

CAR: Cross-Vehicle Kinodynamics Adaptation via Mobility Representation

Tong Xu*, Chenhui Pan*, and Xuesu Xiao

Abstract—Developing autonomous off-road mobility typically requires either extensive, platform-specific data collection or relies on simplified abstractions, such as unicycle or bicycle models, that fail to capture the complex kinodynamics of diverse platforms, ranging from wheeled to tracked vehicles. This limitation hinders scalability across evolving heterogeneous autonomous robot fleets. To address this challenge, we propose Cross-vehicle kinodynamics Adaptation via mobility Representation (CAR), a novel framework that enables rapid mobility transfer to new vehicles. CAR employs a Transformer encoder with Adaptive Layer Normalization to embed vehicle trajectory transitions and physical configurations into a shared mobility latent space. By identifying and extracting commonality from nearest neighbors within this latent space, our approach enables rapid kinodynamics adaptation to novel platforms with minimal data collection and computational overhead. We evaluate CAR using the Verti-Bench simulator, built on the Chrono multi-physics engine, and validate its performance on four distinct physical configurations of the Verti-4-Wheeler platform. With only one minute of new trajectory data, CAR achieves up to 67.2% reduction in prediction error compared to direct neighbor transfer across diverse unseen vehicle configurations, demonstrating the effectiveness of cross-vehicle mobility knowledge transfer in both simulated and real-world environments.

I. INTRODUCTION

Autonomous off-road mobility across heterogeneous vehicle platforms remains a fundamental challenge in field robotics [1], [2]. Modern ground robot fleets encompass vehicles with diverse physical configurations, including variations in suspension stiffness, drivetrain architecture, payload distribution, and actuation type, ranging from wheeled to tracked systems. These structural differences lead to distinct kinodynamics behaviors, even under similar terrain conditions. To address such diversity, prior works have proposed platform-specific mobility solutions through physics-based modeling [3], [4], learning-based kinodynamics models [5]–[8], and end-to-end policies [9], [10]. While these approaches have demonstrated promising results on individual platforms, they are typically designed and validated for a single vehicle type. As a result, they lack a holistic framework that enables different ground vehicles to share and transfer common off-road mobility knowledge across a heterogeneous robot fleet.

In practice, when a new vehicle is introduced or an existing platform is modified, conventional approaches require collecting large-scale, platform-specific datasets and retraining mobility models from scratch [11]. This repetitive engineering process is both time-intensive and computationally expensive, fundamentally constraining the scalability

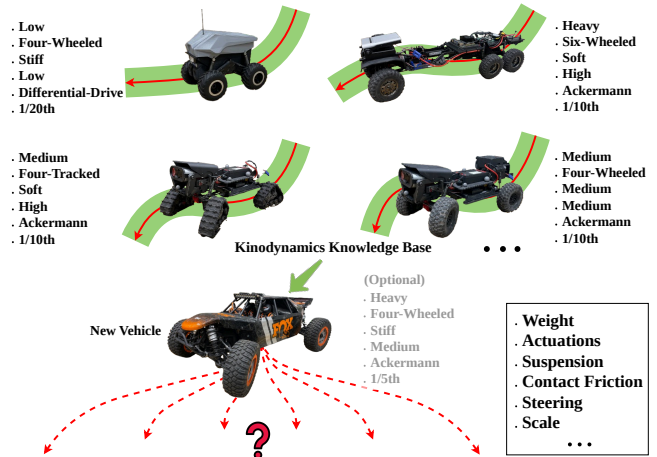


Fig. 1: A heterogeneous fleet of vehicles with diverse physical configurations form a kinodynamics knowledge base. When a new vehicle is introduced, shared mobility representations in this knowledge base enable rapid adaptation with minimal data, without designing or retraining from scratch.

of fleet deployments. Therefore, a scalable framework is required to enable knowledge sharing across the fleet while preserving the flexibility to rapidly adapt to platform-specific kinodynamics.

To overcome this scalability limitation and enable adaptation across heterogeneous vehicle platforms, self-supervised representation learning offers a promising direction by uncovering shared structure in mobility data without explicit supervision. In particular, contrastive learning methods have demonstrated strong capability in constructing structured latent space, where semantically similar samples are encouraged to cluster while dissimilar ones are separated [12], [13]. Such representations provide transferable priors that generalize across diverse tasks and domains. In robotics, learned representations have been leveraged for task transfer [14], [15], sim-to-real adaptation [16], [17], and multi-task or multi-skill policy learning [18], [19]. However, learning cross-vehicle kinodynamics representations that effectively incorporate various physical configurations to adapt to new vehicles with minimal data remains an open challenge.

Motivated by such a challenge, we present Cross-vehicle kinodynamics Adaptation via mobility Representation (CAR), a novel framework that enables rapid kinodynamics adaptation to new vehicle platforms through a shared mobility knowledge base across a heterogeneous fleet, as shown in Fig. 1. Our contributions are summarized as follows:

- A cross-vehicle mobility representation framework that

All authors are with the Department of Computer Science, George Mason University {txu25, cpan7, xiao}@gmu.edu

*Equally contributing authors

employs a Transformer encoder with Adaptive Layer Normalization (AdaLN) [20] to embed trajectory transitions and physical configurations into a structured, shared latent space;

- An efficient selection strategy that identifies relevant mobility neighbors within the shared latent space to extract transferable kinodynamics knowledge in order to minimize data overhead on new vehicles;
- A rapid adaptation mechanism that integrates weighted dataset aggregation, weighted loss optimization, and regulated gradient update to transfer mobility knowledge to new vehicles using only one minute of new trajectory data; and
- Empirical validation of our method in simulation [21] and on four distinct physical configurations of the Verti-4-Wheeler platform [22], demonstrating superior performance over state-of-the-art baselines.

II. RELATED WORK

In this section, we review related work on kinodynamics modeling, self-supervised representation learning, and cross-vehicle knowledge transfer for rapid adaptation.

A. Kinodynamics Modeling

Kinodynamics modeling seeks to capture the coupled relationship among various vehicular components to enable accurate motion prediction and planning. Traditional approaches relied on physics-based models that explicitly parameterize vehicle properties such as size, weight, friction, actuation, suspension, and contact forces [23]. While physically interpretable, these models require careful parameter tuning and often struggle to generalize across different platforms.

More recently, learning-based approaches have emerged as a compelling alternative, leveraging data-driven methods to approximate kinodynamics behavior from onboard sensory observations [8], [24], [25]. End-to-end policy learning methods further bypass explicit kinodynamics modeling by directly learning control policies from demonstration data [26]. Despite these advances, existing approaches remain platform-specific and do not generalize across varying vehicle configurations. The challenge of learning shared mobility representations that transfer across heterogeneous platforms remains unaddressed.

B. Representation Learning in Robotics

Self-supervised representation learning has emerged as a powerful paradigm for extracting structured and transferable representations from high-dimensional robotic data. Contrastive learning methods [12], [13] construct structured latent space by aligning representations of positive augmented views while separating dissimilar samples. Beyond contrastive approaches, cross-embodiment learning has explored the transfer of representations and policies across robots with diverse physical configurations [27], [28]. These approaches typically learn shared representations that abstract embodiment-specific details while preserving task-relevant structure, enabling policy reuse across platforms with varying

kinematics or actuation layouts. However, existing cross-embodiment approaches primarily focus on policy-level behavior learning despite embodiment differences, rather than explicitly modeling distinct kinodynamics grounded in diverse physical configurations. In other words, they *abstract away* embodiment differences, rather than *adapt to* them.

C. Rapid Adaptation and Knowledge Transfer

Few-shot adaptation aims to rapidly generalize a learned model to new tasks or domains using minimal data. Foundational approaches, such as Model-Agnostic Meta-Learning (MAML) [14] and its first-order extensions [29]–[31], achieve this by optimizing parameter initializations that can be efficiently fine-tuned with only a few gradient updates. However, they are usually incapable of knowing when adaptation is not possible due to inimitable differences of a new task. Furthermore, despite their applicability to a wide range of tasks, these purely data-driven methods do not leverage problem-specific knowledge, e.g., vehicle configurations, to efficiently facilitate adaptation. In extremely data-scarce scenarios where even few-shot fine-tuning is not possible, such adaptation efficiency is paramount. One avenue toward adaptation efficiency is to only identify and leverage prior knowledge closely relevant to a new task, while effectively mitigating its difference with existing ones.

Therefore, CAR moves beyond the limitations of few-shot adaptation by explicitly leveraging kinodynamics-specific vehicle configurations and trajectories to identify the most relevant mobility neighbors within a structured latent space, excluding irrelevant vehicles to prevent unnecessary and detrimental adaptation. By combining weighted dataset aggregation and loss optimization with gradient updates regulated by extremely scarce new vehicle data, our framework enables rapid kinodynamics adaptation to new vehicle configurations using only one minute of new trajectory data.

III. METHOD

Our objective is to enable rapid kinodynamics adaptation to new vehicle platforms with minimal new data. We decompose our method into three phases: (1) learning a shared mobility latent space that encodes physical configurations and vehicle trajectories, (2) identifying the most relevant mobility neighbors within this space, and (3) leveraging prior knowledge to enable rapid kinodynamics adaptation. An overview of the framework is illustrated in Fig. 2.

A. Problem Formulation

We define vehicle kinodynamics as a six Degrees of Freedom (DoFs) problem over the state space $S \subset \mathbb{SE}(3)$. The vehicle state at time t is defined as:

$$\mathbf{s}_t = [x_t, y_t, z_t, \phi_t, \varphi_t, \eta_t] \in \mathbb{SE}(3),$$

where x_t, y_t, z_t represent the vehicle position in 3D, and $\phi_t, \varphi_t, \eta_t$ denote roll, pitch, and yaw. Higher order derivatives can be included in \mathbf{s}_t to capture more dynamic maneuvers

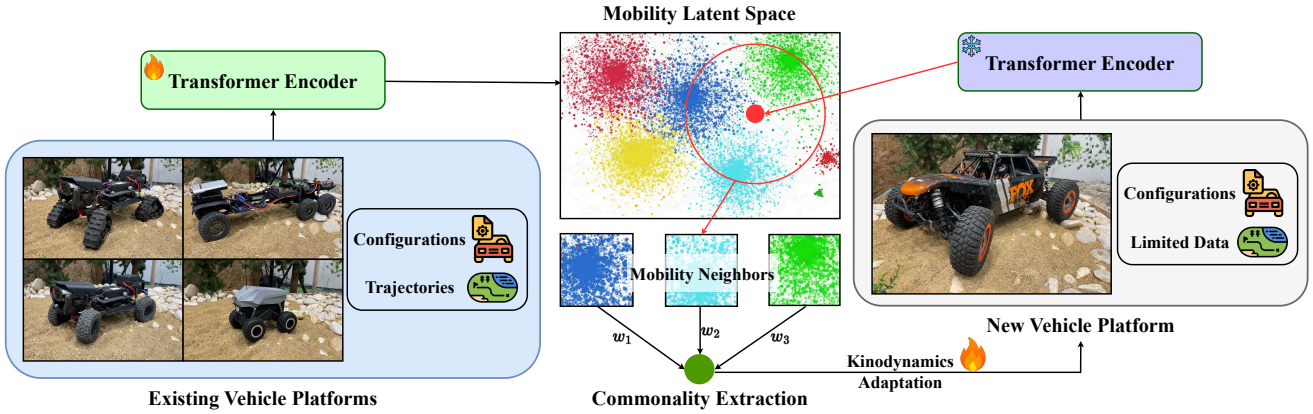


Fig. 2: CAR Overview: Physical configurations and vehicle trajectories from a diverse fleet (left) are embedded into a shared mobility latent space (middle). When a new vehicle is introduced (right), its representation identifies nearest mobility neighbors in this mobility latent space. Inversely proportional to the distances to the mobility neighbors, weights w_1, w_2, w_3 are used to aggregate datasets and bias training loss. Gradient updates regulated by minimal new vehicle data then enables rapid kinodynamics adaptation.

(see experiments for details). The control input $\mathbf{u}_t \in \mathcal{U}$ corresponds to steering and speed. The forward kinodynamics model of a single vehicle is expressed as:

$$\mathbf{s}_{t+1} = f(\mathbf{s}_t, \mathbf{u}_t).$$

Traditional kinodynamics models are typically trained independently for each vehicle platform and do not explicitly account for variations in physical configurations, limiting their transferability within a heterogeneous fleet. We represent the physical configuration of a vehicle v as a parameter descriptor $\mathbf{c} \in \mathbb{R}^{d_c}$, where d_c denotes the number of configurable physical parameters. This descriptor encodes platform-specific properties that govern kinodynamics behavior, such as chassis mass, suspension stiffness, tire friction, actuation layout, and geometric scale.

To enable cross-vehicle knowledge sharing, we assume access to an existing heterogeneous fleet $\mathcal{V}_{\text{train}} = \{v_1, v_2, \dots, v_k\}$. For each vehicle $v_i \in \mathcal{V}_{\text{train}}$, characterized by a distinct physical configuration \mathbf{c}_i , a trajectory dataset $\mathcal{D}_i = \{\tau_1, \tau_2, \dots, \tau_M\}$ is available as prior knowledge, e.g., from previous design effort or deployment experiences. Each trajectory $\tau_j = \{(\mathbf{s}_t, \mathbf{u}_t, \mathbf{s}_{t+1})\}_{t=1}^H$ consists of H discrete kinodynamics transition steps.

Given a new vehicle v_{new} , the problem becomes to derive the forward kinodynamics model f_θ , parameterized by θ , with as little new data as possible, i.e., $\mathcal{D}_{\text{new}} = \{\tau_1, \tau_2, \dots, \tau_{M_{\text{new}}}\}$, $M_{\text{new}} \ll M$, with or without its physical configuration \mathbf{c}_{new} .

B. Cross-Vehicle Mobility Representation

Given trajectory transitions and physical configurations from a heterogeneous training fleet, we aim to learn an encoder that maps each vehicle’s mobility behavior into a shared latent space. This latent space should (i) preserve kinodynamics (dis)similarity across vehicles and (ii) incorporate physically meaningful priors from vehicle configurations to support downstream mobility neighbor identification and rapid kinodynamics adaptation.

1) *Continuous Physical Configuration Embeddings*: To incorporate platform-specific physical priors, we map each configuration descriptor $\mathbf{c} \in \mathbb{R}^{d_c}$ to a continuous embedding via a learnable function $g_\xi(\cdot)$:

$$\mathbf{e}_c = g_\xi(\mathbf{c}) \in \mathbb{R}^d,$$

where d is the embedding dimension. This yields a compact conditioning vector that captures continuous variation in physical configurations.

2) *Transformer Encoder with AdaLN Conditioning*: We employ a Transformer encoder e_ψ to embed trajectories into a shared mobility latent space. Each transition token $\mathbf{x}_t = [\mathbf{s}_t, \mathbf{u}_t, \mathbf{s}_{t+1}]$ is first projected to dimension d , prepended with a learnable CLS token \mathbf{h}_{cls} , and augmented with sinusoidal positional embeddings PE:

$$\mathbf{H}^{(0)} = [\mathbf{h}_{\text{cls}}, \text{Proj}(\mathbf{x}_{t_0}), \dots, \text{Proj}(\mathbf{x}_{t_0+L-1})] + \text{PE}.$$

The sequence is processed through N Transformer blocks. To inject physical configuration priors without overwhelming trajectory encoding, AdaLN [20] conditioning is applied only at the final block. The configuration embedding \mathbf{e}_c produces per-channel affine modulation parameters via a linear projection Φ :

$$[\Delta\beta_1, \Delta\gamma_1, \Delta\beta_2, \Delta\gamma_2] = \Phi(\mathbf{e}_c),$$

which modulate the attention and feed-forward sublayers as:

$$\tilde{\mathbf{H}} = \mathbf{H} \odot (1 + \Delta\gamma_1) + \Delta\beta_1,$$

$$\mathbf{H} \leftarrow \text{LN}(\mathbf{H} + \text{Attn}(\tilde{\mathbf{H}})),$$

$$\tilde{\mathbf{H}} = \mathbf{H} \odot (1 + \Delta\gamma_2) + \Delta\beta_2,$$

$$\mathbf{H} \leftarrow \text{LN}(\mathbf{H} + \text{FFN}(\tilde{\mathbf{H}})),$$

where \odot denotes element-wise multiplication, $\text{LN}(\cdot)$ is Layer Normalization, $\text{Attn}(\cdot)$ is multi-head self-attention, and $\text{FFN}(\cdot)$ is a position-wise feed-forward network. The final mobility embedding is extracted from the CLS output:

$$\mathbf{z} = e_\psi(\tau, \mathbf{c}) = \mathbf{H}_{\text{cls}}^{(N)} \in \mathbb{R}^d.$$

3) *Dual-Path Triplet Training*: To structure the latent space such that kinodynamically similar vehicles cluster together while dissimilar ones are separated, we train the encoder with a dual-path triplet loss. For each anchor trajectory, we sample a positive trajectory from the same vehicle configuration and a negative trajectory from a different one, forming triplets (τ_a, τ_p, τ_n) . The training objective consists of two complementary losses:

(i) **Unconditional Mobility Loss**. We encode trajectories without configuration conditioning to emphasize trajectory-level kinodynamics:

$$\mathbf{z}^u = e_\psi(\tau, \emptyset).$$

We apply a triplet margin loss:

$$\mathcal{L}_u = \max\left(\|\mathbf{z}_a^u - \mathbf{z}_p^u\|_2 - \|\mathbf{z}_a^u - \mathbf{z}_n^u\|_2 + \delta, 0\right),$$

where δ is a predefined threshold.

(ii) **Conditional Mobility Loss**. We additionally encode trajectories with configuration conditioning to encourage physically grounded clustering:

$$\mathbf{z}^c = e_\psi(\tau, \mathbf{c}),$$

and apply the same triplet loss:

$$\mathcal{L}_c = \max\left(\|\mathbf{z}_a^c - \mathbf{z}_p^c\|_2 - \|\mathbf{z}_a^c - \mathbf{z}_n^c\|_2 + \delta, 0\right).$$

The final training objective is $\mathcal{L} = \mathcal{L}_u + \mathcal{L}_c$. This dual-path objective jointly captures intrinsic mobility patterns of the trajectories and the influence of physical vehicle configurations, producing a structured latent space that supports reliable cross-vehicle knowledge transfer.

C. Mobility Neighbor Identification

Given the trained Transformer encoder e_ψ , we project the new vehicle v_{new} into the shared mobility latent space and identify its relevant mobility neighbors from $\mathcal{V}_{\text{train}}$.

Centroid Computation. For each vehicle $v_i \in \mathcal{V}_{\text{train}}$, we encode a set of trajectories using the conditional encoder and compute a centroid $\boldsymbol{\mu}_i$ in the latent space:

$$\boldsymbol{\mu}_i = \frac{1}{|\mathcal{Z}_i|} \sum_{\mathbf{z} \in \mathcal{Z}_i} \mathbf{z}, \quad \mathcal{Z}_i = \{e_\psi(\tau_j, \mathbf{c}_i) \mid \tau_j \in \mathcal{D}_i\}.$$

To reduce noise and preserve the most informative latent dimensions, we apply Principle Component Analysis (PCA) and retain the minimum number of principal components that explain at least 90% of the total variance. Similarly, the new vehicle centroid $\boldsymbol{\mu}_{\text{new}}$ is computed from its limited dataset \mathcal{D}_{new} using the same PCA projection.

Mobility Neighbor Identification. To characterize the inter-cluster structure of the training fleet, we compute the pairwise cosine distances between all training centroids:

$$d_{ij} = d_{\cos}(\boldsymbol{\mu}_i, \boldsymbol{\mu}_j), \forall i < j.$$

We then define an adaptive distance threshold:

$$\epsilon = \bar{d} + 2\sigma_d,$$

where \bar{d} and σ_d denote the mean and standard deviation of the inter-cluster cosine distances, respectively.

Given the centroid $\boldsymbol{\mu}_{\text{new}}$, we evaluate its cosine distance to each training centroid and identify mobility neighbors as:

$$\mathcal{N} = \{v_i \in \mathcal{V}_{\text{train}} \mid d_{\cos}(\boldsymbol{\mu}_{\text{new}}, \boldsymbol{\mu}_i) \leq \epsilon\}.$$

If no training vehicle satisfies this criterion, the new vehicle is classified as out-of-distribution to prevent detrimental adaptation.

Inverse-Distance Weighting. For each identified mobility neighbor $v_i \in \mathcal{N}$, we assign a weight inversely proportional to its squared cosine distance to the new vehicle centroid:

$$\tilde{w}_i = \frac{1}{d_{\cos}(\boldsymbol{\mu}_{\text{new}}, \boldsymbol{\mu}_i)^2}, \quad w_i = \frac{\tilde{w}_i}{\sum_{j=1}^{|\mathcal{N}|} \tilde{w}_j}.$$

To emphasize the most relevant neighbors, we sort neighbors in descending order of their weights and select the top neighbors whose cumulative weight exceeds 0.9:

$$\mathcal{N}_{\text{new}} = \left\{ v_i \mid \sum_{j=1}^i w_j \geq 0.9 \quad \& \quad \sum_{j=1}^{i-1} w_j < 0.9 \right\},$$

The weights corresponding to \mathcal{N}_{new} are then used to guide dataset aggregation and loss optimization during rapid kinodynamics adaptation.

D. Rapid Kinodynamics Adaptation

Given the identified mobility neighbors \mathcal{N}_{new} and their associated weights w_i , CAR computes the kinodynamics model f_θ of the new vehicle through weighted knowledge transfer with gradient regulation.

Weighted Dataset Aggregation. For each neighbor vehicle $v_i \in \mathcal{N}_{\text{new}}$, we sample trajectories from its dataset \mathcal{D}_i , with the number of samples proportional to its weight w_i . This ensures that vehicles with higher mobility similarity contribute more during adaptation. The resulting aggregated training set is

$$\mathcal{D}_{\text{agg}} = \bigcup_{v_i \in \mathcal{N}_{\text{new}}} \mathcal{D}_i^{(w_i)},$$

where $\mathcal{D}_i^{(w_i)}$ denotes the weighted subset sampled from \mathcal{D}_i .

Weighted Loss Optimization. The kinodynamics prediction loss is computed separately on each neighbor dataset and combined as a weighted objective:

$$\mathcal{L}_{\text{train}} = \sum_{v_i \in \mathcal{N}_{\text{new}}} w_i \mathcal{L}(f_\theta; \mathcal{D}_i^{(w_i)}),$$

where \mathcal{L} denotes the Mean Squared Error (MSE) over sampled trajectories. This objective prioritizes model updates from neighbors with the most mobility relevance to the new vehicle.

Gradient Regulation via Limited New Data. To ensure that adaptation remains consistent with the limited trajectory data \mathcal{D}_{new} , we regulate gradient updates using \mathcal{D}_{new} as a constraint. Inspired by Gradient Episodic Memory [32], we compute constraint gradients from \mathcal{D}_{new} and enforce:

$$\nabla_\theta \mathcal{L}(f_\theta; \mathcal{D}_{\text{new}}) \cdot \nabla_\theta \mathcal{L}_{\text{train}} \geq 0.$$

If violations occur, the training gradient is projected onto the nearest gradient direction that does not increase the loss on \mathcal{D}_{new} . The constraint gradients are periodically refreshed as the model evolves.

Collectively, these three components enable CAR to transfer shared knowledge from the most relevant mobility neighbors to the new platform using a small amount of new data.

IV. IMPLEMENTATIONS

In this section, we present implementation details of our approach and experiments.

A. Vehicle Configurations and Datasets

We evaluate CAR using the Verti-Bench simulator [21], where trajectories are generated for a heterogeneous fleet of vehicles. In our experiments, the physical configuration of each vehicle v is represented as: $\mathbf{c} = [\alpha_m, \mu_f, \alpha_s]^\top$, where α_m denotes the chassis mass scaling ratio relative to the nominal mass in the simulator, μ_f is the rigid tire friction coefficient, and α_s represents the suspension spring stiffness scaling ratio relative to the nominal stiffness. The training fleet comprises eight vehicle configurations formed by selecting one out of two values for each parameter: $\alpha_m \in \{0.5, 4.0\}$, $\mu_f \in \{0.6, 0.9\}$, and $\alpha_s \in \{0.6, 1.8\}$.

To collect training trajectories, each vehicle executes sinusoidal random exploration on flat terrain. Steering follows $\mathbf{u}_{\text{steer}}(t) = \sin(\omega_s t)$ with $\omega_s \sim \mathcal{U}(0.1, 0.5)$ Hz, combined with speed $\mathbf{u}_{\text{speed}}(t) = v_c + A \sin(\omega_v t)$, where $\omega_v \sim \mathcal{U}(0.1, 2.5)$ Hz. The speed amplitude A and center velocity v_c are determined such that the minimum speed lies in $\mathcal{U}(3, 4)$ m/s and maximum speed lies in $\mathcal{U}(8, 10)$ m/s.

We adopt a gravity-aligned body frame, in which current state is denoted as $[0, 0, 0, \phi_t, \varphi_t, 0, \dot{\eta}_t, v_t] \in \mathbb{R}^8$, where the position components (x, y, z) and yaw are zeroed (and therefore omitted), while ϕ_t and φ_t retain their world-frame values. To better capture high-speed kinodynamics, we additionally include the yaw rate $\dot{\eta}_t$ and longitudinal velocity v_t . The corresponding next state encodes the relative motion in the current body frame: $\mathbf{s}_{t+1} = [\Delta x, \Delta y, \Delta z, \phi_{t+1}, \varphi_{t+1}, \Delta \eta] \in \mathbb{R}^6$, where $(\Delta x, \Delta y, \Delta z, \Delta \eta)$ denote body-frame pose change, and $(\phi_{t+1}, \varphi_{t+1})$ are the absolute roll and pitch angles at the next timestep.

B. Cross-Vehicle Mobility Representation

The cross-vehicle mobility encoder consists of 4 Transformer blocks, each with embedding dimension $d = 16$ and 4 attention heads. For each trajectory, transition tokens $\mathbf{x}_t = [\mathbf{s}_t, \mathbf{u}_t, \mathbf{s}_{t+1}] \in \mathbb{R}^{14}$ are linearly projected to \mathbb{R}^{16} and prepended with a learnable CLS token augmented by sinusoidal positional embeddings. Vehicle configuration parameters $(\alpha_m, \mu_f, \alpha_s)$ are mapped to continuous embeddings to enable generalization across new configurations. Each parameter is first min-max normalized to $[0, 1]$ using the predefined ranges: $\alpha_m \in [0.5, 4.0]$, $\mu_f \in [0.6, 0.9]$, and $\alpha_s \in [0.6, 1.8]$, and then embedded through a two-layer MLP with $\{3, 8, 16\}$ neurons and Tanh activations into a

continuous \mathbb{R}^{16} embedding space. The trajectory encoder and configuration modulator share the same Transformer backbone, with AdaLN conditioning applied exclusively at the final Transformer block with a scale factor of 0.5.

To enable flexible inference, the configuration modulator adopts a learnable null embedding $\mathbf{e}_\emptyset \in \mathbb{R}^{16}$. During training, configuration embeddings are randomly replaced by \mathbf{e}_\emptyset with a dropout probability of 0.1. This mechanism allows the encoder to operate in both *conditional* mode, incorporating known physical configurations, and *unconditional* mode, relying solely on trajectories when configuration information is unavailable.

The mobility encoder is trained with the dual-path triplet loss with margin $\delta = 4.0$. We use the Adam optimizer with learning rate 1×10^{-4} and batch size 128, training for up to 100K iterations. Early stopping with patience 10 is applied based on conditional embedding separation evaluated every 1,000 iterations on a held-out validation set.

C. Kinodynamics Model

The kinodynamics model f_θ is implemented as a lightweight MLP that predicts the 6-DoF state change from the current state and action. The model processes state and action inputs through separate branches, each implemented as a two-layer MLP with $\{8, 16\}$ hidden neurons, LayerNorm, and Tanh activations. The encoded features are concatenated and passed through two output heads: a translation and yaw head predicting $[\Delta x, \Delta y, \Delta z, \Delta \eta]$, and a roll-pitch head predicting $[\phi_{t+1}, \varphi_{t+1}]$, each consisting of two fully connected layers with dimensions $\{32, 16\}$, BatchNorm, ReLU activations, and dropout.

During rapid kinodynamics adaptation, f_θ is trained using weighted knowledge transfer with gradient regulation described in Section III-D. The training loss is computed via autoregressive rollout over a horizon of $T_{\text{pred}} = 16$ steps. Constraint gradients for gradient regulation are computed from \mathcal{D}_{new} using three sampled trajectories, and refreshed every 100 gradient steps. The kinodynamics model performs 2,000 gradient steps until convergence using the Adam optimizer with learning rate 1×10^{-3} , minimizing the MSE between predicted and ground-truth state changes.

V. EXPERIMENTS

We design our experiments to evaluate CAR from three aspects: (1) cross-vehicle few-shot adaptation via identified mobility neighbors, (2) few-shot generalization compared with Model-Agnostic Meta-Learning (MAML) [14], and (3) ablation studies of both the adaptation mechanisms and mobility encoder design.

A. Cross-Vehicle Few-Shot Adaptation

In this experiment, we evaluate the ability of CAR to identify and leverage the most relevant mobility neighbors within the structured latent space for effective cross-vehicle knowledge transfer. We compare CAR against two baselines:

- **From Scratch:** A model trained on the full set of 400 trajectories collected from the new vehicle, serving as a data-intensive upper-bound; and

- **Mobility Neighbors:** Models trained from the identified mobility neighbors and directly applied to the new vehicle without any adaptation.

TABLE I: Cross-Vehicle Few-Shot Adaptation Performance on New Vehicle Configurations, denoted as $[\alpha_m, \mu_f, \alpha_s]$.

Configuration	Model	MSE \pm Std \downarrow
[3.0, 0.75, 1.6]	CAR (Proposed)	0.632 \pm 0.825
	Neighbor 1	1.261 \pm 1.117
	From Scratch	0.173 \pm 0.251
[0.6, 0.7, 0.8]	CAR (Proposed)	0.734 \pm 0.730
	Neighbor 1	2.241 \pm 1.857
	Neighbor 2	3.305 \pm 2.811
	From Scratch	0.569 \pm 0.775
[0.6, 0.75, 1.2]	CAR (Proposed)	0.696 \pm 0.972
	Neighbor 1	0.894 \pm 0.722
	Neighbor 2	1.938 \pm 3.110
	Neighbor 3	2.495 \pm 2.680
	From Scratch	0.222 \pm 0.241

Only one mobility neighbor is identified for the first new vehicle, which only uses regulated gradients on the only neighbor’s dataset. The second and third new vehicles have two and three neighbors respectively. As shown in Table I, CAR outperforms direct transfer from individual mobility neighbors without adaptation across all new vehicle configurations, achieving up to a 67.2% reduction in prediction error. Although the “From Scratch” upper-bound attains the lowest MSE by leveraging 400 platform-specific trajectories, CAR only uses three trajectories. This result highlights the effectiveness of structured mobility representation in substantially reducing data requirements while maintaining strong adaptation performance.

B. Cross-Vehicle Few-Shot Generalization

We compare CAR against MAML [14] to evaluate cross-vehicle generalization in the few-shot regime. Unlike MAML, which learns a platform-agnostic parameter initialization, CAR explicitly leverages physical vehicle configurations to identify and prioritize the most relevant mobility neighbors within the structured latent space. During adaptation, CAR uses three trajectories from the new vehicle for mobility neighbor and weight identification, as well as gradient regulation, without fine-tuning on platform-specific data.

Table II shows CAR outperforms MAML: Using only three trajectories from the new vehicle, CAR reduces prediction error by 79.1% compared to MAML under the same few-shot setting. Remarkably, even when MAML is fine-tuned on 400 trajectories of the new vehicle, its prediction accuracy remains inferior to CAR adapted with just three trajectories. These results demonstrate that structured mobility representations informed by vehicle configurations enable

TABLE II: Cross-Vehicle Few-Shot Generalization Compared with MAML.

Model	New Vehicle Data	MSE \pm Std \downarrow
CAR (Proposed)	3 Trajectories	0.734 \pm 0.730
MAML	3 Trajectories	3.507 \pm 7.582
	400 Trajectories	1.329 \pm 2.412

more data-efficient and robust cross-vehicle generalization than purely data-driven, initialization-based meta-learning.

C. Ablation Studies

We conduct ablation studies to evaluate the contribution of key components in CAR’s kinodynamics adaptation mechanism and mobility encoder design.

1) *Kinodynamics Adaptation Components:* Each variant removes one component from the full kinodynamics adaptation mechanism:

- **CAR:** The complete adaptation framework incorporating mobility neighbor identification (MN), weighted dataset aggregation (WD), weighted loss optimization (WL), and gradient regulation (GR);
- **CAR w/o MN:** Removes mobility neighbor identification. Instead of selecting the most relevant mobility neighbors in the latent space, all vehicles in the training fleet are treated uniformly during adaptation;
- **CAR w/o WD:** Removes weighted dataset aggregation. All identified mobility neighbor datasets are treated uniformly, eliminating selective emphasis on the most relevant mobility neighbors;
- **CAR w/o WL:** Removes weighted loss optimization. Despite weighted dataset aggregation, their losses are treated equally during optimization without similarity-based weighting; and
- **CAR w/o GR:** Removes gradient regulation via limited new vehicle data. The model is adapted solely using weighted datasets and loss without constraining gradient updates to preserve consistency with the new platform.

TABLE III: Ablation Study for Kinodynamics Adaptation Mechanisms.

Variant	MSE \pm Std \downarrow
CAR (Proposed)	0.734 \pm 0.730
CAR w/o MN	1.674 \pm 2.791
CAR w/o WD	1.729 \pm 2.554
CAR w/o WL	1.034 \pm 1.367
CAR w/o RG	1.173 \pm 1.333

As shown in Table III, each component of the adaptation mechanism contributes to CAR’s performance. Mobility neighbor identification and weighted dataset aggregation have the largest impact: their removal increases prediction

error by 128% and 135% respectively, indicating that effective cross-vehicle adaptation requires both identifying the most relevant mobility neighbors and weighting their contributions proportionally during knowledge transfer. Weighted loss optimization and gradient regulation provide complementary gains, with their removal degrading performance by 41% and 60% respectively, confirming that prioritizing kinodynamically similar sources and constraining parameter updates with limited new vehicle data are both essential for robust adaptation.

2) *Mobility Encoder Architecture*: We compare three encoder variants: our *conditional* encoder, which applies AdaLN to modulate trajectory representations using physical configuration embeddings during training and inference; an *unconditional* encoder, trained with AdaLN modulation but deployed when configuration information is unavailable during inference; and a *simple* encoder that relies solely on trajectories during both training and inference.

TABLE IV: Ablation Study for Encoder Design.

Encoder Variant	MSE \pm Std \downarrow
Conditional (Proposed)	0.734 \pm 0.730
Unconditional (Proposed)	1.043 \pm 1.242
Simple	1.797 \pm 4.033

Table IV validates the architectural design of injecting physical priors via AdaLN conditioning. By explicitly grounding the mobility representation in vehicle configurations, the *conditional* encoder achieves 29.63% reduction in prediction error compared to the *unconditional* counterpart. Furthermore, removing the conditioning architecture entirely from training (*simple*) severely degrades performance.

D. Physical Experiment

For real-world validation, we deploy CAR on different open-source 1/10th-scale Verti-4-Wheeler platforms [22], with state information provided at 100Hz via a motion capture system. The training fleet comprises three distinct configurations: Rear Right Wheel Removed, Four-Wheeled, and Four-Tracked. As the new platform, we introduce a heavy-payload configuration with increased mass and altered inertia distribution, representing an embodiment shift relative to the training fleet.

TABLE V: Physical Experiment Results.

Model	MSE \pm Std \downarrow
CAR (Proposed)	0.144 \pm 0.186
Neighbor 1	0.151 \pm 0.184
Neighbor 2	0.247 \pm 0.188
From Scratch	0.111 \pm 0.195

To construct the kinodynamics knowledge base datasets, each training vehicle is driven using sinusoidal steering and

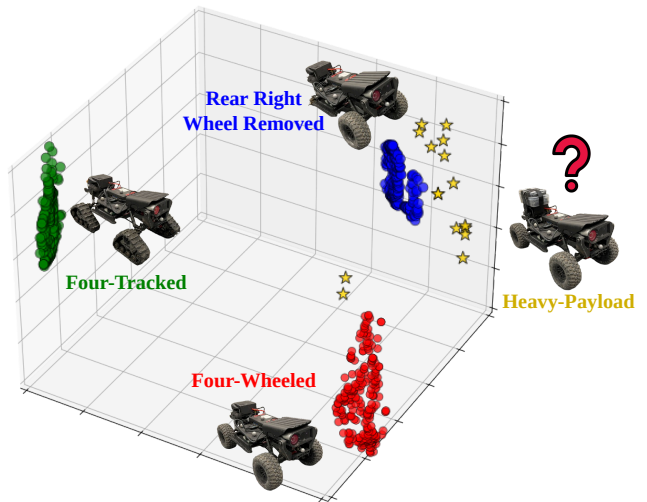


Fig. 3: The physical mobility latent space is learned from three distinct configurations (blue, red, and green). When the new heavy-payload platform (yellow) is introduced, it is projected into this space to identify the most relevant mobility neighbors (blue and red) before rapid adaptation.

velocity commands for 750 seconds per configuration. For rapid adaptation, CAR transfers knowledge to the new heavy-payload platform. The physical experiment results in Table V demonstrate that CAR achieves the lowest prediction error compared with the identified mobility neighbors, reducing prediction error by 4.6% relative to the best neighbor using only 45 s of new trajectory data. Although training from scratch attains a lower error with 750 s of platform-specific data, the small performance gap indicates that CAR recovers most of the heavy-payload kinodynamics with minimal interaction data.

The learned mobility latent space provides an interpretable explanation for the observed transfer behavior. As shown in Fig. 3, the heavy-payload configuration is embedded closer to the Rear Right Wheel Removed and Four-Wheeled configurations than the Four-Tracked platform, indicating an underlying kinodynamics similarity with these two mobility neighbors. This geometric structure is consistent with the performance trends shown in Table V.

VI. CONCLUSIONS AND LIMITATIONS

We present CAR, a novel framework designed to address the scalability challenges of enabling rapid kinodynamics adaptation across heterogeneous fleets. By leveraging a Transformer encoder conditioned with AdaLN, CAR effectively maps vehicle trajectories and physical configurations into a shared, structured mobility latent space. This representation enables reliable identification of the most relevant mobility neighbors for any newly introduced platform. Through weighted dataset aggregation, weighted loss optimization, and gradient regulation, the proposed framework achieves rapid kinodynamics adaptation using limited one minute of new data. Our evaluation demonstrates that CAR reduces

prediction error by up to 67.2% compared to direct neighbor transfer, highlighting its effectiveness for scalable cross-vehicle mobility knowledge transfer.

A key limitation of this work is that it is evaluated only on flat terrain, where kinodynamics variations are primarily driven by physical configurations rather than diverse off-road terrain properties, obstacles, or dynamic environmental conditions. Future work will extend the mobility representation to jointly consider vehicle embodiment together with terrain geometry and semantics, enabling more comprehensive and realistic rapid kinodynamics adaptation.

ACKNOWLEDGMENTS

This work has taken place in the RobotiXX Laboratory at George Mason University. RobotiXX research is supported by National Science Foundation (NSF, 2350352), Army Research Office (ARO, W911NF2320004, W911NF2520011), Google DeepMind (GDM), Clearpath Robotics, FrodoBots Lab, Raytheon Technologies (RTX), Tangenta, Mason Innovation Exchange (MIX), and Walmart.

REFERENCES

- [1] X. Xiao, B. Liu, G. Warnell, and P. Stone, "Motion planning and control for mobile robot navigation using machine learning: a survey," *Autonomous Robots*, vol. 46, no. 5, pp. 569–597, 2022.
- [2] N. Wang, X. Li, K. Zhang, J. Wang, and D. Xie, "A survey on path planning for autonomous ground vehicles in unstructured environments," *Machines*, vol. 12, no. 1, p. 31, 2024.
- [3] X. Cai, J. Queeney, T. Xu, A. Datar, C. Pan, M. Miller, A. Flather, P. R. Osteen, N. Roy, X. Xiao *et al.*, "Pietra: Physics-informed evidential learning for traversing out-of-distribution terrain," *IEEE Robotics and Automation Letters*, vol. 10, no. 3, pp. 2359–2366, 2025.
- [4] Z. Zhao, B. Li, Y. Du, T. Fu, and C. Wang, "Physord: a neuro-symbolic approach for physics-infused motion prediction in off-road driving," in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2024, pp. 11 670–11 677.
- [5] X. Xiao, J. Biswas, and P. Stone, "Learning inverse kinodynamics for accurate high-speed off-road navigation on unstructured terrain," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 6054–6060, 2021.
- [6] A. Datar, C. Pan, M. Nazeri, A. Pokhrel, and X. Xiao, "Terrain-attentive learning for efficient 6-dof kinodynamic modeling on vertically challenging terrain," in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2024, pp. 5438–5443.
- [7] P. Atreya, H. Karnan, K. S. Sikand, X. Xiao, S. Rabiee, and J. Biswas, "High-speed accurate robot control using learned forward kinodynamics and non-linear least squares optimization," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 11 789–11 795.
- [8] A. Pokhrel, M. Nazeri, A. Datar, and X. Xiao, "Cahsor: Competence-aware high-speed off-road ground navigation in se(3)," *IEEE Robotics and Automation Letters*, vol. 9, no. 11, pp. 9653–9660, 2024.
- [9] T. Xu, C. Pan, and X. Xiao, "Vertiselector: Automatic curriculum learning for wheeled mobility on vertically challenging terrain," in *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2025, pp. 11 136–11 143.
- [10] Y. Pan, C.-A. Cheng, K. Saigol, K. Lee, X. Yan, E. A. Theodorou, and B. Boots, "Imitation learning for agile autonomous driving," *The International Journal of Robotics Research*, vol. 39, no. 2-3, pp. 286–302, 2020.
- [11] P. V. Borges, T. Peynot, S. Liang, B. Arain, M. Wildie, M. G. Minareci, S. Lichman, G. Samvedi, I. Sa, N. Hudson *et al.*, "A survey on terrain traversability analysis for autonomous ground vehicles: Methods, sensors, and challenges," *Field Robotics*, vol. 2, pp. 1567–1627, 2022.
- [12] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A simple framework for contrastive learning of visual representations," in *International conference on machine learning*. Pmlr, 2020, pp. 1597–1607.
- [13] K. He, H. Fan, Y. Wu, S. Xie, and R. Girshick, "Momentum contrast for unsupervised visual representation learning," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 9729–9738.
- [14] C. Finn, P. Abbeel, and S. Levine, "Model-agnostic meta-learning for fast adaptation of deep networks," in *International conference on machine learning*. PMLR, 2017, pp. 1126–1135.
- [15] C. Devin, A. Gupta, T. Darrell, P. Abbeel, and S. Levine, "Learning modular neural network policies for multi-task and multi-robot transfer," in *ICRA*, 2017.
- [16] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel, "Domain randomization for transferring deep neural networks from simulation to the real world," in *2017 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE, 2017, pp. 23–30.
- [17] X. B. Peng, M. Andrychowicz, W. Zaremba, and P. Abbeel, "Sim-to-real transfer of robotic control with dynamics randomization," in *2018 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2018, pp. 3803–3810.
- [18] T. Haarnoja, K. Hartikainen, P. Abbeel, and S. Levine, "Latent space policies for hierarchical reinforcement learning," in *International Conference on Machine Learning*. PMLR, 2018, pp. 1851–1860.
- [19] T. Yu, D. Quillen, Z. He, R. Julian, K. Hausman, C. Finn, and S. Levine, "Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning," in *Conference on robot learning*. PMLR, 2020, pp. 1094–1100.
- [20] W. Peebles and S. Xie, "Scalable diffusion models with transformers," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2023, pp. 4195–4205.
- [21] T. Xu, C. Pan, M. B. Rao, A. Datar, A. Pokhrel, Y. Lu, and X. Xiao, "Verti-bench: A general and scalable off-road mobility benchmark for vertically challenging terrain," in *Robotics: Science and Systems (RSS)* 2025, 2025.
- [22] A. Datar, C. Pan, M. Nazeri, and X. Xiao, "Toward wheeled mobility on vertically challenging terrain: Platforms, datasets, and algorithms," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 16 322–16 329.
- [23] C. Pavlov and A. M. Johnson, "Soil displacement terramechanics for wheel-based trenching with a planetary rover," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 4760–4766.
- [24] X. Meng, N. Hatch, S. A. Lambert, A. Li, N. Wagener, M. Schmittle, J. Lee, W. Yuan, Z. Chen, S. Deng *et al.*, "Terrainet: Visual modeling of complex terrain for high-speed, off-road navigation," in *Robotics: Science and Systems (RSS)*, 2023.
- [25] W. Xiao, H. Xue, T. Tao, D. Kalaria, J. M. Dolan, and G. Shi, "Anycar to anywhere: Learning universal dynamics model for agile and adaptive mobility," in *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2025, pp. 8819–8825.
- [26] T. Xu, C. Pan, and X. Xiao, "Reinforcement learning for wheeled mobility on vertically challenging terrain," in *2024 IEEE International Symposium on Safety Security Rescue Robotics (SSRR)*, 2024, pp. 125–130.
- [27] 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," in *Proceedings of Robotics: Science and Systems*, 2024.
- [28] R. Doshi, H. R. Walke, O. Mees, S. Dasari, and S. Levine, "Scaling cross-embodied learning: One policy for manipulation, navigation, locomotion and aviation," in *Conference on Robot Learning*. PMLR, 2025, pp. 496–512.
- [29] A. Nichol, J. Achiam, and J. Schulman, "On first-order meta-learning algorithms," *arXiv preprint arXiv:1803.02999*, 2018.
- [30] A. Fallah, A. Mokhtari, and A. Ozdaglar, "On the convergence theory of gradient-based model-agnostic meta-learning algorithms," in *International conference on artificial intelligence and statistics*. PMLR, 2020, pp. 1082–1092.
- [31] A. Nagabandi, I. Clavera, S. Liu, R. S. Fearing, P. Abbeel, S. Levine, and C. Finn, "Learning to adapt in dynamic, real-world environments through meta-reinforcement learning," in *International Conference on Learning Representations (ICLR)*, 2019.
- [32] B. Liu, X. Xiao, and P. Stone, "A lifelong learning approach to mobile robot navigation," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 1090–1096, 2021.