ICRA2026 Late Breaking Results Poster presented at
2026 IEEE International Conference on Robotics and Automation (ICRA 2026)
June 1-5, 2026. Vienna, Austria

Learning What Matters: Task Tailored Dynamics Models through Differentiable MPC

Jan Węgrzynowski, Piotr Kicki, Grzegorz Czechmanowski, Krzysztof Walas

***Abstract*—In model-based control, dynamics models are typically trained by minimizing open-loop prediction errors uniformly across all states. However, due to finite model capacity, this misallocates representational power, as not all prediction errors impact the downstream closed-loop performance equally. In this extended abstract, we propose a task-aware training methodology for a prediction model used in the context of Model Predictive Control (MPC). By extracting analytical sensitivities via differentiable MPC, we construct a loss function that weights multi-step dynamics model prediction errors based on their impact on the closed-loop task cost. Experimental results on a simulated 7DoF manipulator demonstrate that our sensitivity-weighted loss significantly improves closed-loop tracking performance compared to standard Mean Squared Error (MSE) or variance-based state standardization.**

I. Introduction

The deployment of optimization-based control, particularly Model Predictive Control (MPC), increasingly relies on high-capacity function approximators like neural networks to capture complex, unmodeled phenomena [1]. When integrated into MPC frameworks, learned dynamics models can significantly enhance controller closed-loop performance [2], [3].

Traditionally, these neural dynamics models are trained via regression to minimize the open-loop multi-step or single-step prediction error between the simulated trajectory and the ground truth. Standard loss functions, such as MSE, implicitly assume that an error in any state dimension at any future time step is equally important. However, in the context of closed-loop control, a uniform loss is suboptimal: a small prediction error in a critical state can severely degrade task performance, while a larger error in a less relevant state might be entirely compensated for by the controller's feedback loop. This is especially pronounced in cases where the model's representational capacity is bounded, e.g. due to real-time control requirements.

In this work, we introduce a control-aware loss formulation tailored for MPC. By leveraging the known, differentiable structure of the MPC optimization problem [4], we analytically approximate the sensitivity of the closed-loop task cost with respect to specific state prediction errors along the horizon. We utilize these sensitivities to stronger penalize prediction errors that significantly impact controller performance.

This work was supported by an internal PUT grant 0214/SBAD/0256 and 0214/SBAD/0257. All authors are with the Institute of Robotics and Machine Intelligence, Poznan University of Technology, Poznan, Poland, and with IDEAS Research Institute, Warsaw, Poland
jan.wegrzynowski@put.poznan.pl

II. Task-Tailored Model Training

Let $\hat{\mathbf{x}}_{t:t+t_p} \in \mathbb{R}^{t_p \times n}$ represent a predicted state trajectory of length t_p, rolled out from an initial true state x_t using a sequence of applied controls and a parameterized dynamics model f_θ^{rk4}. Let $\mathbf{x}_{t:t+t_p}^{\text{true}}$ denote the corresponding ground-truth trajectory. We formulate a sensitivity-weighted objective function for training the dynamics model:

$$\mathcal{L}_{\text{sens}}(\theta) = \left\| \mathbf{S}_{t:t+t_p} \odot \left(\hat{\mathbf{x}}_{t:t+t_p} - \mathbf{x}_{t:t+t_p}^{\text{true}} \right) \right\|_2^2 \tag{1}$$

where $\mathbf{S}_{t:t+t_p} \in \mathbb{R}^{t_p \times n}$ is a time- and state-dependent weighting tensor, and $\odot$ is the Hadamard product.

A. Sensitivity of Task Objective to Prediction Errors

Ideally, $\mathbf{S}$ should be proportional to the impact of the open-loop prediction error $\mathbf{e}_{t:t+t_p} = \hat{\mathbf{x}}_{t:t+t_p} - \mathbf{x}_{t:t+t_p}^{\text{true}}$ on the total episodic task cost $J = \sum_{\tau=0}^{T} c(x_\tau, u_\tau)$. Because the controller replans at every step, the open-loop error at step t influences J entirely through its perturbation of the immediately applied action u_t. Thus, we can decompose that impact as follows

$$\mathbf{S}_{t:t+t_p} \propto \frac{\partial J}{\partial \mathbf{e}_{t:t+t_p}} = \underbrace{\frac{\partial J}{\partial u_t}}_{\substack{\text{Task Sens.}\\ \text{to Action}}} \cdot \underbrace{\frac{\partial u_t}{\partial \mathbf{e}_{t:t+t_p}}}_{\substack{\text{Controller Sens.}\\ \text{to Pred. Error}}}. \tag{2}$$

B. Estimating the Weighting Tensor

Computing the exact analytical value of $\partial J / \partial u_t$ is intractable as it requires unrolling Jacobians through the unknown true system dynamics. Instead, we approximate $\mathbf{S}$ using the finite MPC horizon as a proxy:

$$\hat{\mathbf{S}}_{t:t+t_p} = \frac{\partial \hat{J}_t^{\text{MPC}}}{\partial u_0^{\text{MPC}}} \cdot \frac{\partial u_0^{\text{MPC}}}{\partial \mathbf{e}_{t:t+t_p}}. \tag{3}$$

The first term, $\partial \hat{J}_t / \partial u_0$, evaluates how a change in the selected action degrades the estimated MPC cost. The second term, $\partial u_0 / \partial \mathbf{e}_{t:t+t_p}$, is the sensitivity of the controller to dynamics model errors, which is obtained directly via the backward pass of a differentiable MPC solver [4].

III. Experimental Results

We evaluated our methodology using a Context-Aware Deep Lagrangian Model (CaDeLaC) [1] controlling a 7DoF Franka Panda manipulator tracking randomized trajectories in MuJoCo. We fine-tuned a pretrained CaDeLaC model using multi-step rollout losses ($p \in \{1, 4\}$). We compared our sensitivity weighting (`ours`) against baseline uniform

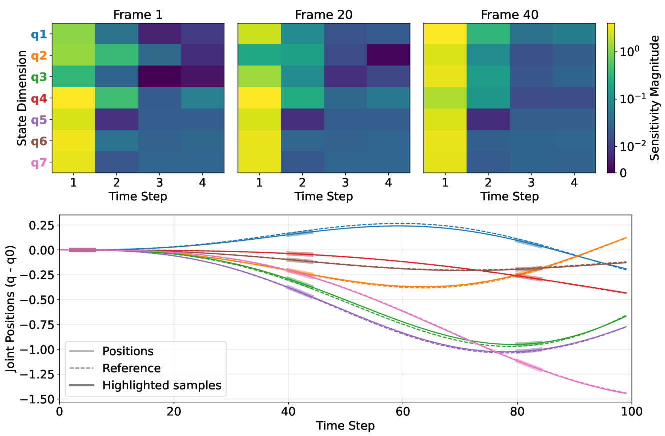

Fig. 1. Cost map showing the influence of dynamics model errors across states and the prediction horizon on the estimated task cost for a 7DoF manipulator.

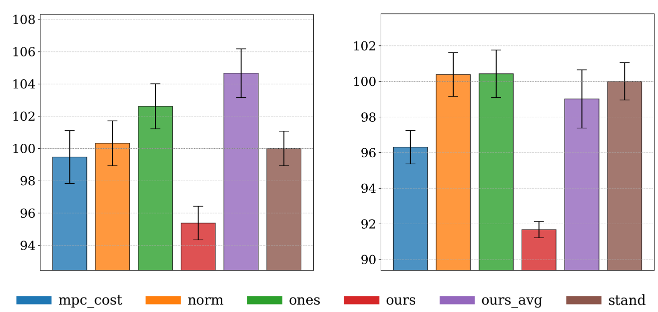

Fig. 2. Normalized closed-loop tracking cost (left $p = 1$ and right $p = 4$ training rollout length). The sensitivity-weighted model (ours) achieves superior tracking performance compared to all baseline scaling methods.

weighting (ones), state standardization (stand), range normalization (norm), and an MPC stage-cost based weighting (mpc_cost). Sensitivities were computed using the acados framework [4].

Figure 1 shows that the estimated cost sensitivity over state dimensions and the prediction horizon is strongly non-uniform, validating the need for control-aware weighting.

In closed-loop evaluation (Fig. 2), models trained with our task-aligned sensitivity weighting achieved the lowest tracking cost. For standard 1-step rollouts ($p = 1$), our method reduced tracking cost to 95.4% of the stand baseline. Crucially, the benefits scaled with rollout length: at $p = 4$, ours reduced the normalized cost to 91.7%, while uniform and naive cost weightings remained near 100%. Moreover, we performed an ablation study by averaging the sensitivities over the dataset (ours_avg), which resulted in a significant performance drop, proving that local, sample-dependent sensitivity is key to allocating model capacity effectively.

IV. Conclusions

By leveraging differentiable MPC to extract analytical sensitivities, we demonstrate that weighting single-step and multi-step prediction errors by their downstream task impact significantly improves closed-loop performance over generic statistical normalizations, without increasing model capacity.

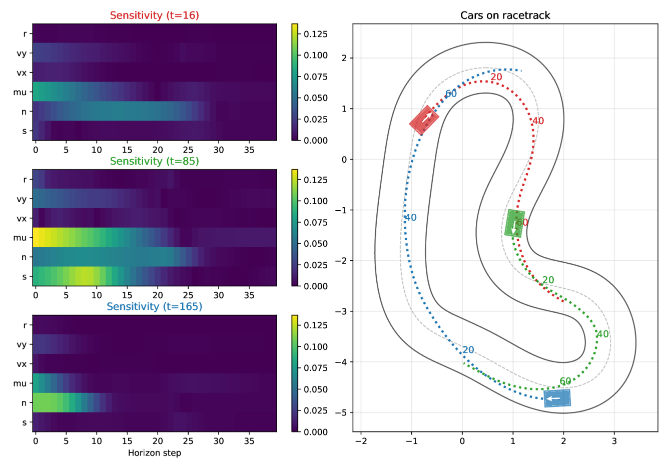

Fig. 3. Example sensitivities for a real-world F1Tenth autonomous racing task. Left: Sensitivity maps for the dynamic states at different points in time. Right: Corresponding car positions on the track. The sensitivity of the task objective to individual states changes heavily depending on current track position and corresponding velocities.

V. Future Work

As future work, we plan to extend this approach to other robotic domains with highly dynamic environments, such as autonomous racing. Figure 3 provides an initial outlook in this direction, illustrating sensitivity evaluations for an autonomous racing vehicle on a track. The importance of predicting specific state dimensions varies dramatically over time; for instance, prediction errors in lateral deviation (n) and track relative heading (μ) are heavily penalized when entering a sharp curve, whereas different states are prominent ahead of straights. Because the importance of specific states is highly contextual, applying our sensitivity-weighted loss represents a promising avenue for advancing high-speed closed-loop control of racing vehicles.

References

[1] L. Schulze, J. Peters, and O. Arenz, "Context-aware deep lagrangian networks for model predictive control," in *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2025, pp. 6939–6946.

[2] J. Kabzan, L. Hewing, A. Liniger, and M. N. Zeilinger, "Learning-based model predictive control for autonomous racing," vol. 4, no. 4, pp. 3363–3370. [Online]. Available: https://ieeexplore.ieee.org/document/8754713/

[3] K. Y. Chee, T. Z. Jiahao, and M. A. Hsieh, "KNODE-MPC: A knowledge-based data-driven predictive control framework for aerial robots." [Online]. Available: http://arxiv.org/abs/2109.04821

[4] R. Verschueren, G. Frison, D. Kouzoupis, J. Frey, N. van Duijkeren, A. Zanelli, B. Novoselnik, T. Albin, R. Quirynen, and M. Diehl, "acados – a modular open-source framework for fast embedded optimal control," *Mathematical Programming Computation*, 2021.