

V-Max: A Reinforcement Learning Framework for Autonomous Driving

Valentin Charraut^{1,†}, Waël Doulazmi^{1,2,†}, Thomas Tournaire^{1,†}, Thibault Buhet¹

{firstname.lastname}@valeo.com

¹ Valeo Brain

² Centre for Robotics, Mines Paris - PSL

[†] Equal contribution

Abstract

Learning-based decision-making has the potential to enable generalizable Autonomous Driving (AD) policies, reducing the engineering overhead of rule-based approaches. Imitation Learning (IL) remains the dominant paradigm, benefiting from large-scale human demonstration datasets, but it suffers from inherent limitations such as distribution shift and imitation gaps. Reinforcement Learning (RL) presents a promising alternative, yet its adoption in AD remains limited due to the lack of standardized and efficient research frameworks. To this end, we introduce V-Max, an open research framework providing all the necessary tools to facilitate RL research for AD. V-Max is built on Waymax (Gulino et al., 2023), a hardware-accelerated AD simulator designed for large-scale experimentation. We extend it using ScenarioNet’s (Li et al., 2023b) approach, enabling the fast simulation of diverse AD datasets. V-Max integrates a set of observation and reward functions, transformer-based encoders, and training pipelines. Additionally, it includes adversarial evaluation settings and an extensive set of evaluation metrics. Through a large-scale benchmark, we investigate how network architectures, observation functions, training data, and reward shaping impact RL performance. Code is available at: github.com/valeoai/v-max

1 Introduction

Reinforcement Learning (RL, Sutton & Barto (2018)) has proven to be a powerful approach for controlling real-world systems, with milestones in dexterous robotic manipulation and industrial process control (Rajeswaran et al., 2018; Degraeve et al., 2022). RL’s ability to learn adaptive policies through closed-loop interaction makes it an appealing framework for Autonomous Driving (AD, Kiran et al. (2022)), where decision-making agents must continuously respond to unseen scenarios and distribution shifts while maintaining high levels of robustness.

However, applying RL to real-world tasks such as AD introduces significant challenges, particularly regarding sample efficiency and training environments. As a result, RL remains underused in AD research due to practical constraints. Imitation Learning (IL, Bansal et al. (2019)) is often favored instead, as it capitalizes on vast driving datasets collected by vehicle fleets and reduces decision-making to a supervised learning task. The absence of RL-compatible environments made RL unusable in the nuPlan challenge (Karnchanachari et al., 2024), which led the organizers to conclude that learning-based methods could not compete with simple rule-based baselines (Dauner et al., 2023).

This gap has motivated recent efforts to improve the accessibility of RL research for AD. Notably, ScenarioNet provides an open-source framework for standardizing and replaying AD datasets in MetaDrive, an RL-compatible simulator that facilitates research on RL generalization in driving (Li et al., 2023a;b). In parallel, Gulino et al. (2023) released Waymax, a hardware-accelerated driving simulator capable of running large-scale simulations at unprecedented speeds, making RL’s sample inefficiency less of a limiting factor for experimentation. Waymax was developed as a high-speed simulation tool, but it lacks essential benchmarking capabilities for RL research, requiring practitioners to build full training pipelines from scratch.

In this work, we introduce V-Max, a framework that extends Waymax with all the necessary tools for RL research in autonomous driving. V-Max provides a set of observation and reward functions, multiple transformer-based encoders, and a complete training pipeline for standard RL algorithms. All these elements are implemented using the JAX framework (Bradbury et al., 2018), enabling training and simulation to be performed within the same computation graph. Additionally, V-Max leverages ScenarioNet’s approach to enable the accelerated simulation of diverse driving datasets, whereas Waymax was originally limited to the Waymo Open Motion Dataset (WOMD, Ettinger et al. (2021)). With these features, V-Max aims to standardize RL experimentation for AD, making algorithm comparisons more reproducible and accelerating progress in learning-based decision-making.

We enhance Waymax’s evaluation metrics by re-implementing nuPlan’s metrics (Karnchanachari et al., 2024) and introducing additional metrics, such as traffic light violations, for a more comprehensive assessment of policy performance. To further evaluate robustness, we integrate *ReGentS* (Yin et al., 2024), enabling evaluation against adversarial agents. We conduct a large-scale benchmark with these tools, systematically analyzing how observation functions, reward shaping, training data selection, network architectures, and learning algorithms impact performance and sample efficiency. These experiments demonstrate V-Max’s versatility, facilitating research and development on decision-making for AD.

Our contributions are as follows:

1. V-Max provides a fully integrated, JAX-based, RL training pipeline, including observation and reward functions, and transformer-based encoders inspired by motion forecasting.
2. V-Max supports multi-dataset accelerated simulation by extending Waymax with ScenarioNet’s approach.
3. V-Max integrates comprehensive evaluation tools, including the reimplementation of nuPlan’s driving quality metrics, and integration of ReGentS for adversarial evaluation.
4. We perform a benchmark on the impact of network architectures, observation choices, reward shaping, and training data on RL performance in AD, resulting in a policy that successfully completes 97.86% of the scenarios in WOMD.

2 Related Work

2.1 Reinforcement Learning for Autonomous Driving

There are two main formulations of the Autonomous Driving (AD) task in the literature. The first category consists of *end-to-end* approaches (Chen et al., 2024), which aim to learn vehicle controls directly from raw sensor data. Kendall et al. (2019) successfully applied End-to-End RL to lane-following in real-world settings, while Toromanoff et al. (2020) won the first CARLA (Dosovitskiy et al., 2017) challenge using Reinforcement Learning (RL) with a supervised pretraining. These works demonstrated RL’s potential in AD, particularly as a way to overcome the limitations of Imitation Learning (IL), such as distribution shift, causal confusion and imitation gap (Walsman et al., 2022). However, methods based solely on RL still fail to perform in the end-to-end setting, the main reason being that RL gradients are insufficient to train the large neural networks needed for perception (Chen et al., 2024). This issue is further compounded by the difficulty of creating realistic and fast simulators for the closed-loop training required for RL. Most works rely on the




Figure 1: **Overview of the V-Max framework.** ScenarioMax standardizes multiple datasets into a Waymax-compatible format. The simulation runs in Waymax (Gulino et al., 2023), which provides the simulator state s_t . An observation o_t is extracted and processed using a JAX-based neural encoder (Bradbury et al., 2018) before being fed into an RL agent implemented with Brax (Freeman et al., 2021). The RL agent selects an action a_t (acceleration, steering), which is executed in the simulator, receiving a reward r_t based on evaluation metrics. JAX enables to run multiple instances of this process in parallel, on the same device.

CARLA simulator, which allows procedurally generated scenarios to be played in the Unreal Engine (Dosovitskiy et al., 2017). While generative world models such as GAIA-1 (Hu et al., 2023) offer photorealistic closed-loop simulation, their computational cost remains a barrier to large-scale RL training.

The parallel approach is to work at mid-level and decouple the decision-making problem from the real-world perception task. In this *mid-to-end* paradigm, agents process post-perception data, i.e. a structured high-level representation of the scene, and output vehicle controls. The release of large post-perception datasets like WOMD, nuScenes and Argoverse 2 (Caesar et al., 2020; Ettinger et al., 2021; Wilson et al., 2021) accelerated mid-to-end research, with a focus on the trajectory prediction sub-task. Closed-loop evaluation and training of mid-level agents was made possible with the appearance of data-driven simulators, that can replay scenarios from real-world driving while taking into account the agent’s actions. Research on mid-level decision-making mainly revolves around IL and methods to improve its robustness, such as data augmentation (Bansal et al., 2019), model-based generative adversarial IL (MGAIL) (Bronstein et al., 2022), policy gradients (Scheel et al., 2022), and curriculum learning (Bronstein et al., 2023). Notably, the *nuPlan Challenge 2023* (Karnchanachari et al., 2024) remains the only public competition for the mid-to-end AD task, and its closed-loop challenge was won by PDM (Dauner et al., 2023), a rule-based approach that significantly outperformed all the other learning-based approaches, which were all variants of imitation learning.

Lu et al. (2023) demonstrated that combining IL and RL with a simple reward signal can improve policy robustness in corner cases underrepresented in the training dataset. Similarly, Grislain et al. (2024) showed that incorporating an RL objective is needed to mitigate the imitation gap, which arises from the discrepancy between the observations of human experts and those of mid-to-end AD agents (e.g. sound, turn signals). Cusumano-Towner et al. (2025) showed that self-play can generate highly robust policies, surpassing all prior approaches on CARLA, nuPlan, and Waymax. Their work heavily relies on a proprietary high-speed simulator, highlighting how accelerated simulation can enable large-scale RL training and significantly impact learning-based decision-making for AD.

The successful application of RL in the simpler simulated racing task ((Jaritz et al., 2018; Wurman et al., 2022)), further motivates the idea that performing simulators are key to unlock the potential of RL.

2.2 Frameworks for mid-to-end Autonomous Driving

V-Max is a framework built on Waymax (Gulino et al., 2023) which is a data-driven, accelerated, mid-to-end AD simulator. Besides Waymax, other frameworks related to V-Max include nuPlan (Caesar et al., 2021), Nocturne (Vinitsky et al., 2023), MetaDrive (Li et al., 2023a), and GPUDrive (Kazemkhani et al., 2024). Below, we compare them to V-Max.

Datasets. All the aforementioned frameworks enable data-driven simulation, where driving scenes are instantiated by replaying real-world data. MetaDrive also integrates procedural generation, allowing to artificially instantiate driving maps and specific situations (e.g. lane merging, roundabouts). Nocturne, GPUDrive and Waymax are limited to the WOMD dataset (Ettinger et al., 2021), while nuPlan uses its own dataset. MetaDrive and V-Max are compatible support both nuPlan and WOMD, as well as other datasets like Argoverse 2 (Wilson et al., 2021), thanks to the use of ScenarioNet’s standardization (Li et al., 2023b).

Hardware-Acceleration. Waymax supports both acceleration on GPUs and TPUs enabling high speed simulation. If additionally the training pipeline is written using the JAX library (Bradbury et al., 2018), which is the case in V-Max, then simulation and training can be performed within the same computation graph, eliminating communication bottlenecks with the host machine. GPUDrive achieves GPU-acceleration through the Madrona game engine (Shacklett et al., 2023). Hardware-acceleration makes V-Max, Waymax and GPUDrive two to three orders of magnitude faster than CPU-based simulators like nuPlan, MetaDrive, and Nocturne.

Multi-Agent Environments. Waymax supports environments with multiple controllable agents, a feature that V-Max uses to perform adversarial evaluation. While multi-agent RL (MARL) can technically be implemented in Waymax, V-Max is designed for traditional single-agent RL and does not include MARL-specific functionalities. In contrast, GPUDrive is explicitly designed and optimized for multi-agent learning, making it the better choice for MARL and self-play applications.

Observation. In the mid-to-end setting, simulators provide perfect perception of the scene, making the first design choice the selection of what the driving agent observes. There are two approaches to this decision. The first approach models partial observability to reduce the sim-to-real gap. Nocturne and GPUDrive use sensor-based observations that replicate camera or LiDAR properties, where vehicles can occlude one another. V-Max also implements these sensor-based observations, along with the noisy observations from IGDrivSim (Grislain et al., 2024), which were designed to highlight the limitations of IL. The second approach, observation shaping, focuses on selecting an observation that maximizes policy performance while minimizing memory usage. V-Max provides tools for observation shaping and includes a comparison of different observation choices in Table 10, a topic not addressed in other frameworks.

Evaluation. MetaDrive, Nocturne, Waymax, and GPUDrive evaluate driving agents using a goal-reaching metric, which measures the percentage of scenarios where an agent successfully reaches its destination without collisions or off-road violations. nuPlan introduces a more sophisticated scoring system that also considers driving quality. V-Max integrates both the goal-reaching metric from Waymax and nuPlan’s scoring system, enabling more comprehensive evaluations and facilitating direct comparisons between agents.

3 The V-Max Framework

Figure 1 provides an overview of the V-Max framework, which formulates mid-to-end AD as a partially observable Markov decision process (POMDP, Spaan (2012)). In this section, we present the core components of V-Max and how they extend Waymax to facilitate RL’s application to AD.




Figure 2: **Illustration of potential limitations when using expert trajectories as learning targets.** The pink rectangle represents the ego, and the blue rectangles are the other vehicles. The blue path is one SDC path and the red cross is the last waypoint of the expert’s ground truth. *Left:* The trajectory terminates before a traffic signal, inadvertently encoding the implicit knowledge that the expert stopped at a red light. *Center:* The trajectory ends in the middle of the intersection, showing that the expert stopped to let other vehicles pass, which unintentionally teaches the policy when to yield. *Right:* The trajectory ends immediately prior to an intersection, which may result in the policy incidentally avoiding collisions by terminating at this location. In each scenario, we overlay the self-driving vehicle (SDC) path in blue, which provides a topologically consistent road representation without encoding such implicit behavioral biases, thus constituting a more appropriate supervisory signal for policy optimization.

3.1 Rules of the Game

Simulation. The simulation process leverages a `simulator_state` that encapsulates data from a bird’s-eye view (BEV) representation, under the assumption that the perception problem is fully resolved. This `simulator_state` includes comprehensive records of real-world scenarios, encompassing logged trajectories and high-definition (HD) maps. The primary objective of the ego vehicle is to predict control outputs, specifically acceleration and steering, to govern the vehicle’s motion from time t to $t + 1$. Vehicle dynamics are modeled using a continuous bicycle model, which forms the basis for motion planning and control. The simulation operates over a 9-second scenario duration, running at a frequency of 10 Hz. The initial second of each scenario is typically simulated using log-replay to establish a historical context for scene perception.

The scenario concludes when the ego vehicle violates critical safety constraints. Critical failures considered include collisions with other objects, deviations from the road, and crossing intersections under a red light. Notably, the latter constraint is not originally present in the Waymax framework and has been introduced within the V-Max framework.

Goal. V-Max does not prescribe a universal goal for the policy; instead, it allows practitioners to define the desired behavior of the ego vehicle thanks to SDC (Self-Driving Car) paths. Waymax defines SDC path as the routes given to an agent by combining the logged future trajectory of the agent with all possible future routes after the logged trajectories. At the time of writing, these paths are not publicly available in WOMD. An alternative is to rely on expert-logged trajectories only as reference paths. However, this approach is problematic because expert trajectories represent privileged information that consistently demonstrates safe behavior, as shown in Figure 2. To overcome this limitation, V-Max enhances the simulation environment by incorporating reconstructed SDC paths. This addition enables researchers to define various tasks such as navigating to specific destinations or following predetermined routes.

3.2 Training RL Agents

ScenarioMax. One of V-Max’s key contributions is ScenarioMax, an extension of ScenarioNet (Li et al., 2023b) that converts multiple open-source driving datasets into a single, compatible TfRecord format. This integration process requires several preprocessing steps to ensure data consistency and quality across different sources.

Our approach includes SDC paths reconstruction by creating drivable area definitions for the ego vehicle using road lane data. Since the original SDC paths are not publicly available, we derive them from the simulator state information. We construct paths by starting at the lane closest to the SDC’s initial position, then following exit lanes. When multiple lane options exist, we create separate paths. Our method generates 10 distinct paths, selected based on their proximity to the SDC’s final position. This approach captures important route options while maintaining diverse targets. Improvements could be made by adding adjacent lanes, allowing for more complex maneuvers such as safe lane changing.

While ScenarioNet proposed a scenario description format, Waymax simulator requires specific data fields to construct the `simulator_state`. To address this gap, we augment the HD map data by adding directional vectors to each map point and defining proper roadgraph types. We also apply proper labeling to match the `tf.Example` format used by the Waymo Open Motion Dataset (WOMD).

Training pipeline. V-Max uses a flexible wrapper system to encapsulate environments, drawing inspiration from the Brax (Freeman et al., 2021) framework’s approach to parallel simulation while extending it for autonomous driving.

Notable wrappers include the `AutoResetWrapper` that restarts scenarios automatically when completed and the `VmapWrapper` that handles batched scenarios during training to accelerate policy development. We significantly modified the `BraxWrapper` to better integrate with our learning processes. We also added a wrapper to reconstruct one SDC path on the fly in a simulator state to support the original WOMD dataset. This wrapper is not fully recommended as it can contain errors due to the difficulty to reconstruct dynamic data in JAX jitted functions.

To support diverse learning paradigms, V-Max provides a standardized training pipeline that creates consistent agent-simulator interactions across different learning approaches (imitation learning, off-policy, and on-policy methods). Observation and feature extraction wrappers provide a flexible mechanism for processing BEV data and state representations. The reward function module is designed for customization, allowing practitioners to define task-specific objectives and shape agent behavior through tailored incentives.

In addition to these foundational components, V-Max includes popular decision-making algorithms implementations, facilitating rapid experimentation with different policy-learning techniques. A dedicated encoder catalog further enhances the system by offering a range of neural network architectures optimized for extracting high-level representations from input features.

Observation function. Selecting the right input features is essential for the performance of learning-based methods in autonomous driving. While Waymax provides a function to transform the simulator state to the self-driving car (SDC) view, it doesn’t offer complete tools to build input features for neural networks. To solve this problem, we developed feature extractors that organize data into input features such as: (1) trajectory features showing how object motion; (2) roadgraph features describing roads and lanes; (3) traffic light features showing signal states; and (4) path target features indicating where the vehicle should go. Figure 3 shows how the data is processed from the `simulator_state` to adequate features for a neural network.

The entire feature extraction system can be customized through `yaml` configurations, giving practitioners flexibility in designing observations.




Figure 3: **Visualization of the observation process.** *Left:* Scene-centered view of a scenario. SDC paths are displayed, in blue paths containing the expert trajectory, and in green showing alternative route options. *Center:* Ego-centric transformation with HD map filtering within a rectangular bounding box (50 meters front, 5 meters back, 20 meters on both sides of the SDC). Optionally, noise and masking can be applied to the perception of the scene. *Right:* Neural network input representation. Road boundaries are highlighted after roadgraph filtering. Only the eight closest objects to the SDC are retained, and the SDC path containing the ground truth is selected and interpolated into 10 points spaced 5 meters apart, providing a compact representation of the environment for decision making.

Network architectures. To process mid-level observations, we leverage architectures developed for motion forecasting challenges (Ettinger et al., 2021; Wilson et al., 2021). These challenges focus on predicting the future trajectories of all agents in a driving scene and use the same structured scene representations as our task. Since most motion forecasting models are built on encoder-decoder architectures, their encoders can be repurposed to extract meaningful features from a mid-level driving scene, making them suitable as value and policy networks in RL algorithms.

The motion forecasting competitions are dominated by transformer-based architectures (Vaswani et al., 2017). The attention mechanism is particularly useful for encoding temporal dependencies in the SDC’s past trajectory, modeling interactions between the SDC and other road users, and capturing relationships between the SDC and road features. These properties make transformers a compelling architecture for our task. While Waymax reports training results with the Wayformer architecture (Nayakanti et al., 2023), no official public implementation is available. We reimplement Wayformer along with other state-of-the-art encoders from the motion forecasting literature using JAX, enabling in-graph training.

Reward function. In Waymax, the reward is defined as a weighted sum of multiple components. We follow this approach and extend it by adding more reward functions based on the metrics defined in subsection 3.3.

3.3 Evaluation and Benchmarking

Metrics. Waymax proposes the following metrics: collision rate, offroad rate, route progress ratio, and average displacement error (ADE, ℓ_2 -distance between the agent’s position and the expert’s position at each timestep, averaged over the trajectory). We re-implemented the metrics used in the nuPlan challenge (Karnchanachari et al., 2024), which provide a more fine-grained assessment of driving quality. Notably, nuPlan distinguishes between the collisions imputable to the agent’s action, and unavoidable incidents, such as rear-end collisions. Additionally, we integrate a red-light violation check, a feature absent from both Waymax and nuPlan.

Episode score. To aggregate multiple metrics into a single score, we adopt the methodology from the nuPlan challenge (Karnchanachari et al., 2024). Each episode is assigned a score based on a

hybrid weighted average of all metric scores. The complete list of metrics and their corresponding weights are provided in the supplementary material.

Evaluation setups. The main evaluation setup used in V-Max is closed-loop non-reactive, where other agents replay their logged trajectories. The advantage of this setup is that all non-controlled agents exhibit human-like behavior, as they follow real-world recorded data. However, a key limitation arises when the agent’s actions deviate from those originally taken by the expert, leading to unrealistic interactions. A common example is when an agent drives slower than the expert, causing other vehicles to collide with it from behind. This issue is partly mitigated by the short duration of scenarios (8s) and nuPlan’s distinction between at-fault and unavoidable collisions.

In order to enable reactive evaluations, where the other vehicles are controlled by a policy rather than replaying their logged trajectory, we introduce a multi-agent version of WOMD. To construct it, we first define *controllable* vehicles as cars that are not offroad, and has data for the whole scenario (excluding cars that appear and disappear from the data because they were too far to be consistently detected). We then extract from the validation dataset (41872 scenarios) all scenarios with at least one controllable agent, and obtain a dataset of 31379 scenarios, with a mean of 8.72 controllable agents per scenario.

For non-ego agents, traffic lights data may be of poor quality, as the data was originally recorded from the ego vehicle, we robustify traffic lights by inferring it from vehicles movement, similar to what was done in the nuPlan dataset (Karnchanachari et al., 2024). Yan et al. (2025) went a step further by adding missing traffic lights in intersections in the WOMD dataset, which is planned to be integrated in V-Max.

Then, we construct SDC paths for the identified controllable agents, resulting in a new dataset, in which any policy developed on V-Max can be used to drive the controllable agents. It enables to perform closed-loop reactive evaluation with IDM policies similar to nuPlan and Waymax, excepted our IDM policies respect traffic lights, and follow the road rather than logged-trajectories.

This also enables us to try putting various policies like PDM or even evaluating how our RL policy drives when other agents use the same policy. We believe this could serve as a foundation for researchers trying to develop adversarial policies for robustness assessments.

Another evaluation setup available in V-Max applies Gaussian perturbations to the first 10 timesteps of the agent’s trajectory, following the methodology of Bansal et al. (2019). This setup assesses the policy’s ability to recover from distribution shifts, as agents in the training data are most often initialized at the center of their lane.

Finally, V-Max also integrates ReGentS (Yin et al., 2024), a methodology for generating adversarial scenarios by modifying real-world driving data. In ReGentS, surrounding objects (e.g., vehicles, cyclists, and pedestrians) are optimized to create challenging situations for the agent while maintaining realistic and physically plausible interactions. The method prevents unrealistic swinging turns and unavoidable rear-end collisions, ensuring that the generated scenarios provide meaningful robustness evaluations.

4 Case Study: RL Design Choices for Autonomous Driving

We now present the functionalities of V-Max through distinct design choices for implementing an RL agent. To isolate the impact of each component, we use the control configuration defined in Table 1 and conduct experiments on observation functions, network encoders, reward shaping, and training datasets.

We replicate each run¹ across 3 random seeds, and report mean and standard deviation of evaluation metrics. Accuracy denotes the percentage of episodes completed without failure conditions (collisions, off-road, or traffic signal violations). Additionally, we report collision and off-road rates,

¹Runs are executed on a single NVIDIA L4 GPU for 12-48 hours per run.

Table 1: Control configuration used in all the experiments

Algorithm	Encoder	Reward	Training Dataset	Evaluation Dataset
SAC	LQ	Navigation	WOMD Training	WOMD Valid

along with Waymax’s progress ratio metric, which measures how far the agent progresses relative to the expert. Those metrics, as well as nuPlan’s metrics are aggregated in the V-Max score defined in [Appendix B](#). Simulations are executed with a maximum of 64 objects per scenario to avoid memory overload.

Observation function. We perform an ablation study ([Table 10](#) in Appendix) to understand which BEV features are most important for a RL agent.

We find that increasing temporal context from 1 to 5 observation steps does not always improves V-Max scores, and adding lane center or road line features does not systematically improve performance. It suggests that road edge with waypoints information alone provides sufficient spatial context for many scenarios in the context of the simulation. The best configuration (5 steps, 16 objects, road edge features, waypoints) achieves 0.88 V-Max score compared to 0.84-0.87 for other variants. All observation functions are illustrated in [Figure 4](#).

Encoders. We implemented the following encoder architectures with the Flax library ([Heek et al., 2024](#)): (1) **Latent-query** (LQ): inspired from ([Jaegle et al., 2021](#)), (2) **Latent-query hierarchical** (LQH) ([Bronstein et al., 2022](#)); (3) **Motion Transformer** (MTR) ([Shi et al., 2024](#)), (4) **Wayformer** ([Nayakanti et al., 2023](#)). For comparison, we also take an architecture that uses one multi-layer perceptron to encode each feature (road, trajectories...) separately: **MLP**. And an architecture that don’t use separate encodings: **None**.

The results are reported in [Table 2](#). The Latent-query (LQ) encoder performs best across all metrics, with other transformer-based architectures (LQH, MTR, Wayformer) showing similar results. The MLP encoder performs significantly worse, confirming our intuition about the relevance of transformers-based architectures for mid-to-end AD.

Table 2: Performance of encoder architectures, evaluated using the control configuration ([Table 1](#)).

Encoder	Accuracy \uparrow	Collision \downarrow	Off-road \downarrow	Progress \uparrow	V-Max Score \uparrow
None	69.95 \pm 1.72	25.13 \pm 1.40	4.51 \pm 0.23	87.26 \pm 3.43	0.53 \pm 0.01
MLP	87.54 \pm 0.48	9.24 \pm 0.31	2.92 \pm 0.24	104.03 \pm 2.93	0.68 \pm 0.01
LQ	97.45\pm0.22	1.73\pm0.15	0.67\pm0.07	155.38 \pm 5.34	0.87\pm0.01
LQH	96.28 \pm 0.35	2.52 \pm 0.12	1.02 \pm 0.20	162.23\pm12.26	0.84 \pm 0.01
MTR	95.94 \pm 0.24	2.42 \pm 0.24	1.42 \pm 0.38	154.88 \pm 2.53	0.84 \pm 0.01
Wayformer	96.08 \pm 0.42	2.70 \pm 0.41	0.99 \pm 0.11	161.94 \pm 7.20	0.84 \pm 0.00

Table 3: Performance achieved with different reward functions, evaluated using the control configuration ([Table 1](#)).

Reward	Accuracy \uparrow	Collision \downarrow	Off-road \downarrow	Progress \uparrow	V-Max Score \uparrow	Comfort \uparrow
Safety	96.73 \pm 0.57	2.21 \pm 0.29	0.90 \pm 0.34	78.66 \pm 3.18	0.67 \pm 0.04	57.24 \pm 1.34
Navigation	97.45\pm0.22	1.73\pm0.15	0.67 \pm 0.07	155.38 \pm 5.34	0.87 \pm 0.01	68.38 \pm 3.23
Behavior	97.25 \pm 0.31	1.92 \pm 0.20	0.64\pm0.05	139.91\pm4.06	0.88\pm0.02	70.86\pm2.66

Reward shaping. We developed a hierarchical reward structure consisting of three levels: Safety (basic collision and violation penalties), Navigation (adding progress incentives), and Behavior (in-

corporating comfort and speed compliance). To optimize reward weights, we conducted a grid search detailed in Table 12 in the Appendix.

$$\begin{aligned}
 r_t^{\text{safety}} &= -\mathbb{I}_{\text{collision}}(t) - \mathbb{I}_{\text{off-road}}(t) - \mathbb{I}_{\text{red-lightviolation}}(t) \\
 r_t^{\text{navigation}} &= r_t^{\text{safety}} - 0.2 \cdot \mathbb{I}_{\text{off-route}}(t) + 0.2 \cdot \mathbb{I}_{\text{progress}(t) > \text{progress}(t-1)} \\
 r_t^{\text{behavior}} &= r_t^{\text{navigation}} + 0.2 \cdot \mathbb{I}_{\text{comfort}}(t) - 0.1 \cdot \mathbb{I}_{v(t) > v_{\text{lim}}}
 \end{aligned}$$

Results in Table 3 show that the simple safety reward achieves high accuracy but results in a conservative policy, as reflected in its low progress score and overall V-Max score. Adding a navigation objective significantly improves progress while maintaining safety. The behavior reward adds an incentive to drive with respects to the speed limits and the comfort of vehicle.

Reward shaping introduces significant optimization challenges due to competing objectives between reward components. Maximizing progress along the route while maintaining driving comfort creates a fundamental trade-off that requires precise speed control. High comfort weights can produce overly cautious policies with reduced speed, while high progression weights lead to aggressive driving that reduces safety. This demonstrates the challenge of multi-objective optimization in autonomous driving: individual reward terms can conflict with each other, which might be inefficient in order to achieve the desired driving behavior.

Cross-dataset experiments. To validate ScenarioMax, we conduct a cross-dataset evaluation using agents trained on WOMD, nuPlan and combination of both. We report the results in Table 4. Agents trained on a single dataset achieve the best performance on their respective dataset but are outperformed when evaluated on an unseen dataset, highlighting the distribution shift between WOMD and nuPlan. In contrast, the model trained on both datasets is able to achieve competitive scores on both datasets, demonstrating improved generalization.

Table 4: **Preliminary study on generalization**, performance of an agent trained with the control configuration (Table 1) across nuPlan and WOMD datasets.

Train Dataset	Accuracy \uparrow	Collision \downarrow	Off-road \downarrow	Progress \uparrow	V-Max Score \uparrow
Evaluated on WOMD					
WOMD	97.45\pm0.22	1.73\pm0.15	0.67\pm0.07	155.38 \pm 5.34	0.87\pm0.01
nuPlan	76.81 \pm 1.53	7.30 \pm 1.25	1.30 \pm 0.28	163.77 \pm 8.15	0.67 \pm 0.01
WOMD + nuPlan	96.24 \pm 0.68	2.49 \pm 0.57	0.98 \pm 0.10	172.29\pm10.28	0.85 \pm 0.02
Evaluated on nuPlan					
WOMD	87.73 \pm 0.64	1.89\pm0.20	2.66 \pm 0.15	308.64 \pm 4.04	0.76 \pm 0.01
nuPlan	95.33 \pm 0.42	2.27 \pm 0.23	1.86 \pm 0.17	319.84\pm14.02	0.82\pm0.00
WOMD + nuPlan	95.38\pm0.71	2.04 \pm 0.40	1.98\pm0.21	315.95 \pm 15.77	0.82\pm0.01
Evaluated on WOMD + nuPlan					
WOMD	94.21 \pm 0.05	1.80\pm0.13	1.42 \pm 0.07	211.43 \pm 3.18	0.83 \pm 0.01
nuPlan	82.76 \pm 1.16	5.69 \pm 0.92	1.47 \pm 0.22	213.79 \pm 9.95	0.72 \pm 0.01
WOMD + nuPlan	95.97\pm0.68	2.35 \pm 0.51	1.29\pm0.13	218.33\pm11.96	0.84\pm0.01

5 Benchmark

We now present a benchmark of various approaches reimplemented in V-Max and report the best performance achieved by each method. Each method is executed across three random seeds, and the policy with the highest V-Max score is selected. The hyperparameters used for training are detailed in Appendix E.

5.1 Standard evaluation

We first compare approaches in the closed-loop non-reactive evaluation setup. PPO (Schulman et al., 2017) and SAC (Haarnoja et al., 2018) represent RL approaches, IDM (Treiber et al., 2000) and PDM (Dauner et al., 2023) represent rule-based approaches, more details on our implementation of PDM can be found in Appendix C. Note that we used *PDM-closed*, as it has the best scores in closed-loop evaluation and is purely rule-based. To represent IL, we implemented BC with a Wayformer architecture, similar to the one used in Waymax’s benchmark. Finally, we also reimplemented BC-SAC from (Lu et al., 2023), trained with the Safety reward.

Notably, PPO and BC train significantly faster than SAC, so we allow them to run for a higher number of timesteps to ensure fair comparisons at equivalent training durations. All the trainings and evaluations are made using WOMD.

Table 5: **V-Max benchmark** - evaluations on WOMD Valid, in the closed-loop non-reactive setting. A complete definition of metrics is provided in Appendix B.

Metric	Expert	IDM	PDM	SAC	BC+SAC	PPO	BC
Collision rate ↓	0.56	7.32	5.14	1.28	1.87	7.81	13.14
Offroad ↓	0.53	2.47	1.05	0.65	0.55	1.14	6.92
Red-light violation ↓	0.65	0.26	0.21	0.21	0.20	0.33	0.55
Progress ↑	98.94	144.44	140.58	170.85	90.98	189.52	86.87
At-fault collisions ↓	0.49	2.78	0.96	0.89	1.25	5.57	4.49
Making progress ↑	99.45	97.11	97.78	99.02	90.49	97.58	87.97
Driving direction compliance ↑	97.82	99.66	99.52	98.11	90.95	99.26	97.47
TTC within bound ↑	99.16	92.86	95.90	96.12	96.96	83.63	94.79
Progress along route ratio ↑	99.18	92.65	92.17	95.13	75.60	94.61	78.64
Speed limit compliance ↑	99.91	99.67	99.70	97.69	99.59	98.35	98.55
Multiple lanes score ↑	90.96	99.25	96.28	90.62	72.65	94.00	89.68
Comfort ↑	96.92	69.75	58.05	67.84	68.95	18.69	93.32
Accuracy ↑	98.26	89.96	93.60	97.86	97.40	90.75	79.42
Accuracy only at-fault ↑	98.33	94.51	97.78	98.25	98.02	88.07	92.99
V-Max score ↑	0.945	0.876	0.887	0.892	0.713	0.784	0.721

Results of the benchmark are reported in Table 5. The expert serves as an upper bound for accuracy, as failures only occur due to data errors. The poor performance of BC and PPO is likely due to limited hyperparameter tuning, as we concentrated our efforts on optimizing the SAC baseline.

We see that V-Max enabled us to train a highly performing driving policy using a standard RL algorithm, achieving an accuracy of 97.86% and strong performance across all sub-metrics of the V-Max score, using a single NVIDIA L4 GPU for 48 hours.

Although a RL planner can achieve expert-like accuracy in simulation, we see that our agents still fail to optimize for the comfort metric and the overall V-Max score, which could be investigate in future works in order to push a RL planner im the real world.

5.2 Reactive Evaluations

Table 6 presents our first experiment on using our multi-agent dataset for reactive evaluations. We see that safer models, like PDM and our SAC policy, allow to reduce collision rates for all policies. We specifically see how we managed to mitigate the weakness of log-replay by looking at the "Other collisions" raw, that account for collisions not imputable to the ego (rear-bumper collisions mainly).

Table 6: **Reactive evaluations** – Results on the first 1000 scenarios of our multi-agent dataset. Expert corresponds to the non-reactive setting.

Metric	IDM vs				PDM vs				SAC vs			
	Expert	IDM	PDM	SAC	Expert	IDM	PDM	SAC	Expert	IDM	PDM	SAC
Accuracy \uparrow	89.7	91.0	96.0	95.5	94.3	95.1	98.4	97.8	97.7	95.9	98.5	98.4
At-Fault Collision \downarrow	2.4	3.4	1.9	1.8	0.2	0.4	0.1	0.3	0.7	1.1	0.6	0.4
Other collisions \downarrow	6.10	4.10	0.20	1.00	4.70	3.20	0.30	0.90	0.60	2.00	0.10	0.50

5.3 Adversarial Evaluation

Table 7 showcases how ReGentS can be used to perform adversarial evaluations. The ReGentS agents are optimized to have a trajectory that aggressively cuts-in the ego. We see that the RL agent, thanks to its closed-loop training, is far more robust than the other counterparts.

However, as it could be expected, most of the collisions created by ReGentS are classified as not imputable to the ego, but they should be taken into account in this setting, as such aggressive behaviors can occur in real life.

Table 7: **Adversarial evaluations** – Results on the first 100 valid scenario for ReGentS.

Metric	IDM		PDM		SAC	
	Non Reactive	ReGentS	Non Reactive	ReGentS	Non Reactive	ReGentS
Accuracy \uparrow	92	43	96	53	98	76
At-Fault Collision \downarrow	1	5	0	11	2	11
Other collisions \downarrow	5	51	4	33	0	10

6 Conclusion

In this work, we introduced V-Max, a framework designed to facilitate Reinforcement Learning research for mid-to-end Autonomous Driving (AD). Built on Waymax, V-Max extends its capabilities with a JAX-based RL training pipeline, multi-dataset accelerated simulation, and comprehensive evaluation tools. Using these tools, we trained high-performing SAC agents, showing how V-Max can help advance RL research for AD. To further support progress in this area, we ensure full reproducibility by publishing our framework and benchmarks.

While V-Max provides a foundation for AD research, rigorously evaluating driving policies remains an open challenge. Current evaluation protocols (in V-Max and the frameworks discussed in [section 2](#)) average scenario metrics across the entire validation dataset. However, driving difficulty follows a long-tail distribution ([Makansi et al., 2021](#); [Bronstein et al., 2023](#)), where most scenarios are easily solvable while a small subset presents significant challenges. Developing benchmarks that explicitly account for this distribution would enable a more rigorous assessment of policy robustness.

Additionally, further research on adversarial scenario generation, could enable deeper robustness assessment of driving policies. ReGentS is a good starting point, diffuser-based methods is an alternative approach ([Pronovost et al., 2023](#)). Similarly, the development of more realistic simulation agents, as explored in the *Waymo Open Sim Agents Challenge* (WOSAC, ([Montali et al., 2023](#))) could improve realism of closed-loop simulators, reducing the reliance on non-reactive evaluation.

References

Mayank Bansal, Alex Krizhevsky, and Abhijit Ogale. ChauffeurNet: Learning to Drive by Imitating the Best and Synthesizing the Worst. In *Proceedings of Robotics: Science*

- and Systems*, volume 15, June 2019. ISBN 978-0-9923747-5-4. URL <https://www.roboticsproceedings.org/rss15/p31.html>.
- James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018. URL <http://github.com/jax-ml/jax>.
- Eli Bronstein, Mark Palatucci, Dominik Notz, Brandyn White, Alex Kuefler, Yiren Lu, Supratik Paul, Payam Nikdel, Paul Mougin, Hongge Chen, Justin Fu, Austin Abrams, Punit Shah, Evan Racah, Benjamin Frenkel, Shimon Whiteson, and Dragomir Anguelov. Hierarchical Model-Based Imitation Learning for Planning in Autonomous Driving. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 8652–8659, October 2022. DOI: 10.1109/IROS47612.2022.9981695. URL <https://ieeexplore.ieee.org/document/9981695>. ISSN: 2153-0866.
- Eli Bronstein, Sirish Srinivasan, Supratik Paul, Aman Sinha, Matthew O’Kelly, Payam Nikdel, and Shimon Whiteson. Embedding Synthetic Off-Policy Experience for Autonomous Driving via Zero-Shot Curricula. In *Proceedings of The 6th Conference on Robot Learning*, pp. 188–198. PMLR, March 2023. URL <https://proceedings.mlr.press/v205/bronstein23a.html>. ISSN: 2640-3498.
- Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuScenes: A Multimodal Dataset for Autonomous Driving. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 11618–11628, Seattle, WA, USA, June 2020. IEEE. ISBN 978-1-72817-168-5. DOI: 10.1109/CVPR42600.2020.01164. URL <https://ieeexplore.ieee.org/document/9156412/>.
- Holger Caesar, Juraj Kabzan, Kok Seang Tan, Whye Kit Fong, Eric Wolff, Alex Lang, Luke Fletcher, Oscar Beijbom, and Sammy Omari. NuPlan: A closed-loop ML-based planning benchmark for autonomous vehicles, June 2021. URL <http://arxiv.org/abs/2106.11810>. arXiv:2106.11810 [cs].
- Li Chen, Penghao Wu, Kashyap Chitta, Bernhard Jaeger, Andreas Geiger, and Hongyang Li. End-to-end Autonomous Driving: Challenges and Frontiers. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 1–20, 2024. ISSN 1939-3539. DOI: 10.1109/TPAMI.2024.3435937. URL <https://ieeexplore.ieee.org/document/10614862/?arnumber=10614862>. Conference Name: IEEE Transactions on Pattern Analysis and Machine Intelligence.
- Marco Cusumano-Towner, David Hafner, Alex Hertzberg, Brody Huval, Aleksei Petrenko, Eugene Vinitsky, Erik Wijmans, Taylor Killian, Stuart Bowers, Ozan Sener, Philipp Krähenbühl, and Vladlen Koltun. Robust Autonomy Emerges from Self-Play, February 2025. URL <http://arxiv.org/abs/2502.03349>. arXiv:2502.03349 [cs].
- Daniel Dauner, Marcel Hallgarten, Andreas Geiger, and Kashyap Chitta. Parting with Misconceptions about Learning-based Vehicle Motion Planning. In *Proceedings of The 7th Conference on Robot Learning*, pp. 1268–1281. PMLR, December 2023. URL <https://proceedings.mlr.press/v229/dauner23a.html>. ISSN: 2640-3498.
- Jonas Degraeve, Federico Felici, Jonas Buchli, Michael Neunert, Brendan Tracey, Francesco Carpanese, Timo Ewalds, Roland Hafner, Abbas Abdolmaleki, Diego de las Casas, Craig Donner, Leslie Fritz, Cristian Galperti, Andrea Huber, James Keeling, Maria Tsimpoukelli, Jackie Kay, Antoine Merle, Jean-Marc Moret, Seb Noury, Federico Pesamosca, David Pfau, Olivier Sauter, Cristian Sommariva, Stefano Coda, Basil Duval, Ambrogio Fasoli, Pushmeet Kohli, Kory Kavukcuoglu, Demis Hassabis, and Martin Riedmiller. Magnetic control of tokamak plasmas

- through deep reinforcement learning. *Nature*, 602(7897):414–419, February 2022. ISSN 1476-4687. DOI: 10.1038/s41586-021-04301-9. URL <https://www.nature.com/articles/s41586-021-04301-9>. Publisher: Nature Publishing Group.
- Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. CARLA: An Open Urban Driving Simulator. In *Proceedings of the 1st Annual Conference on Robot Learning*, pp. 1–16. PMLR, October 2017. URL <https://proceedings.mlr.press/v78/dosovitskiy17a.html>. ISSN: 2640-3498.
- Scott Ettinger, Shuyang Cheng, Benjamin Caine, Chenxi Liu, Hang Zhao, Sabeek Pradhan, Yuning Chai, Ben Sapp, Charles R. Qi, Yin Zhou, Zoey Yang, Aurélien Chouard, Pei Sun, Jiquan Ngiam, Vijay Vasudevan, Alexander McCauley, Jonathon Shlens, and Dragomir Anguelov. Large Scale Interactive Motion Forecasting for Autonomous Driving: The Waymo Open Motion Dataset. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 9710–9719, 2021. URL https://openaccess.thecvf.com/content/ICCV2021/html/Ettinger_Large_Scale_Interactive_Motion_Forecasting_for_Autonomous_Driving_The_Waymo_ICCV_2021_paper.html.
- C. Daniel Freeman, Erik Frey, Anton Raichuk, Sertan Girgin, Igor Mordatch, and Olivier Bachem. Brax - A Differentiable Physics Engine for Large Scale Rigid Body Simulation, 2021. URL <http://github.com/google/brax>.
- Clémence Grislain, Risto Vuorio, Cong Lu, and Shimon Whiteson. IGDriSim: A Benchmark for the Imitation Gap in Autonomous Driving, November 2024. URL <http://arxiv.org/abs/2411.04653>. arXiv:2411.04653 [cs].
- Cole Gulino, Justin Fu, Wenjie Luo, George Tucker, Eli Bronstein, Yiren Lu, Jean Harb, Xinlei Pan, Yan Wang, Xiangyu Chen, John Co-Reyes, Rishabh Agarwal, Rebecca Roelofs, Yao Lu, Nico Montali, Paul Mouglin, Zoey Yang, Brandyn White, Aleksandra Faust, Rowan McAllister, Dragomir Anguelov, and Benjamin Sapp. Waymax: An Accelerated, Data-Driven Simulator for Large-Scale Autonomous Driving Research. *Advances in Neural Information Processing Systems*, 36:7730–7742, December 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/hash/1838feeb71c4b4ea524d0df2f7074245-Abstract-Datasets_and_Benchmarks.html.
- Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor. In *Proceedings of the 35th International Conference on Machine Learning*, pp. 1861–1870. PMLR, July 2018. URL <https://proceedings.mlr.press/v80/haarnoja18b.html>. ISSN: 2640-3498.
- Jonathan Heek, Anselm Levskaya, Avital Oliver, Marvin Ritter, Bertrand Rondepierre, Andreas Steiner, and Marc van Zee. Flax: A neural network library and ecosystem for JAX, 2024. URL <http://github.com/google/flax>.
- Anthony Hu, Lloyd Russell, Hudson Yeo, Zak Murez, George Fedoseev, Alex Kendall, Jamie Shotton, and Gianluca Corrado. GAIA-1: A Generative World Model for Autonomous Driving, September 2023. URL <http://arxiv.org/abs/2309.17080>. arXiv:2309.17080 [cs].
- Andrew Jaegle, Felix Gimeno, Andy Brock, Oriol Vinyals, Andrew Zisserman, and Joao Carreira. Perceiver: General Perception with Iterative Attention. In *Proceedings of the 38th International Conference on Machine Learning*, pp. 4651–4664. PMLR, July 2021. URL <https://proceedings.mlr.press/v139/jaegle21a.html>. ISSN: 2640-3498.
- Maximilian Jaritz, Raoul De Charette, Marin Toromanoff, Etienne Perot, and Fawzi Nashashibi. End-to-End Race Driving with Deep Reinforcement Learning. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 2070–2075, Brisbane, QLD, May 2018. IEEE. DOI: 10.1109/icra.2018.8460934. URL <https://ieeexplore.ieee.org/document/8460934/>.

- Napat Karnchanachari, Dimitris Geromichalos, Kok Seang Tan, Nanxiang Li, Christopher Eriksen, Shakiba Yaghoubi, Noushin Mehdipour, Gianmarco Bernasconi, Whye Kit Fong, Yiluan Guo, and Holger Caesar. Towards learning-based planning: The nuPlan benchmark for real-world autonomous driving. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 629–636, May 2024. DOI: 10.1109/ICRA57147.2024.10610077. URL <https://ieeexplore.ieee.org/document/10610077?arnumber=10610077>.
- Saman Kazemkhani, Aarav Pandya, Daphne Cornelisse, Brennan Shacklett, and Eugene Vinit-sky. GPUDrive: Data-driven, multi-agent driving simulation at 1 million FPS. In *The Thirteenth International Conference on Learning Representations*, October 2024. URL <https://openreview.net/forum?id=ERv8ptegFi>.
- Alex Kendall, Jeffrey Hawke, David Janz, Przemyslaw Mazur, Daniele Reda, John-Mark Allen, Vinh-Dieu Lam, Alex Bewley, and Amar Shah. Learning to Drive in a Day. In *2019 International Conference on Robotics and Automation (ICRA)*, pp. 8248–8254, Montreal, QC, Canada, May 2019. IEEE. ISBN 978-1-5386-6027-0. DOI: 10.1109/ICRA.2019.8793742. URL <https://ieeexplore.ieee.org/document/8793742/>.
- B Ravi Kiran, Ibrahim Sobh, Victor Talpaert, Patrick Mannion, Ahmad A. Al Sallab, Senthil Yogamani, and Patrick Pérez. Deep Reinforcement Learning for Autonomous Driving: A Survey. *IEEE Transactions on Intelligent Transportation Systems*, 23(6):4909–4926, June 2022. ISSN 1558-0016. DOI: 10.1109/TITS.2021.3054625. URL <https://ieeexplore.ieee.org/abstract/document/9351818>. Conference Name: IEEE Transactions on Intelligent Transportation Systems.
- Quanyi Li, Zhenghao Peng, Lan Feng, Qihang Zhang, Zhenghai Xue, and Bolei Zhou. MetaDrive: Composing Diverse Driving Scenarios for Generalizable Reinforcement Learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(3):3461–3475, March 2023a. ISSN 1939-3539. DOI: 10.1109/TPAMI.2022.3190471. URL <https://ieeexplore.ieee.org/document/9829243>. Conference Name: IEEE Transactions on Pattern Analysis and Machine Intelligence.
- Quanyi Li, Zhenghao (Mark) Peng, Lan Feng, Zhizheng Liu, Chenda Duan, Wenjie Mo, and Bolei Zhou. ScenarioNet: Open-Source Platform for Large-Scale Traffic Scenario Simulation and Modeling. *Advances in Neural Information Processing Systems*, 36:3894–3920, December 2023b. URL https://proceedings.neurips.cc/paper_files/paper/2023/hash/0c26a501df8fb919a0350e2df06b5d39-Abstract-Datasets_and_Benchmarks.html.
- Yiren Lu, Justin Fu, George Tucker, Xinlei Pan, Eli Bronstein, Rebecca Roelofs, Benjamin Sapp, Brandyn White, Aleksandra Faust, Shimon Whiteson, Dragomir Anguelov, and Sergey Levine. 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)*, pp. 7553–7560, October 2023. DOI: 10.1109/IROS55552.2023.10342038. ISSN: 2153-0866.
- Osama Makansi, Özgün Çiçek, Yassine Marrakchi, and Thomas Brox. On Exposing the Challenging Long Tail in Future Prediction of Traffic Actors. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 13147–13157, 2021. URL https://openaccess.thecvf.com/content/ICCV2021/html/Makansi_On_Exposing_the_Challenging_Long_Tail_in_Future_Prediction_of_ICCV_2021_paper.html.
- Nico Montali, John Lambert, Paul Mougin, Alex Kuefler, Nicholas Rhinehart, Michelle Li, Cole Gulino, Tristan Emrich, Zoey Zeyu Yang, Shimon Whiteson, Brandyn White, and Dragomir Anguelov. The Waymo Open Sim Agents Challenge. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2023. URL <https://openreview.net/forum?id=5FnttJZQFn>.

- Nigamaa Nayakanti, Rami Al-Rfou, Aurick Zhou, Kratarth Goel, Khaled S. Refaat, and Benjamin Sapp. Wayformer: Motion Forecasting via Simple & Efficient Attention Networks. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 2980–2987, London, United Kingdom, May 2023. IEEE. ISBN 9798350323658. DOI: 10.1109/ICRA48891.2023.10160609. URL <https://ieeexplore.ieee.org/document/10160609/>.
- Ethan Pronovost, Meghana Reddy Ganesina, Noureldin Hendy, Zeyu Wang, Andres Morales, Kai Wang, and Nick Roy. Scenario Diffusion: Controllable Driving Scenario Generation With Diffusion. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neural Information Processing Systems*, volume 36, pp. 68873–68894. Curran Associates, Inc., 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/file/d95cb79a3421e6d9b6c9a9008c4d07c5-Paper-Conference.pdf.
- Aravind Rajeswaran, Vikash Kumar, Abhishek Gupta, Giulia Vezzani, John Schulman, Emanuel Todorov, and Sergey Levine. Learning Complex Dexterous Manipulation with Deep Reinforcement Learning and Demonstrations. In *Robotics: Science and Systems XIV*, June 2018. URL <https://www.roboticsproceedings.org/rss14/p49.html>.
- Oliver Scheel, Luca Bergamini, Maciej Wolczyk, Błażej Osipiński, and Peter Ondruska. Urban Driver: Learning to Drive from Real-world Demonstrations Using Policy Gradients. In *Proceedings of the 5th Conference on Robot Learning*, pp. 718–728. PMLR, January 2022. URL <https://proceedings.mlr.press/v164/scheel22a.html>. ISSN: 2640-3498.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal Policy Optimization Algorithms, August 2017. URL <http://arxiv.org/abs/1707.06347>. arXiv:1707.06347 [cs].
- Brennan Shacklett, Luc Guy Rosenzweig, Zhiqiang Xie, Bidipta Sarkar, Andrew Szot, Erik Wijmans, Vladlen Koltun, Dhruv Batra, and Kayvon Fatahalian. An Extensible, Data-Oriented Architecture for High-Performance, Many-World Simulation. *ACM Trans. Graph.*, 42(4), 2023. DOI: 10.1145/3592427.
- Shaoshuai Shi, Li Jiang, Dengxin Dai, and Bernt Schiele. MTR++: Multi-Agent Motion Prediction With Symmetric Scene Modeling and Guided Intention Querying. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(5):3955–3971, May 2024. ISSN 1939-3539. DOI: 10.1109/TPAMI.2024.3352811. URL <https://ieeexplore-ieee-org.minesparis-psl.idm.oclc.org/document/10398503>. 1 citations (Crossref) [2024-06-06] Conference Name: IEEE Transactions on Pattern Analysis and Machine Intelligence.
- Matthijs T. J. Spaan. Partially Observable Markov Decision Processes. In Marco Wiering and Martijn van Otterlo (eds.), *Reinforcement Learning: State of the Art*, pp. 387–414. Springer Verlag, 2012.
- Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning, second edition: An Introduction*. Bradford Books, Cambridge, Massachusetts, 2nd edition edition, November 2018. ISBN 978-0-262-03924-6. URL <http://incompleteideas.net/book/the-book-2nd.html>.
- Marin Toromanoff, Emilie Wirbel, and Fabien Moutarde. End-to-End Model-Free Reinforcement Learning for Urban Driving Using Implicit Affordances. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7153–7162, 2020. URL https://openaccess.thecvf.com/content_CVPR_2020/html/Toromanoff_End-to-End_Model-Free_Reinforcement_Learning_for_Urban_Driving_Using_Implicit_Affordances_CVPR_2020_paper.html.
- Martin Treiber, Ansgar Hennecke, and Dirk Helbing. Congested traffic states in empirical observations and microscopic simulations. *Physical Review E*, 62(2):1805–1824, August 2000. ISSN 1063-651X, 1095-3787. DOI: 10.1103/PhysRevE.62.1805. URL <https://link.aps.org/doi/10.1103/PhysRevE.62.1805>.

- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is All you Need. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL https://proceedings.neurips.cc/paper_files/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html.
- Eugene Vitisnky, Nathan Lichtlé, Xiaomeng Yang, Brandon Amos, and Jakob Foerster. Nocturne: a scalable driving benchmark for bringing multi-agent learning one step closer to the real world, February 2023. URL <http://arxiv.org/abs/2206.09889>. arXiv:2206.09889 [cs].
- Aaron Walsman, Muru Zhang, Sanjiban Choudhury, Dieter Fox, and Ali Farhadi. Impossibly Good Experts and How to Follow Them. In *The Eleventh International Conference on Learning Representations*, September 2022. URL https://openreview.net/forum?id=sciA_xgYofB.
- Benjamin Wilson, William Qi, Tanmay Agarwal, John Lambert, Jagjeet Singh, Siddhesh Khandelwal, Bowen Pan, Ratnesh Kumar, Andrew Hartnett, Jhony Kaesemodel Pontes, Deva Ramanan, Peter Carr, and James Hays. Argoverse 2: Next Generation Datasets for Self-Driving Perception and Forecasting. *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, 1, December 2021. URL https://datasets-benchmarks-proceedings.neurips.cc/paper_files/paper/2021/hash/4734ba6f3de83d861c3176a6273cac6d-Abstract-round2.html.
- Peter R. Wurman, Samuel Barrett, Kenta Kawamoto, James MacGlashan, Kaushik Subramanian, Thomas J. Walsh, Roberto Capobianco, Alisa Devlic, Franziska Eckert, Florian Fuchs, Leilani Gilpin, Piyush Khandelwal, Varun Kompella, HaoChih Lin, Patrick MacAlpine, Declan Oller, Takuma Seno, Craig Sherstan, Michael D. Thomure, Houmeh Aghabozorgi, Leon Barrett, Rory Douglas, Dion Whitehead, Peter Dürr, Peter Stone, Michael Spranger, and Hiroaki Kitano. Out-racing champion Gran Turismo drivers with deep reinforcement learning. *Nature*, 602(7896): 223–228, February 2022. ISSN 0028-0836, 1476-4687. DOI: 10.1038/s41586-021-04357-7. URL <https://www.nature.com/articles/s41586-021-04357-7>.
- Xintao Yan, Erdao Liang, Jiawei Wang, Haojie Zhu, and Henry X. Liu. Improving Traffic Signal Data Quality for the Waymo Open Motion Dataset, June 2025. URL <http://arxiv.org/abs/2506.07150>. arXiv:2506.07150 [cs].
- Yuan Yin, Pegah Khayatan, Eloi Zablocki, Alexandre Boulch, and Matthieu Cord. ReGentS: Real-World Safety-Critical Driving Scenario Generation Made Stable. In *ECCV 2024 Workshop on Multimodal Perception and Comprehension of Corner Cases in Autonomous Driving*, September 2024. URL <https://openreview.net/forum?id=dJqcdUgEdw>.

Supplementary Materials

The following content was not necessarily subject to peer review.

A Performances

We run a performance comparison between Waymax and V-Max, measuring steps per second (SPS) when replaying expert trajectories. The benchmark task involves converting expert steering and acceleration commands through Waymax’s InvertibleBicycle model while incrementally adding V-Max’s observation extraction and metrics computation within the environment step. All measurements were conducted on a single Nvidia L4 GPU with 16 CPUs and 64GB RAM. For fair comparison, Waymax’s default observation function was overridden to return a dummy array, as the full state observation typically reduces performance by 30-40%.

Table 8: SPS comparison between Waymax and V-Max features.

	Waymax	V-Max
<i>Single scenario, 64 objects</i>		
Env	623.08	615.63
Env + Obs	N/A	543.55
Env + Metrics	N/A	367.46
Env + Metrics + Obs	N/A	342.72
<i>Batch of 32 scenarios, 64 objects</i>		
Env	10109.73	9942.12
Env + Obs	N/A	8210.96
Env + Metrics	N/A	5387.38
Env + Metrics + Obs	N/A	4609.01

B Metrics

Waymax Metrics

- Offroad: Binary flag indicating whether the SDC left the drivable area at any point in the scenario.
- Collision (overlap): Binary flag indicating whether the SDC collided with another object at any point in the scenario.
- Wrongway: Waymax-specific metric based on SDC paths, indicating whether the SDC is more than 3.5 meters away from the closest SDC path.
- Offroute: Similar to wrongway but with respect to on-route paths, which are the SDC paths that contain the expert’s logged trajectory.
- `sdc_progress`: computes how much the SDC progressed along on-route paths, and divides it by the distance the expert did cover on those paths. Can be greater than 1.

For example, if the SDC takes a right turn at an intersection while the expert proceeded straight, the SDC will be considered offroute but not wrongway.

nuPlan Metrics

- Progress along route: same definition as Waymax, but capped to 1.
- Making progress: binary flag indicating if progress along route is superior to 20%.
- At-fault collision: Binary flag following nuPlan’s criteria for assigning collision responsibility:

- Collisions with stopped vehicles are always at-fault.
- If the SDC is stopped, it is never at-fault.
- If the SDC is occupying multiple lanes, it is at-fault.
- Rear-bumper collisions are not at-fault, while front-bumper collisions are at-fault.
- TTC within bound: indicates if the time-to-collision (ttc) with ahead vehicles remain superior to 0.95s.
- Speed limit compliance: defined by nuPlan as:

$$\text{nuplan_speed_compliance} = \max \left(0.0, 1.0 - \frac{\sum_t \text{speed_violation}_t \cdot \Delta t}{\max(\text{threshold}, 1e - 3) \cdot T} \right), \quad (1)$$

where speed_violation_t is the magnitude of overspeeding at timestep t , Δt is the time step duration, and T is the total scenario duration.

- Driving direction compliance: Based on distance traveled into oncoming traffic. We check if the vehicle is effectively driving into oncoming traffic lanes using the road information, rather than SDC paths.
 - Score = 1.0 if wrong-way distance $\leq 2.0\text{m}$.
 - Score = 0.5 if wrong-way distance is between 2.0m and 6.0m.
 - Score = 0 if wrong-way distance $> 6.0\text{m}$.
- Comfort: binary indicating if the trajectory is comfortable based on jerk, acceleration and yaw rates.

V-Max Metrics

- Red-light violation: Binary flag indicating whether the SDC crossed an intersection while the traffic light was red.
- Time spent on multiple lanes: Evaluated based on roadgraph information rather than SDC paths. We added this metric to encourage agent to remain on one lane, to set the thresholds, we looked at the expert trajectories.
 - Score = 1.0 if time spent on multiple lanes $\leq 3.4\text{s}$.
 - Score = 0.5 if time spent on multiple lanes is between 3.4s and 5.7s.
 - Score = 0 if time spent on multiple lanes $> 5.7\text{s}$.

Table 9: Metrics and their weights in nuPlan aggregate score and V-Max aggregate score. [†] indicates metrics that appear only in the V-Max aggregate score.

Metric name	Multiplier weight	Average weight
No at-fault collisions	{0, 1}	-
Offroad	{0, 1}	-
Red-light violation [†]	{0, 1}	-
Making progress	{0, 1}	-
Driving direction compliance	{0, 0.5, 1}	-
TTC within bound	-	5
Progress along route ratio	-	5
Speed limit compliance	-	4
Multiple lanes score [†]	-	3
Comfort	-	2

C Reimplementing PDM in V-Max

We first reimplement IDM (Treiber et al., 2000), using the path target as a reference path, and considering the first object or red light along this path as the leading agent. This baseline is already quite strong. Then PDM-Closed (Dauner et al., 2023), generates multiple trajectories by using PDM with the target path and two offset paths (± 1.0 meter) and five desired speeds, namely $\{20\%, 40\%, 60\%, 80\%, 100\%\}$ of the speed limit.

External objects are forecasted using a constant-speed and constant-heading assumption, and each candidate trajectory is scored based on this forecast. The first action of the highest-scoring trajectory is then executed, and the process repeats at each timestep, which might explain the low comfort metric observed in Table 5.

Following Dauner et al. (2023), we introduce an emergency braking maneuver when the estimated time-to-collision (TTC) falls below 2.0s. While this improves PDM’s evaluation scores, it also makes the behavior less realistic.

D Observation functions

We present and details the observation functions used in the experiments. We tested different observations represented in Table 10. For object features, the module extracts waypoints, velocity, yaw, size, and validity information from the n closest objects to the ego vehicle. Roadgraph features include waypoints, direction, speed limits, and validity data, sampled at 2-meter intervals within a maximum distance of 50 meters. The module selects the top 200 roadgraph elements within a bounding box extending 50 meters forward, 5 meters back, and 20 meters to each side. Traffic light features capture waypoints, state, and validity information for the 5 closest traffic lights. Finally, path target features consist of waypoints for the target trajectory, sampled from the SDC paths.

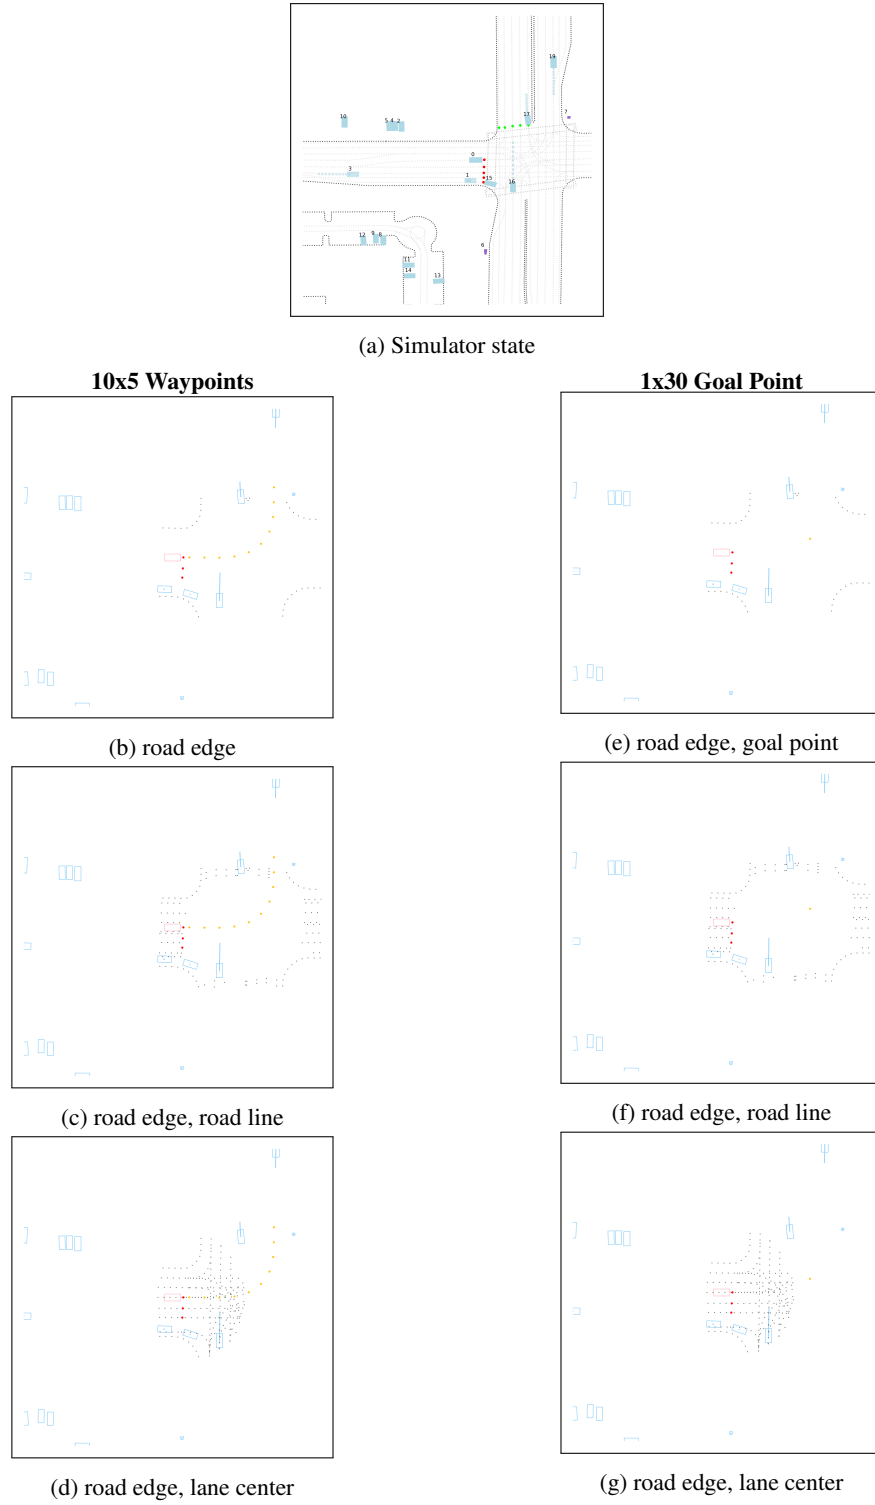


Figure 4: Illustration of the observation functions tested in [section 4](#).

Table 10: Observation configuration hyperparameter sweep results on WOMD Valid dataset. All experiments use SAC with LQ encoder and navigation reward (progression=0.2, off-route=-0.2).

Observation Config				Road Features	Safety Metrics				Performance	
Steps	Obj	Path Pts	Gap		Collision ↓	Offroad ↓	Red-light ↓	Progress ↑	Accuracy ↑	V-Max ↑
1	8	1	30	road_edge	1.54	0.62	0.27	165.65	97.11	0.81
1	8	1	30	road_edge+lane	1.79	1.25	0.43	179.84	96.03	0.82
1	8	1	30	road_edge+line	1.48	1.08	0.35	151.91	96.37	0.79
1	8	10	5	road_edge	1.54	0.63	0.21	169.1	96.97	0.87
1	8	10	5	road_edge+lane	1.46	1.31	0.15	173.9	96.54	0.84
1	8	10	5	road_edge+line	1.41	0.92	0.14	161.0	96.92	0.84
1	16	1	30	road_edge	1.56	0.73	0.20	163.91	96.80	0.79
1	16	1	30	road_edge+lane	1.08	0.71	0.19	142.1	97.53	0.87
1	16	1	30	road_edge+line	1.33	0.81	0.17	134.45	96.88	0.80
1	16	10	5	road_edge	1.28	0.64	0.29	160.5	97.33	0.87
1	16	10	5	road_edge+lane	1.60	0.87	0.15	161.9	96.55	0.87
1	16	10	5	road_edge+line	1.61	0.86	0.25	168.06	96.66	0.84
3	8	1	30	road_edge	1.38	0.61	0.15	161.93	97.30	0.81
3	8	1	30	road_edge+lane	1.09	0.82	0.23	150.5	97.28	0.84
3	8	1	30	road_edge+line	1.64	1.39	0.39	145.67	95.93	0.78
3	8	10	5	road_edge	1.84	0.62	0.13	172.5	96.51	0.86
3	8	10	5	road_edge+lane	1.92	1.33	0.26	174.5	96.05	0.85
3	8	10	5	road_edge+line	1.23	0.84	0.35	170.8	97.13	0.84
3	16	1	30	road_edge	1.00	1.06	0.17	154.22	97.10	0.81
3	16	1	30	road_edge+lane	1.49	1.26	0.20	151.0	96.46	0.85
3	16	1	30	road_edge+line	1.56	1.32	0.11	137.78	96.30	0.79
3	16	10	5	road_edge	1.14	0.69	0.12	139.2	97.09	0.87
3	16	10	5	road_edge+lane	1.03	0.94	0.10	142.2	97.18	0.87
3	16	10	5	road_edge+line	1.17	0.85	0.22	152.5	97.12	0.85
5	8	1	30	road_edge	1.70	0.68	0.28	167.38	96.82	0.80
5	8	1	30	road_edge+lane	1.52	1.33	0.19	169.8	96.38	0.83
5	8	1	30	road_edge+line	1.53	1.17	0.15	162.69	96.64	0.82
5	8	10	5	road_edge	1.23	0.92	0.21	166.4	97.04	0.87
5	8	10	5	road_edge+lane	1.59	0.87	0.15	155.25	96.78	0.81
5	8	10	5	road_edge+line	1.26	1.07	0.15	144.97	96.74	0.83
5	16	1	30	road_edge	1.35	0.62	0.25	158.57	97.26	0.82
5	16	1	30	road_edge+lane	1.45	1.22	0.24	160.5	96.56	0.84
5	16	1	30	road_edge+line	1.15	0.98	0.15	147.78	97.03	0.83
5	16	10	5	road_edge	1.09	0.73	0.17	153.6	97.24	0.88
5	16	10	5	road_edge+lane	1.55	1.20	0.18	181.9	96.49	0.84
5	16	10	5	road_edge+line	1.34	1.09	0.20	153.5	96.80	0.86

E Training hyperparameters

We details hyperparameters used for all learning algorithms in the experiments. Every algorithm are trained in one single NVIDIA L4.

Table 11: Algorithms hyperparameters used in experiments

Parameter	BC	SAC	BC+SAC	PPO
Total Steps	200M	25M	25M	200M
Num Envs	16	16	16	16
Unroll Length	80	1	1(RL)/80(IL)	40
IL frequency	-	-	8	-
Learning Rate	1e-4	1e-4	1e-4/5e-5	1e-4
Batch Size	64	64	64	512
Num Minibatches	-	-	-	16
Discount	-	0.99	0.99	0.99
Entropy α	-	0.2	0.2	0.2
Updates/step	32	4	8	4
Buffer size	-	1000000	1000000	-
Learning start	-	50000	50000	-
GAE factor	-	-	-	0.95
Clip ϵ	-	-	-	0.2
Loss function	Log_prob	-	-	-

Table 12: Reward shaping hyperparameter sweep results on WOMD Valid dataset. All experiments use SAC with LQ encoder, 5 observation steps, 8 objects, and road_edge features.

Prog.	Reward Weights			Safety Metrics			Progress	Performance	
	Off-route	Comfort	Overspeed	Collision	Offroad	Red-light		Accuracy	V-Max
				↓	↓	↓	↑	↑	↑
0.1	-0.3	0	0	1.08	0.57	0.15	113.2	97.64	0.86
0.1	-0.6	0	0	1.22	0.63	0.13	111.1	97.06	0.85
0.1	-1.0	0	0	0.94	0.55	0.08	100.1	97.50	0.83
0.2	-0.2	0	0	1.18	0.67	0.15	155.4	97.45	0.87
0.3	-0.3	0	0	1.50	0.75	0.28	185.6	97.11	0.87
0.3	-0.6	0	0	1.44	0.67	0.32	184.7	96.94	0.87
0.3	-0.6	0	0	1.87	0.74	0.25	187.2	96.49	0.85
0.3	-1.0	0	0	1.94	0.98	0.23	181.3	96.16	0.84
0.5	-0.6	0	0	1.83	0.85	0.24	197.3	96.49	0.85
0.5	-0.6	0	0	2.26	0.71	0.18	190.1	95.42	0.81
0.5	-1.0	0	0	1.69	0.71	0.21	192.8	96.43	0.84
0.2	-0.2	0.1	-0.1	1.49	0.55	0.19	164.3	97.19	0.86
0.2	-0.2	0.1	-0.2	1.37	0.71	0.14	168.5	97.01	0.86
0.2	-0.2	0.2	-0.1	1.06	0.64	0.20	139.9	97.25	0.88
0.2	-0.2	0.2	-0.2	1.26	0.62	0.27	162.2	97.26	0.88
0.3	-0.2	0.1	-0.1	1.86	0.66	0.27	176.2	96.44	0.86
0.3	-0.2	0.2	-0.1	2.22	0.69	0.42	202.5	96.24	0.80
0.3	-0.2	0.2	-0.2	1.93	0.66	0.26	188.0	96.25	0.83

Table 13: Encoders and decoders hyperparameters used in experiments

MLP policy decoder		
Parameter	Value	Description
Layer sizes	[256, 64, 32]	Number and size of layers
MLP value decoder		
Layer sizes	[256, 64, 32]	Number and size of layers
LQH encoder		
Embedding sizes	[256, 256]	Embedding layer sizes
dk	64	Dense encoder dimension
num_latents	16	Learnable latent size
cross_num_heads	2	Attention heads count
cross_head_feat	16	Features per attention head
ff_mult	2	Feedforward layer multiplier
LQ encoder		
Embedding sizes	[256, 256]	Embedding layer sizes
depth	4	Attention layers count
num_latents	16	Learnable latent size
num_self_heads	2	Self-attention heads
self_head_feat	16	Features per self-attention
cross_num_heads	2	Cross-attention heads
cross_head_feat	16	Features per cross-attention
ff_mult	2	Feedforward layer multiplier
MTR encoder		
Embedding sizes	[256, 256]	Embedding layer sizes
dk	64	Dense encoder dimension
num_latents	16	Learnable latent size
num_self_heads	2	Self-attention heads
self_head_feat	16	Features per self-attention
ff_mult	2	Feedforward layer multiplier
k	8	Nearest objects in attention
Wayformer encoder		
Embedding sizes	[256, 256]	Embedding layer sizes
dk	64	Dense encoder dimension
num_latents	16	Learnable latent size
num_self_heads	2	Self-attention heads
self_head_feat	16	Features per self-attention
depth	2	Attention layers count
ff_mult	2	Feedforward layer multiplier
fusion_type	late	Fusion mechanism type



(a) Scene-centered view of the initial scenario.



(b) Scene-centered view after applying ReGentS. The adversarial agent (in red) follows a trajectory designed to collide with the SDC agent (in pink). The RL agent adapts by deviating from its path to avoid the collision and then re-centering itself in the lane once the adversarial object has passed.

Figure 5: Visualization of adversarial evaluation of our best RL agent using ReGentS (Yin et al., 2024)