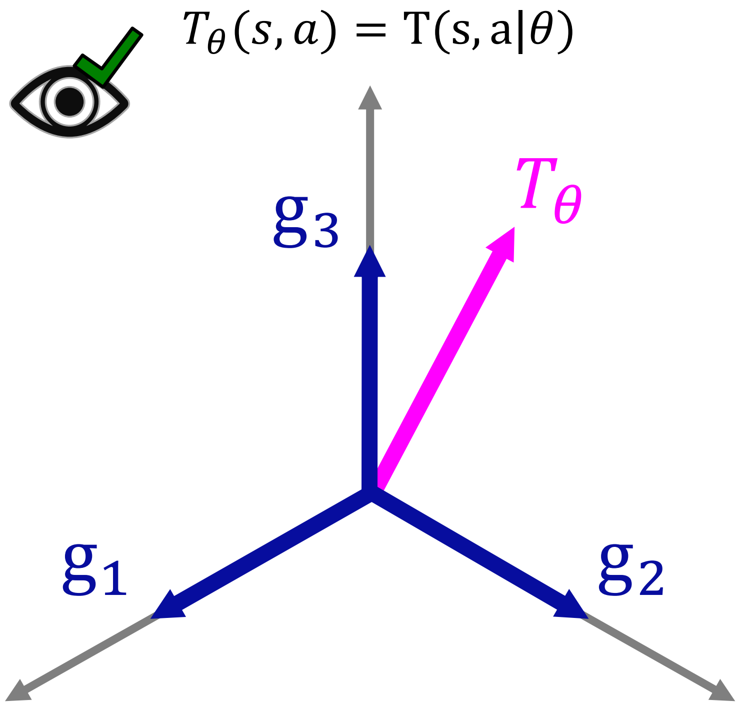

Given data over many episodes, we can learn basis functions over the space of induced dynamics, where $g_i: S \times A \to S$

Hidden-Parameter Dynamics Prediction

$T_\theta(s, a) = \mathrm{T}(\mathrm{s}, \mathrm{a}|\theta)$

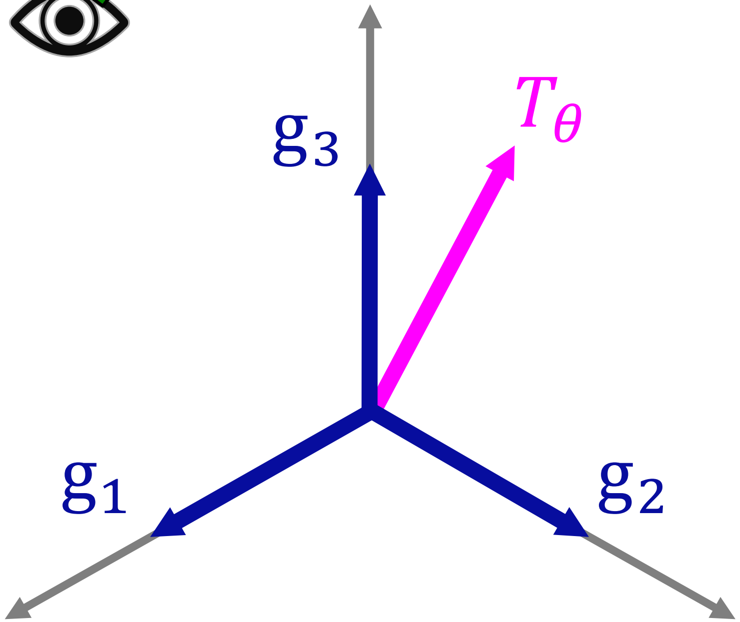

Given data over many episodes, we can learn basis functions over the space of induced dynamics, where $g_i: S \times A \to S$

Small, online dataset

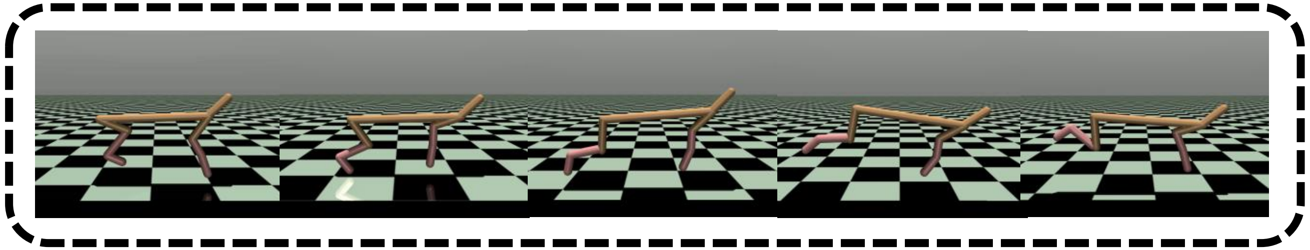

Future states of the episode

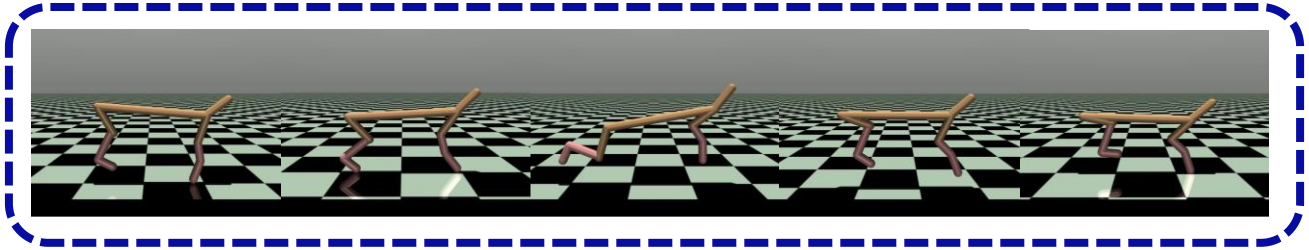

Once we have learned basis functions, can we predict dynamics after collecting a small amount of online data?

Hidden-Parameter Dynamics Prediction w/ Neural ODEs

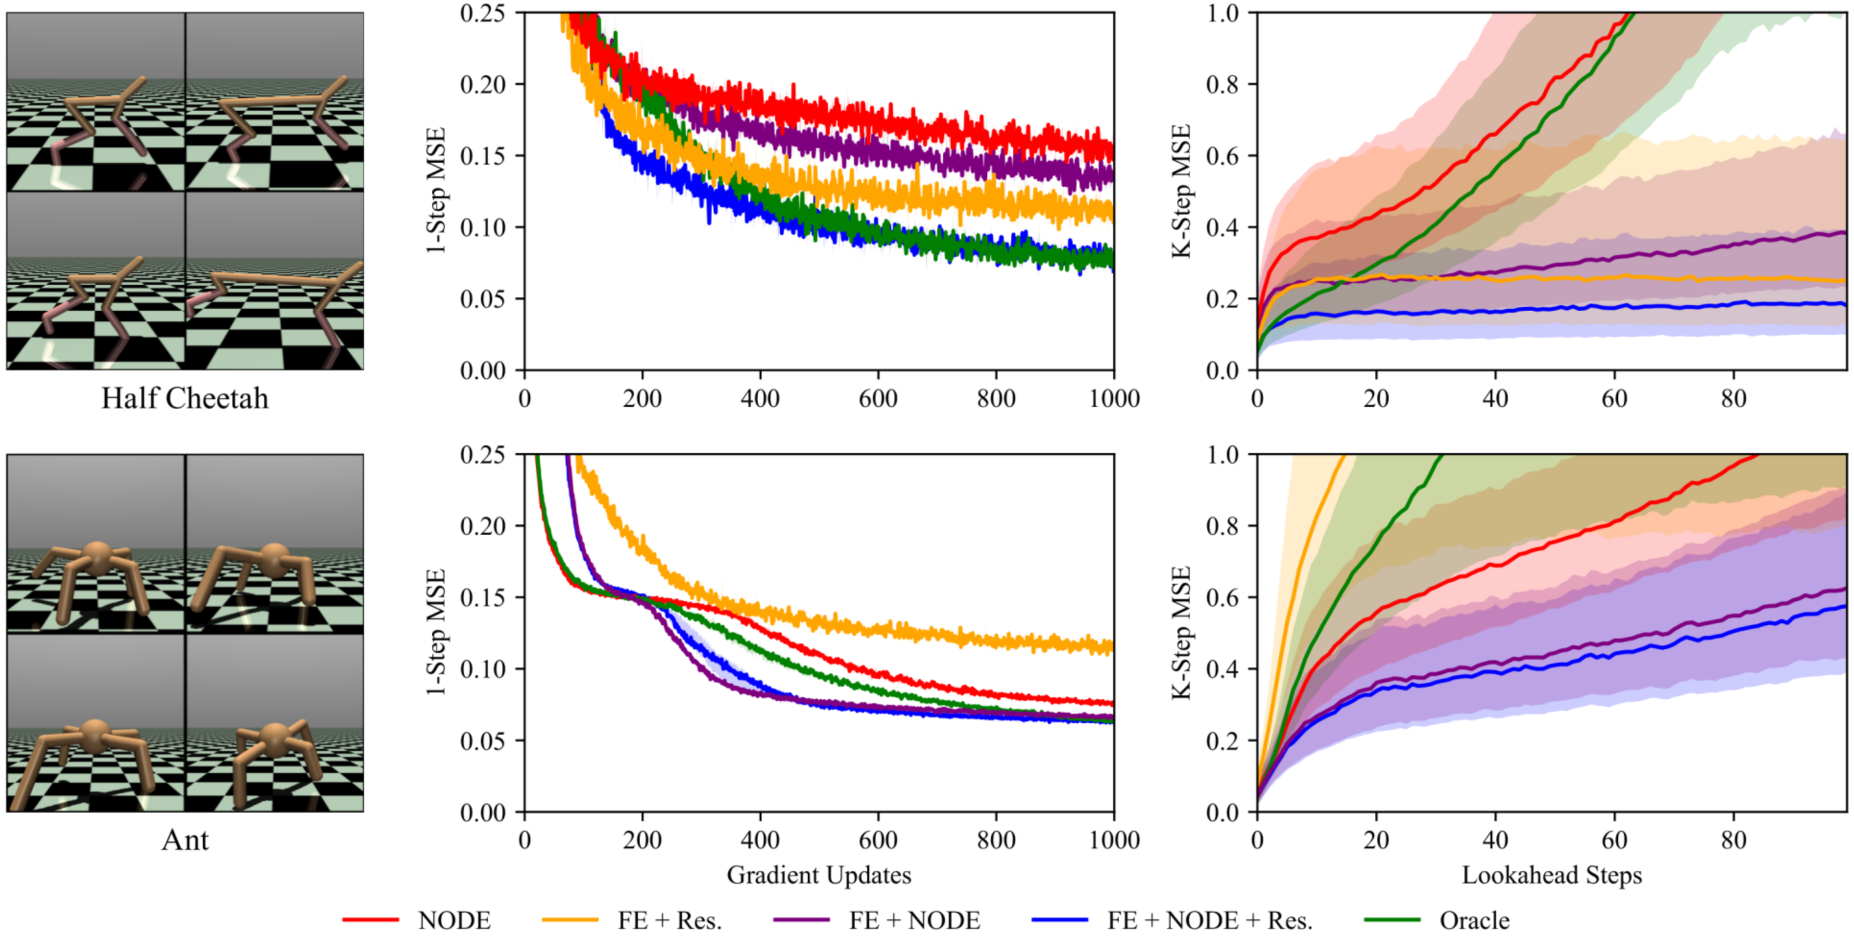

K-Step lookahead is significantly improved

Takeaway: Function encoders inherit the inductive biases of their basis functions

The Impact on Downstream Controllers

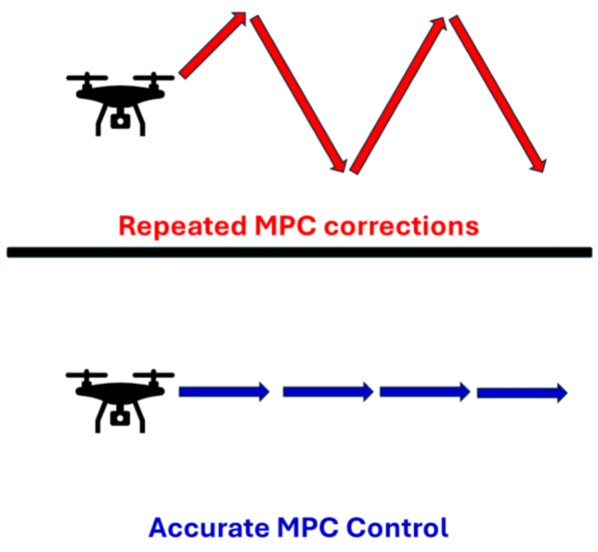

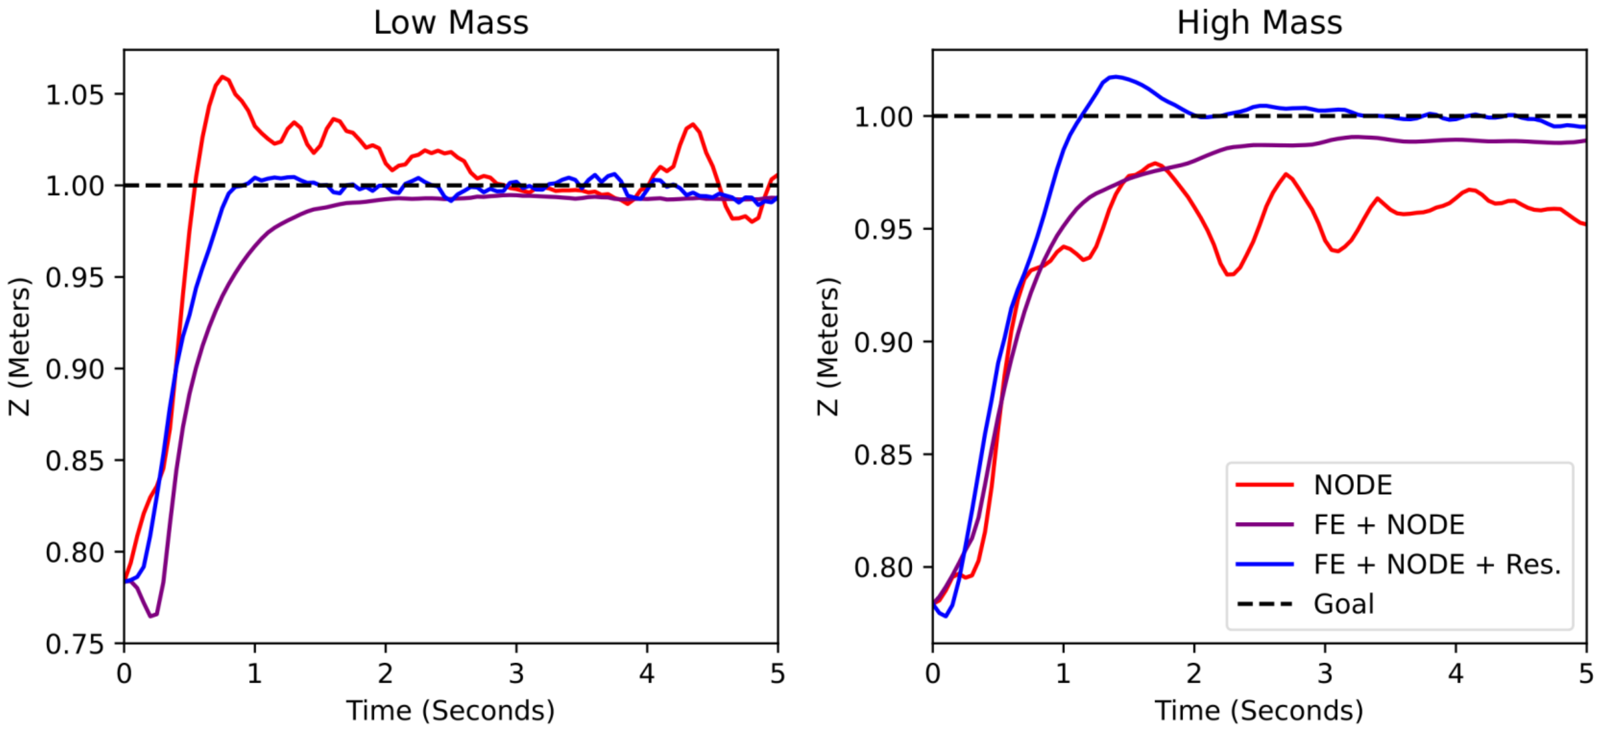

Example: A quadrotor carrying a package, where the weight of the package is a hidden parameter. The controller is using a learned model for model predictive control (MPC).

Learned Basis Functions and Transfer

Online Prediction

Context Representation

Dynamics Prediction

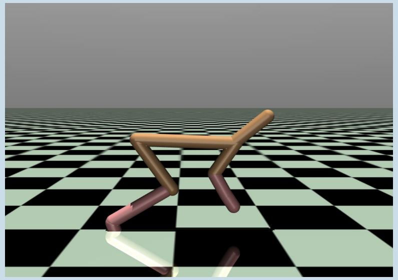

Neural ODEs

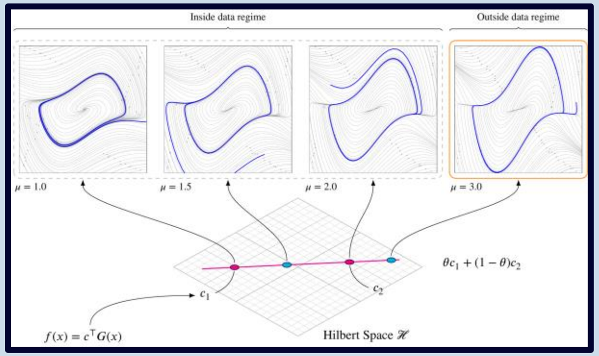

Operator Learning

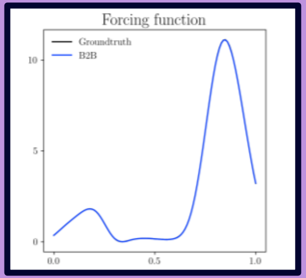

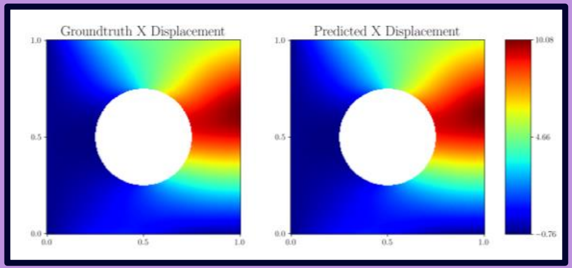

Multi-Task RL

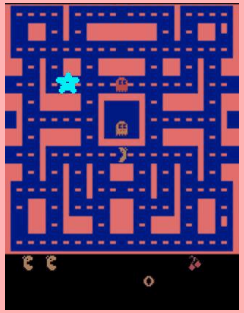

Multi-Agent RL