Model Learning and Online Adaptation for Environmental Disturbances in Autonomous Off-Road Driving

Anoushka Alavilli[1,2], Charles Noren[2], Jason Gibson[2], Guanya Shi[1], Jeff Schneider[1], Shehryar Khattak[2]
[1]Robotics Institute, Carnegie Mellon University, [2]FieldAI

1. Motivation & Introduction

Fig 1. Examples of environmental disturbances encountered during real-world deployments. **Left:** A skid-steer mobile robot (SSMR) transporting a heavy payload in off-road terrain. A human operator may choose a payload with unknown mass that slides freely. In the shown case, water sloshing within the payload introduces additional disturbances. **Right:** An all-terrain vehicle (ATV) crossing a deep stream with strong downstream current. Although we focus on SSMRs in this work, we use this as a motivating example for the numerical studies in our work.

Navigation in the wild requires reasoning about embodiment-specific dynamics that can vary significantly across platforms and environments.

Human-specified tasks can induce new dynamics, such as transporting an unknown payload or maintaining a path through strong environmental disturbances.

These effects create a major adaptation challenge: dynamics may change rapidly, and dynamics learned under one set of conditions may not remain accurate at deployment.

Can structured physics priors and lightweight online adaptation enable robust model-based navigation under changing dynamics?

2. Methods

Structured Dynamics: Rely on known robot physics and learn only the unmodeled external/generalized forces:

$$M(q)\ddot{q} + C(q,\dot{q}) + G(q) = B\tau + \tau_{\text{ext.}}$$

where M, C, G, B come from the robot model and $\tau_{\text{ext.}}$ captures tire forces, friction, and other unmodeled dynamics.

Basis Model Learning & Online Adaptation: Factor the residual dynamics as

$$\tau_{\text{ext.}} = \Phi^{w}(x)\theta,$$

with Φ^{w} trained offline on domain-randomized simulation data, modeled after [1], while θ is adapted online from recent observations:

$$\theta^{\star} = \arg\min_{\theta} \sum_{t} \|\Phi^{w}(x_t)\theta - y_t\|_2^2 + \lambda\|\theta - \theta_0\|_F^2,$$

with training target y_t. This lets the model retain a shared nonlinear basis while adapting only a small set of coefficients at deployment. This model structure is inspired by [2-4].

Closed-loop planning with MPPI: At every control step, re-estimate θ_t from the most recent state transitions, then use the adapted dynamics model to roll out candidate control sequences in MPPI:

$$\theta^{\star} = (\Phi^{\top}\Phi + \lambda I)^{-1}(\Phi^{\top}Y + \lambda\theta_0)$$

The adapted $\theta^{\star}$ is held fixed across the MPPI prediction horizon and shared across all sampled rollouts.

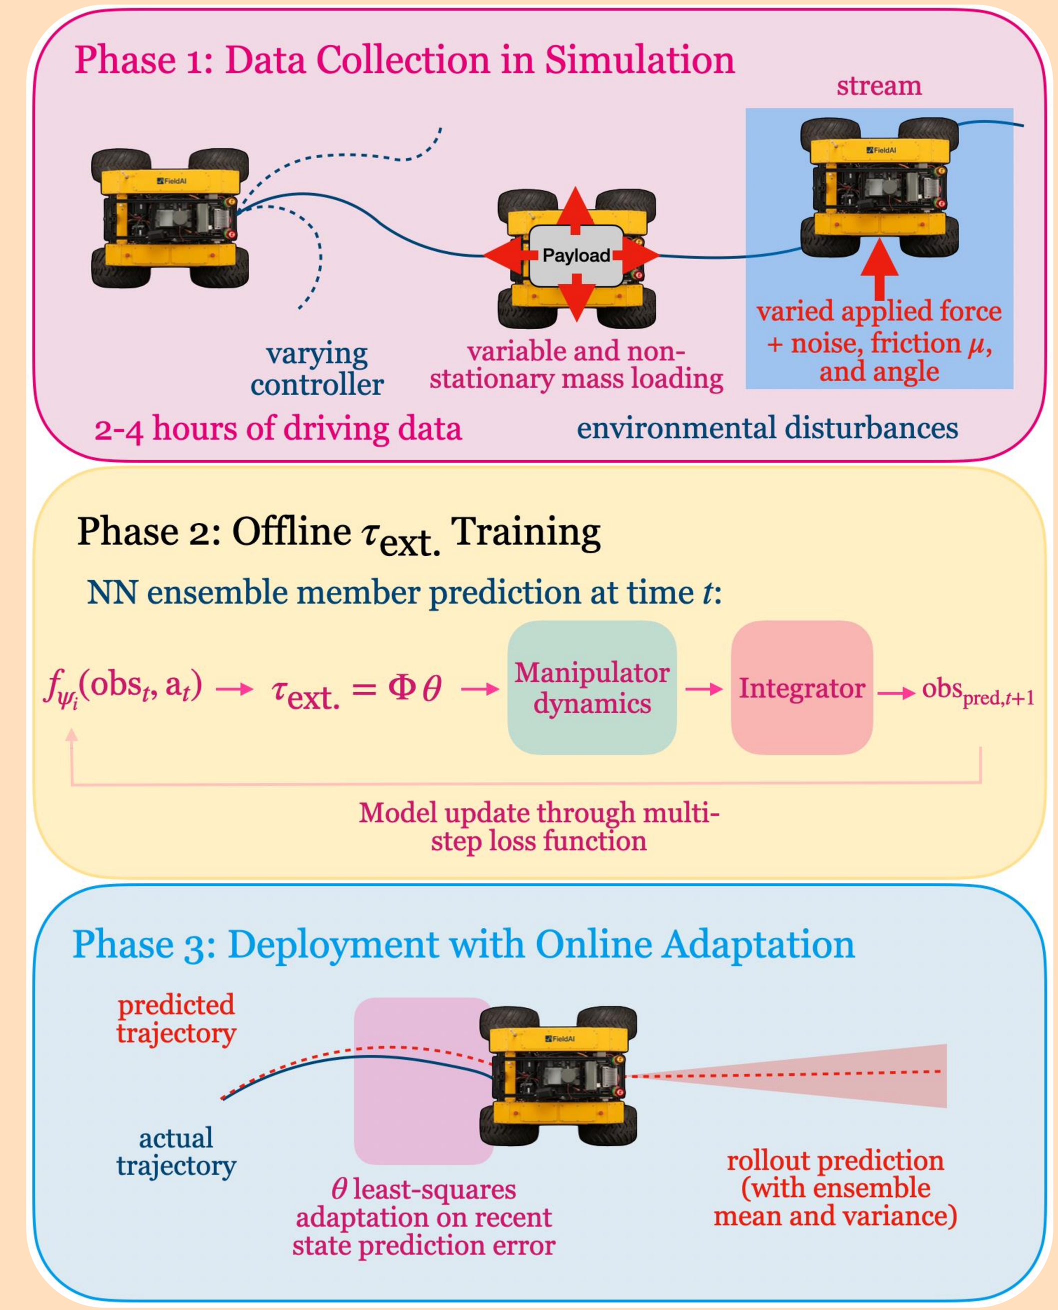

Fig. 2 An overview of the three phases of our proposed method: Data collection in simulation, offline model training, and deployment with online adaptation.

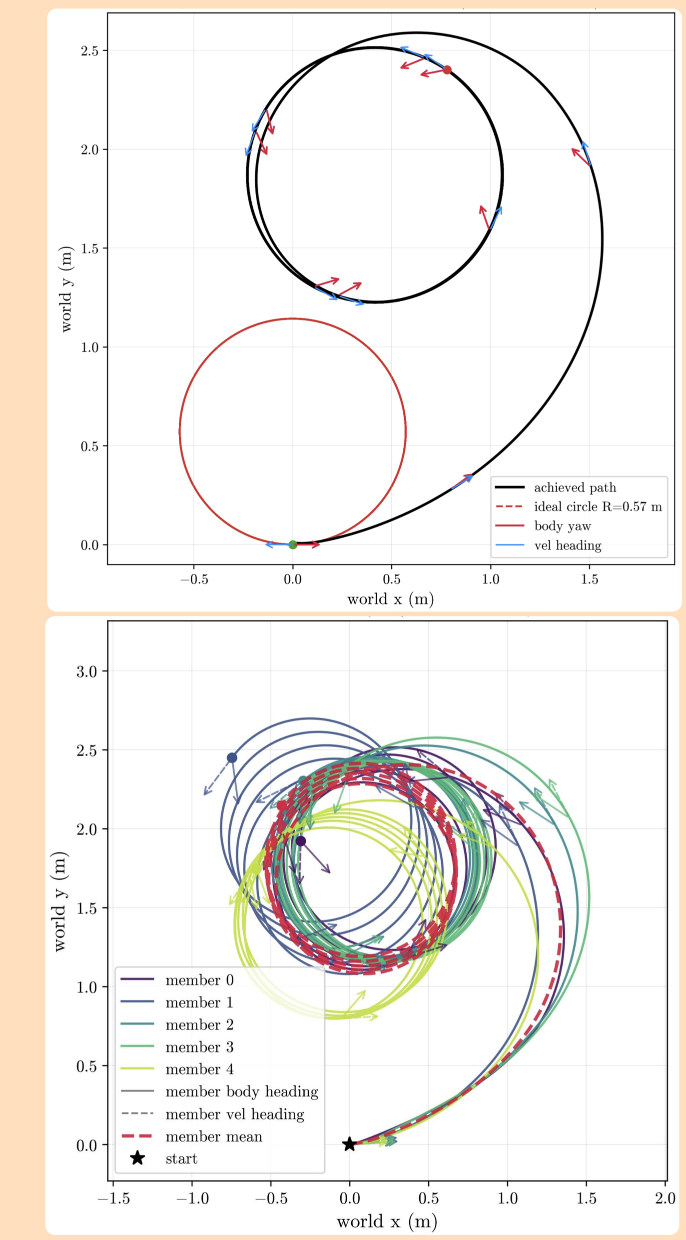

Fig. 3 Illustrative ensemble prediction example. (A) Achieved and ideal trajectories for a high-slip trajectory. (B) Deployment-time trajectory predictions. Ensemble spread represents epistemic uncertainty. Ensemble mean closely matches the achieved trajectory.

3. Results

Numerical Study Examples & Closed-loop Control

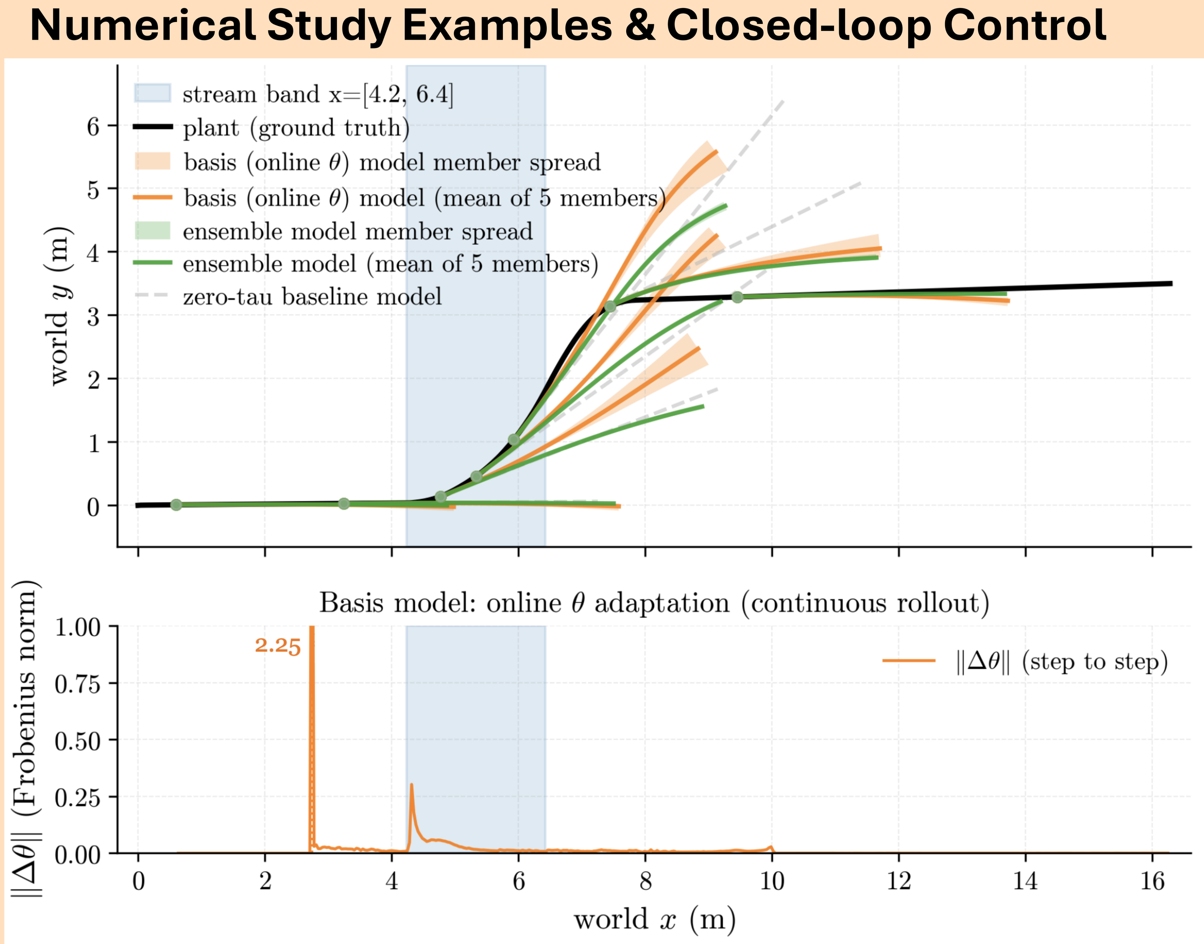

Fig. 4

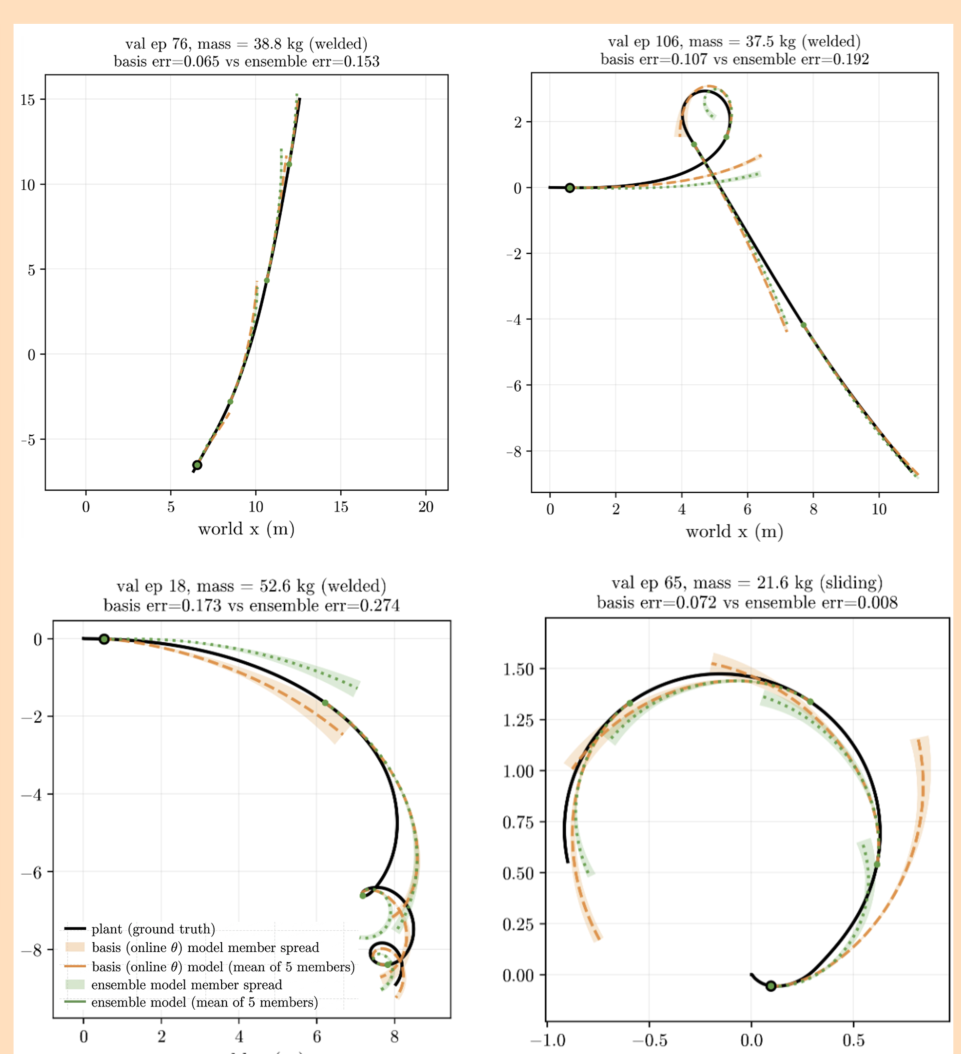

Fig. 5

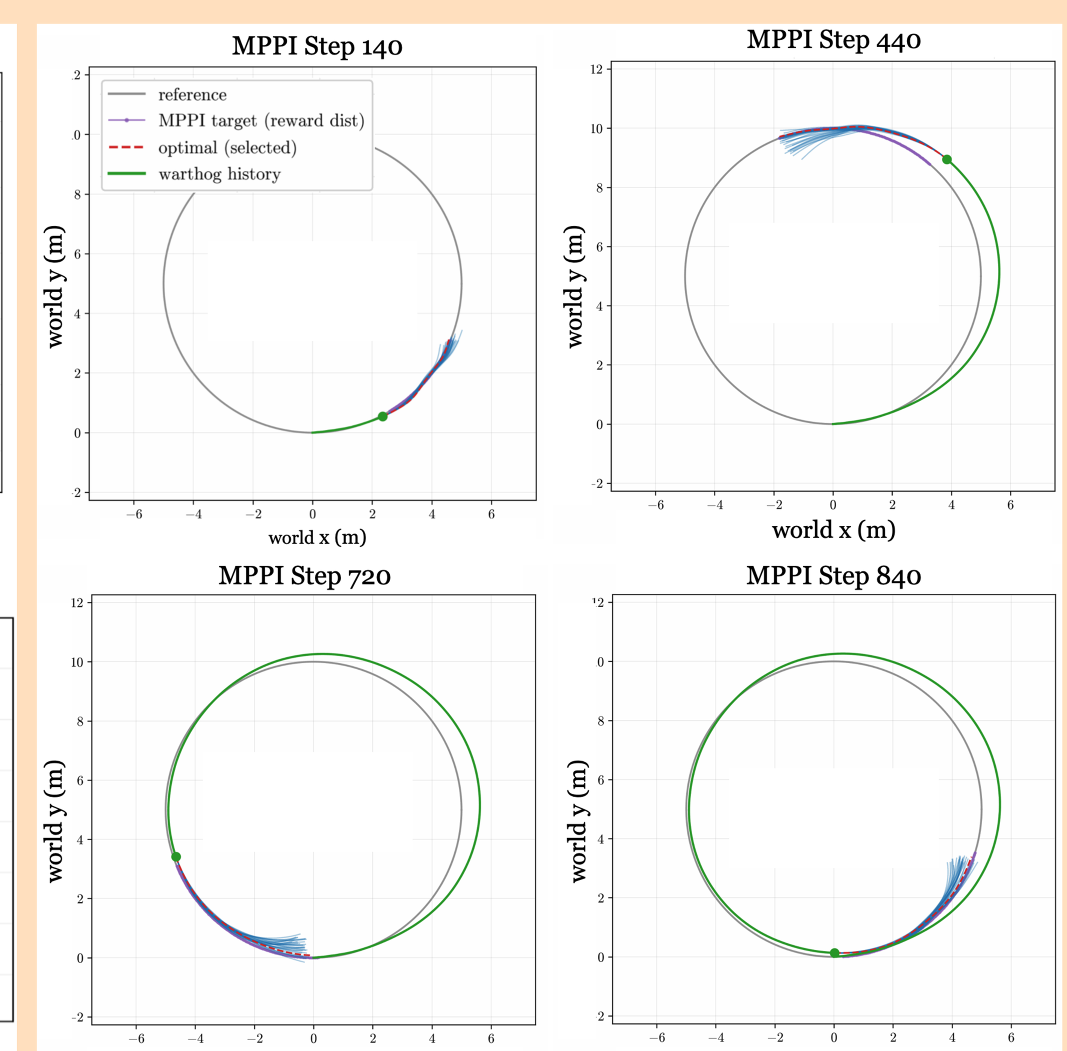

Fig. 6

- Fig. 4: **Stream crossing experiment:** A comparison of model rollouts at selected points in a selected instance of the stream-crossing experiment using a *H=2* s horizon. As expected, the basis model with online adaptation (ours) most closely matches the ground-truth trajectory. We observe adaptation by θ due to start-up transients, stream entry, and stream exit.
- Fig. 5: **Mass loading experiment:** A comparison of model rollouts at selected points in random validation trajectories using a *H = 2* s horizon.
- Fig. 6: SSMR under **closed-loop MPPI** using our learned dynamics model. The controller successfully tracks the reference path despite wheel slip.

Aggregate Numerical Studies Results

Model	Stream	Mass Loading
Mean rollout loss ($h = 1, \ldots, 120$)		
Basis (online θ)	**0.1819**	**0.0507**
Ensemble	0.1958	0.0804
$\tau_{\text{ext.}} = 0$ baseline	0.4670	0.3639
Final-step rollout loss ($h = 120$)		
Basis (online θ)	**0.2664**	**0.1150**
Ensemble	0.2737	0.1702
$\tau_{\text{ext.}} = 0$ baseline	0.7911	0.7234

Table 1: Validation rollout losses for the stream-crossing and mass-loading datasets.

- **Online adaptation consistently improves prediction accuracy.** The structured basis model achieves the lowest rollout loss in both stream-crossing and mass-loading experiments, outperforming the frozen ensemble and nominal-dynamics baselines.
- **The model adapts to qualitatively different dynamics shifts.** It handles both **external disturbances** from stream forces and **changes to vehicle inertial properties** from unknown/freely sliding payloads, despite these perturbing *different aspects* of the underlying dynamics (i.e., $\tau_{\text{ext.}}$ and M, respectively).

4. References & Acknowledgments

[1] Xiao, W., et al., "AnyCar to Anywhere: Learning Universal Dynamics Model for Agile and Adaptive Mobility," in *2025 IEEE International Conference on Robotics and Automation (ICRA)*, 2025, pp. 8819–8825. doi: 10.1109/ICRA55743.2025.11128396.

[2] O'Connell, M., et al., Neural-Fly enables rapid learning for agile flight in strong winds.*Sci. Robot.* **7**, eabm6597(2022).DOI:10.1126/scirobotics.abm6597

[3] Lupu, E.S., et al., "MAGIC VFM—Meta-learning adaptation for ground interaction control with visual foundation models (abstract reprint)," in *Proceedings of the Fortieth AAAI Conference on Artificial Intelligence (AAAI 2026)*, 2026, Art. no. 4491. doi: 10.1609/aaai.v40i47.41392.

[4] Levy, J., et al., "Learning to Walk from Three Minutes of Real-World Data with Semi-structured Dynamics Models," in *Proceedings of the 8th Annual Conference on Robot Learning (CoRL)*, 2024.

This work was done while AA was an intern at FieldAI.