

Building reliable sim driving agents by scaling self-play

Daphne Cornelisse, Aarav Pandya, Kevin Joseph, Joseph Suárez, and Eugene Vinitsky

This is the author's version of a work that can be found published online by arXiv on February 20, 2025, and last revised on May 19, 2025 (v3). <https://doi.org/10.48550/arXiv.2502.14706>.

Cornelisse, D., Pandya, A., Joseph, K., Suárez, J., and Vinitsky, E. (2025). Building reliable sim driving agents by scaling self-play. <https://doi.org/10.48550/arXiv.2502.14706>.

Building reliable sim driving agents by scaling self-play

Daphne Cornelisse^{*1}, Aarav Pandya¹, Kevin Joseph¹, Joseph Suárez², and Eugene Vinitzky¹

¹NYU Tandon School of Engineering

²Puffer.ai

May 21, 2025

Abstract

Simulation agents are essential for designing and testing systems that interact with humans, such as autonomous vehicles (AVs). These agents serve various purposes, from benchmarking AV performance to stress-testing system limits, but all applications share one key requirement: reliability. To enable sound experimentation, a simulation agent must behave as intended. It should minimize actions that may lead to undesired outcomes, such as collisions, which can distort the signal-to-noise ratio in analyses. As a foundation for reliable sim agents, we propose scaling self-play to thousands of scenarios on the Waymo Open Motion Dataset under semi-realistic limits on human perception and control. Training from scratch on a single GPU, our agents solve almost the full training set within a day. They generalize to unseen test scenes, achieving a 99.8% goal completion rate with less than 0.8% combined collision and off-road incidents across 10,000 held-out scenarios. Beyond in-distribution generalization, our agents show partial robustness to out-of-distribution scenes and can be fine-tuned in minutes to reach near-perfect performance in such cases. We open-source the pre-trained agents and integrate them with a batched multi-agent simulator. Demonstrations of agent behaviors can be viewed at <https://sites.google.com/view/reliable-sim-agents>, and we open-source our agents at <https://github.com/Emerge-Lab/gpudrive>.

1 Introduction

Simulation agents are a core part of safely developing and testing systems that interact with humans, such as autonomous vehicles (AVs). In the context of self-driving, these agents, also referred to as road user behavior models, serve two primary purposes: establishing benchmarks for AV behavior [Engström et al., 2024], and representing other road users in simulators to enable statistical safety testing in both nominal and rare, long-tail scenarios [Corso et al., 2021, Montali et al., 2024]. While each use case brings particular requirements, *reliability* is an important one that they share.

A reliable simulation agent consistently behaves as intended by the designer, minimizing unintended actions. For instance, agents designed to stress-test AVs should reliably initiate realistic near-collision events, generating safety-critical scenarios to provide meaningful information about the system’s behavior in edge cases. Conversely, nominal agents must focus on replicating typical road behavior to simplify experiments that vary other environmental factors, such as weather. In either case, unreliable sim agents introduce *noise* into the evaluation process by producing trajectories that crash too infrequently in the stress-test case and too frequently in the nominal case.

How can we build sim agents that are *close enough*¹ to reality while maximizing designer specifications, i.e., reliability? One approach relies on generative models, which have shown remarkable progress in producing diverse, human-like behaviors through imitation learning from demonstrations [Xu et al., 2023, Pillion et al., 2024, Huang et al., 2024]. However, it is uncertain whether they meet the reliability standards of a fully automated AV development pipeline. This is highlighted by the best-performing models in the Waymo Open Sim Agent Challenge [Montali et al., 2024, WOSAC], a well-known benchmark for realistic nominal road user behavior. While state-of-the-art models in the 2024 challenge closely replicate logged human trajectories and achieve high

^{*}Corresponding author: cornelisse.daphne@nyu.edu

¹Here, close enough is emphasized because what constitutes an acceptable model of human behavior depends highly on the use case.

scores on various distributional metrics, they still fall short in critical areas. Ground-truth human trajectories in the dataset rarely or never involve collisions or off-road movements; the best aggregate collision and off-road rates among the top WOSAC 2024 are approximately 4% and 1-2%, respectively [Zhang et al., 2024b, Huang et al., 2024, Zhou et al., 2024, CATK, VBD, BehaviorGPT]. This is well below the capabilities of human drivers.

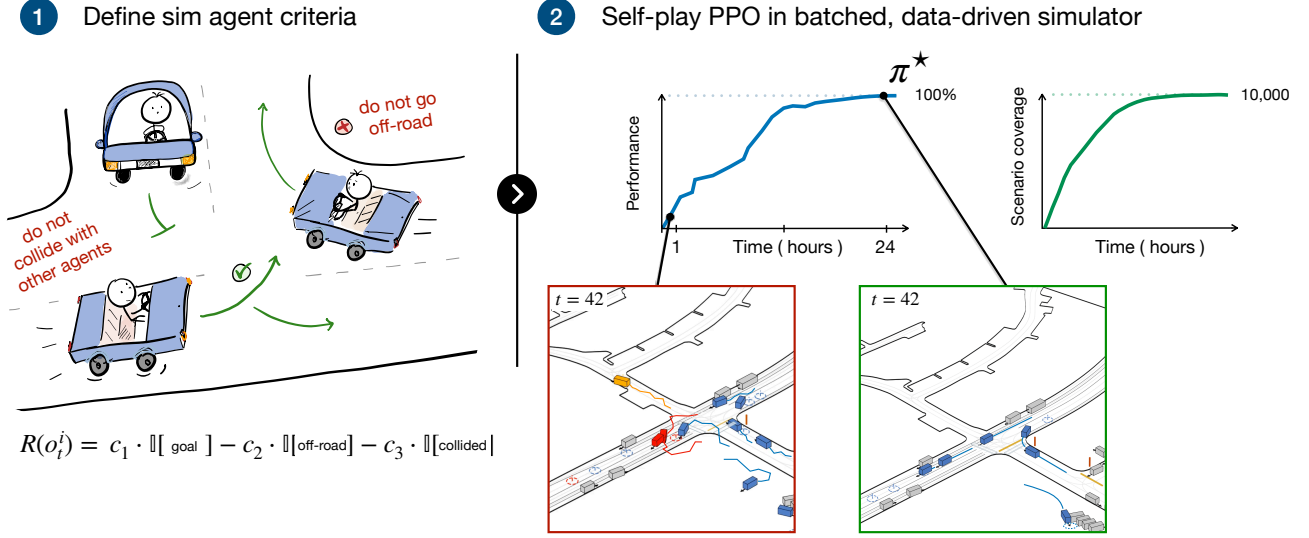


Figure 1: **Overview of approach.** *Left:* We define several criteria to guide the learning of simulation agents through rewards. The reward function is a weighted combination of these criteria: $r(o_t^i) = \sum_i c_i \cdot \mathbb{I}[\text{criteria}_i]$. Here, we focus on achieving goal-directed nominal sim agent behavior—ensuring agents stay on the road and avoid collisions while navigating to a target position. *Right:* Over 24 hours on a single GPU, we iterate through 10,000 scenarios (green curve) from the Waymo Open Motion Dataset in GPUdriVe [Kazemkhani et al., 2024], reaching high performance (blue curve, reliability) on the defined criteria after 2 billion agent steps by self-play PPO. The example scenarios illustrate agent behavior at different stages of training. Initially, agents display random behavior and frequently collide with each other and the road edges (marked in orange and red), but their behavior becomes streamlined over many iterations.

Importantly, these relatively high failure rates may limit the scalability of AV evaluation and development, especially as generative models are increasingly used to create rare safety-critical scenarios that are underrepresented in real-world data [Mahjourian et al., 2024]. When trajectories deviate unpredictably, researchers or engineers must find out: Is the observed outcome a signal or an artifact of simulator noise? For example, if 1 in 10 scenarios reflects unintended behavior, distinguishing meaningful failures from artifacts becomes a time-consuming task. As such, making sim agents more reliable seems to be a key pillar to further scale AV evaluation and development.

The question becomes: How can we close this reliability gap for state-of-the-art sim agents? Assuming we can precisely define what the agent should adhere to (e.g., stay on the road), there is reason to believe that self-play reinforcement learning (RL) is a piece of the puzzle. Evidence from a broad body of recent literature on games shows that self-play RL, combined with well-defined criteria (e.g., maximize score X) can produce agents capable of perfect, superhuman gameplay in the large compute and data regime [Silver et al., 2018, Bakhtin et al., 2023, Berner et al., 2019].

We study whether self-play at scale improves the reliability of sim agents. Specifically, we ask:

1. How does the *reliability* (as measured by performance on the test set of metric X) of sim agents through self-play scale as a function of the data available?
2. How well do these agents generalize to unseen scenarios and out-of-distribution events?

To explore these questions, we train agents via self-play using a semi-realistic human perception framework in a data-driven simulator [Kazemkhani et al., 2024]. We evaluate performance across thousands of scenarios from the Waymo Open Motion Dataset [Ettinger et al., 2021]. We find that self-play PPO scales effectively with on-policy data and compute. After sufficient training, models generalize well to 10,000 unseen test traffic scenarios, effectively closing the train-test gap.

At scale, self-play PPO sim agents consistently achieve the specified criteria (Section 2.2): staying on the road, avoiding collisions, and reaching a target position. This establishes a framework where agents can be tuned to achieve specific collision rates, enabling both nominal and safety-critical traffic simulation. By improving the reliability standards of sim agents, our approach supports the continued scaling and automation of AV development and evaluation pipelines.

Finally, we take a step towards fine-tuning these agents for behaviors underrepresented in the dataset, a useful capability for safety-critical applications. To facilitate further research, we open-source the pre-trained agents at www.github.com/Emerge-Lab/gpudrive, allowing others to reproduce our results and seamlessly use these sim agents.

2 Method

2.1 Dataset and simulator

We conduct our experiments in GPUdriVe, a data-driven, multi-agent, GPU-accelerated simulator [Kazemkhani et al., 2024]. GPUdriVe contains $K = 160,147$ real-world traffic scenarios from the Waymo Open Motion Dataset [Ettinger et al., 2021, WOMD]. Each scenario $k \in K$ comprises a static road graph, R_k , and a time series of *joint* logged human trajectories: $S_k = \{(s_t, A_t)_{t=0}^{T=90}, R_k\}$ where $s_t \in \mathbb{R}^{(1,F)}$ represents the world state represented as F features at time t , and $A_t \in \mathbb{R}^{(N,2)}$ represents the action matrix for all N agents in the scene. The joint agent demonstrations are 9 seconds long and discretized at 10 Hz.

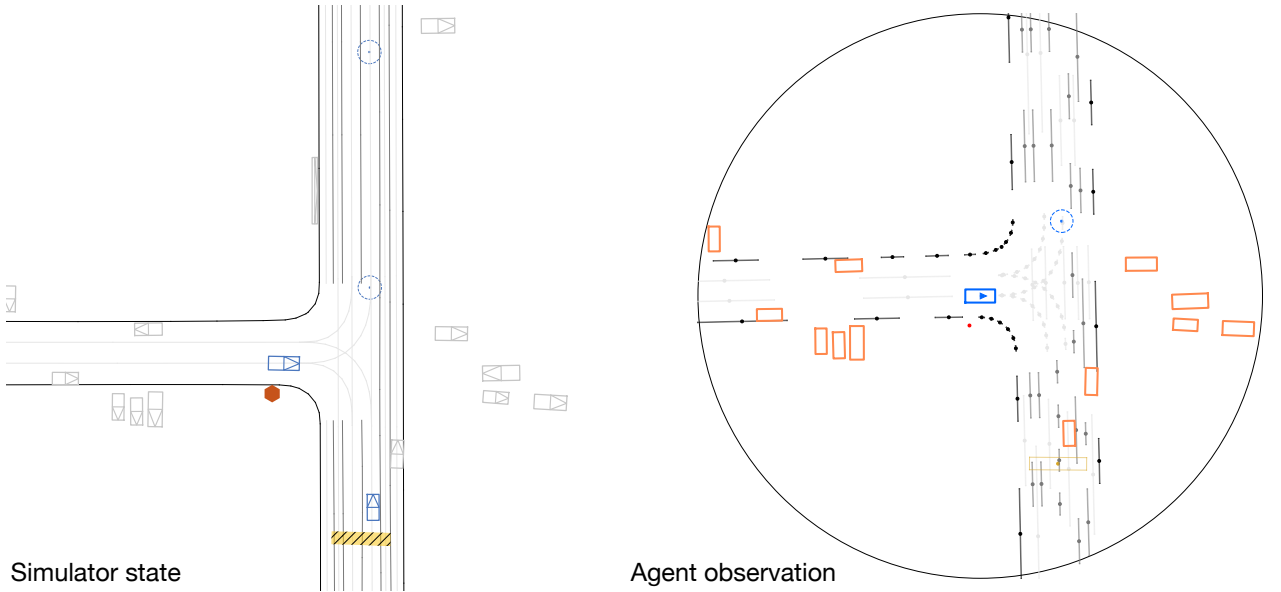


Figure 2: **Sample scenario state with corresponding agent observation.** *Left:* Example scenario from the Waymo Open Motion Dataset rendered in GPUdriVe as shown from a bird’s eye view. The boxes (\square) indicate controlled agents and the circles (\odot) indicate the goal positions for every controlled agent. *Right:* Scene view from the agent in the center (\square). Agents see a subset of the road points within a configurable radius (here $r_o = 50$ meters) and their corresponding types and segment length. Road types are road edges (\bullet) and road lanes (\circ). They can also view the relative position and velocity of the other agents in the scene (\square). Agents in gray are static throughout the episode as they are parked cars but this information is not visible to the agent i.e. the agent does not know that the gray cars are guaranteed not to move and consequently all cars are orange in the agent observation view.

2.2 Task definition and measuring performance

2.2.1 Task definition

We aim to study how the reliability of simulation agents trained via self-play scales with data. To this end, we design a task with well-defined metrics such that experimental results are easy to interpret. Given a traffic scenario S_k with N controlled agents we task every agent to navigate to a designated goal position while satisfying two criteria: (1) avoiding collisions with other agents and (2) staying on the road.

To obtain valid goals, we use the endpoints (x_T^i, y_T^i) (marked by \odot in Figure 2) from the WOMD. Agents are initialized from the starting positions (x_0^i, y_0^i) of the WOMD. Given how the WOMD is collected and processed, we know that the human agents in the dataset must have successfully reached their endpoints within 9 seconds (or 91 steps). We assume that in principle, all agents should be capable of doing the same. To reflect this, a scenario is considered solved when *all controlled agents* reach their target positions within 91 steps while adhering to the specified criteria.

2.2.2 Metrics

We use four *scene-based metrics* to quantify performance:

- Goal achieved \uparrow : Percentage of agents that reached their target position within $T = 91$ steps.
- Collided \downarrow : Percentage per scenario indicating objects that collided, at any point in time, with any other object, i.e. when the agent bounding boxes touch.
- Off-road \downarrow : Percentage of agents per scenario that went off the road or touched a road edge, at any point in time.
- Other \downarrow : Percentage of agents per scenario that did not collide or go off-road but also did not reach the goal position.

The Collided and Off-road metrics align with the Waymo Open Sim Agent Challenge and Waymax [Montali et al., 2024, Gulino et al., 2024]. Specifically, Collided is part of the “object interaction metrics” category and the off-road events are part of the “map-based metrics” category. Under the assumption that human road users have near zero collision and off-road events, we can meaningfully compare our scores to the top submissions [Huang et al., 2024, Zhou et al., 2024]².

The Goal achieved metric is not directly reported in WOSAC, making it less comparable. The most similar metric is the Route Progress Ratio used in Waymax [Gulino et al., 2024], which measures how far an agent travels along the logged trajectory. However, since our focus is not on mimicking logged trajectories but on precisely reaching a particular goal, a binary metric is, in our case, a more meaningful indicator of performance. However, reaching the goal roughly corresponds to a Route Progress Ratio of 100%.

Agent-based metrics: Since the scene-based metrics are biased towards scenes with a small number of agents (one agent colliding in a scene with 2 agents vs. 10 scenes provides a fraction of 1/2 vs 1/10th), we also report the metrics above in *agent-based* way, where we aggregate the counts across the whole dataset and then divide them by the number of total agents.

In both cases, the ceiling for this task is 100% Goal achieved, 0% Collided, and 0% Off-road.

2.3 State and observation space

This section outlines the design choices and parameterization of the observation \mathbf{o}_t^i for agent i at time t . We make these choices to reflect semi-realistic limits on human perception. The observation encodes the agent’s partial view of the scenario state s_t , capturing the information necessary for decision-making. In this work, we model the RL problem as a Partially Observed Stochastic Game [Hansen et al., 2004, POSG], where agents make simultaneous decisions under partial observability. We briefly describe the core aspects of the observation space and give a complete overview in Appendix A. Design choices:

²Technically, WOSAC frames this as a distribution-matching problem: metrics are first computed as event counts, which are then compared to the distribution of log replay trajectories across several rollouts.

- **Relative coordinate frame:** All agent information is presented in an ego-centric coordinate frame to align with human-like perception.
- **Observation radius:** The observation radius r_o determines the visible area around the agent. For our experiments, we set $r_o = 50$ meters, as illustrated in Figure 2.
- **No history:** Agents only receive information from the current timestep.
- **Road graph:** We reduce the full road graph, which consists of up to 10,000 sparsely distributed road points, for computational efficiency. To reduce the number of points corresponding to straight lines, we run the polyline reduction threshold of the polyline decimation algorithm [Visvalingam and Whyatt, 2017] in GPUdrive to 0.1 which roughly cuts the number of points by a factor of 10. We also cap the maximum visible road points at 200, selecting 200 points from those in the view radius in a random order if there are more than 200 points, creating a sparse view of the local road graph. Empirical results show this is sufficient for agents to navigate toward goals without going off the road or causing collisions.
- **Normalization:** Features are normalized to be between -1 and 1 by the minimum and maximum value in their respective category. Details are found in Tables 3, 4, and 5.

2.4 Action space and dynamics model

To align with the control outputs of real human road users more closely, we take the action for every agent i to be a vector of the following discrete random variables: $\mathbf{a}_t^i = (\tilde{a}, \tilde{s})$ where acceleration actions are 7 actions defined over an evenly spaced grid between $[-4, 4]$ and the steering wheel angle are 13 actions defined over an evenly spaced grid between $[-\pi, \pi]$. The bounds are set to reflect the kinematic constraints of real driving. We assume that the random variables \tilde{a}, \tilde{s} are not independent (e.g. sharp turns are less likely at high acceleration) and model the conditional joint probability mass function (pmf) of the two discrete random variables, where we condition on the current observation of agent i at time step t : $\pi_{\tilde{a}, \tilde{s}}(a, s \mid \mathbf{o}_t^i) := P(\tilde{a} = a, \tilde{s} = s \mid \mathbf{o}_t^i)$. The conditional pmf π_θ describes the behavior under the assumption that \mathbf{o}_t^i takes a fixed set of values. The total joint action space contains $7 \times 13 = 91$ actions. With these actions, agents are stepped in the simulator using an Ackermann bicycle model [Rajamani, 2011].

2.5 Reward function

We define the individual agent rewards as follows:

$$r(\mathbf{o}_t^i, \mathbf{a}_t^i) = w_{\text{Goal achieved}} \cdot \mathbb{I}[\text{Goal achieved}] - w_{\text{Collided}} \cdot \mathbb{I}[\text{Collided}] - w_{\text{Offroad}} \cdot \mathbb{I}[\text{Offroad}]$$

2.6 Collision behavior

During training and testing, we allow agents to continue the episode even after going off-road or colliding with another agent in the scene. Agents receive a penalty for each collision or off-road event, allowing them to accrue multiple penalties throughout an episode. See Appendix D.1 for the details.

2.7 Model

We use a neural network with an encoder and a shared embedding, as illustrated in Figure 8 (see Appendix B.2). The flat observation vector is first decomposed into three modalities: the dense ego state, the sparse road graph, and the sparse partner observations. Each modality is processed independently. Inspired by the late fusion approach in Wayformer [Nayakanti et al., 2023], we then concatenate the outputs, apply max pooling, and pass the result through a shared embedding. This hidden embedding is fed into separate actor and critic heads, each implemented as a single feedforward layer. The model only has $\sim 50,000$ trainable parameters.

2.8 Training

Self-play PPO In each scenario, we control up to $N = 64$ agents using a shared, decentralized policy π_θ . Actions are independently sampled from the policy based on the ego views of each agent i during every step in

the rollout: $\mathbf{a}_t^i \sim \pi_\theta(\cdot \mid \mathbf{o}_t^i)$. We train agents using Proximal Policy Optimization [Schulman et al., 2017, PPO] using batches of $S = 800$ distinct scenarios, with the set of training scenarios uniformly resampled every 2 million steps. Initially, agents exhibit random behavior and crash frequently. Over time, the agents’ behavior becomes more streamlined, creating smooth trajectories with high rates of reaching the goals.

3 Related work

Self-play for agents in games: Self-play RL [Samuel, 1959, Tesauro et al., 1995] has been a core ingredient in creating effective agents across a wide range of complex games. Notable examples include superhuman gameplay in two-player zero-sum games like Chess and Go [Silver et al., 2018], expert human-level play in Stratego [Perolat et al., 2022] and Starcraft [Vinyals et al., 2019], as well many-player games that require some level of cooperation like Diplomacy [Bakhtin et al., 2023] and Gran Turismo [Wurman et al., 2022]. These successes have demonstrated the effectiveness of self-play, particularly in the large-data, large-compute regime. However, the majority of its successes are in variants of zero-sum games whereas driving tasks are likely general-sum and feature many-agent interaction.

RL for driving agents: Reinforcement learning has been explored for the design of autonomous driving agents, though state-of-the-art agents are currently far below the human rate of between 800000 km per police-reported traffic crash in the United States [Stewart et al., 2023] or as much as 1 crash per 24800 km in more challenging domains such as San Francisco [Flannagan et al., 2023]. These agents are frequently trained in simulators built atop large open-source driving datasets [Gulino et al., 2024, Vinitsky et al., 2022, Kazemkhani et al., 2024] such as Waymo Open Motion [Ettinger et al., 2021, WOMD], [Caesar et al., 2020, NuScenes], [Zheng et al., 2024, ONE-Drive] though there are also procedurally generated [Li et al., 2022] and non-data-driven simulators [Dosovitskiy et al., 2017]. These datasets collectively add up to tens of thousands of hours of available data and are often used to train RL agents in *log-replay* mode, a setting in which only one agent is learning and the remainder are either replaying human trajectories or executed hand-coded policies.

The complexity of scaling RL in these settings has led to the creation of batched simulators [Kazemkhani et al., 2024, GPUDrive], [Gulino et al., 2024, Waymax], [Cusumano-Towner et al., 2025, Gigaflow] whose high throughput helps ameliorate issues of sample complexity. Many works have explored ways to use these simulators to learn high-quality reinforcement learning agents through RL including uses of self-play [Peng et al., 2021, Vinitsky et al., 2022, Zhang et al., 2022, 2023, 2024a, Jaeger et al., 2025]. Our work is mostly distinct from these by the scale of training and a significantly lower crash and off-road rate than has previously been observed with the exception of [Cusumano-Towner et al., 2025, Gigaflow]. *Gigaflow* achieves a scale and robustness above what is in this paper, crashing between every hundred thousand to million kilometers in self-play, whereas we crash on the order of every 10 kilometers. However, our work provides complementary insights to *Gigaflow* in several ways 1) the results are open-source, creating models and training scripts that can be reused in other works; 2) our training setup draws scenes from log-replay data, creating a curriculum of both easy and challenging tasks as opposed to scattering goals over a large map as is done in *Gigaflow*.

4 Results

Table 1: Aggregate scene-based performance in % across $N = 10,000$ randomly sampled train and test traffic scenarios from the Waymo Open Motion Dataset (mean \pm std). Metrics are defined in section 2.2.2.

Dataset	Goal achieved \uparrow	Collided \downarrow	Off-road \downarrow	Other \downarrow
Train	99.84 \pm 1.27	0.38 \pm 2.91	0.26 \pm 2.17	0.13 \pm 1.14
Test	99.81 \pm 1.53	0.44 \pm 3.17	0.31 \pm 2.59	0.14 \pm 1.16

4.1 Scaling with data

Solving the full Waymo Open Motion Dataset under partial observability We investigate whether agents with a partial view of the environment can solve all scenarios in the Waymo Open Motion Dataset. Our results




Figure 3: **Scaling with data.** Average performance with standard errors on 10,000 unseen scenarios from the WOMB validation set as a function of the training dataset size. The striped lines indicate optimal performance.

show that almost all scenarios can be solved successfully. After 2 billion training steps, agents achieve a goal-reaching rate of 99.84%, a collision rate of 0.38%, and an off-road rate of 0.26% on the training dataset.³ Furthermore, as depicted in Figure 4, zooming in on the final four hours of training suggests that metrics exhibit a continued, albeit gradual, improvement, indicating that performance can be further increased with additional training. This training run took 24 hours on a single NVIDIA A100 GPU.

The agent-based metrics are similar to the scene-based metrics reported above: a goal rate of 99.72%, a collision rate of 0.26%, and an off-road rate of 0.35%. Sample rollouts with the best-trained policy are shown in Figures 6, 7 and on the project page.




Figure 4: **Batch performance throughout training.** *Left:* Average reward per agent (maximum of 1) as a function of wall-clock time. We train agents for at most 24 hours. *Center:* Goal achievement rate per batch as a function of global steps (2 billion steps generated in 24 hours). *Right:* Percentage of agents that collide with another agent (red) or with a road edge (orange). All curves are smoothed using a rolling window of 250 steps. The inset figures show a zoomed-in view of the final four hours of the run, with the y-axes displayed on a logarithmic scale. The red annotations on the insets indicate the minimum and maximum values within the zoomed-in window. Note that the metrics reported during training are by excluding trivial agents, we only control agents that have to drive for more than 2 meters to reach their goal destination.

Effective generalization to unseen scenarios with sufficient data We conduct experiments with 100, 1,000, 10,000, and 100,000 unique training scenarios to assess how self-play performance scales with the diversity of training scenes. Table 1 summarizes the results. We find no significant train-test gap when training with 10,000 scenarios or more, indicating the model generalizes well to new, unseen situations. Figure 3 shows the key metrics as a function of training dataset size. Notably, with 10,000 training scenarios, the model reaches

³These metrics are computed in alignment with the way they are defined in WOSAC, but it should be noted that this is an over-optimistic metric as it includes many agents that simply need to remain in place as they are initialized right next to their goals. Excluding such agents and only controlling agents that must drive more than 2 meters before reaching their goal gives performance metrics of: 99.40% goal-reaching rate, 0.5% collision rate, and 0.6% offroad rate.

nearly the ceiling of our benchmark, achieving a 99.81% goal-reaching rate, 0.44% collision rate, and 0.31% off-road rate on 10,000 held-out test scenarios. Note that these numbers do not add to 1 as a vehicle may collide and go off-road in the same step.

4.2 Extrapolative generalization and fast fine-tuning

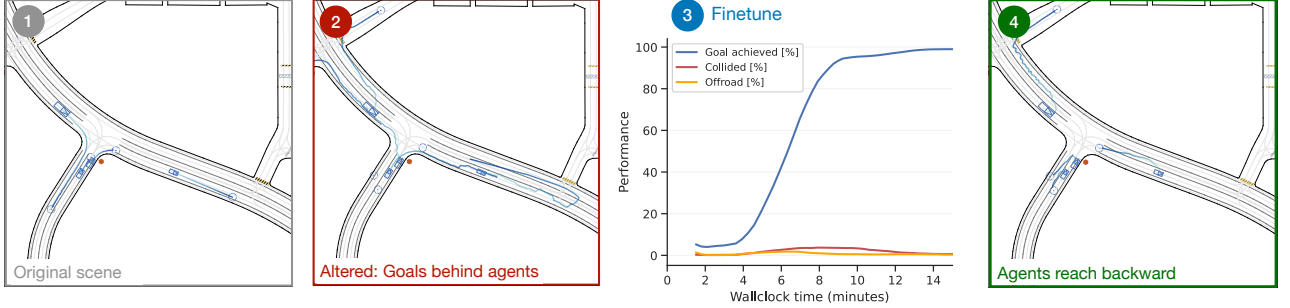


Figure 5: **Fine-tuning agent behaviors** 1: In most scenarios, agent target positions are located in front of them. The figure shows a typical example from the dataset with rollouts from the trained policy. 2: Fewer than 2% of agent goals require backward driving or a U-turn. To evaluate agent performance in such out-of-distribution cases, we create hand-designed scenarios where goals are placed behind agents. As expected, performance drops significantly (by 50%), as agents struggle to reach these goals. In this scene, only a single agent achieves its new goal. 3: To address this, we fine-tune a model pre-trained on 10,000 WOMD scenarios using the 13 hand-designed cases. Within 15 minutes, agents successfully learn to navigate to the goals behind them. 4: A rollout of the fine-tuned model demonstrates its ability to handle the altered scenario. Each agent executes a U-turn to get to its goal.

Navigating backwards Beyond generalization to within distribution scenarios, as reported in Section 4.1, we are interested in agent performance in out-of-distribution events. This is useful to know, as researchers may typically manipulate scenarios or make them harder in some way to test the limits of AV systems. Where do these agents break, and how easily can they be finetuned? We analyzed the complete training data set to find rare behaviors and found that agents rarely have to reach a goal that is behind them (1.86%, Details in Appendix E). We then hand-designed 13 scenarios from the test dataset with a total of 27 agents across all scenes, placing goals behind agents. Figure 5.2 illustrates an example of such a scene. We summarize the results in Table 7. Overall, we observe that the performance drops from 100% to 53.5% goal-reaching rate when placing the targets behind the agents. Unsurprisingly, agents exhibit poor performance on events that are rare or entirely unobserved in the training scenarios.

Fast finetuning As a proof of concept, we demonstrate how self-play reinforcement learning enables rapid fine-tuning of a model to learn new behaviors, such as navigating backward, using only a few samples. Figure 5 provides an overview of our approach. Initially, introducing a rare and possibly out-of-distribution scenario—where goals are positioned behind agents—leads to a drop in performance (1 \rightarrow 2). To address this, we take the 13 hand-designed scenarios and fine-tune the policy that was pre-trained on 10,000 WOMD scenarios (3). The model starts with a low goal-reaching rate but quickly adapts, achieving 100% success within 15 minutes of training. After fine-tuning, agents can reliably reach goals behind them (4). We share an accompanying video of before and after finetuning at [the project page](#).

Realism assessment While our objective in this work is not to develop a *human-like agent*, it is informative to assess how close the resulting agent behavior matches the distributional properties of human driving. To obtain an indication of the agent realism, we evaluate on the Waymo Open Sim Agents Challenge (WOSAC) [Montali et al., 2023], an established benchmark for assessing realism in driving. In a slight deviation from the standard benchmark, we condition the agents on the *end points* of the log trajectories, instead of conditioning on the first second of logs.

Table 2: Comparison of model performance across WOSAC evaluation metrics. Arrows indicate whether higher (\uparrow) or lower (\downarrow) values are better. Note that collision/offroad here is a likelihood, not a measure of collision/offroad rate.

Model	Realism \uparrow	minADE \downarrow	Kinematics \uparrow				Interaction \uparrow			Map-Based \uparrow	
	Meta	[m]	Lin Speed	Lin Accel	Ang Speed	Ang Accel	Dist Obj	Collis	TTC	Dist edge	Offroad
Logged Oracle	-	-	0.476	0.478	0.578	0.694	0.476	1.000	0.883	0.715	1.000
Self-play PPO	0.629	8.205	0.170	0.272	0.302	0.471	0.073	0.970	0.882	0.262	0.910
VBD	0.720	1.474	0.359	0.366	0.420	0.522	0.368	0.934	0.815	0.651	0.879

Table 2 shows the aggregate scores across 100 held-out scenes. Despite not intentionally optimizing for realism, our method gets a realism meta-score of 0.629. The simulated collision rate of the agent in WOSAC is 0.5%, which is comparable to what we report in Table 1. As ranked by the realism meta-metric, it is above a linear extrapolation policy (0.324) [Montali et al., 2024], and competitive with entries on the 2024 WOSAC challenge, such as VPD-BP (0.63), and GigaFlow (0.6190) [Cusumano-Towner et al., 2025].

5 Discussion and conclusions

Our results lead us to three main conclusions:

- **Self-play at scale reliably achieves well-defined criteria in unseen scenarios:** Our findings suggest that self-play RL scales effectively with available data (Section 4.1), achieving low-collision performance on the Waymo Open Motion Dataset (WOMD) with no train-test gap. To the best of our knowledge, this is the first demonstration of this level of performance on WOMD. Compared to state-of-the-art supervised models, such as CATK [Zhang et al., 2024b], VBD [Huang et al., 2024], and BehaviorGPT [Zhou et al., 2024], our approach appears to reduce the collision and off-road rates by approximately 4 \times . Our agents crash on the order of once every 30 minutes, which, while well below human capabilities, represents a meaningful increase over available baselines.
- **Rare events remain a challenge.** Agents struggle with rare or out-of-distribution scenarios, such as goals placed behind them (Section C) or navigating roundabouts. In these cases, performance drops significantly, indicating that performance on uncommon situations remains a key limitation.
- **Fine-tuning quickly improves performance in unseen scenarios:** Fine-tuning on a small subset of hand-designed cases can improve agent performance. In our experiments, fine-tuning a pre-trained model for just a few minutes enables agents to achieve near-perfect goal-reaching rates on previously difficult tasks (Section 4.2).

As our agents may be independently interesting to use as part of other simulators or in autonomous vehicle test cases, we open-source our agents at www.github.com/Emerge-Lab/gpudrive.

We note three limitations of our work that provide some nuance. First, our benchmark, built atop the Waymo Open Motion Dataset, consists of short-horizon scenarios that are only 9 seconds long. Second, we excluded pedestrians, cyclists, and traffic lights. Finally, we train agents to optimize specific criteria above realism, we find that there is still ample room for improvement in making their behavior human-like, as indicated by the WOSAC meta-score of 0.63 (Section 4.2). An interesting direction for future work is to balance reliability with realism, ensuring agents meet performance standards while accurately reflecting human driving behavior across diverse scenarios.

While not explored in this paper, we anticipate that our findings extend to other domains such as neuroscience, where agent-based modeling is gaining momentum [Aldarondo et al., 2024, Johnson-Yu et al., 2024, Castro et al., 2025]. Researchers are increasingly using physics-based simulators to create digital twins of animals, enabling cost-effective and controlled experimentation. For these agents to be useful models of animal behavior, reliability and robustness appear essential. For example, a rodent foraging model, should not exhibit free movement. We hope our work contributes to the improvement of agent-based modeling, helping to enhance controllability and robustness across different domains.

References

- Diego Aldarondo, Josh Merel, Jesse D Marshall, Leonard Hasenclever, Ugne Klibaite, Amanda Gellis, Yuval Tassa, Greg Wayne, Matthew Botvinick, and Bence P Ölveczky. A virtual rodent predicts the structure of neural activity across behaviours. *Nature*, 632(8025):594–602, 2024.
- Anton Bakhtin, David J. Wu, Adam Lerer, Jonathan Gray, Athul Paul Jacob, Gabriele Farina, Alexander H. Miller, and Noam Brown. Mastering the game of no-press diplomacy via human-regularized reinforcement learning and planning. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023. URL <https://openreview.net/forum?id=F61FwJTZhb>.
- Christopher Berner, Greg Brockman, Brooke Chan, Vicki Cheung, Przemyslaw Debiak, Christy Dennison, David Farhi, Quirin Fischer, Shariq Hashme, Christopher Hesse, Rafal Józefowicz, Scott Gray, Catherine Olsson, Jakub Pachocki, Michael Petrov, Henrique Pondé de Oliveira Pinto, Jonathan Raiman, Tim Salimans, Jeremy Schlatter, Jonas Schneider, Szymon Sidor, Ilya Sutskever, Jie Tang, Filip Wolski, and Susan Zhang. Dota 2 with large scale deep reinforcement learning. *CoRR*, abs/1912.06680, 2019. URL <http://arxiv.org/abs/1912.06680>.
- Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nusenes: A multimodal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020.
- Pablo Samuel Castro, Nenad Tomasev, Ankit Anand, Navodita Sharma, Rishika Mohanta, Aparna Dev, Kuba Perlin, Siddhant Jain, Kyle Levin, Noémi Éltető, et al. Discovering symbolic cognitive models from human and animal behavior. *bioRxiv*, pages 2025–02, 2025.
- Anthony Corso, Robert Moss, Mark Koren, Ritchie Lee, and Mykel Kochenderfer. A survey of algorithms for black-box safety validation of cyber-physical systems. *Journal of Artificial Intelligence Research*, 72:377–428, 2021.
- Marco F. Cusumano-Towner, David Hafner, Alexander 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. *CoRR*, abs/2502.03349, 2025. doi: 10.48550/ARXIV.2502.03349. URL <https://doi.org/10.48550/arXiv.2502.03349>.
- Alexey Dosovitskiy, Germán Ros, Felipe Codevilla, Antonio M. López, and Vladlen Koltun. CARLA: an open urban driving simulator. In *1st Annual Conference on Robot Learning, CoRL 2017, Mountain View, California, USA, November 13-15, 2017, Proceedings*, volume 78 of *Proceedings of Machine Learning Research*, pages 1–16. PMLR, 2017. URL <http://proceedings.mlr.press/v78/dosovitskiy17a.html>.
- Johan Engström, Shu-Yuan Liu, Azadeh DinparastDjadid, and Camelia Simoiu. Modeling road user response timing in naturalistic traffic conflicts: a surprise-based framework. *Accident Analysis & Prevention*, 198: 107460, 2024.
- Scott Ettinger, Shuyang Cheng, Benjamin Caine, Chenxi Liu, Hang Zhao, Sabeek Pradhan, Yuning Chai, Ben Sapp, Charles R Qi, Yin Zhou, et al. Large scale interactive motion forecasting for autonomous driving: The waymo open motion dataset. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9710–9719, 2021.
- Carol Flannagan, Andrew Leslie, Raymond Kiefer, Scott Bogard, Geoff Chi-Johnston, Laura Freeman, Rayman Huang, David Walsh, and Joseph Anthony. Establishing a crash rate benchmark using large-scale naturalistic human ridehail data. Technical report, UMTRI, 2023.
- Cole Gulino, Justin Fu, Wenjie Luo, George Tucker, Eli Bronstein, Yiren Lu, Jean Harb, Xinlei Pan, Yan Wang, Xiangyu Chen, et al. Waymax: An accelerated, data-driven simulator for large-scale autonomous driving research. *Advances in Neural Information Processing Systems*, 36, 2024.
- Eric A. Hansen, Daniel S. Bernstein, and Shlomo Zilberstein. Dynamic programming for partially observable stochastic games. In *National Conference on Artificial Intelligence*, 2004.

- Zhiyu Huang, Zixu Zhang, Ameya Vaidya, Yuxiao Chen, Chen Lv, and Jaime Fernández Fisac. Versatile scene-consistent traffic scenario generation as optimization with diffusion. *arXiv preprint arXiv:2404.02524*, 2024.
- Bernhard Jaeger, Daniel Dauner, Jens Beißwenger, Simon Gerstenecker, Kashyap Chitta, and Andreas Geiger. Carl: Learning scalable planning policies with simple rewards. *arXiv preprint arXiv:2504.17838*, 2025.
- Sonja Johnson-Yu, Satpreet Harcharan Singh, Federico Pedraja, Denis Turcu, Pratyusha Sharma, Naomi Saphra, Nathaniel Sawtell, and Kanaka Rajan. Understanding biological active sensing behaviors by interpreting learned artificial agent policies. In *Workshop on Interpretable Policies in Reinforcement Learning@RLC-2024*, 2024.
- Saman Kazemkhani, Aarav Pandya, Daphne Cornelisse, Brennan Shacklett, and Eugene Vinitzky. Gpudrive: Data-driven, multi-agent driving simulation at 1 million fps. *arXiv preprint arXiv:2408.01584*, 2024.
- 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*, 2022.
- Reza Mahjourian, Rongbing Mu, Valerii Likhoshesterov, Paul Mougín, Xiukun Huang, Joao Messias, and Shimon Whiteson. Unigen: Unified modeling of initial agent states and trajectories for generating autonomous driving scenarios. *arXiv preprint arXiv:2405.03807*, 2024.
- Nico Montali, John Lambert, Paul Mougín, Alex Kuefler, Nicholas Rhinehart, Michelle Li, Cole Gulino, Tristan Emrich, Zoey Yang, Shimon Whiteson, Brandyn White, and Dragomir Anguelov. The waymo open sim agents challenge. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine, editors, *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*, 2023. URL http://papers.nips.cc/paper_files/paper/2023/hash/b96ce67b2f2d45e4ab315e13a6b5b9c5-Abstract-Datasets_and_Benchmarks.html.
- Nico Montali, John Lambert, Paul Mougín, Alex Kuefler, Nicholas Rhinehart, Michelle Li, Cole Gulino, Tristan Emrich, Zoey Yang, Shimon Whiteson, et al. The waymo open sim agents challenge. *Advances in Neural Information Processing Systems*, 36, 2024.
- Nigamaa Nayakanti, Rami Al-Rfou, Aurick Zhou, Kratarth Goel, Khaled S. Refaat, and Benjamin Sapp. Wayformer: Motion forecasting via simple & efficient attention networks. In *IEEE International Conference on Robotics and Automation, ICRA 2023, London, UK, May 29 - June 2, 2023*, pages 2980–2987. IEEE, 2023. doi: 10.1109/ICRA48891.2023.10160609. URL <https://doi.org/10.1109/ICRA48891.2023.10160609>.
- Zhenghao Peng, Quanyi Li, Ka-Ming Hui, Chunxiao Liu, and Bolei Zhou. Learning to simulate self-driven particles system with coordinated policy optimization. In Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan, editors, *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 10784–10797, 2021. URL <https://proceedings.neurips.cc/paper/2021/hash/594ca7adb3277c51a998252e2d4c906e-Abstract.html>.
- Julien Perolat, Bart De Vylder, Daniel Hennes, Eugene Tarassov, Florian Strub, Vincent de Boer, Paul Muller, Jerome T Connor, Neil Burch, Thomas Anthony, et al. Mastering the game of stratego with model-free multiagent reinforcement learning. *Science*, 378(6623):990–996, 2022.
- Jonah Philion, Xue Bin Peng, and Sanja Fidler. Trajenglish: Traffic modeling as next-token prediction. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024. URL <https://openreview.net/forum?id=Z59Rb5bPPP>.
- Rajesh Rajamani. *Vehicle dynamics and control*. Springer Science & Business Media, 2011.
- A. L. Samuel. Some studies in machine learning using the game of checkers. *IBM Journal of Research and Development*, 3(3):210–229, 1959. doi: 10.1147/rd.33.0210.

- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *CoRR*, abs/1707.06347, 2017. URL <http://arxiv.org/abs/1707.06347>.
- David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dharmashan Kumaran, Thore Graepel, et al. A general reinforcement learning algorithm that masters chess, shogi, and go through self-play. *Science*, 362(6419):1140–1144, 2018.
- Timothy Stewart et al. Overview of motor vehicle traffic crashes in 2021. Technical report, United States. Department of Transportation. National Highway Traffic Safety ..., 2023.
- Gerald Tesauro et al. Temporal difference learning and td-gammon. *Communications of the ACM*, 38(3):58–68, 1995.
- Eugene Vinitsky, 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. In Sanmi Koyejo, S. Mohamed, A. Agarwal, Danielle Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*, 2022. URL http://papers.nips.cc/paper_files/paper/2022/hash/191e9e721a2748a860714fb23aaf7c5d-Abstract-Datasets_and_Benchmarks.html.
- Oriol Vinyals, Igor Babuschkin, Wojciech M. Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H. Choi, Richard Powell, Timo Ewalds, Petko Georgiev, Junhyuk Oh, Dan Horgan, Manuel Kroiss, Ivo Danihelka, Aja Huang, Laurent Sifre, Trevor Cai, John P. Agapiou, Max Jaderberg, Alexander Sasha Vezhnevets, Rémi Leblond, Tobias Pohlen, Valentin Dalibard, David Budden, Yury Sulsky, James Molloy, Tom Le Paine, Çağlar Gülçehre, Ziyu Wang, Tobias Pfaff, Yuhuai Wu, Roman Ring, Dani Yogatama, Dario Wünsch, Katrina McKinney, Oliver Smith, Tom Schaul, Timothy P. Lillicrap, Koray Kavukcuoglu, Demis Hassabis, Chris Apps, and David Silver. Grandmaster level in starcraft II using multi-agent reinforcement learning. *Nat.*, 575(7782):350–354, 2019. doi: 10.1038/S41586-019-1724-Z. URL <https://doi.org/10.1038/s41586-019-1724-z>.
- Maheswari Visvalingam and James D Whyatt. Line generalization by repeated elimination of points. In *Landmarks in Mapping*, pages 144–155. Routledge, 2017.
- 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 Raj 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. Outracing champion gran turismo drivers with deep reinforcement learning. *Nat.*, 602(7896):223–228, 2022. doi: 10.1038/S41586-021-04357-7. URL <https://doi.org/10.1038/s41586-021-04357-7>.
- Danfei Xu, Yuxiao Chen, Boris Ivanovic, and Marco Pavone. Bits: Bi-level imitation for traffic simulation. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2929–2936. IEEE, 2023.
- Chris Zhang, Runsheng Guo, Wenyuan Zeng, Yuwen Xiong, Binbin Dai, Rui Hu, Mengye Ren, and Raquel Urtasun. Rethinking closed-loop training for autonomous driving. In Shai Avidan, Gabriel J. Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner, editors, *Computer Vision - ECCV 2022 - 17th European Conference, Tel Aviv, Israel, October 23-27, 2022, Proceedings, Part XXXIX*, volume 13699 of *Lecture Notes in Computer Science*, pages 264–282. Springer, 2022. doi: 10.1007/978-3-031-19842-7_16. URL https://doi.org/10.1007/978-3-031-19842-7_16.
- Chris Zhang, James Tu, Lunjun Zhang, Kelvin Wong, Simon Suo, and Raquel Urtasun. Learning realistic traffic agents in closed-loop. In Jie Tan, Marc Toussaint, and Kourosh Darvish, editors, *Conference on Robot Learning, CoRL 2023, 6-9 November 2023, Atlanta, GA, USA*, volume 229 of *Proceedings of Machine Learning Research*, pages 800–821. PMLR, 2023. URL <https://proceedings.mlr.press/v229/zhang23b.html>.
- Chris Zhang, Sourav Biswas, Kelvin Wong, Kion Fallah, Lunjun Zhang, Dian Chen, Sergio Casas, and Raquel Urtasun. Learning to drive via asymmetric self-play. In Ales Leonardis, Elisa Ricci, Stefan Roth, Olga Russakovsky, Torsten Sattler, and Gül Varol, editors, *Computer Vision - ECCV 2024 - 18th European*

Conference, Milan, Italy, September 29-October 4, 2024, Proceedings, Part LXII, volume 15120 of *Lecture Notes in Computer Science*, pages 149–168. Springer, 2024a. doi: 10.1007/978-3-031-73033-7_9. URL https://doi.org/10.1007/978-3-031-73033-7_9.

Zhejun Zhang, Peter Karkus, Maximilian Igl, Wenhao Ding, Yuxiao Chen, Boris Ivanovic, and Marco Pavone. Closed-loop supervised fine-tuning of tokenized traffic models. *arXiv preprint arXiv:2412.05334*, 2024b.

Yupeng Zheng, Zhongpu Xia, Qichao Zhang, Teng Zhang, Ben Lu, Xiaochuang Huo, Chao Han, Yixian Li, Mengjie Yu, Bu Jin, Pengxuan Yang, Yuhang Zheng, Haifeng Yuan, Ke Jiang, Peng Jia, Xianpeng Lang, and Dongbin Zhao. Preliminary investigation into data scaling laws for imitation learning-based end-to-end autonomous driving. *CoRR*, abs/2412.02689, 2024. doi: 10.48550/ARXIV.2412.02689. URL <https://doi.org/10.48550/arXiv.2412.02689>.

Zikang Zhou, Haibo Hu, Xinhong Chen, Jianping Wang, Nan Guan, Kui Wu, Yung-Hui Li, Yu-Kai Huang, and Chun Jason Xue. Behaviorgpt: Smart agent simulation for autonomous driving with next-patch prediction. *arXiv preprint arXiv:2405.17372*, 2024.

A Observation features and design choices

The observation at time step t for agent i , \mathbf{o}_i^t , is multi-modal and consists of three types of information: the ego state, the visible view of the scene, and the partner observation. We set the maximum number of agents per scenario throughout the experiments, $N = 64$. We limit agents to vehicles. A given agent’s observation is provided as a flattened vector of ~ 3000 elements.

Table 3: Ego state features and dimensions provided in the observation \mathbf{o}_i^t .

Feature	Dimension	Description
Speed	1	The speed of the agent
Vehicle length	1	Length of the agents’ bounding box
Vehicle width	1	Width of the agents’ bounding box
Relative goal position	2	Distance from agent to the target position in the x and y axis
Collision state	1	Whether the agent is in collision (1) or not (0)

Table 4: Visible view or road graph features and dimensions provided in the observation \mathbf{o}_i^t . The road graph consists of a sampled set of R nearest road points, where R is set to 200 in the experiments.

Feature	Dimension	Description
x	$1 \cdot R$	Relative x coordinate of the road point
y	$1 \cdot R$	Relative y coordinate of the road point
Segment length	$1 \cdot R$	Length of the road segment associated with the (x, y) coordinate
Segment width	$1 \cdot R$	Width of the road segment associated with the (x, y) coordinate
Segment height	$1 \cdot R$	Height of the road segment associated with the (x, y) coordinate
Segment orientation	$1 \cdot R$	Angle between the segment midpoint and the ego agent
Type	$1 \cdot R$	Integer indicating the type of the road point. Existing types are: Road edge (impassable; boundary of the road), road lane, road line, stop sign, crosswalk, and speed bump. Integers are one-hot encoded during training, which multiplies the feature dimension by the total number of classes.

Table 5: Partner (“the other”) agent features and dimensions provided in the observation \mathbf{o}_i^t . Partner information is visible within the observation radius.

Feature	Dimension	Description
Speed	$1 \cdot N - 1$	The speed of the other agents
(\mathbf{x}, \mathbf{y})	$2 \cdot N - 1$	Relative positions of the other $N - 1$ agents in the scene. Information is only provided if the partner agents are within the observation radius of the ego agent, and are left as zero otherwise.
(θ_x, θ_y)	$2 \cdot N - 1$	Relative orientation of the other $N - 1$ agents in the scene. Information is only provided if the partner agents are within the observation radius of the ego agent, and are left as zero otherwise.
(w, l, h)	$3 \cdot N - 1$	The width, length, and height of the bounding boxes of the other agents.

B Additional figures

B.1 Sample rollouts




Figure 6: Example rollouts with the best-trained policy. Agents controlled by the trained policy are shown in blue, while static agents are colored in grey.




Figure 7: Example rollouts with the best-trained policy. Agents controlled by the trained policy are shown in blue, while static agents are colored in grey.

B.2 Neural Network Architecture

C Distribution of errors and remaining failure modes

We analyze scenarios that are not perfectly solved, defined as those with a collision rate or off-road rate greater than 0, or where at least one agent fails to reach its goal. A selection of failure modes can be viewed on the project page. Together, these account for 8.95% of the test dataset (896 out of 10,000 scenarios). Figure 9 shows the histogram of error distributions, revealing that most unsolved scenarios have only a small error rate. We compute the Pearson correlation between off-road fractions and collision rates to examine potential relationships between failure modes. The result, $\rho = 0.0135$, is not significant at $\alpha = 0.05$, indicating no meaningful correlation between these two metrics in the unsolved scenarios and suggesting that errors are spread across scenarios.

Additionally, we analyze the top 0.5% failure modes in each category (collision rates, off-road rates, and agents that did not reach the goal position) of the test set. This analysis provides information about challenging aspects of these scenarios. The key takeaways are as follows.

Rare map layouts and objects: High off-road rates occur in scenarios with rarely occurring road structures. One example of this is roundabouts. A large fraction (15%) of the top fraction of collision rates was in roundabout

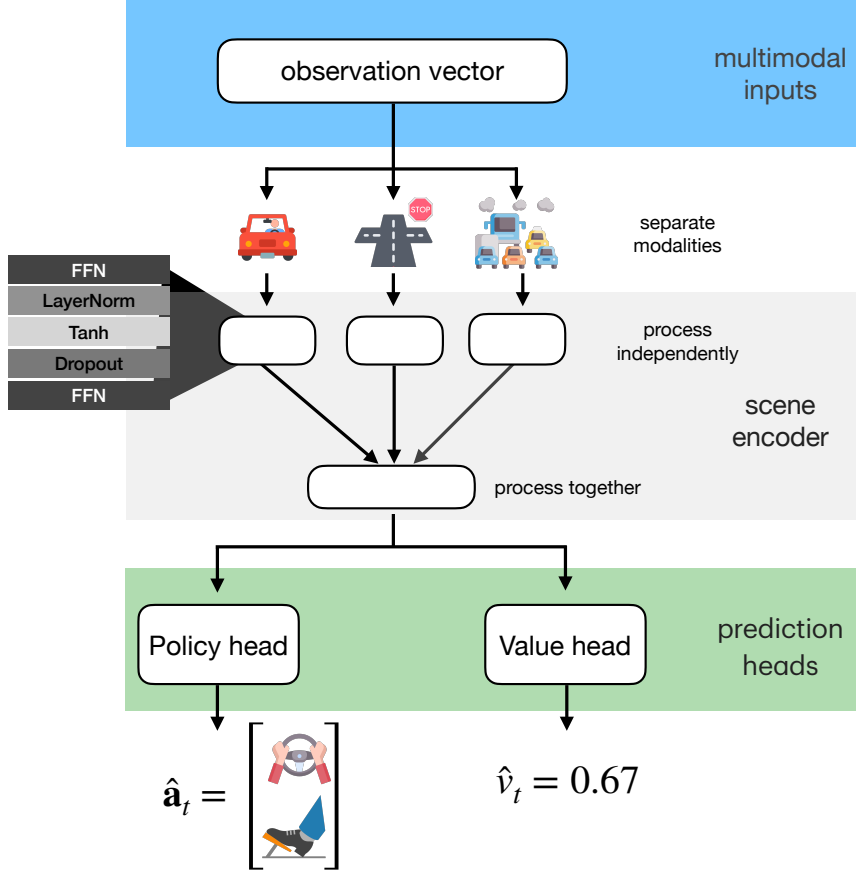


Figure 8: **Network architecture.** The relative observation vector o_t^j is first decomposed into its separate modalities: the ego state (i.e. the agent’s information about itself and its goals), the visible portion of the road graph, and the speeds, yaws, and relative positions of the other agents in the scene. These modalities are first processed separately. Their outputs are combined and max pooled, then processed together. The hidden layer is finally fed into an actor and a critic head.

scenes. The rest included road layouts that are simply harder to navigate, such as tight corners, narrow lane entries, parking lots, etc. Larger vehicles especially struggle with such maps. This coupled with multiple vehicles crowding leads to some of them going off-road.

Coordination: High collision rates occur in intersections, speedy highways, and crowded scenes where sophisticated interaction is required (eg: letting another agent pass before you, making space for another agent to overtake, etc). Crowding and interaction coupled with rare map layouts compound the difficulty of the scene and lead to a higher collision and off-road rate.

Out of time: Some agents have goals further away than others. Having a finite horizon of 91 steps means trying to squeeze past agents and narrow lanes when it is very hard to. This leads to a higher collision and off-road rate compared to scenes with closer goals. This can also compound difficulty in scenes with the aforementioned properties.

D Considerations for learning sim agents through self-play PPO

D.1 Collision behavior

GPUDrive supports three types of collision behaviors: *ignore*, *remove*, and *stop*. Each of these has different effects on the types of behaviors agents learn over time. We briefly outline some things to be aware of below, which might be useful for future experiments.




Figure 9: **Probability distribution function for each type of error for scenes that are not fully solved..** *Left:* Percentage of agents that collided. *Middle:* Percentage of agents that went off-road. *Right:* Percentage of agents that neither failed nor reached their goal. Note that almost all scenes contain just a single failure.

Ignoring collisions When collision behavior is ignored, the agent is not terminated when it collides with another agent or touches a road edge. As such, it can proceed to the goal and collide within a single episode. To discourage collisions, it seems reasonable to give agents a penalty. However, since, in most scenarios, the probability of getting negative signals in an episode with random behavior (e.g. hitting a road edge) is significantly larger than the probability of receiving a positive signal (getting to the goal), the value function may become overly pessimistic because the majority of the advantages the agent is receiving will be negative, and as such the probability of actions that lead to these negative advantages, such as higher acceleration, will be decreased. This can lead to a behavior where agents freeze (they learn to stay on the road) and do not head towards the goal. This can be avoided by ensuring that agents receive enough positive signals along with negative ones, especially early on during learning. This can be achieved by sufficient exploration through a large enough entropy coefficient.

Removing agents at collision Another option is to simply terminate agents whenever they do something that is not desired (in our case colliding) without assigning penalties (giving negative rewards). This means that the goal can only be achieved if the agent does not do something bad. Since the penalty in this case is implicit, the value function can not become overly pessimistic. Instead, the advantages will be 0 most of the time early on in training. Once the first positive signals are achieved by accident (which is inevitable given the small maps of the WOMB and a high enough entropy coefficient), the probability of the right action sequences will be increased until all agents hit their goals without colliding or going off-road.

E Analyses.

E.1 Detecting out-of-distribution events

To find rare goals, we analyzed the full training dataset (~ 129,000 scenes) or about 4.2 million controllable agents. Of these, we found approximately 30,000 agents (0.73%) making a U-turn, and 47,000 agents (1.13%) driving in reverse (see Appendix E for the exact definition of these events). Further, most agents driving in reverse were simply pulling out of park, with goals immediately behind them. We observed a distinct lack of goals where the agent needs to execute a complex U-turn, making it plausibly out of distribution. We then hand-designed 13 scenarios from the test dataset with a total of 27 agents across all scenes, placing goals behind agents. This was done by setting the new goal for each agent to $(x_f - x_i, y_f - y_i)$, where (x_i, y_i) is the initial position, and (x_f, y_f) is the original goal. We chose the scenes in such a way that doing this process for all controlled agents results in valid and reachable goals. Figure 5.

1. **U-turn:** For each time step t where the agent is valid, we check the condition: $\text{abs}(\text{heading}[t] - \text{heading}[\text{initial}]) > 150^\circ$.
2. **Driving in reverse:** For each time step t where the agent is valid, calculate the direction of its velocity vector and subtract it from its heading angle. If the absolute difference is greater than a threshold (150°), it

Table 6: Overview of collision behaviors

Collision behavior	Pro's	Caveats
Ignore	1) Agents receive a diverse range of observations	1) Value function can become overly pessimistic; 2) Large exploration space: A large set of possible states leads to agents seeing lots of useless observations during exploration
Stop	1) Closest to the real-world effect of collisions	1) Introduces extra challenge during learning: drive around other stopped agents
Remove	1) Simple	1) Might lead to unrealistic behavior when used and agents are not removed from the scene; 2) Might be difficult to reach certain states because the agents are always removed upon collisions; 3) Number of completed episodes is large in the beginning since most agents are terminated within 10 steps and subsequently a lot of <code>reset()</code> calls early on in training, which decreases the SPS.

is driving in reverse. Note: We only detect driving in reverse if it occurs for more than a threshold (10) consecutive steps, and above a minimum magnitude velocity (0.5 km/hr).

E.2 Placing goals behind agents

Table 7: Aggregate performance comparison between Altered and Original goal positions (mean \pm std).

Class	Goal achieved \uparrow	Collided \downarrow	Off-road \downarrow	Other \downarrow
Altered	53.5 \pm 38.4	10.0 \pm 8.3	6.7 \pm 16.1	41.7 \pm 32.0
Original	100.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0

F PPO implementation details.

F.1 Hyperparameters

Table 8 reports the hyperparameters used for the results in our experiments.

G Compute resources

Experiments were run on either a single NVIDIA A100 or an RTX4080 device for 12-36 hours per experiment. Including hyperparameter tuning and experimentation, all runs combined for this paper took approximately 5 GPU days.

Table 8: PPO Algorithm Hyperparameters

Parameter	Value	Description
total_timesteps	1,000,000,000	Total number of timesteps for training.
batch_size	524,288	Number of timesteps collected in each rollout.
minibatch_size	16,384	Number of timesteps in each minibatch for gradient updates.
learning_rate	3e-4	Initial learning rate for the optimizer.
anneal_lr	false	Whether to anneal the learning rate over time.
gamma	0.99	Discount factor for future rewards.
gae_lambda	0.95	Lambda parameter for Generalized Advantage Estimation.
update_epochs	2	Number of epochs to update the policy network per rollout.
norm_adv	true	Whether to normalize advantages during training.
clip_coef	0.2	PPO clipping coefficient for policy updates.
clip_vloss	false	Whether to clip the value loss.
vf_clip_coef	0.2	Clipping coefficient for value function updates.
ent_coef	0.0001	Entropy regularization coefficient to encourage exploration.
vf_coef	0.5	Coefficient for the value function loss in the total loss.
max_grad_norm	0.5	Maximum norm for gradient clipping.
target_kl	null	Target KL divergence for policy updates (unused if null).
collision_weight	-0.75	Penalty weight for collision events.
off_road_weight	-0.75	Penalty weight for off-road events.
goal_achieved_weight	1.0	Reward weight for achieving the goal.