# SplineTM: B-Spline Tire Modeling for Autonomous Racing

Piotr Kicki, Jan Węgrzynowski, Grzegorz Czechmanowski, Krzysztof Walas

***Abstract*— High-performance autonomous racing requires accurate modeling of vehicle dynamics, particularly in limit-handling regimes where tire-road interactions dominate vehicle behavior. While traditional semi-empirical tire models provide interpretability, they often lack the flexibility to capture complex tire characteristics, especially when combined with simplified vehicle dynamics models. Conversely, purely data-driven methods offer high expressiveness and flexibility, but may suffer from physical inconsistencies and poor extrapolation in extreme regimes, leading to unsafe exploitation by control policies. To address these issues, we introduce SplineTM, a novel tire modeling framework that represents tire forces using B-splines. SplineTM bridges the gap between rigid empirical structures and flexible data-driven approaches by maintaining physical grounding while providing scalable representational capacity. We evaluate the model against strong baselines across three tasks: trajectory prediction, Model Predictive Control, and Reinforcement Learning based sim-to-real transfer. Results on small-scale and full-scale racing platforms demonstrate that SplineTM achieves superior prediction accuracy and significantly faster lap times, providing a robust, differentiable, and interpretable alternative for safe, high-speed autonomous vehicle control.**

## I. INTRODUCTION

Accurate vehicle dynamics modeling is a cornerstone of high-performance autonomous racing. When operating at the limits of handling, small modeling inaccuracies can lead to large prediction errors, degraded control performance, or catastrophic failures. To achieve competitive lap times safely during extreme maneuvers, an autonomous agent must operate within the nonlinear regime of vehicle dynamics, where the interaction between the tires and the track surface becomes the dominant factor.

The landscape of tire modeling spans a spectrum from semi-empirical to purely data-driven approaches. Classic semi-empirical models, such as the Pacejka "Magic Formula"[1], TMEasy [2], or more analytical ones, such as the Dugoff [3], [4] or Fiala [5], utilize structured mathematical formulations to capture the characteristic S-curve of tire forces. While interpretable and physically grounded, their fixed structures lack the flexibility to adapt to peculiar behaviors of specific tire-surface combinations. On the other end, data-driven methods [6], [7] offer unparalleled expressiveness. However, this comes at the cost of poor extrapolation and physical inconsistencies. This is particularly problematic in Reinforcement Learning (RL) and optimal control, where optimization schemes actively exploit these irregularities, leading to unsafe deployment.

This work was supported by an internal PUT grant 0214/SBAD/0256.

All authors are with the Institute of Robotics and Machine Intelligence, Poznan University of Technology, Poznan, Poland, and with IDEAS Research Institute, Warsaw, Poland piotr.kicki@put.poznan.pl

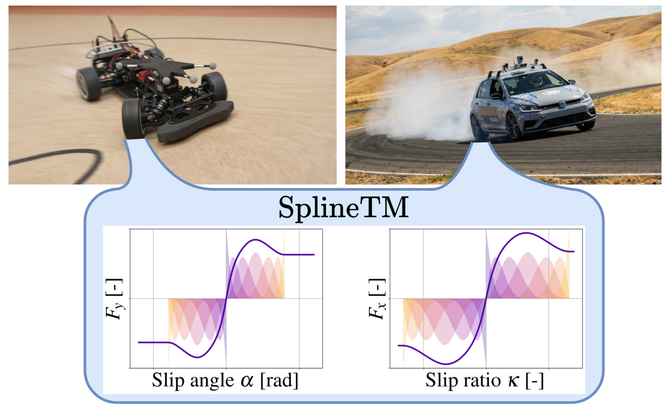


Fig. 1. We propose SplineTM, which utilizes B-splines to model tire-road interactions, combining expressiveness with physical structure.

Recent advancements have sought a middle ground, using neural networks to determine the parameters of semi-empirical models [8], [9]. Nevertheless, the underlying semi-empirical structure is often still constrained by strong parametric coupling and limited expressivity, and some rely on history-dependent features [10] that complicate standard Model Predictive Control (MPC).

To provide a safe, flexible, and interpretable alternative, we propose **SplineTM**, a novel tire modeling framework based on B-splines. Our approach represents longitudinal and lateral tire forces as smooth, piecewise polynomial curves (Fig. 1). By utilizing clamped uniform B-splines [11], we ensure physically consistent behavior while allowing the model to smoothly adapt to complex characteristics.

The contributions of this work are twofold:

- We introduce SplineTM, a B-spline-based tire model that offers a physically-grounded, smooth, differentiable, and interpretable alternative to existing tire models.
- We provide a rigorous evaluation of various tire models in the context of autonomous racing, benchmarking performance across three domains: (i) trajectory prediction on both F1TENTH and full-scale racecars, (ii) safe MPC applied to real-world racing, and (iii) RL policy training and sim-to-real transfer[1].

## II. METHOD

### A. Problem Statement

The problems we consider are trajectory prediction and closed-loop control in autonomous racing. Our goal is to

[1]Supplementary video can be found at https://youtu.be/aa6GFuCYnJM

investigate how the choice of the dynamics model influences safety and performance in these tasks. In trajectory prediction, we aim to develop a dynamics model $\dot{x} = f(x, u)$ that minimizes the Mean Squared Error (MSE) between the predicted state trajectory $\hat{\mathbf{x}}$ and the true system state trajectory $\mathbf{x_{GT}}$ over a prediction horizon. In the control task, we exploit the dynamics model $f$ to control the vehicle through a time trial using either gradient-based MPC or as a simulated environment to train RL policies for zero-shot real-world deployment.

### B. Vehicle Dynamics Model

We adopt the dynamic single-track model, widely used in robotics for its balance of predictive fidelity and computational tractability. The vehicle state is $x = \begin{bmatrix} v_x & v_y & r \end{bmatrix}^T$, representing longitudinal velocity, lateral velocity, and yaw rate. The dynamics are:

$$\begin{bmatrix} \dot{v}_x \\ \dot{v}_y \\ \dot{r} \end{bmatrix} = \begin{bmatrix} \frac{1}{m}\left(F_{xr} + F_{xf}\cos\delta - F_{yf}\sin\delta - F_{drag}\right) + v_y r \\ \frac{1}{m}\left(F_{xf}\sin\delta + F_{yf}\cos\delta + F_{yr}\right) - v_x r \\ \frac{1}{I_z}\left(\left(F_{xf}\sin\delta + F_{yf}\cos\delta\right) l_f - F_{yr} l_r\right) \end{bmatrix},$$

where $\delta$ is the front steering angle, $F_{\square}$ represents lumped tire and drag forces, $m$ is mass, $I_z$ is yaw inertia, and $l_f, l_r$ are distances to the axles. Accurately modeling tire forces $(F_{xf}, F_{xr}, F_{yf}, F_{yr})$ is critical as they are the primary source of energy for a racing vehicle.

### C. B-Spline Tire Model

We model tire forces as functions of the slip angle $\alpha$ and slip ratio $\sigma$, representing the relative velocity between the tire contact patch and road surface. Following standard conventions [12], we define:

$$\alpha_f = \frac{v_y + l_f r}{v_x} - \delta, \quad \alpha_r = \frac{v_y - l_r r}{v_x},$$
$$\sigma_r = \frac{\omega R - v_x}{\max(\omega R, v_x)},$$
$$\sigma_f = \frac{\omega R - (v_x\cos\delta + (v_y + l_f r)\sin\delta)}{\max(\omega R, v_x\cos\delta + (v_y + l_f r)\sin\delta)},$$

where $\omega R$ is the linear velocity of the tire contact patch.

To prevent non-physical boundary exploitation by control policies, we explicitly enforce the coupling between longitudinal and lateral forces via a friction ellipse constraint:

$$F_x = F_z\mu\frac{\sigma_{norm}}{\kappa}T_x(\sigma_{eff}), \quad F_y = F_z\mu\frac{\alpha_{norm}}{\kappa}T_y(\alpha_{eff}),$$
$$\sigma_{eff} = \kappa\bar{\sigma}, \quad \alpha_{eff} = \kappa\bar{\alpha}, \quad \kappa = \sqrt{\sigma_{norm}^2 + \alpha_{norm}^2},$$
$$\sigma_{norm} = \frac{\sigma}{\bar{\sigma}}, \quad \alpha_{norm} = \frac{\alpha}{\bar{\alpha}},$$

where $\bar{\sigma}, \bar{\alpha}$ represent learnable slip scaling factors, and $T_x, T_y$ model normalized tire forces in the pure slip regimes.

Empirical data [1] suggests that $T_x$ and $T_y$ possess characteristic shapes: an initial linear region, smooth saturation, and steady-state behavior at higher slips. We model these using clamped uniform B-splines [11], which natively scale with the number of control points and allow us to impose structural priors. For the longitudinal force:

$$T_x(\sigma_{eff}) = \text{BSpline}(\min(\sigma_{eff}k_\sigma, 1)),$$

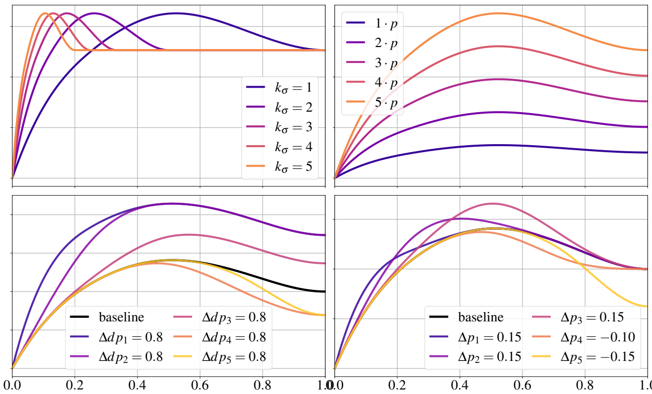


Fig. 2. SplineTM longitudinal tire force $T_x(\sigma)$ for varying parameters. It provides fine-grained shape control while naturally retaining a physically plausible S-curve structure.

where $k_\sigma$ maps the effective slip to the $[0, 1]$ B-spline domain. Although the min operator can introduce non-smoothness, we mitigate this by carefully selecting boundary control points to ensure $\mathcal{C}^1$ continuity at the saturation transition. Specifically, we constrain the last two control points to be equal, ensuring $T'(1) = 0$, guaranteeing a smooth transition to a bounded steady-state value after $\sigma_{eff}k_\sigma$ exceeds 1. By fixing the first control point to 0, we obtain physically grounded behavior at zero slip, though we optionally allow horizontal and vertical offsets $(S_h, S_v)$ to model the effects of camber or internal tire structures [1].

The interior control points dictate the shape of the force curve. To guarantee a physically plausible shape (rising to a peak, then potentially degrading), we parameterize control points $p$ with a vector of non-negative increments $dp_{\uparrow} \in \mathbb{R}_+^{N_\uparrow}$ and another vector of non-negative decrements $dp_{\downarrow} \in \mathbb{R}_+^{N_\downarrow}$. The integers $N_\uparrow$ and $N_\downarrow$ control the capacity of the model before and after the peak. Figure 2 illustrates the diverse yet structurally safe family of characteristics this setup generates. SplineTM integrates physically plausible inductive biases with scalable representational power, bridging the gap between fixed empirical models and unstructured neural network-based ones.

### D. Identification of Model Parameters

We optimize model parameters $\theta$ end-to-end to minimize multi-step Mean Squared Error over a dataset of trajectories $\mathcal{D}$. Each trajectory $\tau$ consists of ground-truth states $\mathbf{x}_{\mathbf{GT}}^{\tau}$ and applied actions $\mathbf{u}^\tau$. The predicted state $\hat{\mathbf{x}}^\tau$ is obtained by integrating the model dynamics $\dot{\hat{x}} = f_\theta(\hat{x}, u)$ from an initial state $x_0^\tau$. We optimize

$$\min_\theta \sum_{t \in T} \|x_t^\tau - \hat{x}_t^\tau\|_2^2 \quad \text{s.t.} \quad \dot{\hat{x}}_t^\tau = f_\theta(\hat{x}_t^\tau, u_t^\tau), \qquad (1)$$

using back-propagation through time via gradient descent.

### E. Model Predictive Control (MPC)

Autonomous racing is a challenging benchmark for evaluating vehicle dynamics models, as discrepancies between the model used in MPC and real-world dynamics directly degrade closed-loop performance, leading to suboptimal lap times or constraint violations.

To evaluate the dynamics models in optimal control, we use the *acados* framework [13]. The objective is to maximize progress subject to soft constraints on track boundaries and hard constraints on actuators limits. We extend the state with Frenet coordinates $(s, n, \mu)$ and control inputs The MPC optimizes over an 80-step (2.64 s-long) horizon. To compensate for real-world delays, we forward-integrate the model from the current state estimate to seed the solver.

### F. Reinforcement Learning (RL)

Similarly, learning racing policies using RL serves as a rigorous benchmark for vehicle dynamics models, as RL agents aggressively maximize reward, inevitably exploiting simulator inaccuracies. We train policies entirely in simulation using PPO [14]. The policy operates at $20\,\mathrm{Hz}$ while the vehicle dynamics are integrated at $100\,\mathrm{Hz}$. To enforce robust behavior without domain randomization, we inject Gaussian observation noise with standard deviation $Q$, which is provided as an input to the network. Higher values of $Q$ force the policy into a conservative driving style, enabling us to test the sim-to-real gap required for safe track completion.

## III. EXPERIMENTS

### A. Experimental Setup

Real-world experiments utilized an Xray GTXE'22 1/8th scale F1TENTH RC car. State estimation was provided by an Optitrack motion capture system. We tested on two virtual tracks: an oval and a technical L-shaped track (Fig. 3).

We evaluate SplineTM against a comprehensive set of baselines: **Dugoff** [4], **TMEasy** [2], **Pacejka** (with and without offsets) [1], **Neural ExpTanh** [9], **Neural Tires** (a 3-layer MLP predicting forces) [7], and an **MLP** approximating full vehicle dynamics [6].

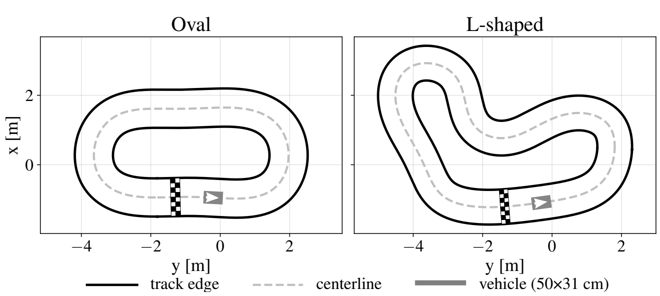


Fig. 3. Layout of the F1TENTH racetracks utilized in our evaluations.

### B. Trajectory Prediction

We evaluated open-loop prediction accuracy over $1\,\mathrm{s}$ horizons on two datasets: (i) an F1TENTH dataset featuring aggressive MPC driving and drifting, and (ii) the Stanford Vehicle Dynamics dataset [15] featuring a full-scale VW Golf operating at handling limits.

As shown in Figure 4, unstructured neural networks (Neural Tires, MLP) achieved the lowest MSE, benefiting from immense representational capacity to fit complex, high-frequency disturbances in the dataset. However, SplineTM with offsets successfully bridged this gap, outperforming the classical Pacejka baseline by 9.5% on F1TENTH and over 30% on the full-scale VW Golf, while utilizing only a fraction of the parameters of the neural models.

The significant performance gap observed in the full-scale dataset highlights the importance of the vertical and horizontal offsets $(S_h, S_v)$. Full-scale vehicles possess complex suspension geometry and rolling resistance that shift the zero-crossing of the force curve. SplineTM captures these structural phenomena natively while maintaining the flexibility to shape the saturation regions using interior control points. In contrast, the highly constrained Neural ExpTanh struggled to fit the data using our end-to-end multi-step loss formulation. SplineTM provides natural scalability while maintaining the interpretable, compact structure necessary for stable downstream control.

### C. Control Performance: MPC and RL

**MPC:** We deployed the identified models as the internal dynamics constraint in a real-time $33\,\mathrm{Hz}$ MPC on the physical F1TENTH car on both oval and L-shaped tracks in a 30-lap-long time trial. Among baselines, only Dugoff, TMEasy, and Pacejka without offsets were able to safely complete 30 laps. In fact, neural network-based models failed to complete even a single lap, which we presume can be caused by the lack of a model structure that makes them inherently easier to exploit. In turn, the models with offsets introduced a slight chattering to the control signal, leading to significantly degraded lap times. The lap times for the solver-friendly baselines and our SplineTM are presented in Figure 5. Although in general, the lap time differences are not big nominally, in the context of autonomous racing, they show a clear advantage. To ensure that the achieved results are statistically significant, we performed the Mann-Whitney U test [16]. For both tracks, we see that the proposed approach achieves the lowest median lap times, significantly outperforming Pacejka and TMEasy models (Mann-Whitney U test $p < 10^{-6}$). In turn, compared with the Dugoff tire model, our approach is comparable on the simpler oval racetrack, but on more technical L-shaped one, there is a significant advantage of $0.09\,\mathrm{s}$ in the mean lap time and the Mann-Whitney U test confirms that SplineTM achieves lower lap times with $p < 10^{-4}$.

**RL Sim-to-Real Transfer:** We trained PPO policies in simulation using the identified models and evaluated zero-shot transfer on the L-shaped track. To ensure safety and prevent track boundary violations, we incrementally scaled up observation noise $Q$ during evaluation until the policy drove conservatively enough to complete 30 laps safely.

As shown in Figure 6, policies trained with SplineTM achieved significantly lower lap times ($p < 10^{-15}$) than all baselines. Moreover, the performance gap increases with the growing number of control points used to represent the tire model. Interestingly, models offering higher prediction accuracy, such as MLP and Neural Tires, did not translate into better racing performance. We suppose that this may be caused by the training process in which the policy learns

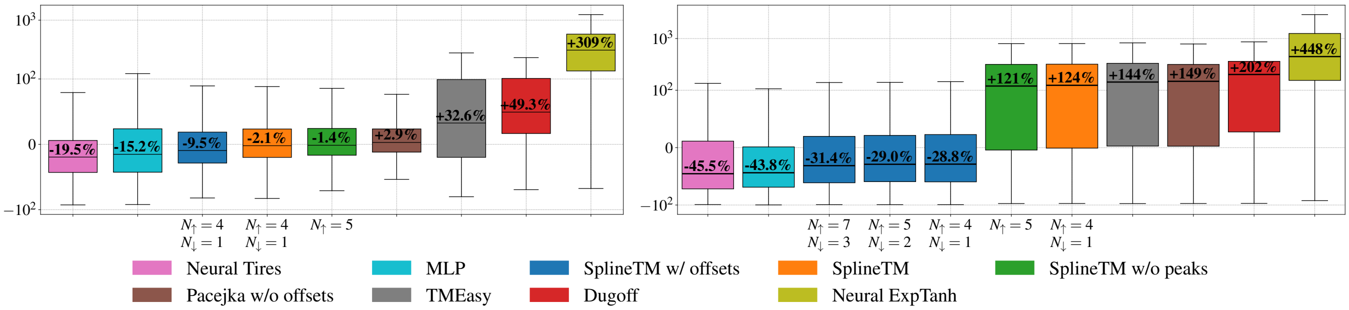


Fig. 4. Relative MSE [%] (↓ lower better) between the considered tire models and the Pacejka baseline for the F1TENTH dataset (left) and Stanford Vehicle Dynamics dataset [15] (right).

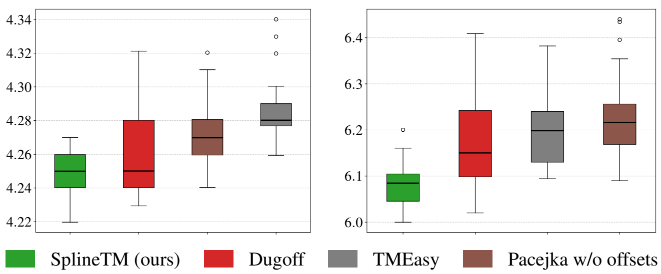


Fig. 5. Lap times [s] for models used in real-world MPC (oval - left, L-shaped - right). Neural baselines failed to deploy safely. SplineTM achieves significantly shorter lap times than Pacejka and TMEasy.

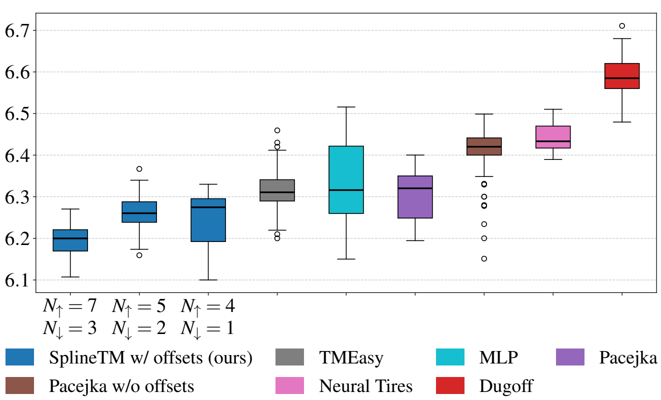


Fig. 6. Lap times [s] achieved on L-shaped track by zero-shot RL policies. SplineTM allows learning safely transferable policies that are significantly faster than all considered baselines.

how to exploit the dynamics model. Thus, even though they are slightly more accurate in the domain of the collected dataset, they may lose accuracy in the region explored by the RL policy. Moreover, unlike the semi-empirical models, purely data-driven ones do not directly impose combined slip constraints.

## IV. CONCLUSIONS

We introduced SplineTM, a B-spline-based tire modeling framework that captures nonlinear limit-handling interactions for autonomous racing. By utilizing piecewise polynomial curves with structural friction-ellipse constraints, SplineTM bridges the gap between rigid semi-empirical models and unconstrained neural networks. Our extensive evaluations demonstrated that while pure neural networks yield the lowest offline prediction errors, their lack of physical structure renders them unsafe or suboptimal in optimization-based control (MPC) and RL. SplineTM effectively prevents policy exploitation and solver divergence while offering superior trajectory prediction and significantly faster real-world lap times compared to classical baselines. Future work will explore online adaptation of the B-spline control points to accommodate real-time tire degradation.

## REFERENCES

[1] H. B. Pacejka and I. Besselink, *Tire and vehicle dynamics*, 3rd ed. Amsterdam ;: Elsevier Sci. & Technology, 2012.

[2] W. Hirschberg, G. Rill, and H. Weinfurter, "User-appropriate tyre-modelling for vehicle dynamics in standard and limit situations," *Vehicle System Dynamics*, vol. 38, no. 2, pp. 103–125, 2002.

[3] H. Dugoff, P. S. Fancher, and L. Segel, "Tire performance characteristics affecting vehicle response to steering and braking control inputs," in *Technical Report Contract CST-460, Office of Vehicle Syst. Research*, 1969.

[4] N. Ding and S. Taheri, "A modified dugoff tire model for combined-slip forces," *Tire Sci. And Technol.*, vol. 38, no. 3, pp. 228–244, 2010.

[5] E. Fiala, "Seitenkräfte am rollenden luftreifen," *Zeitschrift VDI*, vol. 96, no. 29, pp. 1114–1119, 1954.

[6] N. Ding, M. Thompson, J. Dallas, J. Y. Goh, and J. Subosits, "Drifting with unknown tires: Learning vehicle models online with neural networks and model predictive control," in *2024 IEEE Intell. Vehicles Symposium (IV)*, 2024, pp. 2545–2552.

[7] J. Węgrzynowski, G. Czechmanowski, P. Kicki, and K. Walas, "Learning dynamics models for velocity estimation in autonomous racing," in *IEEE/RSJ Int. Conf. on Intell. Rob. and Syst. (IROS)*, pp. 972–979, ISSN: 2153-0866.

[8] J. Chrosniak, J. Ning, and M. Behl, "Deep dynamics: Vehicle dynamics modeling with a physics-constrained neural network for autonomous racing," *IEEE Robot. and Autom. Lett.*, vol. 9, no. 6, pp. 5292–5297, 2024.

[9] F. Djeumou, J. Y. Goh, U. Topcu, and A. Balachandran, "Autonomous drifting with 3 minutes of data via learned tire models," in *2023 IEEE Int. Conf. on Robot. and Autom. (ICRA)*, 2023, pp. 968–974.

[10] Y. Tsuchiya, T. Balch, P. Drews, and G. Rosman, "Online adaptation of learned vehicle dynamics model with meta-learning approach," in *IEEE/RSJ Int. Conf. on Intell. Rob. and Syst. (IROS)*, 2024, pp. 802–809.

[11] C. d. Boor, *A Practical Guide to Splines*. New York: Springer Verlag, 1978.

[12] R. Rajamani, *Vehicle Dynamics and Control*, 2nd ed. Springer, 2012.

[13] R. Verschueren *et al.*, "acados – a modular open-source framework for fast embedded optimal control," *Mathematical Programming Computation*, 2021.

[14] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal Policy Optimization Algorithms," Aug. 2017, arXiv:1707.06347 [cs].

[15] D. Mori, R. K. Aggarwal, N. D. Broadbent, T. Kobayashi, and J. C. Gerdes, "Vehicle dynamics dataset for highly dynamic automated driving," 2025.

[16] H. B. Mann and D. R. Whitney, "On a test of whether one of two random variables is stochastically larger than the other," *The Annals of Mathematical Statistics*, vol. 18, no. 1, pp. 50–60, 1947.