

## 571 Appendix

572 In this appendix, we provide more details about the four experiments and some scenario examples  
573 from the three databases used in the experiments. Section A is on the scenario generation model;  
574 Section B is on the single-agent cross-dataset generalization experiment; Section C is on the multi-  
575 agent reinforcement learning and imitation learning; Section D shows the interface of ROS bridge and  
576 the qualitative results of running Openpilot test. Section E shows more rendered scenario examples  
577 from each database.

578 Codebase, documentation, videos, and a scenario gallery are available at [https://metadriverse.  
579 github.io/scenarionet](https://metadriverse.github.io/scenarionet).

## 580 A Scenario Generation Model

581 We use all the databases (nuPlan, Waymo, PG) in scenarioNet for cross-dataset traffic scenario  
582 generation experiments. The goal is to show that scenarioNet contains diverse traffic scenario data  
583 for training deep neural networks. We conduct our experiments based on TrafficGen [18], a neural  
584 generative model previously developed for reactive traffic scenario generation.

### 585 A.1 Task Definition

586 A traffic scenario is denoted as  $\tau = (\mathbf{m}, \mathbf{s}_{1:T})$ , which lasts  $T$  time steps and contains the High-  
587 Definition (HD) road map  $\mathbf{m}$  and the state series of traffic vehicles  $\mathbf{s}_{1:T} = [s_1, \dots, s_T]$ . Each element  
588  $s_t = \{s_t^1, \dots, s_t^N\}$  is a set of states of  $N$  traffic vehicles at time step  $t$ . Given an existing scenario  
589  $\tau = (\mathbf{m}, \mathbf{s}_{1:T})$ , the goal of TrafficGen is to learn to generate **new** traffic scenarios  $\tau' = (\mathbf{m}, \mathbf{s}'_{1:T'})$   
590 that have similar distribution with  $\tau$  and different states  $\mathbf{s}'$  and longer time steps  $T'$ . After training,  
591 TrafficGen takes  $\mathbf{m}$  as input and generates  $\tau'$  as a totally different scenario.

### 592 A.2 Model Architecture

593 **Traffic Scenario Encoder.** TrafficGen uses vectorization to encode map and vehicle information,  
594 representing lanes as sets of vectors, each vector representing a small region. A vector-based  
595 coordinate system is established for each small region, with each vector comprising a start point  $p_i^s$ ,  
596 endpoint  $p_i^e$ , and information about vehicles in this region. Thus, a traffic snapshot  $\tau$  at time step  $t$  is  
597 represented as  $\tau_t = \mathbf{v} = v_{i=1}^I$ . Cross attention mechanism is applied to the unordered set  $\mathbf{v}$  to fuse  
598 information from different regions into a single context vector.

599 **Decoder for Scenario Generation.** TrafficGen places vehicles by generating a set of weights for all  
600 regions, which is then turned into a categorical distribution. The local position of a tentative vehicle  
601 is modeled by a mixture of  $K$  bivariate normal distributions, as are heading, speed, and size of the  
602 vehicle. Autoregressive sampling is used to create a traffic snapshot. A motion forecasting model is  
603 used as the trajectory generator.

### 604 A.3 Experiment Setting

605 We train a scenario generation model TrafficGen with mixed data. Specifically, we randomly sample  
606 30% data from nuPlan, Waymo, and PG. We filter out the scenarios with less than 8 cars and crop a  
607 rectangular area with a side length of 120m centered on the ego vehicle. Each of the 20s scenarios is  
608 split into 10 traffic snapshots with 2s intervals. The training is executed on servers with 8 x **Nvidia**  
609 2080TI and 256 G memory in Distributed Data-Parallel (ddp) mode. We set the feature size to be  
610 1024, and use 3-layer MLPs with hidden dimensions of [1024, 512, 256] for attribute modeling. The  
611 training takes 16 hours for 100 training epochs and the learning rate decays by 20% at every 30  
612 epochs. The detailed hyperparameters are shown in Table 4.

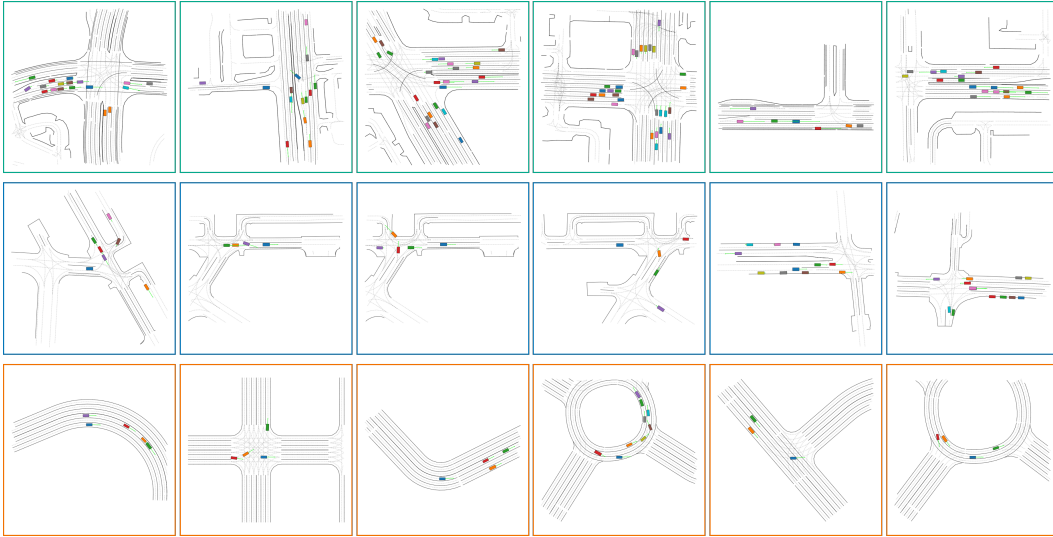


Figure 6: Traffic scenarios generated from the model trained on different databases: Waymo (□), nuPlan (□), and PG (□).

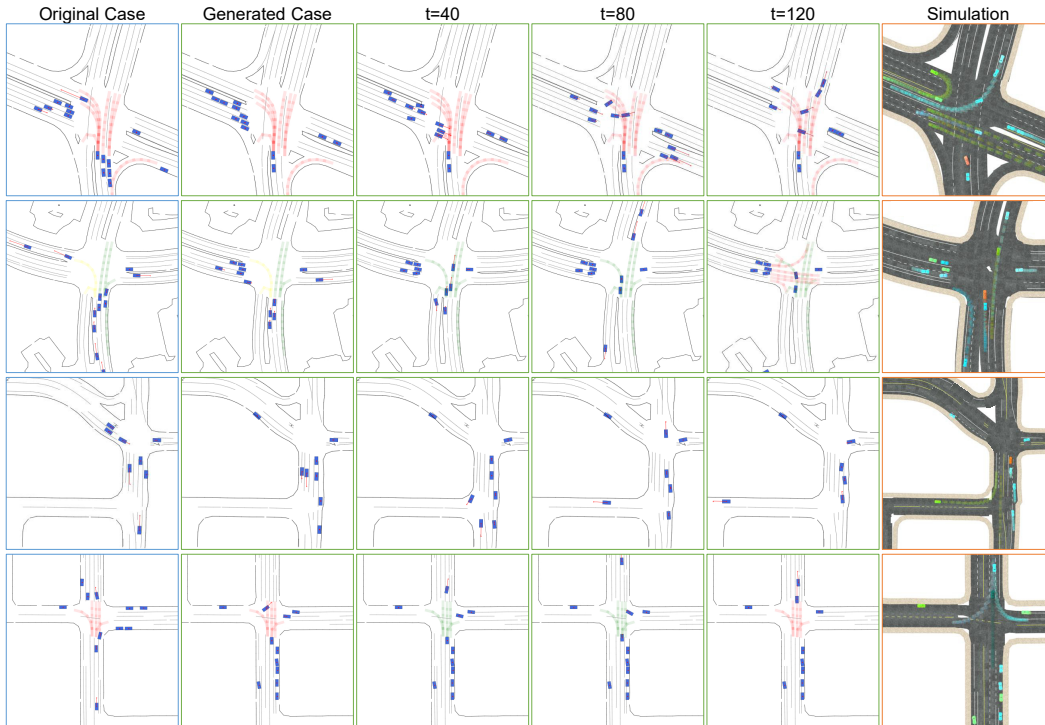


Figure 7: Dynamics of the generated traffic scenarios. The first column is the original case. The middle columns show the generated scenarios at different timesteps. The last column shows the corresponding scenarios imported in the simulation. The green and red dashed lines indicate the traffic light status of this lane at the intersection.



Figure 8: t-SNE visualizations of 3000 scenarios. PG scenarios (□) are located mostly on the left side, while the scenarios from nuPlan (□) and the Waymo scenarios (□) are scattered on the right side.

613 We plot some generated traffic scenarios in Fig. 6 and the dynamics in Fig. 7. It shows that the models  
 614 trained on the processed scenario data from ScenarioNet can generate realistic and diverse traffic  
 615 behaviors and interactions.

#### 616 A.4 t-SNE Visualizations

617 The trained model’s encoder is used to extract feature embeddings of a given scenario sample. The  
 618 embeddings are then visualized with t-SNE method to show the similarities and differences. The  
 619 detailed hyperparameters of t-SNE are shown in Table 5.

620 The t-SNE result is plotted in Fig. 8. The clustering results show that there is a large domain  
 621 gap between real-world scenarios (Waymo and nuPlan) and synthetic scenarios (PG). Besides,  
 622 domain gap exists even in real-world datasets. One prominent feature of the Waymo dataset is the  
 623 complex crossroad (shown in the lower right corner), which the nuPlan dataset is lacking. The t-SNE  
 624 visualization reveals the differences between different traffic datasets.

Table 4: TrafficGen

Hyper-parameter	Value
Batch Size	256
Feature Size	1024
Training epochs	100
Learning Rate	$3e-4$
Activation Function	“relu”
MCG Layers	5

Table 5: t-SNE

Hyper-parameter	Value
Components Number	2
Init	pca
Learning Rate	auto
Perplexity	30
Early Exaggeration	12

## 625 B Single-agent Cross-dataset Generalization Experiment

626 We use the PG database and nuPlan database for cross-dataset generalization experiments. The goal  
 627 is to investigate how the sim-to-real gap affects the generalizability of the learning-based vehicle  
 628 controller. To this end, we train agents on synthetic PG scenarios [28] and nuPlan [6] scenarios  
 629 respectively, and test them on the same held-out real-world test set.

### 630 B.1 Task Setup

631 To be specific, the task is to follow the trajectory of the data collection car and drive as fast as possible  
 632 while avoiding collisions.

633 **Observation.** The observation of the RL agents is as follows:

- 634 • A 120-dimensional vector denoting the Lidar-like point clouds with  $50m$  maximum detecting  
 635 distance centering at the target vehicle. Each entry is in  $[0, 1]$  with Gaussian noise and  
 636 represents the relative distance of the nearest obstacle in the specified direction.
- 637 • A vector containing the data that summarizes the target vehicle’s state such as the steering,  
 638 heading, velocity, and relative distance to the trajectory to follow.
- 639 • The navigation information that guides the target vehicle toward the destination. Concretely,  
 640 it consists of 10 points sampled on the future trajectory and the distance between two  
 641 consecutive points is  $2m$ . The points will be projected to the vehicle coordinates.
- 642 • A 12-dimensional vector denoting the Lidar-like point clouds with  $50m$  maximum detecting  
 643 the boundary of the drivable area, like the solid lines or sidewalks. (Optional)

644 As the vehicle for collecting nuPlan data in Boston sometimes drives out of the drivable area for  
 645 bypassing the cones or barriers, crossing the drivable area boundary usually happens. Therefore, we  
 646 didn’t use the boundary detector in our experiments, and crossing the boundaries like solid lines  
 647 won’t terminate the episode nor penalize the agent.

648 **Action.** The driving policy is a fully end-to-end model and directly controls the low-level throttle  
 649 and steering angle. The action  $a$  is a continuous two-dimensional vector with entries in  $[-1, 1]$ . By  
 650 multiplying coefficients and clipping the extreme value, the action will be converted into the engine  
 651 force and steering angle for changing the vehicle states.

652 **Reward and Cost Scheme.** The reward function is composed of four parts as follows:

$$R = c_1 R_{disp} + c_2 P_{smooth} + c_3 P_{collision} + R_{term}. \quad (1)$$

653 The *displacement reward*  $R_{disp} = d_t - d_{t-1}$ , wherein the  $d_t$  and  $d_{t-1}$  denotes the longitudinal  
 654 movement of the target vehicle in Frenet coordinates of the target trajectory between two consecutive  
 655 time steps, providing a dense reward to encourage the agent to move forward. The *smooth penalty*  
 656  $P_{smooth} = \min(0, 1/v_t - |a[0]|)$  incentivizes the agent to drive smoothly and avoid a large steering  
 657 value change between two timesteps, especially, when the velocity is high.  $v_t$  and  $a[0]$  denote the  
 658 current velocity and the steering value respectively. In addition, if a collision with a vehicle, human,  
 659 or object happens at timestep  $t$ , the agent will receive a *collision penalty*. The penalty is set to  
 660  $P_{collision} = 2$  when colliding with a human or vehicle and  $P_{collision} = 0.5$  for colliding with an  
 661 object like cones and barriers. We also define a sparse *terminal reward*  $R_{term}$ , which is non-zero

662 only at the last time step. At that step, we set  $R_{disp} = R_{speed} = 0$  and assign  $R_{term}$  according to the  
663 terminal state.  $R_{term}$  is set to +10 if the vehicle reaches the destination,  $-5$  for being  $2.5m$  away  
664 from the reference trajectory. We set  $c_1 = 2$ ,  $c_2 = 1$  and  $c_3 = 1$ .

665 **Termination Conditions and Evaluation Metrics.** The episode will be terminated only when: 1)  
666 the agent drives  $2.5m$  away from the reference trajectory 2) the agent arrives at the destination and  
667 3) the agent can not finish the episode in  $recorded\_episode\_length + 50$  steps. For each trained  
668 agent, we evaluate it in the held-out test environments and define the ratio of episodes where the  
669 agent arrives at the destination as the *success rate*. The definition is the same for *out of road* and  
670 *Timeout*. For evaluating the driving behavior, the mean velocity in each scenario is collected and  
671 then averaged across all scenarios. Also, a metric similar to the *success rate* is measured and called  
672 *route completion*, which is the ratio of moving distance to the length of the whole reference trajectory.  
673 Since each agent are trained across 5 random seeds, this evaluation process will be executed for 5  
674 agent which has the same training setting but different random seeds. We report the average and std  
675 on the metrics mentioned above.

## 676 B.2 Curriculum Training System

677 In the single-agent experiments, 40,000 scenarios are used for training agents in both PG scenarios  
678 and nuPlan scenarios. Loading each scenario from scratch costs a significant amount of time, and thus  
679 we would like to buffer the scenarios in RAM for repeated use. However, the memory consumption  
680 is nonnegligible, especially when we train 5 policies concurrently and launch 20 workers to collect  
681 rollout for each policy. Assuming each scenario consumes about 10MB of memory, buffering 40,000  
682 training scenarios in each worker process consumes  $40,000 \times 100 \times 10MB = 40,000GB = 40TB$  of  
683 memory, which requests a powerful and expensive cluster to train agents in large-scale scenarios. We  
684 adopt **curriculum training** scheme for overcoming this issue, which reduces the memory by 99.9%.

685 **Curriculum Training.** We first sort the training scenarios according to the difficulty score calculated  
686 by:  $track\_length \times cumulative\_curvature$ .  $track\_length$  is the moving distance of the data  
687 collection car. A greater moving distance corresponds to a higher velocity, which in turn indicates  
688 a higher difficulty score. This value then will be multiplied by a weight  $cumulative\_curvature$   
689 which quantifies the level of bending in the trajectory. After this, scenarios can be divided into 100  
690 levels with 400 scenarios in each level. Only when the *Success Rate* reaches 75%, the worker will  
691 move to the next level and release the memory used to store scenarios of the previous level. To further  
692 reduce memory usage, we split the scenarios in each level into 20 subsets, and each worker only loads  
693 scenarios from the corresponding subset. Finally, only 0.05% scenarios (20 scenarios) are actually  
694 loaded in each worker, which saves memory by a large margin.

695 This curriculum training scheme not only makes it possible to finish training on a single server but  
696 boosts training efficiency and performance. We conduct an ablation study without the curriculum  
697 training, which costs a long time to train. This training scheme releases every scenario after using and  
698 reloading it when needed. The inferior performance highlights the necessity of having curriculum  
699 training.

## 700 B.3 Results

701 As shown in Fig. 9, we present experiment results for three training settings: *nuPlan with curriculum*  
702 *training*, *nuPlan without curriculum training*, *PG with curriculum training*. The three used databases  
703 are *PG-train*, *nuPlan-train*, and *nuPlan-test*, comprising 40,000, 40,000, and 5,000 scenarios re-  
704 spectively. Considering *nuPlan with curriculum training* as the baseline, our analysis focuses on  
705 two factors: the data source and the inclusion of curriculum training. The evaluation of real-world  
706 scenarios underscores the significance of training with in-distribution real-world data and employing  
707 a curriculum training system.

708 Actually, the poor performance of *nuPlan without curriculum training* is also exposed at the training  
709 stage. In spite of increasing the *data coverage*, the agents under this setting always have a poor  
710 *training success rate* and *training route completion* throughout the whole training process. However,

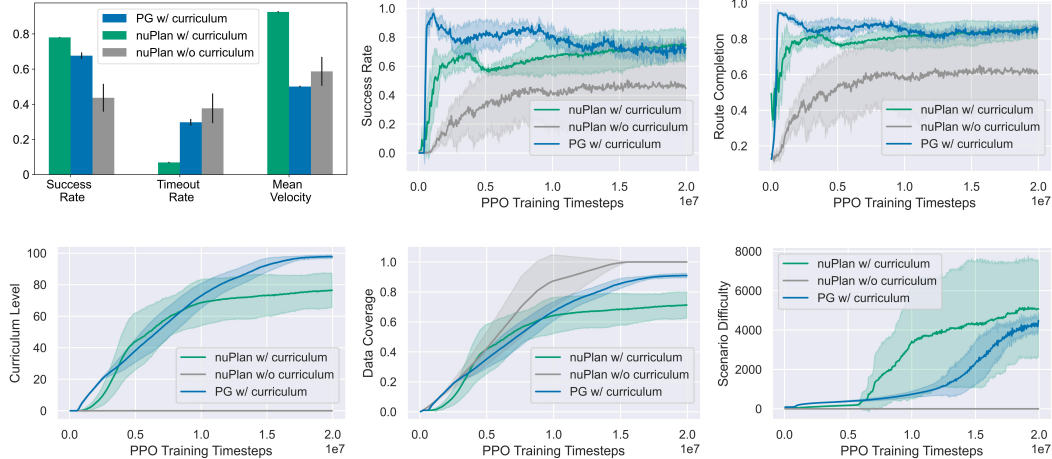


Figure 9: The upper left figure shows the evaluation results for three sets of agents trained in different settings. The other two upper figures show the learning dynamics including *training success rate* and *training route completion*. The bottom figures show the status of the curriculum training system. The *curriculum level* is calculated as the average worker level across 5 seeds, with 20 workers per seed. *Data coverage* refers to the proportion of scenarios in the training set that the agent encounters throughout the entire training process. A value of 1.0 indicates that the agent has visited all scenarios. *Scenario difficulty* is determined by calculating the average difficulty score across all scenarios collected during a PPO optimization epoch.

711 agents trained in *nuPlan with curriculum training* setting are more stable and have an increasing  
 712 *training success rate* as they can fully exploit the scenarios which properly match their current driving  
 713 ability.

714 For investigating the influence of data sources, the conclusion can be drawn only from the test results  
 715 which show that agents trained under *PG with curriculum training* can not generalize to real-world  
 716 settings well. And the failure mode is that agents trained in synthetic scenarios can not learn to  
 717 drive at high speed. And they don't move until other vehicles move far from them. This behavior is  
 718 reflected on the *test mean velocity* as well.

719 Another interesting phenomenon is the drop of *training success rate* when the *scenario difficulty* and  
 720 *curriculum level* increase, which is also reported in [8] as well. And this value finally coverage to the  
 721 test *success rate*. Therefore, we infer that given enough training scenarios, the test performance of  
 722 learning-based agents can be roughly reflected by the training-time performance.

723 The hyper-parameters of the RL training are listed in Table 6

Table 6: PPO

Hyper-parameter	Value
KL Coefficient	0.2
$\lambda$ for GAE [39]	0.95
Discounted Factor $\gamma$	0.99
Number of SGD epochs	20
Train Batch Size	50,000
SGD mini batch size	200
Learning Rate	$1e-4$
Clip Parameter $\epsilon$	0.2
Activation Function	"tanh"
MLP Hidden Units	[512, 256, 128]
MLP Layers	3

## 724 C Multi-agent Policy Learning

### 725 C.1 Experiment Setting

726 We can load the real-world dataset into MetaDrive simulator and create multi-agent interactive policy  
727 environment. We first instantiate agents in the environment where their initial states are loaded from  
728 the real-world dataset. The initial states include position, heading, velocity and the size. Then, we  
729 assign `EnvInputPolicy` to all agents and allow them to be controlled by external RL policies.

730 The ground-truth (GT) trajectories are not accessible to the learning agents but serve as the supervision  
731 via reward function (in RL) or observations (in IL).

732 Concretely, we create the reward function similar to Eq. 1 with slightly different weights:

$$R = R_{disp} + P_{collision} + R_{term}. \quad (2)$$

733 The *displacement reward*  $R_{disp} = d_t - d_{t-1}$ , wherein the  $d_t$  and  $d_{t-1}$  denotes the longitudinal  
734 movement of the target vehicle in Frenet coordinates of the **target trajectory** between two consecutive  
735 time steps, provides a dense reward to encourage the agent to move forward. In addition, if a collision  
736 with a vehicle, human, or object happens at timestep  $t$ , the agent will receive a *collision penalty*.  
737 The penalty is set to  $P_{collision} = 1$  when colliding with a human or vehicle. We also define a  
738 sparse *terminal reward*  $R_{term}$ , which is non-zero only at the last time step. At that step, we set  
739  $R_{disp} = R_{speed} = 0$  and assign  $R_{term}$  according to the terminal state.  $R_{term}$  is set to +10 if the  
740 vehicle reaches the destination, -1 for being 10m away from the reference trajectory. We transform  
741 the GT trajectory into the simulator to create the target trajectory and use the displacement reward as  
742 the primary supervision from the dataset.

743 On the other hand, in multi-agent imitation learning we use the GT trajectory to form a dataset of  
744 observations and follow the setting of learning from observations or say action-free imitation learning.  
745 Specifically, for each frame in the scenario, we load the states of all actors at this frame into the  
746 simulator and utilize the sensor simulation functionality of MetaDrive to simulate the observations of  
747 each actor. The observation follows Sec. B.1. We form the dataset of observation sequences of all  
748 actors in the dataset.

749 The ground-truth trajectory can be used to measure the learned behaviors. The the route completion  
750 rate is the ratio between the length of projected agent trajectory and the length of GT trajectory. The  
751 average distance between agent trajectory and the GT trajectory is computed as follows:

$$\frac{1}{T} \sum_{t=1}^T \|pos_{agent,t} - pos_{GT,t}\|. \quad (3)$$

752 where  $T$  is the minimum length of agent trajectory and the length of GT trajectory. Therefore, for  
753 the agents that terminate quickly after spawning, the average distance will be quite small. The final  
754 distance is computed as distance between the last position of GT trajectory and the last position of  
755 agent trajectory. The cost is the number of crashes of an agent in one episode. There are two terminal  
756 conditions. If the route completion rate exceeds 95%, the agent is marked successful. If the agent  
757 moves out of 10m away from the reference trajectory, the agent is marked out of road and failed.  
758 We don't terminate agent's episode if it crash with other objects.

### 759 C.2 Baseline Details

760 **MA-GAIL.** We use GAIL [23] in multi-agent setting [43] but the discriminator distinguishes state-  
761 next state pair, instead of state-action pair [47]. We use PPO as the underlying RL algorithm in GAIL.  
762 The hyper-parameters are listed in Table 7.

763 **MA-AIRL.** We also train multi-agent Adversarial Inverse RL [59] (AIRL) with an additional  
764 inverse dynamics model for estimating the expert actions. The inverse dynamics model is trained  
765 concurrently with the AIRL policies and learns the action given state-next state pairs from the  
766 environment interactions. AIRL will learn a reward function and use the reward function to build a

767 discriminator as GAIL. We also use PPO as the underlying RL algorithm. The hyper-parameters are  
 768 listed in Table 8.

769 **MARL baselines.** We train independent PPO [40, 38] and TD3 [21] agents as well as the Coordinated  
 770 Policy Optimization (CoPO) agents [36] as MARL baselines. The hyper parameters are given in the  
 771 next section.

### 772 C.3 Hyper-parameters

Table 7: Action-free Multi-agent GAIL		Table 8: Action-free Multi-agent AIRL	
Hyper-parameter	Value	Hyper-parameter	Value
Discriminator LR	5e-5	Discriminator Loss LR	3e-4
Discriminator L2 Norm	1e-5	Discriminator L2 Norm	1e-5
Discriminator SGD Num Iters	1	Discriminator SGD Num Iters	5
Discriminator SGD Minibatch Size	1024	Discriminator SGD Minibatch Size	1024
PPO SGD Minibatch Size	512	Inverse Dynamics SGD Num Iters	100
PPO Batch Size	2000	Inverse Dynamics SGD Minibatch Size	512
PPO LR	1e-4	Inverse Dynamics LR	1e-4
PPO SGD Num Iters	10	PPO SGD Minibatch Size	512
PPO Clip Parameter	0.2	PPO Batch Size	2000
PPO Lambda	0.95	PPO LR	1e-4
PPO Gamma	0.99	PPO SGD Num Iters	10
		PPO Clip Parameter	0.2
		PPO Lambda	0.95
		PPO Gamma	0.99

Table 9: PPO	
Hyper-parameter	Value
PPO SGD Minibatch Size	512
PPO Batch Size	2000
PPO LR	1e-4
PPO SGD Num Iters	10
PPO Clip Parameter	0.2
PPO Lambda	0.95
PPO Gamma	0.99

Table 10: CoPO	
Hyper-parameter	Value
LCF LR	1e-4
LCF Num Iters	5
Neighborhood Distance	40 m
PPO SGD Minibatch Size	512
PPO Batch Size	2000
PPO LR	1e-4
PPO SGD Num Iters	10
PPO Clip Parameter	0.2
PPO Lambda	0.95
PPO Gamma	0.99

Table 11: TD3	
Hyper-parameter	Value
Critic LR	1e-4
Actor LR	1e-4
Tau for Target Update	5e-3
Train Batch Size	100



## 773 D AD stack testing

### 774 D.1 ROS bridge

775 As shown in Fig. 10, We provide a ROS bridge along with the *ScenarioNet*, allowing users from the  
776 ROS community to develop and test their systems with massive real-world data. As shown in Fig. 11,  
777 information like camera, lidar, and mid-level representations can be retrieved from the simulation.




Figure 10: The interface of ROS bridge allowing connecting *ScenarioNet* and ROS community.




Figure 11: Multimodal sensory data provided by *ScenarioNet*.

778 *ScenarioNet* provides not only mid-level scenario representations but multiple sensor outputs like  
779 top-down view, RGB camera, depth camera, and cloud points. The deluxe RGB camera is supported  
780 by the deferred rendering pipeline. This figure shows the `scene-0061` from nuPlan-mini-split and  
781 its digital twin counterparts. It is worth investigating how to reconstruct meshes from the recorded  
782 lidar cloud points and images so that we can simulate the sensor output from new views different  
783 from the recorded ones. Besides, it is a promising topic to study the **closed-loop** sensor fusing and  
784 learning-based control together, which can be supported by *ScenarioNet*.

785 **D.2 Qualitative Results of Openpilot test**

786 Openpilot [9] is an end-to-end solution for driver assistance. Therefore, a platform providing 3D  
787 rendering is necessary, if one would like to study or test the system. *ScenarioNet* is the only one  
788 that provides both real-world scenarios and 3D graphics support. As its navigation module is still a  
789 beta version, we only test Openpilot in scenarios without diverged roads. We build a small database  
790 containing mainly lane-keeping scenarios for conceptually demonstrating that *ScenarioNet* can  
791 connect with commercial AD stack. As shown in Fig. 12, the Openpilot system is robust to common  
792 scenarios like turning right, lane-keeping, stopping at traffic lights, and passing traffic lights. The  
793 demo video is available at <https://youtu.be/Kj1PB0nCTvg>



Figure 12: Openpilot manages to overcome four representative scenarios.

794 **E Scenario Databases**

795 In the database statistics Table, the *Intersection Ratio* is the ratio of scenarios having traffic lights for  
796 real-world data, and the ratio of scenarios containing intersections and roundabouts for synthetic data.

797 **E.1 Waymo Database**

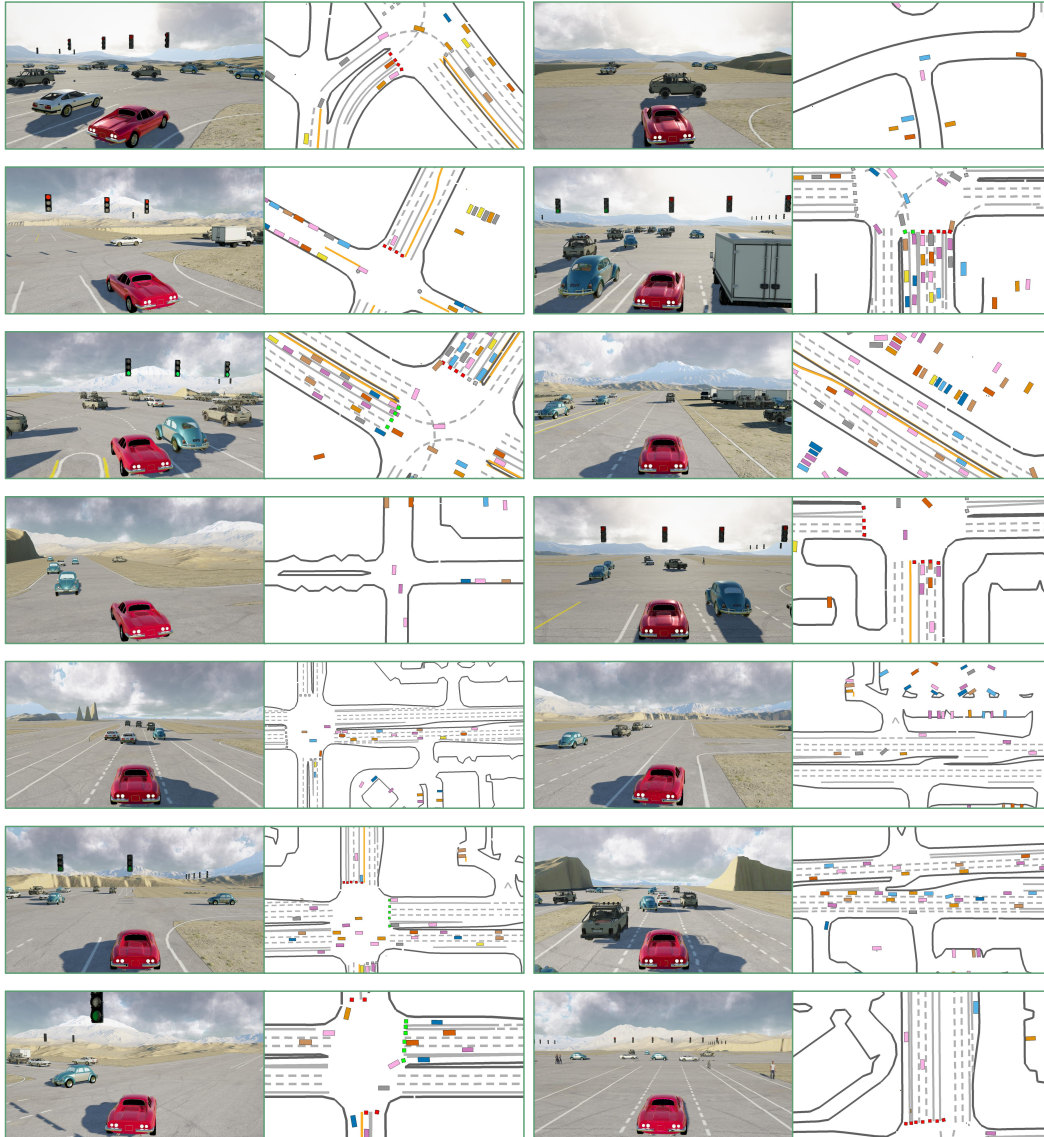


Figure 13: Rendered traffic scenarios and their top-down views from the Waymo database.

798 We construct the Waymo database from the **20 seconds** version [44] motion data with Google Cloud  
799 path at `waymo_open_dataset_motion_v_1_2_0/uncompressed/scenario/training_20s`.  
800 As Waymo data is collected with the altitude calibrated, we can filter the overpass scenarios by  
801 excluding the scenarios with significant changes in height. In addition, we exclude the scenarios  
802 where the ego car waits at the red light for a long time by selecting scenarios where the ego car  
803 moving distances are greater than 10 meters. We provide the rendered scenario examples from this  
804 dataset in Fig. 13. Top-down views show that Waymo scenarios contain diverse road structures and a  
805 large number of vehicles.

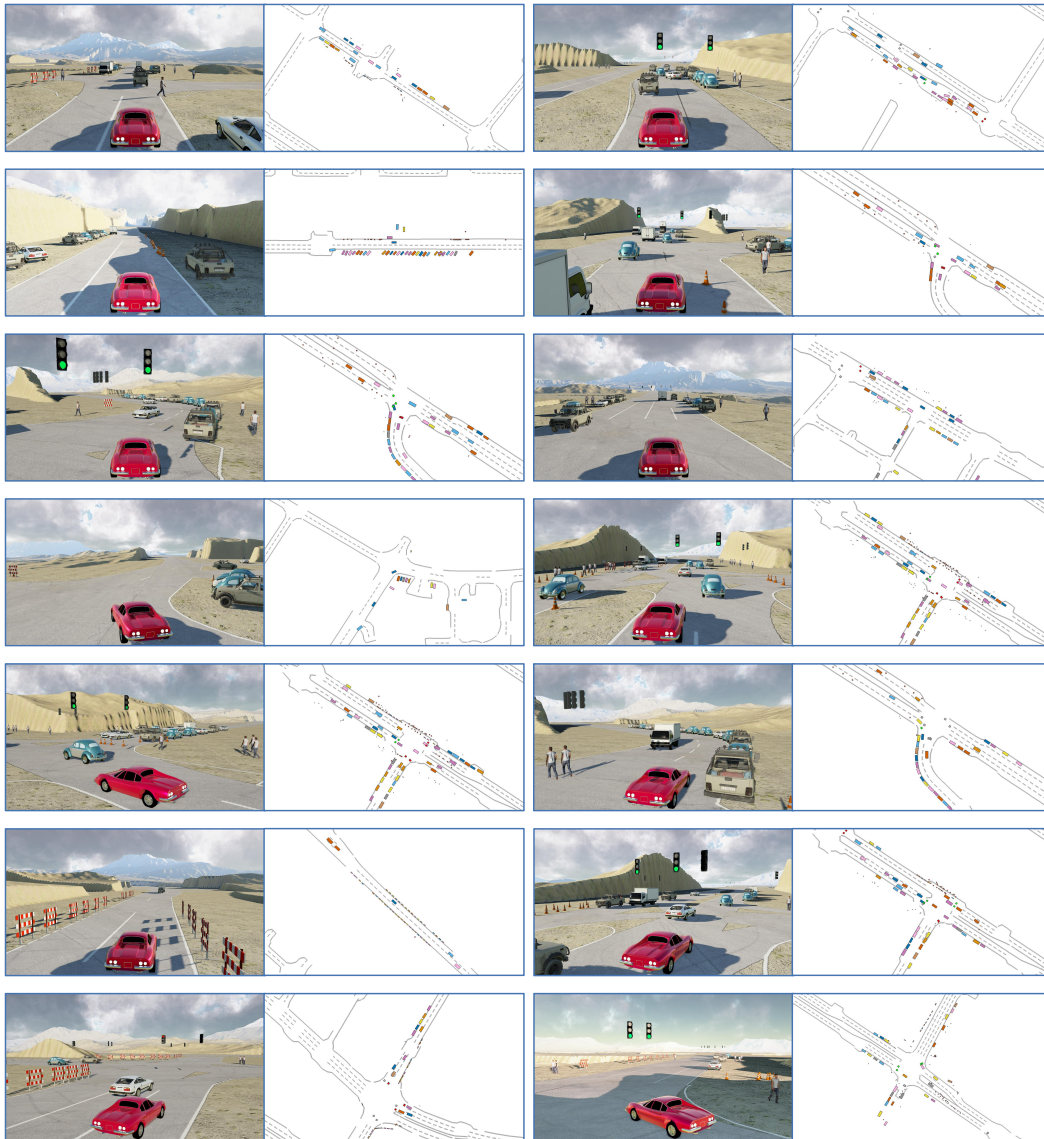


Figure 14: Rendered traffic scenarios and their top-down views from the nuPlan database.

807 According to <https://www.nuscenes.org/nuplan>, nuPlan has more than 1500 hours of driving  
 808 data collected in 4 different cities: Las Vegas, Singapore, Pittsburgh, and Boston, here we only use  
 809 the data collected in Boston as we want to keep all databases in experiments to have similar sizes.  
 810 Data collected in other cities can also be converted to our scenario format.

811 We use the V.1.0 version data, which contains approximately 50,000 scenarios after excluding  
 812 scenarios where the ego car moving distances are less than 10 meters. It is noticeable that nuPlan  
 813 updated the data to version v1.1 recently (one week before the NeurIPS 2023 dataset track deadline),  
 814 which may induce some differences if building a database from this new version. As shown in Fig. 14,  
 815 the scenario examples reveal that *nuPlan Boston split* has cluttered scenes and contains many traffic  
 816 cones and barriers besides vehicles. We additionally find that in some nuPlan scenarios, the ego car  
 817 trajectory is out of the drivable area for sidestepping the barriers or cones.

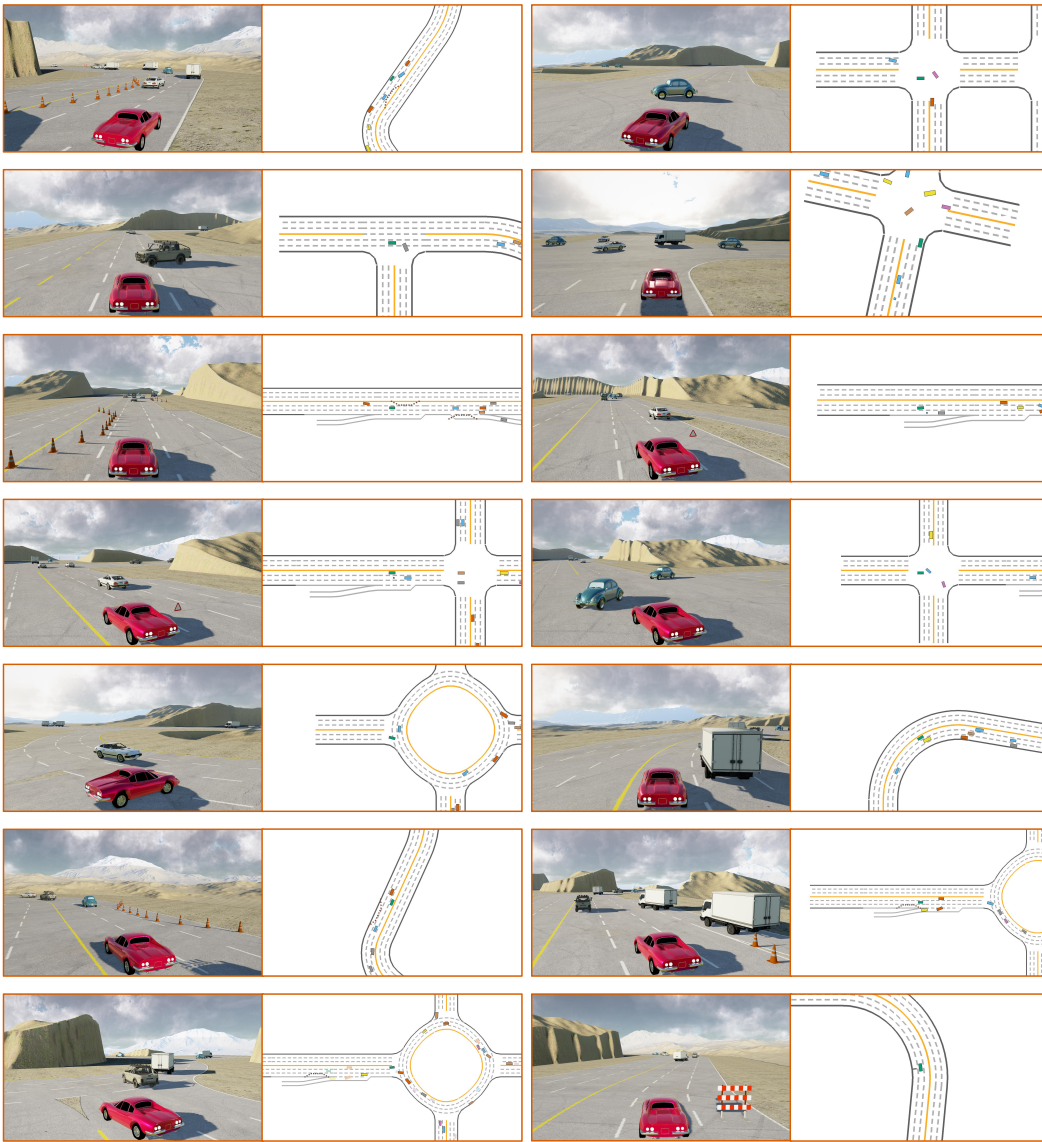


Figure 15: Rendered traffic scenarios and their top-down views from the PG database.

819 Unlike the previous two datasets collected in the real world, PG scenarios are synthesized according to  
 820 a set of rules. For map generation, two blocks are sampled from a set of candidate roadblocks  
 821 and connected to form a map. Those blocks include intersections, roundabouts, straight roads,  
 822 curved roads, Ramp, and so on. Once the map is defined, a traffic generation rule will be  
 823 applied to scatter vehicles and road objects like traffic cones on the map. The detailed scenario  
 824 generation config such as the block distribution for sampling can be found at [https://github.com/metadriverse/metadrive/blob/0a929f8130b34e4428067390f20f872d1d6d224a/metadrive/component/algorithm/blocks\\_prob\\_dist.py#L4](https://github.com/metadriverse/metadrive/blob/0a929f8130b34e4428067390f20f872d1d6d224a/metadrive/component/algorithm/blocks_prob_dist.py#L4).  
 825 All vehicles choose a destination automatically and be actuated by IDM policy which can keep a proper distance from the front  
 826 vehicle and perform a lane change when the front object is static. The scenario data is collected by an  
 827 IDM policy as well. Some scenario examples are shown in Fig. 15.  
 828  
 829

830 **F Discussion: how to improve visual fidelity?**

831 The visual fidelity of the closed-loop simulation can be further improved with the collected raw  
832 sensor data. We already have some plans to improve it and will include this discussion in the paper  
833 for sharing our thoughts regarding realistic sensor simulation.

834 Currently, the sensor simulation including the camera and lidar is achieved in a Computer Graphics  
835 (CG) way, where people try restoring the mesh and texture for objects from real-world data and  
836 shading these models based on lights and materials to make them visually realistic. To improve the  
837 rendering results of *ScenarioNet*, we do plan to reconstruct the geometry data for traffic participants  
838 and objects from the driving videos. Besides, we also consider connecting *ScenarioNet* with Unreal  
839 Engine or Nvidia Omniverse in the future as they could provide better shading results.

840 On the other hand, NeRF is an alternative to improve the quality of sensor simulation. Through  
841 volume rendering, it can directly synthesize new camera views and point clouds from the driving  
842 videos when traffic participants and the ego car move with different trajectories and poses in the  
843 closed-loop training. This way is purely data-driven and can exempt the need for restoring 3D assets  
844 like objects and buildings. Recent results [55, 58, 54] already demonstrate its potential in terms of  
845 camera and lidar simulation. However, how to make the NeRF scene editable is still an open problem;  
846 hence, we plan to investigate this in the future.