

AnyCar to Anywhere: Learning Universal Dynamics Model for Agile and Adaptive Mobility

Wenli Xiao[†]

CMU

wxiao2@andrew.cmu.edu

[†] Equal Contribution

Haoru Xue[†]

UC Berkeley

haoru-xue@berkeley.edu

[†] Equal Contribution

Tony Tao

CMU

Dvij Kalaria

UC Berkeley

John Dolan

CMU

Guanya Shi

CMU

Abstract: Recent works in the robot learning community have successfully introduced generalist models capable of controlling various robot embodiments across a wide range of tasks, such as navigation and locomotion. However, achieving agile control, which pushes the limits of robotic performance, still relies on specialist models that require extensive parameter tuning. To leverage generalist-model adaptability and flexibility while achieving specialist-level agility, we propose AnyCar, a transformer-based generalist dynamics model designed for agile control of various wheeled robots. To collect training data, we unify multiple simulators and leverage different physics backends to simulate vehicles with diverse sizes, scales, and physical properties across various terrains. With robust training and real-world fine-tuning, our model enables precise adaptation to different vehicles, even in the wild and under large state estimation errors. In real-world experiments, AnyCar shows both few-shot and zero-shot generalization across a wide range of vehicles and environments, where our model, combined with a sampling-based MPC, outperforms specialist models by up to 54%. These results represent a key step toward building a foundation model for agile wheeled robot control. We will also open-source our framework to support further research.

Keywords: Dynamics Learning, Foundation Model, Agile Control

1 Introduction

The use of the transformer [1] architecture in contemporary robot learning is ubiquitous across perception [2, 3], planning [4] and control [5] tasks. Reinforcement learning with transformers, such as [6] and [7], is used in many downstream applications in bi-manual manipulation [8], navigation [9], humanoid locomotion [10], and whole-body tele-operation [11, 12]. Vision-language-action (VLA) models such as OpenVLA [13], RT-1 [14], and RT-2 [15] demonstrate the scalability of employing transformers in robotics. These models trained on internet-scale data can generalize knowledge and skills to different complex tasks.

Recent advances in robot learning have also introduced more “specialist” systems that can perform highly agile locomotion tasks [16, 17, 18, 19]. In particular, on wheeled robots, previous works have achieved high-speed autonomous driving on racetracks [20], grass fields [21], loose sand [22, 23], and off-road terrain [24, 25]. However, most of these efforts [24, 25] are optimized for specific car models and environments, requiring extensive system identification and model training [22, 21], which is costly to fine-tune and difficult to transfer to other wheeled platforms.

On the other hand, agile wheeled control for safety-critical applications requires precise dynamics modeling when running at high speed, since small errors can lead to catastrophic failures such as



Figure 1: Performance of AnyCar and baselines in the wild under state estimation errors. Above: A 10 cm tolerance corridor is set as a checkpoint. Below: each row represents the true trajectory of one method, and each column corresponds to a specific setting for the 1/16 scale car: high speed (2 m/s), towing a box, and replacing the front left tire with a plastic wheel. All settings significantly alter the vehicle dynamics.

crashes [26, 21]. There are works that attempt to mitigate this issue by applying neural system identification [5, 27, 26] to adapt the model to different environmental factors such as tire degradation, ground surface imperfections [27], and towed objects [26]. Nevertheless, these methods still require assumptions about the specific car setup (e.g., size and wheelbase of the car).

A key question that arises is: *Can we train a generalist wheeled-robot dynamics model that achieves the performance of a specialist model for each setup?* In this work, we propose AnyCar (depicted in Figure 2), an initial effort to train a vehicle dynamics transformer that can predict the trajectory of various cars in various settings through in-context adaptation. Our contributions are three-fold:

- We build a universal synthetic data generator that collects data across diverse vehicles and environments, using physics engines with varying levels of fidelity (e.g., DBM, MuJoCo, Isaac Sim, Assetto Corsa Gym).
- We propose a two-phase robust vehicle dynamics transformer training method that leverages simulation pre-training and real-world fine-tuning to handle sim2real mismatches and state-estimation errors.
- We integrate the dynamics transformer with a sampling-based MPC and demonstrate real-world performance on different car platforms and in different environments, both indoor and outdoor, achieving up to 54% performance improvement over the baseline methods.

2 Related Work

2.1 Neural Dynamics Model and Adaptive Control

Neural networks, especially transformers, can be used to learn the dynamics of any arbitrary systems [28, 29] or a residual on top of a nominal model, which is expressed in traditional state-space equations [5, 30, 31, 32, 33]. Data-driven techniques can help robots adapt to time-varying dynamics. In particular, open-loop adaptation based on teacher-student training can effectively bridge the Sim2Real gap in RL, such as rapid motor adaptation (RMA) [34]. When real-world ground truth is available, the adaptation can be learned in a more supervised fashion. Safe deep policy adaptation (SafeDPA) [26] performs self-supervised real-world fine-tuning. Neural-Fly [30] trains a residual dynamics model to learn a good representation of environment disturbances with a small amount of real-world data. Learning model predictive control (LMPC) [35, 20] directly regresses a local linear

approximation of the state-space dynamics based on a neighborhood of nearest historical states collected in the real world. However, adaptation-at-scale remains an open challenge, particularly when dynamics differ significantly, such as with robots varying in parameters or embodiments.

2.2 Cross Embodiment

Recent research shows that deep learning models trained at scale can control a variety of robots using the same policy [36, 37]. CrossFormer [36] highlights that training a single transformer on tasks across six embodiments (wheeled robot, quadruped, manipulator, etc.) works even without aligning the action space of these embodiments. The open X-Embodiment dataset [38, 39] focuses on robot manipulation across different embodiments. However, no existing generalist performs agile control. To reconcile generalist adaptability and specialist agility, AnyCar addresses generalization across the same form of wheeled embodiment, which is a scope reduction from cross-embodiment.

2.3 Transformer for Low-Level Control

Transformers have shown potential in low-level control, with attention patterns in Trajectory Transformer effectively capturing properties of Markov decision processes (MDPs) [6]. This raises the question of whether transformers, with their inductive bias on pairwise attention between sequence elements, can achieve efficient training and representation of Markovian dynamics for low-level control. Recent work has begun exploring this area. [12] uses a transformer for humanoid whole-body control, though without addressing cross-embodiment generalization. [5] applies a transformer for online system identification in vehicular robots, using historical states to generate a context vector for a neural dynamics model. Despite these advancements, current literature lacks examples of transformers excelling in both specialized tasks (e.g., agile locomotion or dexterous manipulation) and generalization across different robots or tasks.

3 Overview

Notation. x_t, a_t are state and action at time step t . We use $x_{1:t}$ to denote a sequence $\{x_1, \dots, x_t\}$. \hat{x} denotes estimation of x (e.g., from VIO) and \hat{a} is a with noise.

To demonstrate the practical application of our method, we focus on trajectory tracking in unstructured environments for a range of vehicles, which can be formulated as follows:

$$\begin{aligned} & \underset{a_{0:T}}{\text{maximize}} && \sum_{t=0}^T R(x_t, a_t) \\ & \text{subject to} && x_{t+1} = f(x_t, a_t, c_t), \quad \forall t = 0, 1, \dots, T \end{aligned}$$

where $R(x_t, a_t)$ is the reward function and $x_{t+1} = f(x_t, a_t, c_t)$ represents the system dynamics with c_t representing all physics characteristics related to car dynamics and the environment conditions such as terrain, payload, .etc. The state is defined by $x_t \triangleq [p_t^x, p_t^y, \psi_t, \dot{p}_t^x, \dot{p}_t^y, \omega]$, which contains position (p_t^x, p_t^y) , heading angle ψ_t , linear velocity $(\dot{p}_t^x, \dot{p}_t^y)$, and angular velocity ω . The action $a_t \triangleq [\mathcal{T}, \delta]$ contains throttle \mathcal{T} and steering angle δ . Our model is a seq2seq model that can predict the future state sequence from *imperfect* state and action history sequence and future action sequence:

$$x_{t+1:t+H} \approx f_{\theta}^{\text{AnyCar}} \left(\underbrace{\hat{x}_{t-K:t}, \hat{a}_{t-K:t-1}}_{\text{noisy state and action history}}, \underbrace{a_{t:t+H-1}}_{\text{future actions}} \right), \quad (1)$$

where K denotes history length and H denotes prediction horizon (illustrated in Figure 2 Phase 2). We train a transformer to approximate Equation (1) for various cars and terrains via its in-context adaptation capability. Our model not only can learn adaptive dynamics but also a filter that can handle noisy state estimation \hat{x} and action \hat{a} . Compared with previous works that assumes one specific car model with limited adaptable parameters [27, 20, 26, 22, 5], AnyCar can adapt to various types of car with or without assumptions about the environment. To learn θ , we design two-stage training pipeline. In the first stage, in Section 4, we generate large scale dataset which

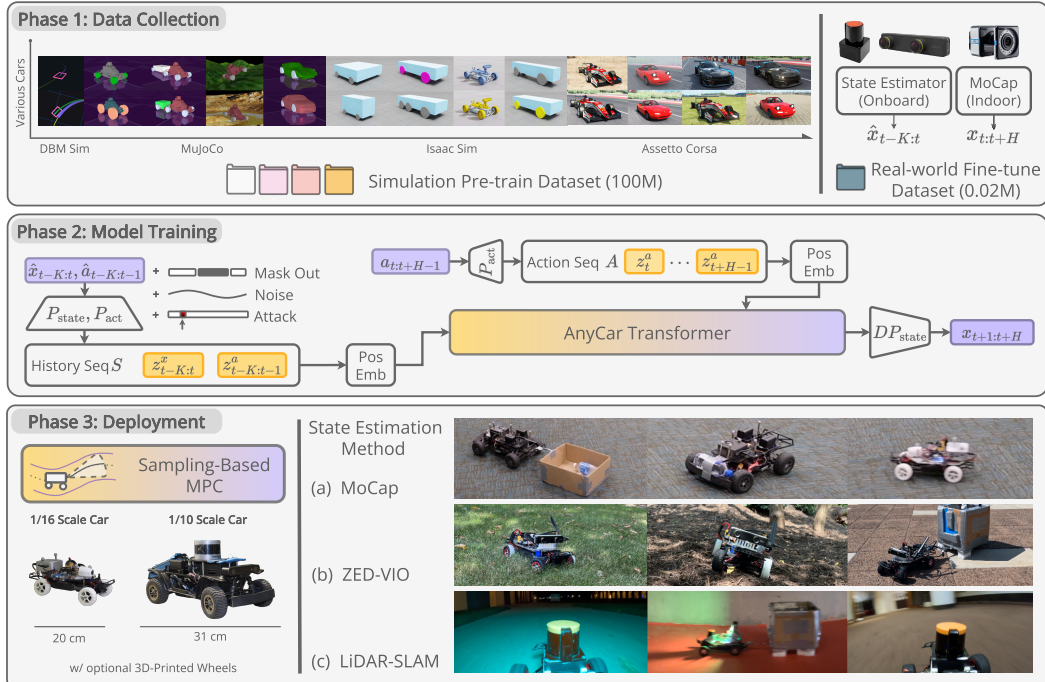


Figure 2: AnyCar System Pipeline. **Phase 1:** We collect 100M data in 4 different simulations for pre-training and 0.02M few-shot real-world data for fine-tuning the model. **Phase 2:** We pre-train the model with the simulation dataset and enhance prediction robustness through masking, adding noise, and attacking the inputs. We also fine-tune using the fine-tuning dataset. **Phase 3:** We deploy AnyCar in the wild under state estimation error (using SLAM [40, 41, 42] and VIO [43]) to control different vehicles (1/10 scale, 1/16 scale) with different settings (tow object, 3D-printed wheels) on different terrains.

contains trajectories of various cars in different terrains using different physics simulations; we then pre-train the AnyCar transformer $f_{\theta_{\text{sim}}}^{\text{AnyCar}}$ (Figure 2 Phase 2) in Section 4.2. In the second stage, we collect few-shot real-world data and fine-tune the transformer $f_{\theta_{\text{fine-tune}}}^{\text{AnyCar}}$ in Section 4.3. After these two stages, the model $f_{\theta_{\text{fine-tune}}}^{\text{AnyCar}}$ can accurately predict future trajectory for different cars even under state-estimation error.

With an accurate dynamics model, we apply the Model Predictive Path Integral (MPPI) to perform trajectory tracking. MPPI is a sampling-based MPC approach that minimizes the cost-to-go for sampled trajectories and selects optimal controls by weighting multiple candidate trajectories based on their performance. We choose MPPI for its ability to leverage parallel computation and its training-free nature, with no need for a reduced-order model. In Section 5.1, we describe the detailed system design for trajectory tracking using AnyCar transformer with MPPI to achieve control at 50Hz in the real world. Finally, in Section 6, we showcase the deployment of AnyCar in various real-world scenarios, demonstrating its versatility and robustness across different vehicles and environments.

4 MODEL AND DATA

In this section, we describe the dataset collection and model training strategies of AnyCar. As shown in Figure 2 Phase 1 and Phase 2, we highlight our data collection in massive simulation and few-shot real-world data, and our robust training and finetuning pipeline.

4.1 Pre-Training with Massive Simulated Data

Scene Generation. To collect a diverse dataset, we leverage the low-cost nature of simulation and generate a large amount of simulated data. Our data generation has three sources of diversity: 1)

dynamics, 2) scenario, 3) controller. As illustrated in Figure 2 Phase 1, we synthesize trajectories of various cars and terrains via four distinct simulators: Dynamic Bicycle Model [27] (DBM)-based numeric simulation, MuJoCo, IsaacSim, and Assetto Corsa Gym [44].

Curriculum Model Training. To ensure data distribution coverage, we diversify between on-policy and off-policy data via a two-stage curriculum learning method. In stage one, we collect high-volume off-policy data as warm-up, by building a hybrid controller where we use a pure pursuit controller for steering δ and a PD controller for throttle \mathcal{T} , to track randomly synthesized reference tracks. After collecting 200M timesteps that are usable for training a general model for non-agile tasks (in which the target velocity is smaller than the physical limit), we switch to stage two, where we deploy an on-policy NN-MPPI controller (described in Section 5.1) to track agile trajectories, and periodically update the network with the collected on-policy trajectory.

4.2 Robust In-Context Adaptive Dynamics Model

Model Structure. As illustrated in Figure 2 Phase 2, the historical states in Equation (1) are linearly projected into a 64-dimensional latent space with an encoder layer $P_{\text{state}}(x) : \mathbb{R}^6 \rightarrow \mathbb{R}^{64}$. The historical actions pass through a different encoder $P_{\text{act}}(a) : \mathbb{R}^2 \rightarrow \mathbb{R}^{64}$ of similar size. We then interleave and stack them to make a complete history token sequence $S^{2K-1 \times 64}$, which is supplied to the transformer as the context. The future actions are tokenized with $P_{\text{act}}(a)$ and stacked as a sequence $A^{H \times 64}$. Both S and A sequences are then summed with two learnable positional encoders. After passing the context S and input A to the transformer decoder, we obtain $S'^{H \times 64}$, which is then down-projected into the state-space with $DP_{\text{state}}(x) : \mathbb{R}^{64} \rightarrow \mathbb{R}^6$. This becomes the final state sequence prediction.

Robust Training. To learn a robust model f_θ for Equation (1) under various disturbances, we propose to add three techniques in pre-training, i.e., mask out, add noise, and attack (Figure 2). We implement *mask out* by applying randomized cross-attention mask to transformer, then *add noise* $\epsilon \sim \mathcal{N}(0, \epsilon^{\max})$. We also add *attack*: unreasonably large or small values to random dimensions in history to simulate state estimation errors in the real-world. As discussed in Section 4, these techniques collaboratively improves robustness for successful real-world deployment.

4.3 Fine-Tuning with state-estimation error reduction

After pre-training $f_{\theta_{\text{sim}}}^{\text{AnyCar}}$, we fine-tune the model with 10 minutes of real-world data (0.02M) to reduce the sim2real gap and mitigate state-estimation error. To collect real-world data, we control the car via joystick to follow a few random curved trajectories in a motion capture field and collect both the state from on-board state-estimators and the ground truth states provided by the Vicon system, as shown in Figure 2 Phase 1. We then fine-tune $f_{\theta_{\text{sim}}}^{\text{AnyCar}}$ on collected data under the constraint of $\|\theta_{\text{fine-tune}} - \theta_{\text{sim}}\|_2 \leq \epsilon_{\text{tune}}$ to prevent catastrophic forgetting. In practice, we also apply data rehearsal to assist. The model after fine-tuning $f_{\theta_{\text{fine-tune}}}^{\text{AnyCar}}$ is ready to deploy at zero-shot.

5 SYSTEM DESIGN

5.1 MPPI Controller

As introduced in Section 3, we choose to evaluate our method with MPPI instead of RL and MPC, because RL would have required optimizing additional policy and value function parameters, which is outside the scope of this work. Likewise, neural MPC [45] would have resulted in a reduced-order approximation of our neural dynamics model that is undesired for a fair evaluation.

We employ a variant of MPPI called Covariance-Optimal MPC (CoVO-MPC) [46], which improves upon vanilla MPPI with adaptive sampling covariance to achieve an optimal convergence rate on the cost function.




Figure 3: Visualization of AnyCar’s transformer attention in three real-world settings: (a) low speed at 0.5 m/s, (b) high speed at 2 m/s, and (c) towing an object at 2 m/s, all tracking the same reference trajectory. AnyCar’s transformer consistently focuses on the nearest 50 steps across all settings and adaptively attends to different sections of the track. For example, it attends to the first corner in setting (b) and the second corner in setting (c).

5.1.1 Trajectory Sampling

Let x_t, a_t be state and action at time t . For given control horizon H , we randomly sample N action sequences $\{a_{t:t+H-1}^i\}_i^N$ in a normal distribution, whose mean and covariance are computed with [46]. We roll out each action sequence with the AnyCar model and compute their cumulative rewards. In order to generate smooth action sequences, we re-parameterize the control sequence $a_{0:T}$ with a set of time-indexed knots represented by $\theta_{0:k}$ [47]. Given query point τ , the control can be evaluated by $a_\tau = \text{spline}(\tau; (\tau_{0:k}, \theta_{0:k}))$.

5.1.2 Reward function

To motivate tracking reference waypoints, we use the single-step reward function from [24], which is given by $r(x, a, \hat{x}) = w_1 \|p - \hat{p}\|^2 + w_2 \|\psi - \hat{\psi}\| + w_3 \|v_x - \hat{v}_x\| + w_4 \|\delta a\|$, where p denotes position, ψ is heading angle, v_x is longitude velocity, and δa is action increment.

5.1.3 Computation

We implement the MPPI in JAX [48] and the transformer model in TransformerEngine. We evaluate 600 action sequence samples on the model in parallel, achieving 20 ms (50 Hz) real-time performance on a RTX 4090 GPU. We note that further optimizations, such as KV caching, sharing history token attentions between samples, and TensorRT conversion can enable edge deployment.

5.2 State Estimation

In our experiments with ground truth data, we use a motion capture system to directly observe the position and velocity of the vehicle. For in-the-wild experiments, we estimate the linear and angular velocity of the vehicle by fusing motor odometry and IMU data using an Extended Kalman Filter (EKF) [49]. We then correct for the odometry drifting using 2D LiDAR SLAM [40, 41, 42], or 3D SLAM with visual-inertial odometry (VIO). We will demonstrate AnyCar’s fine-tuning pipeline can adapt to state estimation error in Section 6.

5.3 Low-Level Controller

Upon receiving a throttle (acceleration) and steering (angle) command, a low-level module maps the throttle to motor current with a simple linear mapping, and maps the steering angle to servo angle. Adapting to the discrepancies in these rough mappings and the latency in actuator responses is also a goal of our fine-tuning pipeline.

6 Experiment

In this section, we aim to demonstrate the capability of the proposed AnyCar by addressing the following questions:

Table 1: Indoor Results using Ground Truth State-Estimation (Motion Capture)

Setting	Method	$E^{\text{Prediction}} \downarrow$	$E^{\text{Tracking}} \downarrow$
Few-shot performance in fine-tuned scenarios			
1/10 Scale Car + Tow Box	PP	-	0.45 ± 0.02
	AnyCar (Ours)	0.27 ± 0.24	0.35 ± 0.04
	AnyCar w/o FT	0.54 ± 0.55	0.47 ± 0.23
	Specialist	0.14 ± 0.09	0.43 ± 0.05
1/10 Scale Car + Payloads	PP	-	0.54 ± 0.03
	AnyCar (Ours)	0.11 ± 0.08	0.41 ± 0.03
	AnyCar w/o FT	0.21 ± 0.11	0.47 ± 0.17
	Specialist	0.26 ± 0.11	0.51 ± 0.03
Zero-shot generalization in unseen scenarios			
1/10 Scale Car + Plastic wheels (All 4 wheels)	PP	-	0.79 ± 0.62
	AnyCar (Ours)	0.26 ± 0.32	0.335 ± 0.03
	AnyCar w/o FT	0.32 ± 0.21	0.45 ± 0.15
	Specialist	0.35 ± 0.15	0.334 ± 0.06
1/10 Scale Car + Plastic wheels (Front 2 wheels)	PP	-	0.52 ± 0.05
	AnyCar (Ours)	0.12 ± 0.08	0.39 ± 0.05
	AnyCar w/o FT	0.20 ± 0.11	0.49 ± 0.09
	Specialist	0.27 ± 0.11	0.41 ± 0.04
1/10 Scale Car + Plastic wheels + Tow box	PP	-	0.57 ± 0.03
	AnyCar (Ours)	0.09 ± 0.06	0.49 ± 0.09
	AnyCar w/o FT	0.18 ± 0.11	0.52 ± 0.09
	Specialist	0.14 ± 0.05	0.60 ± 0.04
1/16 Scale Car + Plastic wheels	PP	-	0.37 ± 0.16
	AnyCar (Ours)	0.17 ± 0.08	0.31 ± 0.06
	AnyCar w/o FT	0.25 ± 0.14	0.44 ± 0.08
	Specialist	0.26 ± 0.11	0.55 ± 0.07

Table 2: Indoor Results under State-Estimation Error

Setting	Method	$E^{\text{Prediction}} \downarrow$	$E^{\text{Tracking}} \downarrow$
Few-shot performance in fine-tuned scenarios			
1/16 Scale Car + Low Speed	AnyCar (Ours)	0.30 ± 0.06	0.35 ± 0.02
	AnyCar w/o FT	0.32 ± 0.07	0.48 ± 0.06
	Specialist	0.33 ± 0.05	0.42 ± 0.01
Zero-shot generalization in unseen scenarios			
1/16 Scale Car + High Speed	AnyCar (Ours)	0.44 ± 0.11	0.57 ± 0.03
	AnyCar w/o FT	0.52 ± 0.15	1.30 ± 0.11
	Specialist	0.48 ± 0.10	1.26 ± 0.89
1/16 Scale Car + Tow 2 Box	AnyCar (Ours)	0.34 ± 0.10	0.57 ± 0.02
	AnyCar w/o FT	0.41 ± 0.14	1.41 ± 0.09
	Specialist	0.39 ± 0.09	0.93 ± 0.62

- **Q1:** Can our model generalize to various cars and terrains, and outperform specialist models?
- **Q2:** Can our model maintain its adaptation capability even with imperfect state estimation?
- **Q3:** Why does the proposed robust vehicle dynamics transformer outperform other baseline models?

Baselines. We compare AnyCar with three baseline methods: 1) *AnyCar w/o FT*: AnyCar without the real-world fine-tuning phase, 2) *PP*: Pure Pursuit controller for steering and PID controller for velocity tracking, and 3) *Specialist*: a DBM model with system identification. We also consider two types of state estimator setups: 1) motion capture (indoor), which can be treated as ground truth, and 2) SLAM [40, 41, 42] (indoor or outdoor), which is less accurate than motion capture and much prone to drifting issues.

Metrics. We use the model prediction error $E^{\text{Prediction}} \triangleq \|x_{t+1:t+H}^{\text{pred}} - x_{t+1:t+H}^{\text{gt}}\|_2$ to assess prediction accuracy. We also define $E^{\text{Tracking}} \triangleq w_2 \|p_t - \hat{p}_t\|_2 + w_3 \|v_t - \hat{v}_t\|_2$, where w_2 and w_3 are the same weights defined for position and velocity tracking rewards in Section 5.1, representing the weighted sum of lateral error and velocity tracking error, to evaluate trajectory tracking performance.

6.1 Evaluate Model In-context Adaptation Capability

To answer **Q1**, we isolate the estimation error and use motion capture to provide ground truth for state estimation in the real world. An 1/10 Scale Car was employed to tow objects and carry payloads, to create a fine-tuning dataset collected by human teleoperation. This dataset was then used to fine-tune our dynamics model. The model was evaluated across two categories: **few-shot performance** (e.g., towing objects and varying payloads, both present in the fine-tuning dataset) and **zero-shot generalization** (e.g., changing 3D-printed wheels, towing objects with modified wheels, and using a smaller car model with altered wheels). We use a trajectory optimization method [20] to compute an on-the-edge reference trajectory on a raceline. Our method, along with other baselines (AnyCar w/o FT, PP, Specialist), was then deployed to track this trajectory. The evaluation considered both model prediction error and closed-loop tracking error with the full controller. Results in Table 1 show AnyCar reaching and outperforming baselines in terms of prediction error and tracking error in both few-shot and zero-shot cases. We observed the same “emergent skill” seen in RT-X[38] that after fine-tuning the model with data of one robot (1/10 scale car) performing a




Figure 4: Comparison with baselines in the wild.

certain task (track curves at low speed), a different robot (1/16 scale car) also gets a performance boost in doing a different task (agile raceline tracking).

6.2 Evaluate Model Capability in the Wild

To address **Q2**, we set up a 2D LiDAR-based SLAM stack on an 1/16 scale car. The car was teleoperated to follow an S-shaped path at low speed, collecting odometry data from both the SLAM system and motion capture to create a small-scale dataset of 24,000 timesteps. The model was then fine-tuned using the method in Section 4.3. Using the fine-tuned model (AnyCar), we first evaluated our method in an indoor environment under different conditions (low speed, high speed, towing two boxes) with SLAM, and computed metrics using ground truth data. The results in Table 2 show AnyCar outperforms baselines with peak improvement of 54%. Based on Table 1 and Table 2, we prove the necessity of fine-tuning that aligning the dynamics model to handle state estimation error.

Next, the system was moved to an outdoor environment without further modification. To validate the robustness to state-estimation capabilities of AnyCar, we set up a narrow corridor along the reference trajectory, allowing the car to pass with a 10 cm tolerance (demonstrated in Figure 1). Poor tracking performance would result in the car colliding with the walls. We compared our method with AnyCar w/o FT and Specialist by calculating the percentage of times the car successfully passed all checkpoints without collision (success rate) for each method. The results in Table 2 show AnyCar achieves the highest success rate in all settings, with the specialist failing consistently due to the state estimation error. We also visualize the transformer attention across different settings in Figure 3, which demonstrates AnyCar’s in-context adaptation in various settings. In addition to SLAM-based state-estimator, we also show the deployment with ZED-VIO [43] on the website¹.

6.3 Interpret the expressiveness and robustness of AnyCar

To address **Q3**, we apply ablation studies on model architecture and training pipeline. Due to the limit of page, we put the detailed results in Appendix .1. The combination of model selection, large-scale data, and robust training explains why the AnyCar performs better than the baseline models.

7 Limitations And Future Work

In this paper, we propose AnyCar, a first step towards foundation model for agile wheeled control. In the future, there are three interesting research directions. One is to use KV caching in the transformer for full onboard computation. The second is to optimize the MPPI control to be aware of model uncertainty and safety. The third is to integrate with existing foundation models for visual navigation [9] to achieve fully agile autonomy in the wild.

¹website: <https://lecar-lab.github.io/anycar/>

References

- [1] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, u. Kaiser, and I. Polosukhin. Attention is All you Need. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL https://papers.nips.cc/paper_files/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html.
- [2] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. Oct. 2020. URL <https://openreview.net/forum?id=YicbFdNTTy>.
- [3] H. Liu, C. Li, Q. Wu, and Y. J. Lee. Visual Instruction Tuning. Nov. 2023. URL <https://openreview.net/forum?id=w0H2xGHlkw>.
- [4] L. Lehnert, S. Sukhbaatar, D. Su, Q. Zheng, P. McVay, M. Rabbat, and Y. Tian. Beyond A*: Better Planning with Transformers via Search Dynamics Bootstrapping. Aug. 2024. URL <https://openreview.net/forum?id=SGoVIC0u0f#discussion>.
- [5] S. J. Wang, H. Zhu, and A. M. Johnson. Pay Attention to How You Drive: Safe and Adaptive Model-Based Reinforcement Learning for Off-Road Driving, Oct. 2023. URL <http://arxiv.org/abs/2310.08674>. arXiv:2310.08674 [cs].
- [6] M. Janner, Q. Li, and S. Levine. Offline Reinforcement Learning as One Big Sequence Modeling Problem. In *Advances in Neural Information Processing Systems*, volume 34, pages 1273–1286. Curran Associates, Inc., 2021. URL <https://proceedings.neurips.cc/paper/2021/hash/099fe6b0b444c23836c4a5d07346082b-Abstract.html>.
- [7] L. Chen, K. Lu, A. Rajeswaran, K. Lee, A. Grover, M. Laskin, P. Abbeel, A. Srinivas, and I. Mordatch. Decision transformer: reinforcement learning via sequence modeling. In *Proceedings of the 35th International Conference on Neural Information Processing Systems, NIPS '21*, pages 15084–15097, Red Hook, NY, USA, June 2024. Curran Associates Inc. ISBN 978-1-71384-539-3.
- [8] T. Zhao, V. Kumar, S. Levine, and C. Finn. Learning Fine-Grained Bimanual Manipulation with Low-Cost Hardware. In *Robotics: Science and Systems XIX*. Robotics: Science and Systems Foundation, July 2023. ISBN 978-0-9923747-9-2. doi:10.15607/RSS.2023.XIX.016. URL <http://www.roboticsproceedings.org/rss19/p016.pdf>.
- [9] D. Shah, A. Sridhar, N. Dashora, K. Stachowicz, K. Black, N. Hirose, and S. Levine. ViNT: A Foundation Model for Visual Navigation, June 2023. URL <https://arxiv.org/abs/2306.14846v2>.
- [10] I. Radosavovic, B. Zhang, B. Shi, J. Rajasegaran, S. Kamat, T. Darrell, K. Sreenath, and J. Malik. Humanoid Locomotion as Next Token Prediction, Feb. 2024. URL <http://arxiv.org/abs/2402.19469>. arXiv:2402.19469 [cs].
- [11] X. Cheng, J. Li, S. Yang, G. Yang, and X. Wang. Open-TeleVision: Teleoperation with Immersive Active Visual Feedback, July 2024. URL <http://arxiv.org/abs/2407.01512>. arXiv:2407.01512 [cs].
- [12] Z. Fu, Q. Zhao, Q. Wu, G. Wetzstein, and C. Finn. HumanPlus: Humanoid Shadowing and Imitation from Humans.
- [13] M. J. Kim, K. Pertsch, S. Karamcheti, T. Xiao, A. Balakrishna, S. Nair, R. Rafailov, E. Foster, G. Lam, P. Sanketi, Q. Vuong, T. Kollar, B. Burchfiel, R. Tedrake, D. Sadigh, S. Levine, P. Liang, and C. Finn. OpenVLA: An Open-Source Vision-Language-Action Model, June 2024. URL <http://arxiv.org/abs/2406.09246>. arXiv:2406.09246 [cs].
- [14] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, J. Dabis, C. Finn, K. Gopalakrishnan, K. Hausman, A. Herzog, J. Hsu, J. Ibarz, B. Ichter, A. Irpan, T. Jackson, S. Jesmonth, N. Joshi, R. Julian, D. Kalashnikov, Y. Kuang, I. Leal, K.-H. Lee, S. Levine, Y. Lu, U. Malla, D. Manjunath, I. Mordatch, O. Nachum, C. Parada, J. Peralta, E. Perez, K. Pertsch, J. Quiambao, K. Rao, M. Ryoo, G. Salazar, P. Sanketi, K. Sayed, J. Singh, S. Sontakke, A. Stone, C. Tan, H. Tran, V. Vanhoucke, S. Vega, Q. Vuong, F. Xia, T. Xiao, P. Xu, S. Xu, T. Yu, and B. Zitkovich. RT-1: Robotics Transformer for Real-World Control at Scale. In *Robotics: Science and Systems XIX*. Robotics: Science and Systems Foundation, July 2023. ISBN 978-0-9923747-9-2. doi:10.15607/RSS.2023.XIX.025. URL <http://www.roboticsproceedings.org/rss19/p025.pdf>.

- [15] B. Zitkovich, T. Yu, S. Xu, P. Xu, T. Xiao, F. Xia, J. Wu, P. Wohlhart, S. Welker, A. Wahid, Q. Vuong, V. Vanhoucke, H. Tran, R. Soricut, A. Singh, J. Singh, P. Sermanet, P. R. Sanketi, G. Salazar, M. S. Ryoo, K. Reymann, K. Rao, K. Pertsch, I. Mordatch, H. Michalewski, Y. Lu, S. Levine, L. Lee, T.-W. E. Lee, I. Leal, Y. Kuang, D. Kalashnikov, R. Julian, N. J. Joshi, A. Irpan, B. Ichter, J. Hsu, A. Herzog, K. Hausman, K. Gopalakrishnan, C. Fu, P. Florence, C. Finn, K. A. Dubey, D. Driess, T. Ding, K. M. Choromanski, X. Chen, Y. Chebotar, J. Carbajal, N. Brown, A. Brohan, M. G. Arenas, and K. Han. RT-2: Vision-Language-Action Models Transfer Web Knowledge to Robotic Control. In *Proceedings of The 7th Conference on Robot Learning*, pages 2165–2183. PMLR, Dec. 2023. URL <https://proceedings.mlr.press/v229/zitkovich23a.html>. ISSN: 2640-3498.
- [16] T. He, C. Zhang, W. Xiao, G. He, C. Liu, and G. Shi. Agile But Safe: Learning Collision-Free High-Speed Legged Locomotion.
- [17] C. Zhang, W. Xiao, T. He, and G. Shi. WoCoCo: Learning Whole-Body Humanoid Control with Sequential Contacts.
- [18] T. He, Z. Luo, X. He, W. Xiao, C. Zhang, W. Zhang, K. Kitani, C. Liu, and G. Shi. OmniH2O: Universal and Dexterous Human-to- Humanoid Whole-Body Teleoperation and Learning.
- [19] T. He, Z. Luo, W. Xiao, C. Zhang, K. Kitani, C. Liu, and G. Shi. Learning Human-to-Humanoid Real-Time Whole-Body Teleoperation, Mar. 2024. URL <http://arxiv.org/abs/2403.04436>. arXiv:2403.04436 [cs, eess].
- [20] H. Xue, E. L. Zhu, J. M. Dolan, and F. Borrelli. Learning Model Predictive Control with Error Dynamics Regression for Autonomous Racing. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 13250–13256, May 2024. doi:10.1109/ICRA57147.2024.10611628. URL <https://ieeexplore.ieee.org/abstract/document/10611628>.
- [21] K. Stachowicz and S. Levine. RACER: Epistemic Risk-Sensitive RL Enables Fast Driving with Fewer Crashes, May 2024. URL <http://arxiv.org/abs/2405.04714>. arXiv:2405.04714 [cs].
- [22] G. Williams, P. Drews, B. Goldfain, J. M. Rehg, and E. A. Theodorou. Aggressive driving with model predictive path integral control. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1433–1440, May 2016. doi:10.1109/ICRA.2016.7487277. URL https://ieeexplore.ieee.org/abstract/document/7487277?casa_token=8gyB7V08MKcAAAAA:Nx9HyqUqpCdvCKEY2xpr8U8Wp5GufuxvfB0SCJRFw0vY9sIIhQ8zajNwtMUKf0E6_fyCjA_XSQ.
- [23] G. Williams, N. Wagener, B. Goldfain, P. Drews, J. M. Rehg, B. Boots, and E. A. Theodorou. Information theoretic MPC for model-based reinforcement learning. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1714–1721, May 2017. doi:10.1109/ICRA.2017.7989202. URL <https://ieeexplore.ieee.org/document/7989202/?arnumber=7989202>.
- [24] T. Han, A. Liu, A. Li, A. Spitzer, G. Shi, and B. Boots. Model Predictive Control for Aggressive Driving Over Uneven Terrain.
- [25] M. Sivaprakasam, P. Maheshwari, M. G. Castro, S. Triest, M. Nye, S. Willits, A. Saba, W. Wang, and S. Scherer. TartanDrive 2.0: More Modalities and Better Infrastructure to Further Self-Supervised Learning Research in Off-Road Driving Tasks, Feb. 2024. URL <https://arxiv.org/abs/2402.01913v1>.
- [26] W. Xiao, T. He, J. Dolan, and G. Shi. Safe Deep Policy Adaptation. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 17286–17292, May 2024. doi:10.1109/ICRA57147.2024.10611340. URL <https://ieeexplore.ieee.org/abstract/document/10611340>.
- [27] D. Kalaria, Q. Lin, and J. M. Dolan. Adaptive Planning and Control with Time-Varying Tire Models for Autonomous Racing Using Extreme Learning Machine. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 10443–10449, May 2024. doi:10.1109/ICRA57147.2024.10610848. URL <https://ieeexplore.ieee.org/abstract/document/10610848>.
- [28] N. Geneva and N. Zabararas. Transformers for modeling physical systems. *Neural Networks*, 146:272–289, Feb. 2022. ISSN 0893-6080. doi:10.1016/j.neunet.2021.11.022. URL <https://www.sciencedirect.com/science/article/pii/S0893608021004500>.
- [29] T. M. Moerland, J. Broekens, A. Plaat, C. M. Jonker, et al. Model-based reinforcement learning: A survey. *Foundations and Trends® in Machine Learning*, 16(1):1–118, 2023.
- [30] M. O’Connell, G. Shi, X. Shi, K. Azzizadenesheli, A. Anandkumar, Y. Yue, and S.-J. Chung. Neural-Fly enables rapid learning for agile flight in strong winds. *Science Robotics*, 7(66):eabm6597, May 2022. doi:10.1126/scirobotics.abm6597. URL <https://www.science.org/doi/10.1126/scirobotics.abm6597>. Publisher: American Association for the Advancement of Science.

- [31] G. Shi, X. Shi, M. O’Connell, R. Yu, K. Azizzadenesheli, A. Anandkumar, Y. Yue, and S.-J. Chung. Neural lander: Stable drone landing control using learned dynamics. In *2019 international conference on robotics and automation (icra)*, pages 9784–9790. IEEE, 2019.
- [32] G. Shi, W. Hönig, X. Shi, Y. Yue, and S.-J. Chung. Neural-swarm2: Planning and control of heterogeneous multirotor swarms using learned interactions. *IEEE Transactions on Robotics*, 38(2):1063–1079, 2021.
- [33] D. Nguyen-Tuong and J. Peters. Model learning for robot control: a survey. *Cognitive processing*, 12: 319–340, 2011.
- [34] A. Kumar, Z. Fu, D. Pathak, and J. Malik. RMA: Rapid Motor Adaptation for Legged Robots, July 2021. URL <http://arxiv.org/abs/2107.04034>. arXiv:2107.04034 [cs].
- [35] U. Rosolia and F. Borrelli. Learning Model Predictive Control for Iterative Tasks. A Data-Driven Control Framework. *IEEE Transactions on Automatic Control*, 63(7):1883–1896, July 2018. ISSN 1558-2523. doi:10.1109/TAC.2017.2753460. URL <https://ieeexplore.ieee.org/abstract/document/8039204>. Conference Name: IEEE Transactions on Automatic Control.
- [36] R. Doshi, H. Walke, O. Mees, S. Dasari, and S. Levine. Scaling Cross-Embodied Learning: One Policy for Manipulation, Navigation, Locomotion and Aviation, Aug. 2024. URL <http://arxiv.org/abs/2408.11812>. arXiv:2408.11812 [cs].
- [37] J. Yang, C. Glossop, A. Bhorkar, D. Shah, Q. Vuong, C. Finn, D. Sadigh, and S. Levine. Pushing the Limits of Cross-Embodiment Learning for Manipulation and Navigation, Feb. 2024. URL <http://arxiv.org/abs/2402.19432>. arXiv:2402.19432 [cs].
- [38] A. O’Neill, A. Rehman, A. Maddukuri, A. Gupta, A. Padalkar, A. Lee, A. Pooley, A. Gupta, A. Mandlekar, A. Jain, A. Tung, A. Bewley, A. Herzog, A. Irpan, A. Khazatsky, A. Rai, A. Gupta, A. Wang, A. Singh, A. Garg, A. Kembhavi, A. Xie, A. Brohan, A. Raffin, A. Sharma, A. Yavary, A. Jain, A. Balakrishna, A. Wahid, B. Burgess-Limerick, B. Kim, B. Schölkopf, B. Wulfe, B. Ichter, C. Lu, C. Xu, C. Le, C. Finn, C. Wang, C. Xu, C. Chi, C. Huang, C. Chan, C. Agia, C. Pan, C. Fu, C. Devin, D. Xu, D. Morton, D. Driess, D. Chen, D. Pathak, D. Shah, D. Büchler, D. Jayaraman, D. Kalashnikov, D. Sadigh, E. Johns, E. Foster, F. Liu, F. Ceola, F. Xia, F. Zhao, F. Stulp, G. Zhou, G. S. Sukhatme, G. Salhotra, G. Yan, G. Feng, G. Schiavi, G. Berseth, G. Kahn, G. Wang, H. Su, H.-S. Fang, H. Shi, H. Bao, H. Ben Amor, H. I. Christensen, H. Furuta, H. Walke, H. Fang, H. Ha, I. Mordatch, I. Radosavovic, I. Leal, J. Liang, J. Abou-Chakra, J. Kim, J. Drake, J. Peters, J. Schneider, J. Hsu, J. Bohg, J. Bingham, J. Wu, J. Gao, J. Hu, J. Wu, J. Wu, J. Sun, J. Luo, J. Gu, J. Tan, J. Oh, J. Wu, J. Lu, J. Yang, J. Malik, J. Silvério, J. Hejna, J. Booher, J. Tompson, J. Yang, J. Salvador, J. J. Lim, J. Han, K. Wang, K. Rao, K. Pertsch, K. Hausman, K. Go, K. Gopalakrishnan, K. Goldberg, K. Byrne, K. Oslund, K. Kawaharazuka, K. Black, K. Lin, K. Zhang, K. Ehsani, K. Lekkala, K. Ellis, K. Rana, K. Srinivasan, K. Fang, K. P. Singh, K.-H. Zeng, K. Hatch, K. Hsu, L. Itti, L. Y. Chen, L. Pinto, L. Fei-Fei, L. Tan, L. J. Fan, L. Ott, L. Lee, L. Weihs, M. Chen, M. Lepert, M. Memmel, M. Tomizuka, M. Itkina, M. G. Castro, M. Spero, M. Du, M. Ahn, M. C. Yip, M. Zhang, M. Ding, M. Heo, M. K. Srirama, M. Sharma, M. J. Kim, N. Kanazawa, N. Hansen, N. Heess, N. J. Joshi, N. Suenderhauf, N. Liu, N. Di Palo, N. M. M. Shafiullah, O. Mees, O. Kroemer, O. Bastani, P. R. Sanketi, P. T. Miller, P. Yin, P. Wohlhart, P. Xu, P. D. Fagan, P. Mitrano, P. Sermanet, P. Abbeel, P. Sundareshan, Q. Chen, Q. Vuong, R. Rafailov, R. Tian, R. Doshi, R. Martín-Martín, R. Bajjal, R. Scalise, R. Hendrix, R. Lin, R. Qian, R. Zhang, R. Mendonca, R. Shah, R. Hoque, R. Julian, S. Bustamante, S. Kirmani, S. Levine, S. Lin, S. Moore, S. Bahl, S. Dass, S. Sonawani, S. Song, S. Xu, S. Haldar, S. Karamcheti, S. Adebola, S. Guist, S. Nasiriany, S. Schaal, S. Welker, S. Tian, S. Ramamoorthy, S. Dasari, S. Belkhale, S. Park, S. Nair, S. Mirchandani, T. Osa, T. Gupta, T. Harada, T. Matsushima, T. Xiao, T. Kollar, T. Yu, T. Ding, T. Davchev, T. Z. Zhao, T. Armstrong, T. Darrell, T. Chung, V. Jain, V. Vanhoucke, W. Zhan, W. Zhou, W. Burgard, X. Chen, X. Wang, X. Zhu, X. Geng, X. Liu, X. Liangwei, X. Li, Y. Lu, Y. J. Ma, Y. Kim, Y. Chebotar, Y. Zhou, Y. Zhu, Y. Wu, Y. Xu, Y. Wang, Y. Bisk, Y. Cho, Y. Lee, Y. Cui, Y. Cao, Y.-H. Wu, Y. Tang, Y. Zhu, Y. Zhang, Y. Jiang, Y. Li, Y. Li, Y. Iwasawa, Y. Matsuo, Z. Ma, Z. Xu, Z. J. Cui, Z. Zhang, and Z. Lin. Open X-Embodiment: Robotic Learning Datasets and RT-X Models : Open X-Embodiment Collaboration0. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6892–6903, May 2024. doi:10.1109/ICRA57147.2024.10611477. URL <https://ieeexplore.ieee.org/document/10611477>.
- [39] O. M. Team, D. Ghosh, H. Walke, K. Pertsch, K. Black, O. Mees, S. Dasari, J. Hejna, T. Kreiman, C. Xu, J. Luo, Y. L. Tan, L. Y. Chen, P. Sanketi, Q. Vuong, T. Xiao, D. Sadigh, C. Finn, and S. Levine. Octo: An Open-Source Generalist Robot Policy, May 2024. URL <http://arxiv.org/abs/2405.12213>. arXiv:2405.12213 [cs].
- [40] S. Macenski and I. Jambrecic. SLAM Toolbox: SLAM for the dynamic world. *Journal of Open Source Software*, 6(61):2783, May 2021. ISSN 2475-9066. doi:10.21105/joss.02783. URL <https://joss.theoj.org/papers/10.21105/joss.02783>.

- [41] S. Macenski, F. Martín, R. White, and J. G. Clavero. The Marathon 2: A Navigation System. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2718–2725, Oct. 2020. doi:10.1109/IROS45743.2020.9341207. URL <http://arxiv.org/abs/2003.00368>. arXiv:2003.00368 [cs].
- [42] D. Fox. KLD-Sampling: Adaptive Particle Filters. In *Advances in Neural Information Processing Systems*, volume 14. MIT Press, 2001. URL <https://proceedings.neurips.cc/paper/2001/hash/c5b2cebf15b205503560c4e8e6d1ea78-Abstract.html>.
- [43] {Stereolabs}. stereolabs/zed-ros2-wrapper, Sept. 2024. URL <https://github.com/stereolabs/zed-ros2-wrapper>. original-date: 2019-04-12T14:23:46Z.
- [44] A. Remonda, N. Hansen, A. Raji, N. Musiu, M. Bertogna, E. Veas, and X. Wang. A Simulation Benchmark for Autonomous Racing with Large-Scale Human Data, July 2024. URL <http://arxiv.org/abs/2407.16680>. arXiv:2407.16680 [cs].
- [45] T. Salzmann, E. Kaufmann, J. Arrizabalaga, M. Pavone, D. Scaramuzza, and M. Ryll. Real-Time Neural MPC: Deep Learning Model Predictive Control for Quadrotors and Agile Robotic Platforms. *IEEE Robotics and Automation Letters*, 8(4):2397–2404, Apr. 2023. ISSN 2377-3766. doi:10.1109/LRA.2023.3246839. URL <https://ieeexplore.ieee.org/document/10049101>. Conference Name: IEEE Robotics and Automation Letters.
- [46] Z. Yi, C. Pan, G. He, G. Qu, and G. Shi. CoVO-MPC: Theoretical Analysis of Sampling-based MPC and Optimal Covariance Design, Jan. 2024. URL <http://arxiv.org/abs/2401.07369>. arXiv:2401.07369 [cs].
- [47] T. Howell, N. Gileadi, S. Tunyasuvunakool, K. Zakka, T. Erez, and Y. Tassa. Predictive Sampling: Real-time Behaviour Synthesis with MuJoCo, Dec. 2022. URL <http://arxiv.org/abs/2212.00541>. arXiv:2212.00541 [cs, eess].
- [48] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. Vander{P}las, S. Wanderman-{M}ilne, and Q. Zhang. google/jax, Sept. 2024. URL <https://github.com/google/jax>. original-date: 2018-10-25T21:25:02Z.
- [49] T. Moore and D. Stouch. A Generalized Extended Kalman Filter Implementation for the Robot Operating System. In E. Menegatti, N. Michael, K. Berns, and H. Yamaguchi, editors, *Intelligent Autonomous Systems 13*, pages 335–348, Cham, 2016. Springer International Publishing. ISBN 978-3-319-08338-4. doi:10.1007/978-3-319-08338-4_25.

Appendix

.1 Ablation Study on Model Design

We conducted experiments on datasets of varying sizes, encompassing trajectories of different cars running in different terrains with timesteps ranging from 1 million (1M) to 100 million (100M). Various model architectures, including transformer, LSTM, GRU, CNN, and MLP, were evaluated, and the normalized prediction error was calculated for each. To ensure the models could be used with MPPI for 50Hz control, we limited the maximum number of parameters for each model to 200K. We create a testing dataset with 1M data points sampled i.i.d. from the simulation, independent of the training dataset, and evaluate all trained models on the same testing dataset. The results shown in Figure 6 demonstrate that as the training dataset size increased from 10M to 100M timesteps, the prediction errors decreased significantly, with the transformer model performing best at the 100M scale. This highlights that the transformer structure is the most effective for modeling diverse car dynamics and environments compared with baseline models.

However, scaling the data and using the appropriate model structure alone are insufficient. Under the optimal data scale and model configuration, we found it crucial to apply the proposed robust training methods (including attack, noise, and mask-out strategies), as discussed in Section 4.2. We systematically evaluate all combinations of these components, resulting in a total of 8 pre-trained models. Each model is fine-tuned using the same dataset and evaluated on real trajectory. The results, shown in Figure 5(a), demonstrate that the model achieves the highest prediction accuracy and stability only when all robust training components are activated. For instance, Figure 5(b)

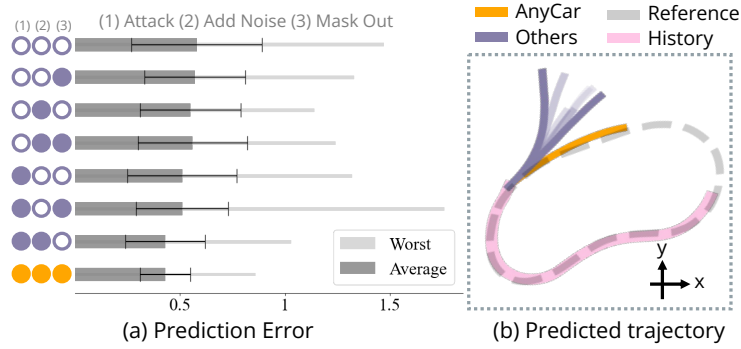


Figure 5: Comparison of AnyCar robust training methods. (a) Evaluate prediction error in real-world trajectories. (b) Demonstrate predicted trajectory of different methods.




Figure 6: Comparison of different model structures and data scales. The reported testing error is normalized using the mean and standard deviation of the evaluation dataset.

shows that without full robust training, the transformer’s predictions are vulnerable to noisy state estimation, leading to significant errors. The combination of model selection, large-scale data, and robust training explains why the AnyCar performs better than the baseline models.