

# Toward Efficient and Robust Behavior Models for Multi-Agent Driving Simulation

Fabian Konstantinidis, Moritz Sackmann, Ulrich Hofmann, Christoph Stiller

**Abstract**—Scalable multi-agent driving simulation requires behavior models that are both realistic and computationally efficient. We address this by optimizing the behavior model that controls individual traffic participants. To improve efficiency, we adopt an instance-centric scene representation, where each traffic participant and map element is modeled in its own local coordinate frame. This design enables efficient, viewpoint-invariant scene encoding and allows static map tokens to be reused across simulation steps. To model interactions, we employ a query-centric symmetric context encoder with relative positional encodings between local frames. We use Adversarial Inverse Reinforcement Learning to learn the behavior model and propose an adaptive reward transformation that automatically balances robustness and realism during training. Experiments demonstrate that our approach scales efficiently with the number of tokens, significantly reducing training and inference times, while outperforming several agent-centric baselines in terms of positional accuracy and robustness.

## I. INTRODUCTION

Given the high cost and potential risks of deploying automated vehicles in real-world environments, simulation is indispensable for the research and development of advanced driver assistant systems. Key applications include generating diverse scenarios that may be difficult or dangerous to recreate in the real world [1]–[3], predicting human driving behavior [4]–[6], and training or evaluating behavior models [7], [8].

A key requirement for multi-agent driving simulation is exhibiting realistic behavior to reduce the sim-to-real gap. For that, behavior models map observations to actions. They are typically trained via Imitation Learning (IL), e.g., [9], [10], which learns directly from real data but yields less robust policies, or via Reinforcement Learning (RL), e.g., [7], [11]–[13], which improves robustness through exploration in simulation but depends on a reward function, which is unknown for real-world drivers.

Instead, we employ Adversarial Inverse Reinforcement Learning (AIRL) [14] to reconstruct a reward signal from real traffic data and simultaneously learn a behavior model that maximizes this reward signal using standard RL. Since the reconstructed reward signal can differ across experiments, it may unintentionally affect the robustness of the model. To address this, we introduce an adaptive reward transformation that balances robustness and realism during training.

Fabian Konstantinidis is with CARIAD SE, Germany, and also with Karlsruhe Institute of Technology, Germany. Moritz Sackmann is with CARIAD SE, Germany. Christoph Stiller is with Karlsruhe Institute of Technology, Germany. This work is a result of the joint research project STADT:up (19A22006E). The project is supported by the German Federal Ministry for Economic Affairs and Climate Action (BMWK), based on a decision of the German Bundestag. The author is solely responsible for the content of this publication.

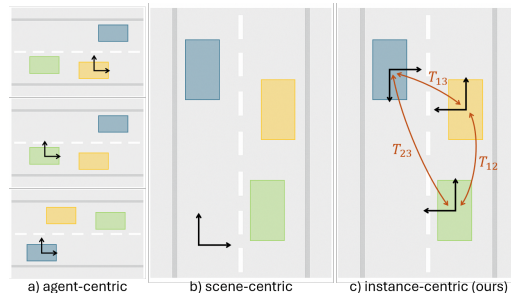


Fig. 1. Illustration of different scene representations. Our instance-centric representation encodes each instance in its own coordinate frame, enabling shared feature extraction.

The second requirement is achieving high execution speed to enable efficient scaling. While prior works, e.g., [7], [15]–[17], have focused on accelerating the simulation backbone, our focus is on optimizing the executed behavior model.

Commonly, agent-centric observations (see Figure 1a) are used, e.g., [6]–[8], [16], where nearby instances, such as agents and static map elements, are represented from each agent’s perspective. However, generating and encoding these local observations for every agent at each simulation step imposes a significant computational burden. Alternatively, scene-centric observations (see Figure 1b) represent all instances in a global coordinate frame, e.g., [18], [19]. While this offers efficient encoding, it comes at the potential sacrifice of pose-invariance, requiring the model to implicitly learn variations in agent positions and orientations.

Instead, to improve the efficiency of the behavior model, we employ instance-centric observations (see Figure 1c), where each instance, i.e., an individual agent or map element, is represented in its own local coordinate frame. Features are encoded using shared context encoders, allowing the encoded instance tokens to be processed across multiple agents. This approach models relationships between tokens in a symmetric manner, independent of any global coordinate frame. As a result, static instances need to be encoded only once, and dynamic instances only once per simulation step, rather than from every agent’s perspective at each step.

**Our Contributions** are four-fold: (1) Our main contribution is the introduction of an instance-centric scene representation for behavior modeling, which enables efficient, viewpoint-invariant encoding of the scene. Instance tokens are shared across all agents. Static instances need to be encoded only once, enabling the reuse of their tokens in subsequent simulation steps. (2) We employ a query-centric symmetric context encoder with relative positional encodings between pairwise instance-centric frames, and evaluate its performance across different model design choices. (3) We propose an adaptive

reward transformation that automatically balances robustness and realism during AIRL training. (4) We demonstrate that our approach yields realistic behavior models, outperforming a diverse set of baselines in robustness and scalability across two automated driving datasets.

## II. RELATED WORK

Achieving both realism and high execution speed is crucial for modern traffic simulations. While high-fidelity behavior modeling ensures the accuracy of the simulated traffic flow, the ability to execute such models efficiently remains equally critical, especially in large-scale or real-time applications.

**Realistic Behavior Simulation:** Early approaches to simulating vehicle behavior include log-replay [15], [20], where trajectories recorded in the real world are simply replayed, rule-based models [15], [21], [22], as well as Behavior Cloning (BC) [9], [10], an open-loop IL method, where a mapping from observations to actions is learned from a dataset using supervised learning. Due to its simplicity and the availability of large-scale datasets, the latter remains widely used. Each approach, however, has its limitations: Replayed trajectories are non-reactive, thus prohibiting interactions with a simulated ego vehicle. Rule-based models often fail to capture the complexity of real-world behaviors. BC performs poorly in rare or unseen scenarios (e. g., near-collisions or very aggressive merges). Even worse, when deploying such a model in a closed-loop simulation, small deviations from expert driving accumulate, leading the model into unfamiliar states - known as *covariate shift* [23].

For learning more robust behavior models, RL is commonly employed [7], [8], [11]–[13]. Since realistically simulating traffic scenarios is inherently a multi-agent problem, self-play RL has emerged as an effective solution. In [12], all agents are controlled by a single shared policy, which is trained jointly through self-play. To enable applying the model to a diverse set of traffic scenarios, a flexible, graph-based input representation with vectorized map features is adopted. Similarly, in [7] and [8], self-play RL has been applied on a large scale, demonstrating that robust and naturalistic behavior emerges from it. The resulting behavior model generalizes well, even in out-of-distribution scenarios. Notably, in [7], the behavior model is trained in randomly initialized traffic scenarios, driving more than 1.6 billion km during training. The resulting behavior model is zero-shot evaluated on three autonomous driving benchmarks, achieving state-of-the-art performance without ever seeing human data during training.

However, achieving realistic behavior with RL requires a carefully tuned reward function that precisely guides the agents toward human-like driving. Recent works [5], [6], [24], [25] address this issue by reconstructing a reward signal from real-world data. In [6], AIRL is used, where a discriminator is trained to distinguish real from simulated behavior, assigning higher scores to more realistic samples. As the goal is to drive as realistically as possible, the output of the discriminator is then used as a reward signal for RL training. Similarly, we train our model using AIRL, but introduce an adaptive reward transformation to balance robustness and realism.

**Enhanced inference speed:** Many existing works on multi-agent driving simulation focus on optimizing the simulation backbone, i. e., the engine responsible for executing actions and producing next-step observations. This is commonly done by implementing the simulator in PyTorch [7], [17], JAX [15], or C++ using CUDA [16] to support in-graph compilation for hardware (GPU/TPU) acceleration. However, we argue that often, not the simulation engine, but the executed behavior model is the limiting factor. We therefore focus on optimizing the executed behavior model for multi-agent simulation.

All the discussed approaches to behavior modeling employ agent-centric observations, where each observation is expressed relative to the simulated agent’s frame of reference. While straightforward, this becomes computationally expensive when simulating multiple agents, as the complexity scales linearly with the number of agents and quadratically with the number of pairwise interactions between agents.

Although coordinated scene-centric agent behavior has already been investigated for single-shot motion forecasting, e. g., [19], [26], to the best of our knowledge, TrafficSim [18] is the only work on learning a scene-centric multi-agent behavior model for closed-loop simulation. In [18], a global map is rasterized and encoded using a CNN. Then, local map features are extracted via Rotated Region of Interest Align and fused with the agent features. Lastly, a joint decoder model, realized as a message passing network, processes all agent features jointly, enabling coordination between them. This scene-centric representation effectively leads to more efficient simulations, as the static environment needs to be encoded only once, allowing the reuse of its encoded tokens in subsequent simulation steps. However, due to the global coordinate frame, this comes with the potential sacrifice of pose-invariance, often leading to performance degradation when the global coordinate frame changes.

To address this, recent single-shot motion prediction methods, e. g., [27]–[30], adopt instance-centric representations, encoding each agent and map element in its own local frame. These pose-invariant features can be reused across all target agents, with relative positions enabling symmetric encoding from any agent’s perspective.

To the best of our knowledge, we are the first work to employ instance-centric representations for closed-loop traffic simulation, enabling us to accelerate the simulation by reusing encoded tokens of static map elements across simulation steps.

## III. METHOD

For learning and evaluating behavior models, a simulation framework is required. In our in-house simulation framework, each vehicle is assigned a route to follow. Similarly, existing works use waypoints [1], [15] or goal points [7], [16]. Vehicles terminate upon colliding or leaving the road.

As illustrated in Figure 2, simulations evolve recursively. First, an observation model encodes the traffic situation, either in an agent-centric or instance-centric manner. Next, a learned behavior model maps the observation to actions, i. e., acceleration and steering angle for vehicles, and acceleration and heading rate for pedestrians. Finally, these actions are

executed via a kinematics model, updating the agent’s state. Repeating these steps for all agents at fine-grained time intervals yields continuous traffic simulation. Our in-house simulation backbone is implemented in C++, thus achieving a low runtime. Figure 3 shows an example scenario.

This sequential decision-making process is formulated as a Partially-Observable Markov Decision Process (POMDP), characterized by the tuple  $(S, O, A, T, R, \Omega, \gamma)$ . In this framework, the agent does not have direct access to the true state  $s \in S$  of the environment. Instead, it receives a noisy observation  $o \in O$ , provided by the observation model  $\Omega : S \rightarrow O$ . Upon executing an action  $a \in A$ , the state of the environment is updated stochastically according to the transition probability density  $T : S \times A \times S \rightarrow [0, \infty)$ . In addition, the agent receives a numerical reward defined by the reward function  $R : S \times A \rightarrow \mathbb{R}$ , as well as a new observation of the updated environment state. The discount factor  $\gamma \in [0, 1)$  balances the trade-off between immediate and future rewards. The solution to a POMDP is the optimal policy  $\pi^* : O \times A \rightarrow [0, \infty)$ , which maps observations to action distributions, maximizing the expected cumulative reward  $J(\pi) = \mathbb{E}_\pi [\sum_{k=0}^{\infty} \gamma^k r_k]$ , where  $r_k = R(s_k, a_k)$ .

### A. Reinforcement Learning

We maximize  $J(\pi)$  using RL, where an agent learns effective actions by interacting with a simulated environment. The process alternates between 1) collecting a set of experiences  $E = \{e_1, \dots, e_M\}$ , with each  $e_k = (o_k, a_k, r_k)$  representing a single-step experience sample, and 2) updating the policy to reinforce actions leading to higher rewards. These steps are alternated until a sufficiently good policy is found.

Commonly, the parameterized policy  $\pi_\theta$  is updated via:

$$\theta \leftarrow \theta + \alpha \frac{1}{M} \sum_{e_k \in E} A(o_k, a_k) \nabla_\theta \log \pi_\theta(a_k | o_k), \quad (1)$$

where  $\alpha$  is the learning rate and the gradients of the policy  $\nabla_\theta \log \pi_\theta(a_k | o_k)$  are weighted by the advantage  $A(o_k, a_k)$ .

In our multi-agent setting, we apply a shared policy across all agents, trained via self-play RL. Specifically, we use Generalized Advantage Estimation (GAE) [31] for estimating  $A(o_k, a_k)$  and Proximal Policy Optimization (PPO) [32] for optimizing the policy. We refer to [7] for a detailed description of how self-play RL is applied to the task of learning behavior models for simulating driver behavior.

### B. Adversarial Inverse Reinforcement Learning

RL requires defining a reward function that accurately captures the incentive structure of real-world driving. Instead of manually defining such a reward function, we use AIRL [14] to reconstruct a surrogate reward signal from real data, given as  $\mathcal{D} = \{(o_1, a_1), (o_2, a_2), \dots\}$ . In AIRL, an additional discriminator model  $D_\phi$  is trained to distinguish generated from real samples, producing the probability  $D_\phi(o, a) \in [0, 1]$  for the observation-action pair being real, i. e., stemming from  $\mathcal{D}$ . The policy is trained via RL using the surrogate reward

$$\tilde{r}(o, a) = \log D_\phi(o, a) - \log(1 - D_\phi(o, a)), \quad (2)$$




Fig. 2. Single simulation step: The behavior model maps observations to actions, which are executed via a kinematic bicycle model for vehicles and kinematic unicycle model for pedestrians.




Fig. 3. Illustration of an example scenario using instance-centric observations. Simulated vehicles and their counterparts from the real-world dataset are depicted as solid and outlined rectangles, respectively.

which guides the policy toward taking realistic actions. Following [6], to smooth the decision boundary during discriminator training, we add random noise to the actions in  $\mathcal{D}$ , matching the standard deviation of the policy.

**Adaptive Reward Transformation:** Typically, the discriminator classifies samples correctly, leading to a negative expected value for  $\tilde{r}(o, a)$ . This often drives agents to terminate as fast as possible to avoid constantly accumulating penalties. To mitigate this, in [6] a positive constant  $c = 5$  is added to  $\tilde{r}$ , shifting its expected value into the positive range. Intuitively, this adjustment encourages survival without altering the policy’s optimal behavior w. r. t. the discriminator.

However, using a fixed  $c$  across experiments, e. g., when comparing different discriminator architectures or scene representations, may result in varying expected values of  $\tilde{r}$ , thereby influencing the agent’s incentive to survive. For instance, a weaker discriminator tends to learn a less distinct decision boundary, yielding higher surrogate rewards, whereas a stronger discriminator learns sharper boundaries, producing lower surrogate rewards. Consequently, policies trained with stronger discriminators have less incentive to avoid termination, i. e., collisions or off-track driving, despite in both cases the optimal behavior w. r. t. the discriminator is to match the distribution of the real data  $\mathcal{D}$ .

To decouple the agent’s survival incentive from the discriminator’s accuracy, we propose an adaptive reward offset

$$c(n) = \tilde{r}_{\text{target}} - \tilde{r}_{\text{mean}}(n), \quad (3)$$

where  $\tilde{r}_{\text{mean}}(n)$  is the average surrogate reward for the generated samples in epoch  $n$ , and  $\tilde{r}_{\text{target}}$  is a hyperparameter, specifying the desired average reward. In other words, our dynamic reward offset can be interpreted as an additional reward term that encourages the policy to survive longer, especially when the discriminator is confident that the given samples are generated. This adaptive approach keeps the expected value of  $\tilde{r}$  consistent, thus enabling fair comparisons across experiments, while still allowing the discriminator to guide the policy toward more realistic behavior by computing the appropriate offset to  $\tilde{r}_{\text{target}}$ .

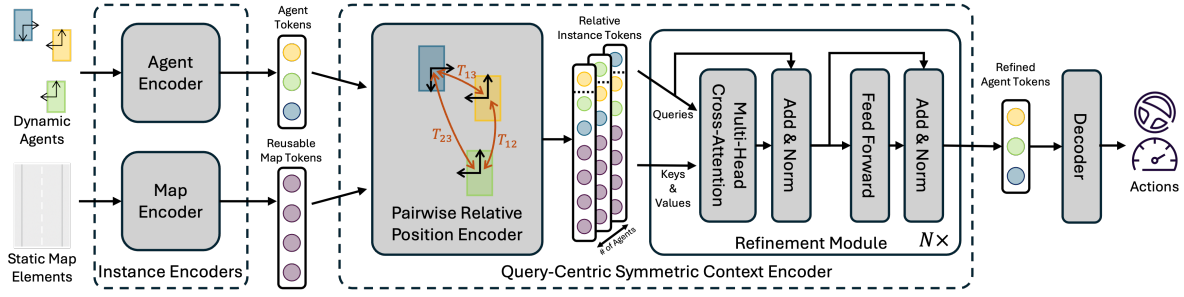


Fig. 4. Illustration of the proposed instance-centric behavior model mapping observations to actions. Instance encoders convert observations into latent tokens. These tokens are then augmented with positional encodings relative to the target agent before being passed through multiple layers of our refinement module. Lastly, each refined agent token is decoded into its corresponding actions. Static map tokens can be reused across simulation steps.

Technically, our adaptive reward transformation is a form of reward shaping. However, unlike conventional shaping methods that add a fixed term to guide the policy toward desirable behavior, our offset adapts dynamically to the discriminator’s performance, automatically balancing realism and robustness during training.

### C. Instance-Centric Observations

Rather than adopting agent-centric observations, as prior works do, e.g., [5], [7], [8], we employ instance-centric observations, representing each instance, i.e., the  $N_a$  agents and  $N_p$  map elements, in its respective local coordinate frame.

For static map elements, such as pedestrian crossings or lane boundaries, we use a vectorized representation [33]. Each map element is approximated by a polyline with a maximum total length of 10 m. A polyline consists of multiple vectors, each described by a set of features  $\mathbf{v} = [\mathbf{v}_{\text{start}}, \mathbf{v}_{\text{end}}, \mathbf{v}_{\text{type}}]^T$ , which includes the vector’s start and end points, as well as a one-hot encoding for the polyline type. For each polyline, the origin of its instance-centric coordinate frame is defined as the mean position of all polyline points. The x-axis is aligned with the mean direction of the vectors forming the polyline.

The local coordinate frame of an agent is determined by its last position and moving direction. Agents are described by their size (width and length), velocity, current speed limit, and a binary indicator for Vulnerable Road Users (VRUs):  $\mathbf{x} = [x_{\text{size}}, x_{\text{vel}}, x_{\text{limit}}, x_{\text{vru}}]^T$ . Note that, compared to existing works utilizing agent-centric observations, here the agent’s position and orientation are not included in the agent features, but are now captured by the position and orientation of its local coordinate frame.

Intuitively, the local coordinate frames serve as anchor poses for the individual instances, with  $\mathbf{p}_i$  being the origin and  $\alpha_i$  being the heading of the  $i$ -th instance relative to an arbitrary global coordinate frame. Following [27], [30], to capture the spatial relationship between two instances  $i$  and  $j$ , we define a coordinate transformation from the local frame of  $i$  to that of  $j$ . Specifically, the relative pose of instance  $j$  in the frame of instance  $i$  is represented as

$$\mathbf{r}_{i \rightarrow j} = [\Delta\alpha_{i \rightarrow j}, \psi_{i \rightarrow j}, \|\mathbf{p}_{i \rightarrow j}\|]^T, \quad (4)$$

where  $\Delta\alpha_{i \rightarrow j} = \alpha_j - \alpha_i$  denotes the heading difference,  $\psi_{i \rightarrow j}$  is the relative azimuth, i.e., the angle of  $\mathbf{p}_{i \rightarrow j} = \mathbf{p}_j - \mathbf{p}_i$  in the reference frame of  $i$ , and  $\|\mathbf{p}_{j \rightarrow i}\|$  is the distance between

the two anchor poses. Both angles are provided as unit-circle embeddings  $f(x) = (\cos x, \sin x)$ .

### D. Behavior Model Details

After obtaining a new observation, it needs to be mapped to the actions that should be executed upon making that observation. Figure 4 illustrates our proposed model.

**Instance Feature Encoders:** We begin by processing the observation into a set of latent instance-tokens. Following the approach introduced in [33], the  $L$  vectors of a polyline are encoded through multiple layers of message-passing:

$$\mathbf{v} \leftarrow f_{\text{rel}}(g_{\text{enc}}(\mathbf{v}), f_{\text{agg}}(\{g_{\text{enc}}(\mathbf{v}_l)\}_{l=1}^L)), \quad (5)$$

followed by an element-wise max-pooling operation over  $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_L\}$  to produce the final polyline token  $\mathbf{z}^{(\text{map})}$ . We implement  $g_{\text{enc}}(\cdot)$  as a multi-layer perceptron (MLP),  $f_{\text{agg}}(\cdot)$  as element-wise max-pooling, and  $f_{\text{rel}}(\cdot)$  as concatenation. Conversely, each agent feature vector  $\mathbf{x}$  is encoded into the same latent space using another MLP, producing the agent token  $\mathbf{z}^{(\text{agent})} = \text{MLP}(\mathbf{x})$ .

Note that using agent-centric observations, requires each instance to be encoded  $N_a \times H$  times — once per agent per simulation step, where  $H$  denotes the total number of steps. In contrast, with our instance-centric representation, the polyline tokens  $Z^{(\text{map})} = \{\mathbf{z}_i^{(\text{map})} \mid i = 1, \dots, N_p\}$  remain unchanged throughout the rollout. This enables reusing the tokens in subsequent simulation steps, reducing the number of encoding a polyline-instance from  $N_a \times H$  times to just one. However, agent states evolve during the simulation rollouts, requiring agent tokens  $Z^{(\text{agent})} = \{\mathbf{z}_i^{(\text{agent})} \mid i = 1, \dots, N_a\}$  to be regenerated at each simulation step. Nonetheless, since agents are encoded in a viewpoint-invariant manner, they only must be encoded  $H$  times instead of  $N_a \times H$  times.

**Query-Centric Symmetric Context Encoder:** After obtaining the instance tokens, we use a symmetric context encoder to refine the queried agent instance tokens in a viewpoint-invariant manner. The idea is to model the interactions between the target agent and nearby instance tokens.

For that, we define the relative positional encoding between a target agent  $i \in \{1, \dots, N_a\}$  and any of the combined instance tokens  $Z = [Z^{(\text{agent})}, Z^{(\text{map})}]$ , indexed by  $j$ , as  $\mathbf{c}_{i \rightarrow j} = \mathbf{r}_{i \rightarrow j} \oplus \mathbb{1}_j^{(\text{agent})} \oplus \mathbb{1}_j^{(\text{route})}$ , where  $\oplus$  denotes concatenation. The binary indicators  $\mathbb{1}_j^{(\text{agent})}$  and  $\mathbb{1}_j^{(\text{route})}$  specify whether instance  $j$  is an agent, or a map polyline that

is part of the route that agent  $i$  should follow, respectively. These indicators are appended here rather than to the vector features  $\mathbf{v}$ , as each agent may follow a different route, making a global encoding at the vector level infeasible. Finally, following the approach proposed in [34], we compute pairwise encodings by applying a feature-wise linear transformation:

$$\mathbf{z}_{i \rightarrow j} = \zeta(\mathbf{c}_{i \rightarrow j}) \odot \mathbf{z}_j + \beta(\mathbf{c}_{i \rightarrow j}), \quad (6)$$

where  $\zeta(\cdot)$  and  $\beta(\cdot)$  are both realized as MLPs. This operation yields the representation of token  $\mathbf{z}_j$  relative to instance  $i$ .

Similarly to [27], [35], we then refine the instance tokens of simulated target agents through multiple Perceiver layers [36]. Each layer consists of a Multi-Head Cross Attention (MHCA) module with a skip connection, followed by layer normalization, an MLP, also with a skip connection, and another layer normalization. The MHCA is expressed as:

$$\mathbf{z}_i^{(k)} = \text{MHCA}^{(k)} \left( \text{Q: } \mathbf{z}_i^{(k-1)}, \text{KV: } \{\mathbf{z}_{i \rightarrow j}\}_{j \in \Gamma(i)} \right), \quad (7)$$

where  $k \in \{1, \dots, K\}$  denotes the  $k$ -th layer, and  $\Gamma(i)$  is the set of neighboring token indices for agent  $i$ , limited to those with anchor points within a set observation radius around agent  $i$  to ensure efficiency. We initialize the first layer with  $\mathbf{z}_i^{(0)} = \mathbf{z}_{i \rightarrow i}$  and use the output of layer  $k$ , denoted  $\mathbf{z}_i^{(k)}$ , as input to the subsequent Perceiver layer.

**Decoder:** The decoder processes the refined instance token  $\mathbf{z}_i^{(K)}$  from the last Perceiver layer. In the behavior model, it outputs the mean and standard deviation of the next action. In the discriminator, which evaluates observation-action pairs, the action is appended to  $\mathbf{z}_i^{(K)}$  to then compute the classification score  $D_\phi$ .

We use three layers of the message passing mechanism defined in (5) and 16 channels per head in (7). Each MLP consists of a linear layer, followed by layer normalization, ReLU activation, and a second linear layer.

#### IV. EXPERIMENTS

**Datasets:** Evaluation is conducted on two datasets: 1) the publicly available INTERACTION dataset [37], containing 38 255 trajectories (4.15% VRUs) across 11 locations; and 2) an in-house dataset commercially licensed by DeepScenario, comprising 88 540 trajectories (43.59% VRUs) across 8 locations. Both datasets were recorded using drones, ensuring high-quality data, and focus on interactive scenarios such as merging lanes, unsignalized intersections, and roundabouts. We use 20% of the recordings per location for validation and 30% for testing. A more compact version of the DeepScenario dataset [38] is publicly available and yields similar results.

**Baselines:** We compare against several compact baselines:

- 1) *Constant Velocity (CV)*: A learning-free baseline that assumes agents continue forward at a constant velocity.
- 2) *LateFusionMLP* [8]: Following [7], [8], [16], this compact agent-centric model consists solely of MLPs and max-pooling operations. We adopt the public implementation [8], replacing its discrete action decoder with ours to support continuous actions and training it within our framework for realistic behavior modeling.

- 3) *GraphAIRL* [6]: A more sophisticated agent-centric model that uses a vectorized scene representation [33] and attention-based interaction modeling. We evaluate two variants: 1) trained with  $c = 5$ , as done in [6], and 2) trained with our proposed adaptive reward offset (3).
- 4) *Behavior Cloning (BC)*: A supervised learning variant of our instance-centric approach, trained for 600 epochs by minimizing the negative log-likelihood of expert actions under the predicted action distribution.

Our agent-centric observations include both nearby agents and map elements within the observation radius. The start and end points of a vector  $\mathbf{v}$  are expressed in the target agent's coordinate frame, with the route indicator  $\mathbb{1}_j^{(\text{route})}$  concatenated to  $\mathbf{v}$ . Positions and orientations of surrounding agents, relative to the target agent, are appended to  $\mathbf{x}$ .

**Training Details:** All models are trained using the AdamW optimizer [39] with a batch size of 1024. Learning rates are set to 0.0002 for the policy model and 0.0001 for the discriminator, both decayed by a factor of 10 during the final 30% of training. For RL, we use advantage normalization. The models are trained for 10 000 epochs with a discount factor  $\gamma = 0.95$ , GAE parameter  $\lambda = 0.95$ , and target rewards  $\tilde{r}_{\text{target}} = 33$  and  $\tilde{r}_{\text{target}} = 9$  for INTERACTION and DeepScenario, respectively. For the policy, the observation radius is set to 50 m. Since the number of observed instances grows quadratically with the observation range, we reduce the discriminator's observation radius to 30 m, allowing it to focus primarily on the target agent and its direct surroundings.

In each epoch, scenarios with approximately 880 agents are sampled from the training set. Validation scenarios contain roughly 1650 agents, while test scenarios feature 5109 and 19 632 agents for INTERACTION and DeepScenario, respectively. Agents are uniformly distributed across locations. Each scenario is simulated for 10 s at 5 Hz. For final evaluation, we select the model that achieves the lowest aggregated score

$$S(\theta) = \frac{\sqrt{\frac{1}{N_a} \sum_{n=1}^{N_a} \left( \text{FDE}_n^{(\pi_\theta)} \right)^2}}{\max(1 - \text{off}^{(\pi_\theta)} - \text{col}^{(\pi_\theta)}, \epsilon)}, \quad (8)$$

i. e., the root mean squared final displacement error normalized by  $(1 - \text{Off-Track Rate} - \text{Collision Rate})$ . To ensure numerical stability, we set  $\epsilon = 10^{-6}$ . The Final Displacement Error (FDE) is the Euclidean distance to the ground truth agent at the final simulation step. This score balances positional accuracy and robustness. During evaluation, terminated vehicles are not removed from simulation.

We use three Perceiver layers and a hidden dimension of 128. Additionally, we evaluate a smaller variant with one Perceiver layer and a hidden dimension of 64, denoted *Ours (small)*. All baseline models use a hidden dimension of 128.

#### A. Results

We evaluate our proposed approach against a diverse set of baselines on two automated driving datasets. Table I presents the quantitative results on the DeepScenario dataset. The proposed model achieves the lowest FDE for the overall and non-VRU subset, as well as the most robust behavior in terms

TABLE I

MODEL PERFORMANCE ON THE DEEPSCEENARIO DATASET.

Model	Observation Type	Num Parameter	FDE			Collision <sup>†</sup> w/		Off- <sup>†</sup>		Aggregated Score	Training Time	
			All	non-VRU	VRU	non-VRU	VRU	Route	Track		T4 16GB	A100 80GB
Expert Demo	-	-	0.0	0.0	0.0	0.16	0.05	3.75	0.83	0.0	-	-
CV	-	-	16.73	29.18	<b>3.99</b>	26.16	3.76	57.52	49.17	123.72	-	-
LateFusionMLP <sup>‡</sup> [8]	AC	105k	9.28±0.14	13.10±0.23	5.33±0.20	1.61±0.23	0.55±0.09	3.89±0.56	0.61±0.02	14.59±0.22	86.85	40.75
GraphAIRL [6]	AC	146k	7.86±0.15	10.78±0.18	4.83±0.24	1.09±0.38	0.79±0.26	3.96±0.33	0.70±0.13	12.58±0.28	125.94	45.95
GraphAIRL <sup>‡</sup> [6]	AC	146k	7.79±0.12	10.84±0.19	4.60±0.14	0.67±0.27	0.42±0.29	4.16±0.55	0.61±0.03	12.41±0.24	129.03	47.10
BC	IC	430k	7.75±0.53	11.40±0.97	4.03±0.25	14.71±2.92	2.90±0.28	16.72±10.43	10.51±9.80	17.35±5.22	12.67	3.07
Ours (small)	IC	60k	7.59±0.23	10.31±0.35	4.74±0.19	0.66±0.33	0.27±0.09	3.85±0.70	0.57±0.03	12.01±0.34	45.53	25.13
Ours	IC	430k	<b>7.09</b> ±0.23	<b>9.69</b> ±0.24	4.43±0.26	<b>0.38</b> ±0.13	<b>0.14</b> ±0.05	<b>3.27</b> ±0.23	<b>0.57</b> ±0.03	<b>11.29</b> ±0.35	125.75	44.07

Reported as mean  $\pm$  std across 7 random seeds. <sup>†</sup> indicates metric reported for the non-VRU subset. <sup>‡</sup> indicates baseline model trained with the proposed target reward. FDE in [m]. Collision, off-route and off-track rates in [%]. Training times in [h]. AC = agent-centric. IC = instance-centric.

TABLE II

MODEL PERFORMANCE ON THE INTERACTION DATASET.

Model	FDE	Collision <sup>†</sup>	Off-Track <sup>†</sup>	Aggregated Score	Training Time	
					T4 16GB	A100 80GB
Expert Demo	0.0	0.0	0.51	0	-	-
CV	17.44	22.59	30.65	49.55	-	-
LateFusionMLP <sup>‡</sup> [8]	10.12±0.14	0.57±0.24	<b>0.30</b> ±0.02	14.20±0.30	97.57	44.48
GraphAIRL [6]	8.51±0.35	1.38±0.48	0.36±0.04	12.49±0.48	152.35	52.26
GraphAIRL <sup>‡</sup> [6]	9.05±0.47	0.70±0.25	<b>0.30</b> ±0.03	12.94±0.50	157.19	55.14
BC	10.28±0.51	13.98±2.15	9.65±7.02	19.09±2.22	2.37	0.68
Ours (small)	8.78±0.39	0.40±0.22	0.32±0.04	12.56±0.58	42.47	24.17
Ours	<b>8.34</b> ±0.29	<b>0.38</b> ±0.18	<b>0.30</b> ±0.02	<b>12.08</b> ±0.26	126.08	50.43

Reported as mean  $\pm$  std across 7 random seeds.

TABLE III

CROSS-DATASET PERFORMANCE ON THE INTERACTION DATASET.

Model	FDE	Collision <sup>†</sup>	Off-Track <sup>†</sup>	Aggregated Score
Expert Demo	0.0	0.0	0.51	0.0
CV	17.44	22.59	30.65	49.55
LateFusionMLP <sup>‡</sup> [8]	15.22±1.34	6.47±1.97	6.89±2.85	22.82±3.21
GraphAIRL [6]	16.05±1.14	5.54±2.57	4.12±2.13	22.80±2.62
GraphAIRL <sup>‡</sup> [6]	14.44±0.49	3.24±0.84	3.89±3.06	20.43±1.44
BC	15.62±1.31	19.69±1.46	13.95±3.65	30.25±2.15
Ours (small)	14.26±1.48	0.81±0.28	1.49±0.79	19.01±1.72
Ours	<b>12.62</b> ±0.69	<b>0.51</b> ±0.16	<b>1.11</b> ±0.16	<b>16.82</b> ±0.71

Reported as mean  $\pm$  std across 7 random seeds. Models are trained on DeepScenario dataset and evaluated on INTERACTION dataset.

of collision, off-route and off-track rates. Overall, our model outperforms all baselines on seven out of eight metrics, and even our small variant with only 60k parameters surpasses all baselines on the same seven metrics. The only exception is VRU FDE, where two baselines achieve marginally better positional accuracy. Notably, our approach achieves lower off-route and off-track rates than ground truth agents, whose errors largely stem from annotation noise.

The learning-free *CV* baseline achieves reasonable positional accuracy for VRUs, as most do not abruptly change speed or direction, but performs poorly on other metrics, underscoring the need to model interactions. *LateFusionMLP* captures more realistic behaviors but still suffers from low robustness and comparably high FDE, likely due to limited model capacity. Both *GraphAIRL* variants demonstrate high positional accuracy and robustness. Notably, training the model with our adaptive reward offset significantly improves its robustness and yields a small gain in positional accuracy. Finally, the *BC* baseline achieves low FDE but incurs extremely high off-track and collision rates due to the aforementioned covariate shift when executed in simulation.

Table II presents the evaluation results on the INTERACTION dataset. Similarly, the proposed model achieves the best performance on all four evaluation metrics. Its compact variant *Ours (small)* remains competitive, exhibiting comparable FDE and outperforming all baselines in robustness.

**Effects on Training Time:** As shown in Tables I and II, training time scales with model size and architectural complex-

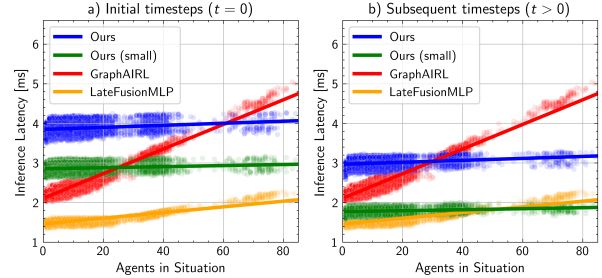


Fig. 5. Regressed inference latency of a single policy-network forward pass w.r.t. the number of agents. a) initial simulation step. b) subsequent simulation steps. Measured on an A100 GPU. Dots denote individual scenes.

ity. Despite having more parameters, *Ours* outperforms both *GraphAIRL* baselines in terms of training times. *Ours (small)* consistently trains faster than all agent-centric baselines, reducing training time by up to 64.71% on DeepScenario and 72.98% on INTERACTION without compromising performance. Moreover, while agent-centric models typically require more computational power, our compact model variant achieves training times on a T4 16 GB GPU comparable to the baselines trained on an A100 80 GB GPU.

**Effects on Inference Time:** Figure 5 presents a quantitative comparison of single-step inference times across our models and baseline approaches. For the agent-centric models, inference time scales linearly with the number of agents, as observations must be encoded separately from each agent’s point of view. Nevertheless, the *LateFusionMLP* consistently exhibits low inference latency due to its compact MLP-based architecture. In contrast, our instance-centric approach decomposes inference into two stages: an initial simulation step, in which the entire map and all agents are encoded, and subsequent steps, in which the encoded map tokens are reused, substantially reducing the inference time for later simulation steps. As a result, simulating  $n$  steps incurs the latency of the initial step only once, followed by the latency of subsequent steps  $n - 1$  times. For scenarios with few agents, the initial forward pass of our proposed small model is slightly slower than that of *GraphAIRL*, since it encodes the entire map rather than only the regions visible to each agent. However, this overhead quickly diminishes as the number of simulation steps increases, yielding greater overall efficiency for longer simulations. Additionally, due to the symmetric scene context encoding, inference times in our instance-centric models stay nearly constant, regardless of the number of simulated agents.

Commonly, multiple simulations are run in parallel. To evaluate the sample throughput of the behavior model under parallel simulation, we introduce the metric *Inference*




Fig. 6. Peak throughput of the behavior model. a) Inference Steps per Second (ISPS) w.r.t. number of parallel simulation environments. b) Distribution of simulated agents across 10 000 randomly sampled scenarios from the INTERACTION dataset. Measured on an A100 GPU.

*Steps per Second (ISPS).* This metric quantifies the number of agents across all simulation environments that can be processed by the behavior model per second. Formally, it is defined as  $ISPS = H \times \sum_{k=1}^N N_A^{(k)} / \Delta T$ , where  $N_A^{(k)}$  denotes the number of agents in the  $k$ -th simulation environment,  $H$  is the number of inference steps taken, and  $\Delta T$  is the elapsed time in seconds. Figure 6 illustrates how the behavior model throughput scales with increasing numbers of parallel simulation environments as well as the distribution of simulated agents across  $10^4$  randomly sampled scenarios. At smaller scales, all models benefit from parallelism, showing substantial gains in ISPS. Agent-centric models saturate early at around 28.7k and 60.3k ISPS, and cannot be scaled beyond 500 environments due to memory limits. In contrast, instance-centric models achieve higher throughput: *Ours (small)* and *Ours* exceed both agent-centric models with only five and twenty parallel environments, respectively, and continue scaling up to 5000 parallel environments. They reach peak throughput of about 93.4k and 340.7k ISPS, representing an up to 11.9-fold improvement over the agent-centric baselines.

*Effects on Cross-Dataset Performance:* Next, we evaluate the models’ performance in unseen traffic scenarios. For that, the models are trained on the DeepScenario dataset and subsequently evaluated on the INTERACTION dataset. As depicted in Table III, all models trained on the DeepScenario dataset, which predominantly contains scenarios from Germany, experience notable performance drops when evaluated on the INTERACTION dataset, which includes diverse scenarios from Germany, China, and the USA. This effect is particularly pronounced in the baseline models. In contrast, our proposed instance-centric model demonstrates greater robustness, maintaining relatively low collision and off-track rates while exhibiting smaller increases in FDE compared to the baselines. We attribute this improved generalization to the viewpoint-invariant encoding of instances, which produces more transferable scene representations. By comparison, agent-centric approaches excel at modeling an individual agent’s local context but remain brittle to changes in viewpoint or environmental conditions, limiting applicability to unseen scenarios. Notably, the collision rate of our models increases only marginally, highlighting their capacity to better generalize to diverse interactions with other agents.

*Effects of Average Reward Target:* Figure 7 presents the results of our ablation study on different choices for the average reward target  $\tilde{r}_{\text{target}}$ , evaluated across values from 1




Fig. 7. Ablation on the validation set w.r.t. the average reward target defined in (3) for a) INTERACTION dataset and b) DeepScenario dataset. Collision and off-track rates are calculated on the non-VRU subset. Results averaged over 5 seeds. Whiskers indicate  $\pm 1$  standard deviation.

to 39 in increments of 2. In general, increasing  $\tilde{r}_{\text{target}}$  leads to more robust behavior, with larger gains being achieved at lower values. Initially, the positional accuracy, measured by FDE, increases significantly, demonstrating that the model benefits from the improved robustness. However, increasing  $\tilde{r}_{\text{target}}$  too much, leads to the positional error to rise again. For instance, this turning point can easily be seen for  $\tilde{r}_{\text{target}} = 9$  on the DeepScenario dataset. This indicates a shift in the model’s behavior, where it starts to prioritize robustness over realism. Based on these results, we selected the values 9 for DeepScenario and 33 for INTERACTION, representing the optimal trade-off between positional accuracy and robustness.

*Example Scenario:* Figure 3 shows a highly interactive example scenario featuring an all-way-stop intersection. All agents are controlled by our learned behavior model. Notably, the simulated and corresponding ground truth trajectories closely match, even over long horizons. The policy captures realistic driving dynamics, including similar speeds, accelerations, and subtle behaviors such as cooperative merging.

## V. CONCLUSION

This work presents a novel approach to robust and efficient RL-based behavior modeling for multi-agent driving simulation. Our instance-centric representation captures agents and map elements in their respective local coordinate frames, enabling viewpoint-invariant scene encoding. A Transformer-based refinement module captures interactions using relative positional encodings between reference frames. Static map elements need to be encoded only once per simulation, reducing redundant computation. As a result, both training and inference times are substantially reduced, and our model scales efficiently with the number of agents, making it particularly suitable for large-scale simulations. We obtain realistic behavior models via AIRL and propose an adaptive reward transformation that decouples the agent’s survival incentive from the discriminator performance, allowing fair comparisons across experiments. Across two automated driving datasets, our approach outperforms diverse baseline models in positional accuracy and robustness, and demonstrates significantly more generalizable behavior in cross-dataset evaluations.

## ACKNOWLEDGMENT

The authors would like to thank Matthias Dingwerth and Matthias Steiner for their support and contributions to the development of the simulation framework.

## REFERENCES

- [1] S. Suo, K. Wong, J. Xu, J. Tu, A. Cui, S. Casas, and R. Urtasun, "Mixsim: A hierarchical framework for mixed reality traffic simulation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 9622–9631.
- [2] K. Chitta, D. Dauner, and A. Geiger, "Sledge: Synthesizing driving environments with generative models and rule-based traffic," in *European Conference on Computer Vision*. Springer, 2024, pp. 57–74.
- [3] A. Amini, I. Gilitschenski, J. Phillips, J. Moseyko, R. Banerjee, S. Karaman, and D. Rus, "Learning robust control policies for end-to-end autonomous driving from data-driven simulation," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 1143–1150, 2020.
- [4] Z. Zhang, A. Liniger, D. Dai, F. Yu, and L. Van Gool, "Trafficbots: Towards world models for autonomous driving simulation and motion prediction," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 1522–1529.
- [5] R. Bhattacharyya, B. Wulfe, D. J. Phillips, A. Kuefler, J. Morton, R. Senanayake, and M. J. Kochenderfer, "Modeling human driving behavior through generative adversarial imitation learning," *IEEE Transactions on Intelligent Transportation Systems*, vol. 24, no. 3, pp. 2874–2887, 2022.
- [6] F. Konstantinidis, M. Sackmann, U. Hofmann, and C. Stiller, "Graph-based adversarial imitation learning for predicting human driving behavior," in *2024 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2024, pp. 857–864.
- [7] M. Cusumano-Towner, D. Hafner, A. Hertzberg, B. Huval, A. Petrenko, E. Vinitsky, E. Wijmans, T. W. Killian, S. Bowers, O. Sener *et al.*, "Robust autonomy emerges from self-play," in *International Conference on Machine Learning*. PMLR, 2025, pp. 11 710–11 737.
- [8] D. Cornelisse, A. Pandya, K. Joseph, J. Suárez, and E. Vinitsky, "Building reliable sim driving agents by scaling self-play," *arXiv preprint arXiv:2502.14706*, 2025.
- [9] M. Bansal, A. Krizhevsky, and A. Ogale, "Chauffeurmet: Learning to drive by imitating the best and synthesizing the worst," *arXiv preprint arXiv:1812.03079*, 2018.
- [10] Y. Lu, J. Fu, G. Tucker, X. Pan, E. Bronstein, R. Roelofs, B. Sapp, B. White, A. Faust, S. Whiteson *et al.*, "Imitation is not enough: Robustifying imitation with reinforcement learning for challenging driving scenarios," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2023, pp. 7553–7560.
- [11] J. Chen, B. Yuan, and M. Tomizuka, "Model-free deep reinforcement learning for urban autonomous driving," in *2019 IEEE intelligent transportation systems conference (ITSC)*. IEEE, 2019, pp. 2765–2771.
- [12] F. Konstantinidis, M. Sackmann, U. Hofmann, and C. Stiller, "Modeling interaction-aware driving behavior using graph-based representations and multi-agent reinforcement learning," in *2023 IEEE 26th International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2023, pp. 1643–1650.
- [13] M. Arief, M. Timmerman, J. Li, D. Isele, and M. J. Kochenderfer, "Importance sampling-guided meta-training for intelligent agents in highly interactive environments," *IEEE Robotics and Automation Letters*, 2024.
- [14] J. Fu, K. Luo, and S. Levine, "Learning robust rewards with adversarial inverse reinforcement learning," in *International Conference on Learning Representations*, 2018.
- [15] C. Gulino, J. Fu, W. Luo, G. Tucker, E. Bronstein, Y. Lu, J. Harb, X. Pan, Y. Wang, X. Chen *et al.*, "Waymax: An accelerated, data-driven simulator for large-scale autonomous driving research," *Advances in Neural Information Processing Systems*, vol. 36, pp. 7730–7742, 2023.
- [16] S. Kazemkhani, A. Pandya, D. Cornelisse, B. Shacklett, and E. Vinitsky, "Gpudrive: Data-driven, multi-agent driving simulation at 1 million fps," *arXiv preprint arXiv:2408.01584*, 2024.
- [17] A. Ścibior, V. Lioutas, D. Reda, P. Bateni, and F. Wood, "Imagining the road ahead: Multi-agent trajectory prediction via differentiable simulation," in *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*, 2021, pp. 720–725.
- [18] S. Suo, S. Regalado, S. Casas, and R. Urtasun, "Trafficssim: Learning to simulate realistic multi-agent behaviors," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 10 400–10 409.
- [19] J. Ngiam, B. Caine, V. Vasudevan, Z. Zhang, H.-T. L. Chiang, J. Ling, R. Roelofs, A. Bewley, C. Liu, A. Venugopal *et al.*, "Scene transformer: A unified architecture for predicting multiple agent trajectories," *arXiv preprint arXiv:2106.08417*, 2021.
- [20] L. Bergamini, Y. Ye, O. Scheel, L. Chen, C. Hu, L. Del Pero, B. Osifski, H. Grimmer, and P. Ondruska, "Simmet: Learning reactive self-driving simulations from real-world observations," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 5119–5125.
- [21] A. Kesting, M. Treiber, and D. Helbing, "General lane-changing model mobil for car-following models," *Transportation Research Record*, vol. 1999, no. 1, pp. 86–94, 2007.
- [22] —, "Enhanced intelligent driver model to access the impact of driving strategies on traffic capacity," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 368, no. 1928, pp. 4585–4605, 2010.
- [23] J. Spencer, S. Choudhury, A. Venkatraman, B. Ziebart, and J. A. Bagnell, "Feedback in imitation learning: The three regimes of covariate shift," *arXiv preprint arXiv:2102.02872*, 2021.
- [24] J. Sun and J. Kim, "Modelling two-dimensional driving behaviours at unsignalised intersection using multi-agent imitation learning," *Transportation Research Part C: Emerging Technologies*, vol. 165, p. 104702, 2024.
- [25] C. Weaver, C. Tang, C. Hao, K. Kawamoto, M. Tomizuka, and W. Zhan, "Betail: Behavior transformer adversarial imitation learning from human racing gameplay," *IEEE Robotics and Automation Letters*, 2024.
- [26] D. A. Su, B. Douillard, R. Al-Rfou, C. Park, and B. Sapp, "Narrowing the coordinate-frame gap in behavior prediction models: Distillation for efficient and accurate scene-centric motion forecasting," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 653–659.
- [27] L. Zhang, P. Li, S. Liu, and S. Shen, "Simpl: A simple and efficient multi-agent motion prediction baseline for autonomous driving," *IEEE Robotics and Automation Letters (RA-L)*, 2024.
- [28] Z. Zhou, J. Wang, Y.-H. Li, and Y.-K. Huang, "Query-centric trajectory prediction," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 17 863–17 873.
- [29] Z. Zhang, A. Liniger, C. Sakaridis, F. Yu, and L. V. Gool, "Real-time motion prediction via heterogeneous polyline transformer with relative pose encoding," *Advances in Neural Information Processing Systems*, vol. 36, pp. 57 481–57 499, 2023.
- [30] S. Shi, L. Jiang, D. Dai, and B. Schiele, "Mtr++: Multi-agent motion prediction with symmetric scene modeling and guided intention querying," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 46, no. 5, pp. 3955–3971, 2024.
- [31] J. Schulman, P. Moritz, S. Levine, M. Jordan, and P. Abbeel, "High-dimensional continuous control using generalized advantage estimation," *arXiv preprint arXiv:1506.02438*, 2015.
- [32] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [33] J. Gao, C. Sun, H. Zhao, Y. Shen, D. Anguelov, C. Li, and C. Schmid, "Vectornet: Encoding hd maps and agent dynamics from vectorized representation," in *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020, pp. 11 522–11 530.
- [34] E. Perez, F. Strub, H. De Vries, V. Dumoulin, and A. Courville, "Film: Visual reasoning with a general conditioning layer," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 32, no. 1, 2018.
- [35] L. Chen, O. Sinavski, J. Hünemann, A. Karnsund, A. J. Willmott, D. Birch, D. Maund, and J. Shotton, "Driving with llms: Fusing object-level vector modality for explainable autonomous driving," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 2024.
- [36] A. Jaegle, F. Gimeno, A. Brock, O. Vinyals, A. Zisserman, and J. Carreira, "Perceiver: General perception with iterative attention," in *International conference on machine learning*. PMLR, 2021, pp. 4651–4664.
- [37] W. Zhan, L. Sun, D. Wang, H. Shi, A. Clausse, M. Naumann, J. Kummerle, H. Königshof, C. Stiller, A. de La Fortelle *et al.*, "Interaction dataset: An international, adversarial and cooperative motion dataset in interactive driving scenarios with semantic maps," *arXiv preprint arXiv:1910.03088*, 2019.
- [38] O. Dhaouadi, J. Meier, L. Wahl, J. Kaiser, L. Scalerandi, N. Wandelburg, Z. Zhou, N. Berinpanathan, H. Banzhaf, and D. Cremers, "Highly accurate and diverse traffic data: The deepscenario open 3d dataset," in *2025 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2025, pp. 377–384.
- [39] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," *arXiv preprint arXiv:1711.05101*, 2017.