How to Learn and Generalize From Three Minutes of Data: Physics-Constrained and Uncertainty-Aware Neural Stochastic Differential Equations

Franck Djeumou, Cyrus Neary, and Ufuk Topcu
The University of Texas at Austin

TEXAS The University of Texas at Austin — CENTER FOR autonomy

Objectives: Learn **data-efficient** dynamics models from **noisy state observations,** while providing estimates of the model's **epistemic uncertainty**.

Approach: Train neural stochastic differential equations (SDE) that **leverage a priori physics knowledge**, and that use the diffusion term to **capture model uncertainty**.

Physics-Informed Structure

$\dot{x} = \mathrm{F}(x, u, \mathrm{g}_1(\cdot), \dots, \mathrm{g}_d(\cdot))$

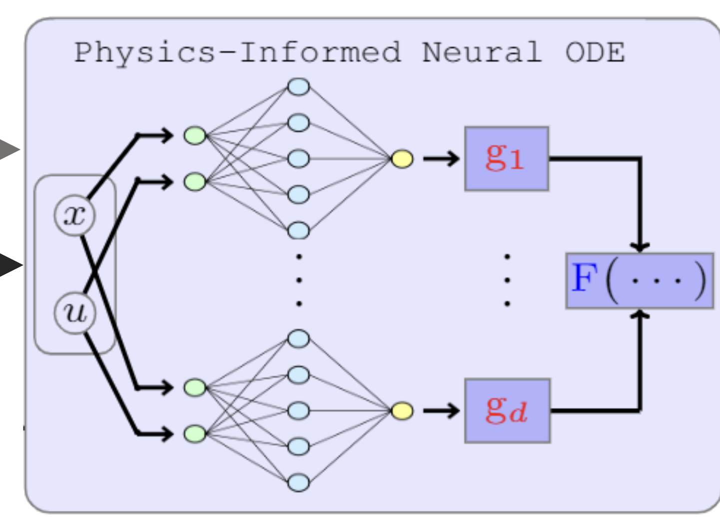

Physics-Aware Drift Term

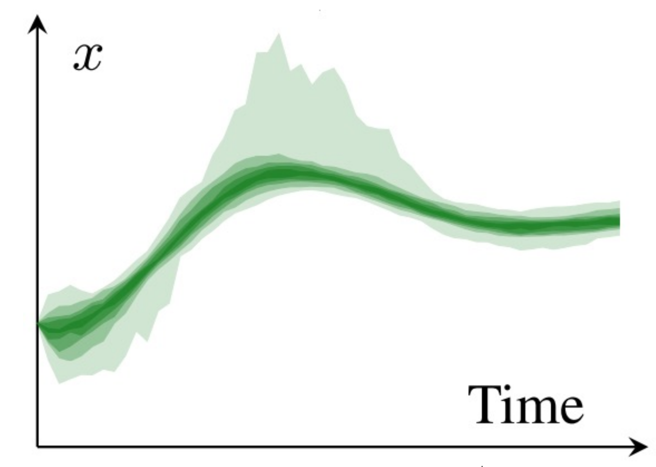

Model predictions are obtained by numerically solving the parametrized stochastic differential equation.

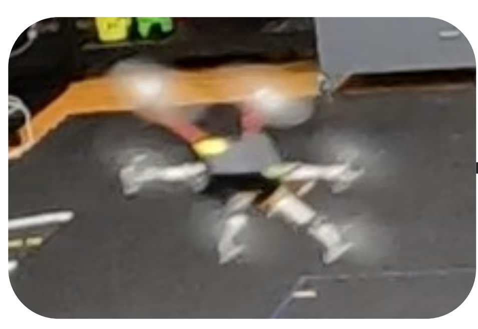

State x

Control Policy

Input u

$$dx = \underbrace{f_\theta(x,u)dt}_{} + \underbrace{\Sigma_\phi(x,u)}_{} * dW$$

SDE Solver

Sampled state trajectories

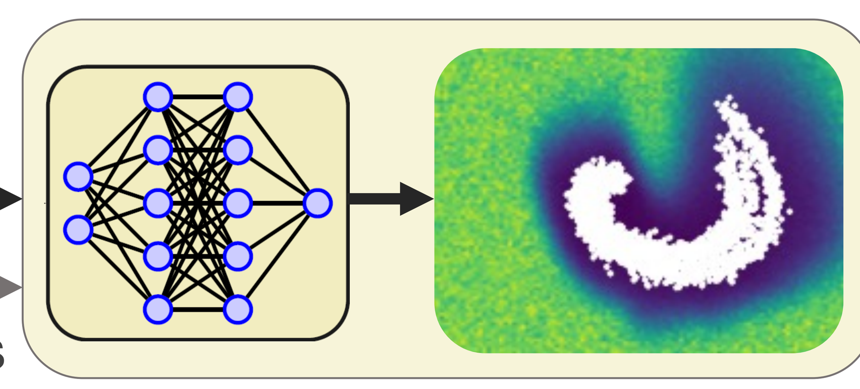

Coverage informs uncertainty estimate

Learned Distance-Aware Uncertainty Estimate

Model predictions are **highly stochastic** on points "far" from the training dataset.

$$J_{total} = \underbrace{\lambda_{data} J_{data}}_{\text{fit training data}} + \underbrace{\lambda_{grad} J_{grad}}_{\text{low stochasticity near training dataset}} + \underbrace{\lambda_{convex} J_{convex} + \lambda_\mu J_\mu}_{\text{high stochasticity far from training dataset}}$$

Illustration of model uncertainty

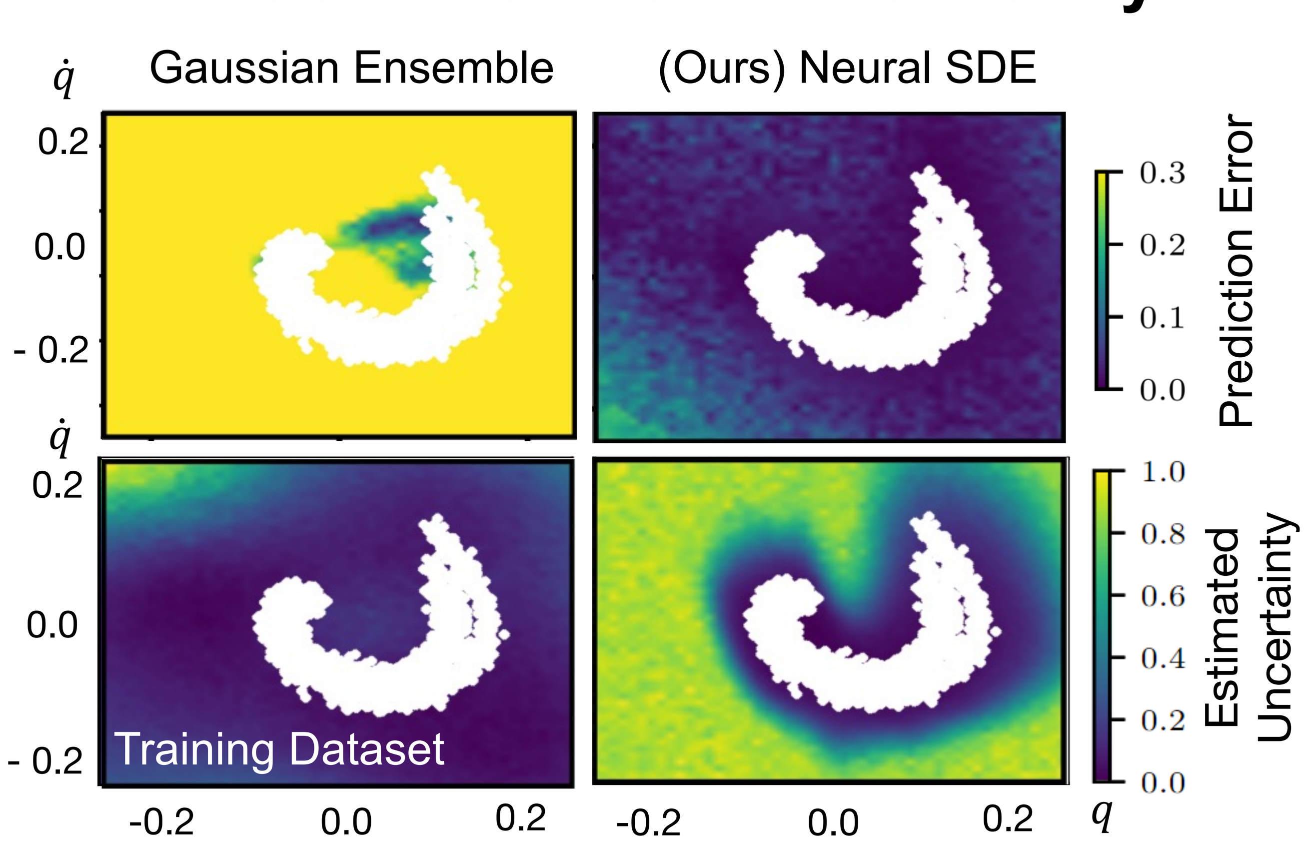

Offline model-based reinforcement learning for cartpole swingup problem

Methodology: use learned dynamics models as simulators for model-free RL algorithm (PPO).

Result: Model-based policies using Neural SDEs are **just as performant as model-free RL**, while requiring **×30 fewer system interactions**. Neural SDE policies outperform policies trained using baseline models.

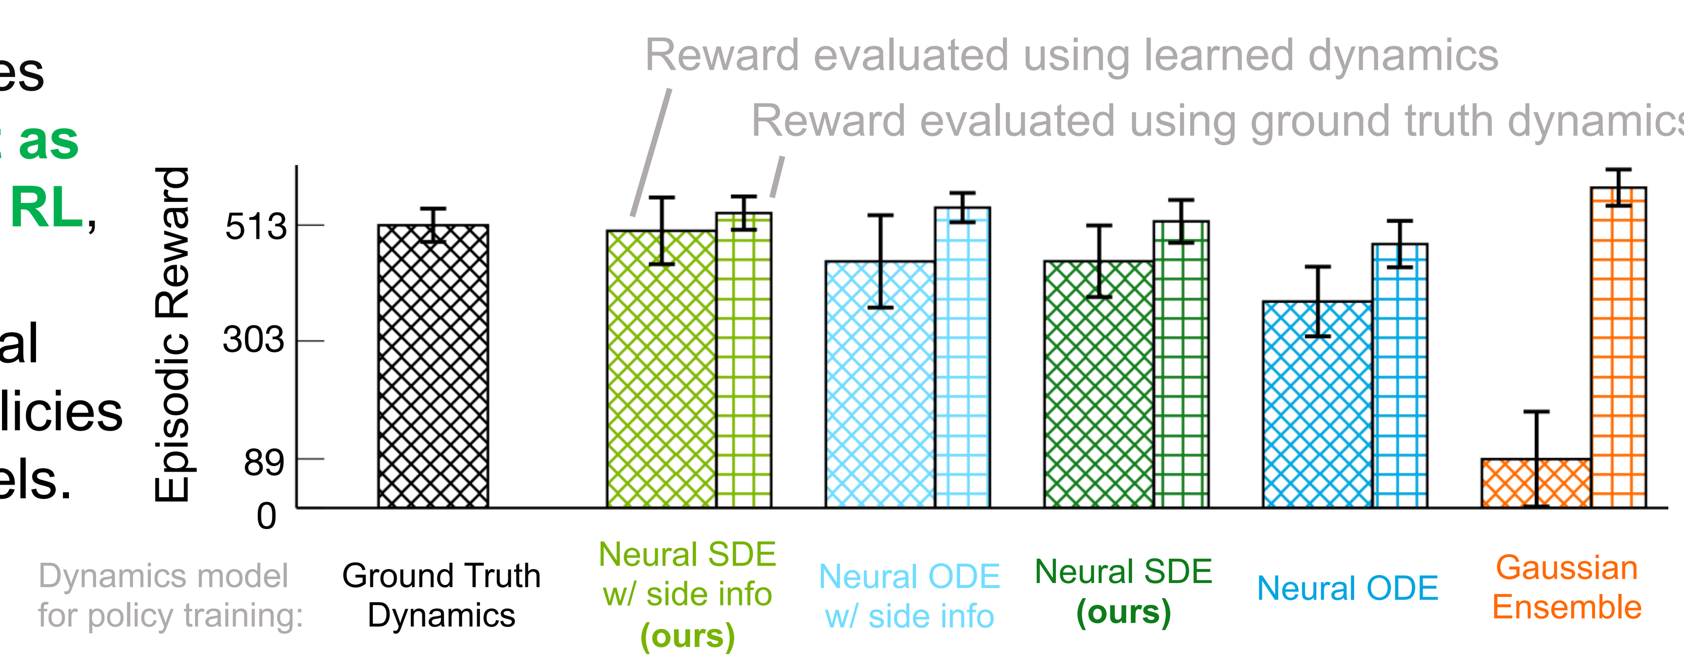

Three minutes of hexacopter data yields an accurate neural SDE

Platform

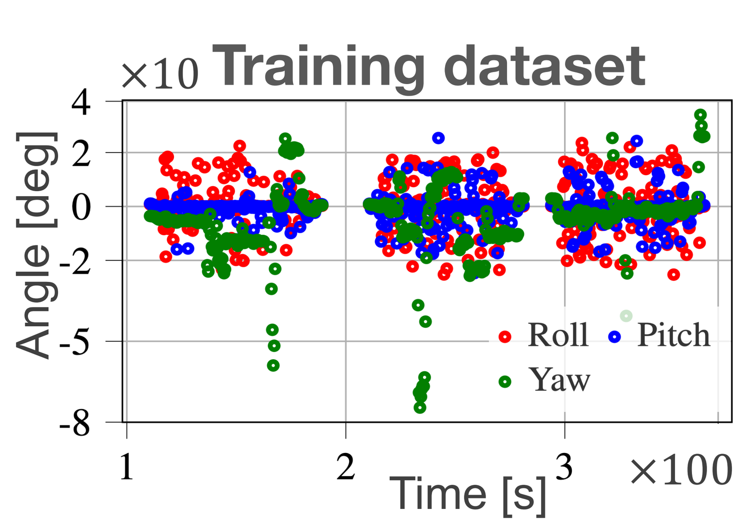

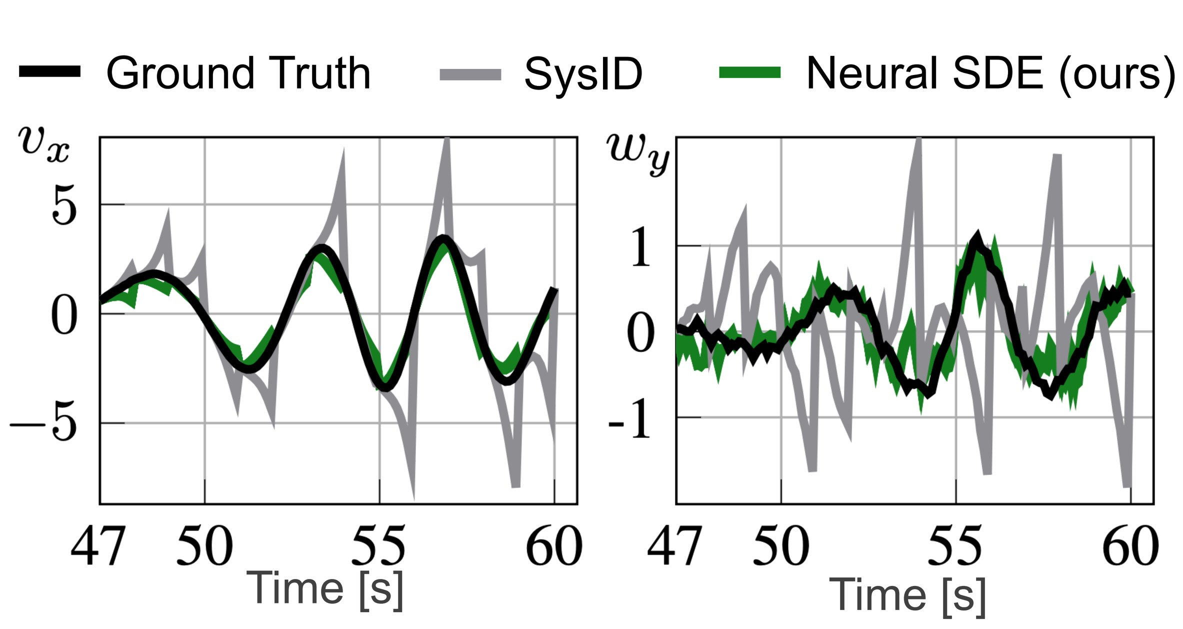

Neural SDE + MPC

Neural SDEs yield performant controllers while operating beyond the training dataset

Lemniscate trajectory

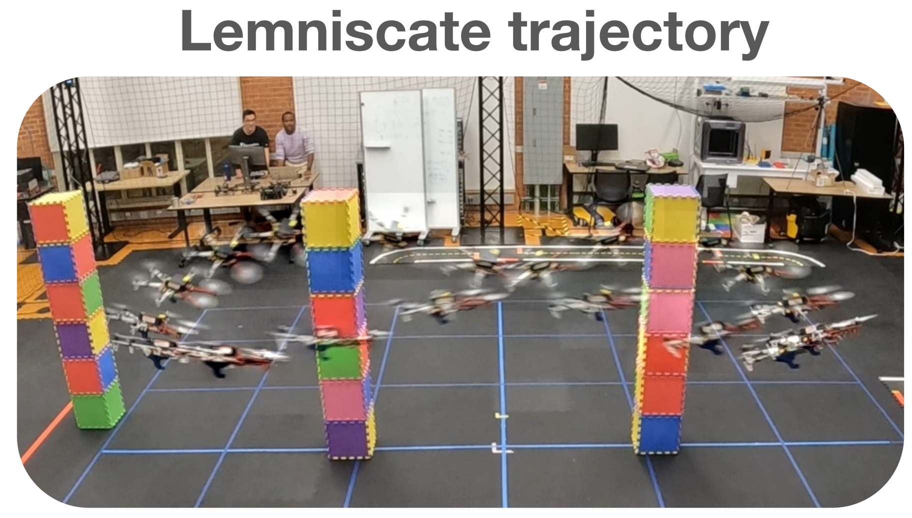

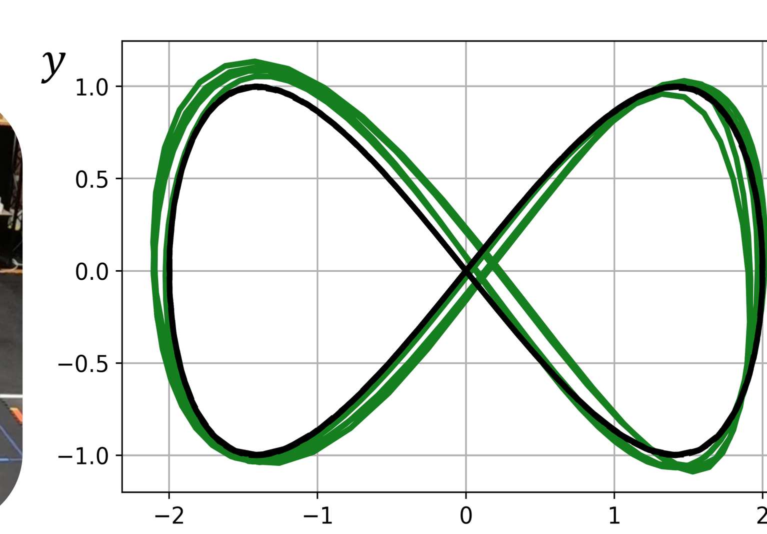

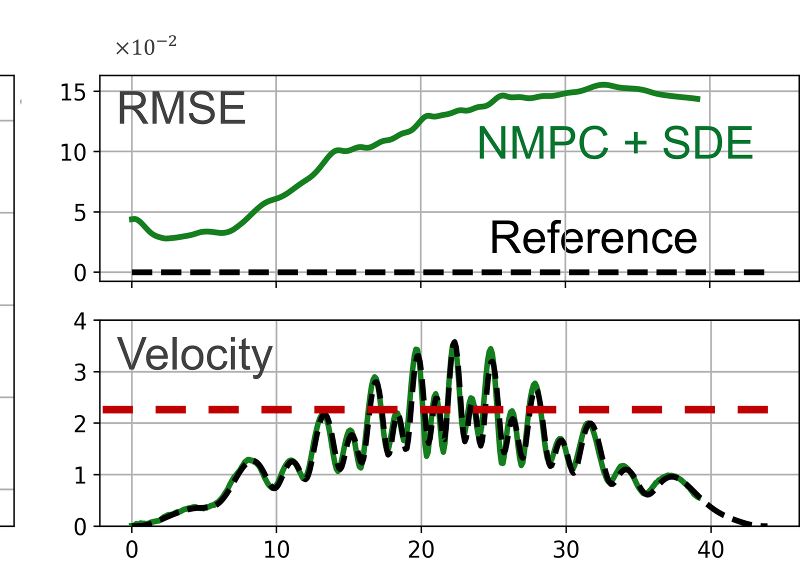

Circular trajectory

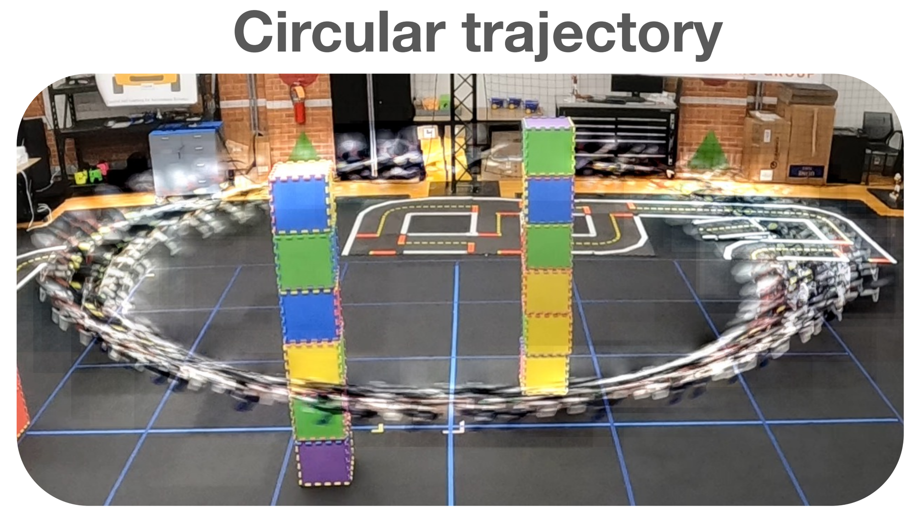

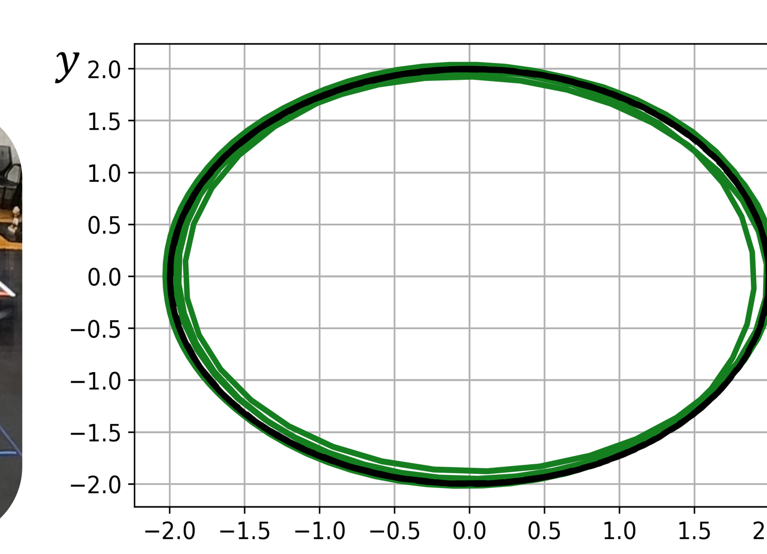

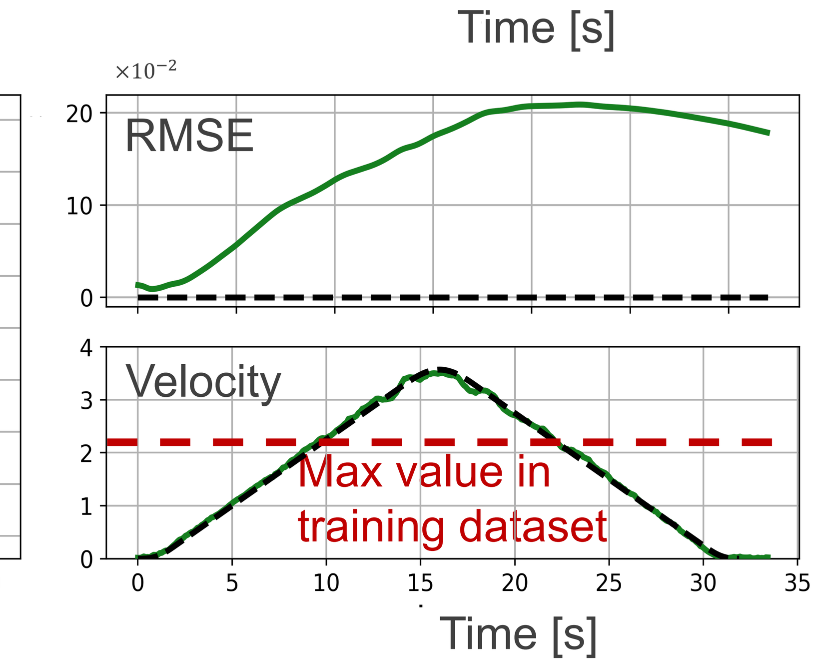

Contact: cneary@utexas.edu