TrajGen: Generating Realistic and Diverse Trajectories with Reactive and Feasible Agent Behaviors for Autonomous Driving

Qichao Zhang[1*], Yinfeng Gao[2*], Yikang Zhang[1*], Youtian Guo[1], Dawei Ding[2],
Yunpeng Wang[3], Peng Sun[3], Dongbin Zhao[1], *IEEE Fellow*

Abstract—Realistic and diverse simulation scenarios with reactive and feasible agent behaviors can be used for validation and verification of self-driving system performance without relying on expensive and time-consuming real-world testing. Existing simulators rely on heuristic-based behavior models for background vehicles, which cannot capture the complex interactive behaviors in real-world scenarios. To bridge the gap between simulation and the real world, we propose TrajGen, a two-stage trajectory generation framework, which can capture more realistic behaviors directly from human demonstration. In particular, TrajGen consists of the multi-modal trajectory prediction stage and the reinforcement learning based trajectory modification stage. In the first stage, we propose a novel auxiliary RouteLoss for the trajectory prediction model to generate multi-modal diverse trajectories in the drivable area. In the second stage, reinforcement learning is used to track the predicted trajectories while avoiding collisions, which can improve the feasibility of generated trajectories. In addition, we develop a data-driven simulator I-Sim that can be used to train reinforcement learning models in parallel based on naturalistic driving data. The vehicle model in I-Sim can guarantee that the generated trajectories by TrajGen satisfy vehicle kinematic constraints. Finally, we give comprehensive metrics to evaluate generated trajectories for simulation scenarios, which shows that TrajGen outperforms either trajectory prediction or inverse reinforcement learning in terms of fidelity, reactivity, feasibility, and diversity.

Index Terms—Simulation scenarios, trajectory prediction, reinforcement learning

I. INTRODUCTION

Recently, self-driving has achieved widespread attention in academic and industry communities [1], [2]. However, how to validate and verify the performance of the decision-making and planning module is still unsolved. In industry, real-world testing with a fleet of autonomous vehicles (AVs) is a common approach. Unfortunately, such an evaluation process is both expensive and time-consuming, since valuable interactive scenarios are rare and safety-critical in the physical world. As a result, simulation has been an important evaluation tool enabling AVs' rapid development and safe deployment.

[1] is with The State Key Laboratory of Management and Control for Complex Systems, Institute of Automation, Chinese Academy of Sciences, Beijing, 100190, China, and also with the College of Artificial Intelligence, University of Chinese Academy of Sciences, Beijing, 100049, China. (email:zhangqichao2014@ia.ac.cn, Dongbin.zhao@ia.ac.cn)
[2] is with the School of Automation and Electrical Engineering, University of Science and Technology Beijing, Beijing, 100083, China.
[3] is with Baidu Inc., Beijing 100085, China.
* means the authors contribute equally.
The I-Sim simulator will be open-sourced soon.

A common scenario generation approach in simulator environments is log replay [3], where the behaviors of traffic participants are replayed according to the collected log trajectories [4]. However, once the behavior of AV is different from the one when the log was collected, those replayed evaluation scenarios may be unrealistic since the replayed traffic participants don't have any reactive capabilities. In other words, a lot of high-quality interactive scenario data collected by the AVs in real-world testing may be valueless to evaluate the AVs with upgraded system version based on the log replay, as shown in Fig. 1. The widely used reactive vehicle behavior model in existing simulators is based on the heuristic method such as Intelligent Driver Model (IDM) [5], MOBIL [6], and so on. However, those models are difficult to model the fidelity and diversity in interactive behaviors of road testing. Feng *et al.* [7] show that the IDM model even calibrated by real-world datasets cannot generate realistic and diverse driving trajectories for highway scenarios.
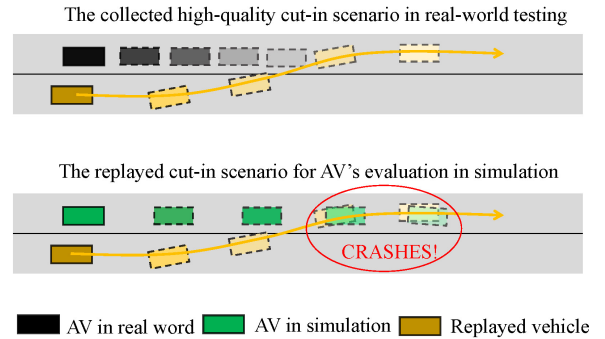

Fig. 1: The issue of log replay. The yellow vehicle is a replayed cut-in vehicle, and the black AV takes a yield behavior in real-world testing. However, when the green AV with an upgraded system version takes an acceleration behavior in this replayed simulation scenario, the yellow replayed vehicle will collide with the green AV. The responsibility of this unrealistic accident is due to the non-reactive behavior of the replayed vehicle, since the green AV has the path priority of starting. Note that the vehicle color fades as the timestep increases.

To bridge the behavior gap between simulation and real-world agents, Wang *et al.* [8] try to generate realistic and diverse trajectories with data-driven and learning models based on real-world data. The common approaches are trajectory prediction or motion forecasting [9], [10] and reinforcement learning (RL) related methods [11], [12]. For the trajectory prediction model, the stochastic multi-modality behaviors can

be captured based on the real-world data [13], [14]. However, the feasibility of generated trajectories cannot be guaranteed such as avoiding the unrealistic collision and satisfying vehicle kinematic constraints, as the agent is considered as a particle model. On the other hand, for the RL related methods [15], [16], the collision rate can be reduced by adding the collision penalty. Considering the vehicle model in the simulator, the vehicle kinematic constraints can be also satisfied. Unfortunately, the generated trajectories lack diversity since the nearly-optimal policy is trained to maximize the cost function, which is difficult to obtain diverse human-like behaviors [17].

Motivated by those, we present TrajGen to generate realistic and diverse scenarios with reactive and feasible agent behaviors based on naturalistic driving data, which is a two-stage trajectory generation framework shown in Fig. 2. The main contributions of this paper are as follows.

- A novel two-stage trajectory generation framework is proposed for simulation scenarios, by combining the trajectory prediction model to guarantee the fidelity and diversity of trajectories, and the RL model to improve the reactivity and feasibility;
- To improve the performance of trajectory prediction, an auxiliary RouteLoss is proposed by utilizing the prior knowledge of High-Definition Map (HD Map). In addition, the social attention mechanism is introduced in the RL model to better capture interaction features between the ego vehicle and nearby vehicles.
- We develop a data-driven simulator I-Sim with an OpenAI Gym environment based on the INTERACTION dataset [18], which can be used for closed-loop RL training and evaluation based on the real-world dataset. Experimental results in I-Sim show that TrajGen outperforms either trajectory prediction or inverse RL in terms of comprehensive metrics.

II. RELATED WORK

In this section, we will discuss relevant trajectory generation methods about multi-modal trajectory prediction, RL based driving behaviors, and existing driving simulators.

A. Multi-modal Trajectory Prediction

Trajectory prediction is used to predict the agent's future trajectories based on given past trajectory and context information. Recently, deep learning based methods have been developed to capture the multi-modal trajectories, which can be divided into two categories: anchor-based and anchor-free prediction. The anchor-based prediction can be considered as conditional forecasting, where the agent's prediction model is conditioned on different types of anchors. [19] and [20] condition the agent's future trajectories on a pre-defined set of anchor trajectories, while [21] and [22] propose goal-conditional multi-modal trajectory prediction based on sparse anchor goals with heuristic goal selection algorithms. [23] proposes the trajectory prediction with lane anchors, which uses existing lane centerlines as anchors. For anchor-free prediction, [24] proposes a sequential probabilistic latent variable

generative model to generate multi-modal interactive trajectories jointly for variable agents. [25] introduces an efficient vectorized representation with graph attention networks for multi-modal prediction. To capture the complex interactions between agents and maps, LaneGCN [26] is proposed with an interaction fusion network, which exploits HD Map with vectorized representation explicitly.

Recently, TrafficSim [13] uses a scene-consistent trajectory predictor as the multi-agent behavior model to simulate realistic and diverse traffic scenarios. To avoid generating collision trajectories, a common-sense loss is designed. Similarly, SimNet [14] uses a multi-modal trajectory predictor as the agent reactive behavior model to generate trajectories for background vehicles. RouteGAN [27] can generate diverse interaction behaviors by controlling the agents separately with desired styles and given final goals. Although the diversity and reactivity are improved based on trajectory prediction models, the feasibility of those produced unconstrained trajectories, such as the collision-free behaviors and vehicle kinematic feasibility, can not be guaranteed. However, the feasibility of generated trajectory is essential for realistic simulation scenarios.

B. RL based Driving Behaviors

Inverse RL is widely used to simulate driving behaviors from the demonstration, which aims to infer the reward function and learn the control policy [28] . Generative adversarial imitation learning (GAIL) [29] is used to simulate multi-agent driving behaviors with the parameter sharing technique in [30]. Considering the complexity of the real-world traffic, [11] proposes a reward augmented imitation learning (RAIL) to learn human driving behavior by embedding the prior knowledge in reward augmentation. [12] uses the adversarial inverse RL (AIRL) to train traffic simulating the agent's car-following behavior. [31] utilizes the model-based GAIL to generate rare-event scenarios for vision-based end-to-end autonomous driving algorithms. However, these methods are difficult to capture the multi-modal properties of the traffic since inverse RL aims to obtain a nearly-optimal policy by maximizing the learned cost function.

Recently, some researchers try to apply RL [32] in generating diverse simulation scenarios. [33] uses deep deterministic policy gradient without exploration noises to generate diverse lane-change scenarios. [34] designs a distinct policy set selector to balance diversity and driving skills. However, Without relying on real-world data, realistic behaviors are difficult to be generated for complex interactive scenarios. The comparison between existing methods and our proposed TrajGen is shown in Table I.

C. Existing Driving Simulator

The autonomous driving simulator is essential for RL based policy training, as well as the closed-loop evaluation. Recently, researchers have leveraged the racing game TORCS [35], microscopic traffic simulator SUMO [36], large-scale city-level traffic simulator CityFlow [37], and hand-coded simulation environment Highway-env [38] to train and evaluate
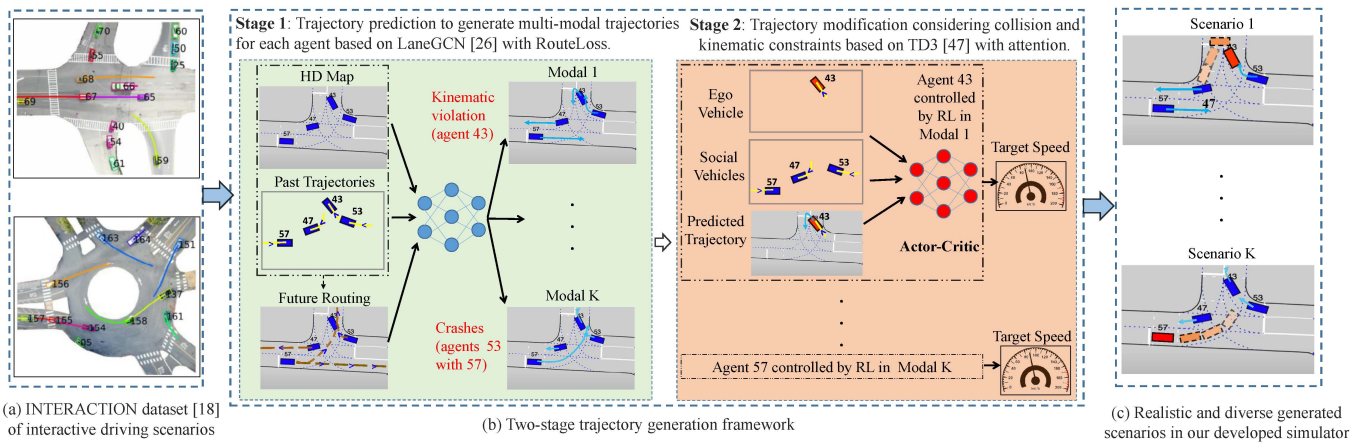
Fig. 2: The proposed TrajGen with a two-stage trajectory generation framework. The blue vehicles are controlled by the trajectory prediction model. As shown in sub-figure(b), the vehicles with kinematic violation (agent 43) or collisions (agent 57) in stage 1 are controlled by the RL model to modify their predicted trajectories in stage 2. For the proposed TrajGen, the fidelity and diversity of trajectories are provided by the trajectory prediction, and the reactivity and feasibility of trajectories are improved by RL. Finally, some realistic and diverse scenarios can be generated in I-Sim as shown in sub-figure(c).

TABLE I: Comparison of related methods

Methods	Fidelity	Reactivity	Feasibility	Diversity
Log replay[3]	✓	✓	✗	✓
Heuristic method[5], [6]	✗	✓	✓	✗
Trajectory prediction[13], [26]	✓	✓	✗	✓
RL[33], [34]	✗	✓	✓	✓
TrajGen (ours)	✓	✓	✓	✓

RL based driving policy. In particular, the provided maps and behavior models for background vehicles in these simulators are simplistic, and can not reflect the complex multi-agent interaction in the real world. SUMMIT [39] and CARLA [40] are open high-fidelity simulators with realistic visual appearance. Different from CARLA, SUMMIT provides many real-world maps, and the agent behavior model is based on a deterministic trajectory prediction model considering the physical constraints. BARK [41] provides a behavior benchmark with the heuristic model, RL model, and dataset tracking model with a few scenarios from the INTERACTION dataset [18]. The newly developed simulator SMARTS [42] and MetaDrive [43] provide excellent interaction environments for multi-agent RL agents in atomic traffic scenes. Different from SMARTS and MetaDrive, the HD Maps and vehicle behaviors in I-Sim are completely based on real-world complex interactive scenarios to guarantee the fidelity, which can facilitate the generalizable RL research in autonomous driving.

III. GENERATING REALISTIC AND DIVERSE SCENARIOS

A. Overview

The simulation scenarios are usually generated as follows: 1) scenarios initialization: specifying the initial scene layout with the road topology and agents' initial states such as position and velocity. 2) forward simulation: simulating the motion of dynamic agents forward, and 3) scenarios termination: ending the scenario when the designed termination conditions are met. In this paper, the initial scene is provided by the INTERACTION dataset shown in Fig. 2(a). And the termination conditions are unrolling for $8s$ during forward simulation, or all agents reach their log trajectories' endpoints. We focus on the motion of dynamic agents for forward simulation by using TrajGen.

Given an initial scene with HD Map M and historical joint states of N dynamic agents, our goal is to design the agent behavior trajectories for their forward simulation. The process of forward simulation can be described as a sequence of joint state $h_t = \{z_t^1, z_t^2, ..., z_t^N\}$ with the unrolling horizon T, where z_t^n means the state of the nth agent including its 2D position and speed information at time t. Given a sequence of historical joint states derived from the perception system, the multi-modal future trajectories should be captured for each agent. For example, agent 57 in Fig. 2(b) can go straight in modal 1 or turn left in modal K. In addition, its behavior should satisfy the vehicle kinematic constraint and the common sense such as avoiding collisions. In the following, we describe the two-stage trajectory generation framework combing trajectory prediction and RL tracking, which can generate realistic and diverse trajectories with reactive and feasible behaviors.

B. Stage 1: Trajectory prediction to generate multi-modal trajectories

For the multi-modal trajectory prediction stage, we aim to generate a joint distribution $P(Y_{pred}|Y_{hist}, M)$ of multi-modal trajectories $Y_{pred} = \{h_{t+1}, h_{t+2}, ..., h_{t+T_{pred}}\}$ for N agents in parallel with a sequence of historical joint states $Y_{hist} = \{h_{t-T_{hist}+1}, h_{t-T_{hist}+2}, ..., h_t\}$, and HD Map context M. T_{pred} denotes the total prediction steps during the unrolling horizon T, and T_{hist} means the historical steps. To capture the complex and diverse interaction behaviors, we employ the LaneGCN model described in [26] with vectorized map

representation on the INTERACTION dataset. A FusionNet with multiple interaction blocks stacked provided in [26] is used to extract interaction features. Based on the extracted features, multiple prediction headers can predict multi-modal trajectories with the following max-margin classification loss and smooth L1 regression loss.

$$L_{cls} = \frac{1}{N(K-1)} \sum_{n=1}^{N} \sum_{k \neq \hat{k}} \max(0, c^{n,k} + \epsilon - c^{n,\hat{k}}) \quad (1)$$

$$L_{reg} = \frac{1}{NT} \sum_{n=1}^{N} \sum_{t=1}^{T_{pred}} reg(z_t^{n,\hat{k}} - z_t^{n,*}) \quad (2)$$

where $\hat{k}$ means the best modal with *Minimum Final Displacement Error* (minFDE), $c^{n,k}$ means the confidence score of the kth modal trajectory for the nth agent, ϵ is the margin, and z^* means ground truth trajectory in prediction horizon. There are N agents and K modals.

INTERACTION is a dataset that provides a fixed Bird-Eye-View (BEV) above complex traffic scenes including intersections, roundabouts, and merging roads. For the INTERACTION dataset, we find that the trajectories generated by LaneGCN shown in Fig. 3(a-c) have a high probability of off-road and collision. To utilize the prior knowledge of HD Map, [44] propose an auxiliary off-road loss to penalize off-road trajectories. In [45], an auxiliary LaneLoss is introduced to encourage diverse trajectories enough to cover every possible maneuver. Note that LaneLoss is designed based on the normal distance of the endpoint of the predicted trajectory from its nearest lane waypoint in the Frenet-Serret Frame. Obviously, this LaneLoss is not helpful to improve the quality of off-road trajectory with an endpoint in the lane like 3(c). Motivated by this, we design a novel auxiliary RouteLoss as follows:

$$L_{route} = \sum_{n=1}^{N} \sum_{k=1}^{K} \min_{f \in L} \sum_{j \in Tseq} \max(0, d(z_j^{n,k}, f(z_j^{n,k})) - m) \quad (3)$$

where L contains all available routes according to agents' history and the HD Map M using a deep limited depth first search (DFS) algorithm over the map graph. $Tseq$ is the set of sampled timestamps from the predicted trajectory including startpoint and endpoint. The deviation $d(\cdot)$ represents the L2 distance between the nth agent's prediction trajectory waypoint $z_j^{n,k}$ for modal k at timestamp j and its corresponding projected route point $f(z_j^{n,k})$ on the nearest route, as shown in Fig. 3(d). The lane width m is used to avoid forcing agents to stay exactly in the lane. RouteLoss is designed to make the generated trajectory close to the nearest available route, which is helpful to avoid the off-road trajectory without sacrificing diversity.

So the total loss of trajectory prediction is

$$L_{total} = L_{cls} + \alpha L_{reg} + \beta L_{route} \quad (4)$$

where L_{cls} is the max-margin loss and L_{reg} is smooth L1 loss defined in Eq.(1) and Eq.(2), α and β are normalization parameters. To avoid the collisions like Fig. 3(a,b), a common-sense auxiliary loss provided in [13] is used in [46], the

experimental results show that the acceleration over the scope phenomenon happens because the vehicle kinematic constraints are not taken into account. We address this challenge by using trajectory modification based on RL and considering the vehicle kinematic model in the simulator.
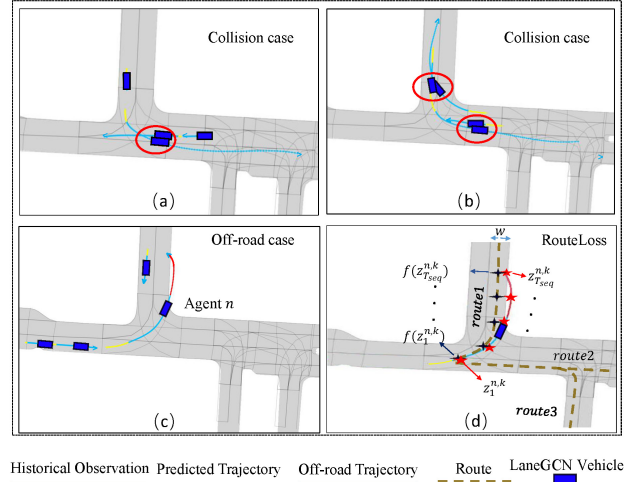

Fig. 3: Illustration of the collision (a,b) and off-road (c) behaviors generated by the LaneGCN model in INTERACTION dataset. To penalize the off-road behaviors, an auxiliary RouteLoss (d) is designed in stage 1, where the red stars denote the sampled waypoints $z_j^{n,k}$ from the kth predicted trajectory of agent n, and black diamonds denote the projected route points $f(z_j^{n,k})$ from the nearest route. The collision behaviors are addressed in stage 2.

C. Stage 2: Trajectory modification to avoid collisions with RL

RL aims to improve the quality of the predicted trajectories generated in stage 1, especially for collision trajectories. It can be considered as the collision avoidance and trajectory tracking task. Considering the continuous action space, the specific RL algorithm used in this paper is Twin Delayed Deep Deterministic Policy Gradient (TD3) [47] with the social attention mechanism [48].

1) RL method:

Definition: In stage 2, RL agents interact with the environment and learn to track prediction trajectories generated by LaneGCN while avoiding collisions by maximizing their cumulative reward. Here we define such tasks as Markov Decision Processes (MDP), which can be represented by the tuple $(\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \gamma)$, where $\mathcal{S}$ is the state space, $\mathcal{A}$ is the action space, $\mathcal{P} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0,1]$ is the transition function which maps the state s_t and action a_t to a probability distribution of the next states s_{t+1}, $\mathcal{R} : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is the reward function which maps the state and its corresponding action to a scalar reward value, and $\gamma \in (0,1)$ is the discount factor. The goal of agents is to find a policy $\pi : \mathcal{S} \rightarrow \mathcal{A}$ which maximizes the expected discount cumulative reward,

i.e., return:

$$\max E[\sum_{t=0}^{\infty} \gamma^t R(s_t, \pi(s_t))] \quad (5)$$

We also define the value function $V^\pi : \mathcal{S} \to \mathbb{R}$, which maps the state to the expected return with running policy π, and the Q-value function $Q^\pi : \mathcal{S} \times \mathcal{A} \to \mathbb{R}$, which maps the state-action pair to the expected return with running π.

TD3 algorithm: TD3 [47] is a popular deep RL method with the Actor-Critic architecture, where both actor and critics are represented using deep networks. In TD3, the actor μ is a policy network parameterized by θ, which takes state as input and outputs an action, the critic Q are composed of two Q-value approximation networks parameterized by ϕ_1 and ϕ_2, which take state and its corresponding action as input and output the Q-value. We use θ', ϕ_1', and ϕ_2' to represent the parameters of the target networks of actor and critics. As an off-policy method, TD3 needs a replay buffer $\mathcal{B}$ to store experiences. Every time a batch of experiences are sampled from the replay buffer, and the critic will be optimized based on the loss $L(\phi_i), i = 1, 2$:

$$L(\phi_i) = E_{(s_t, a_t, r_t, s_{t+1}) \sim \mathcal{B}}[q(s_t, a_t) - Q_{\phi_i}(s_t, a_t)]^2 \quad (6)$$

where $q(s_t, a_t)$ is the target value:

$$q(s_t, a_t) = r_t + \gamma \min_{i=1,2} Q_{\phi_i'}(s_{t+1}, \tilde{a}) \quad (7)$$

with $\tilde{a} = \mu_{\theta'}(s_{t+1}) + \delta, \delta \sim clip(\mathcal{N}(0, \sigma), -l, l)$ the target smoothing action, δ is a random noise clipped by $\pm l$.

After updating ϕ_1 and ϕ_2, the actor's parameter θ is updated by the deterministic policy gradient:

$$\nabla_\theta J(\theta) = E_{s \sim \mathcal{B}}[\nabla_a Q_{\phi_1}(s, a)|_{a=\mu_\theta(s)} \nabla_\theta \mu_\theta(s)] \quad (8)$$

Note that the actor network is updated once after updating a fixed number of d times of the critic networks in TD3 due to the "delayed policy updates" setting.

The parameters of the target networks, ϕ_1', ϕ_2' and θ', are updated softly to ensure the learning process is stable:

$$\phi_i' = \tau \phi_i + (1 - \tau)\phi_i' \quad (9)$$

$$\theta' = \tau \theta + (1 - \tau)\theta' \quad (10)$$

with $\tau \ll 1$. τ is a normalized hyper-parameter that controls the change speed of the target networks.

Social attention mechanism: Attention mechanism is useful to find the internal dependencies between different input instances for deep networks, which has been proved effective in many research fields. To help RL agents better capture dependencies between the ego vehicle and nearby vehicles while making a decision, social attention mechanism [48] is used in our actor network. Note that the ego vehicle means the agent controlled by the RL model.

The social attention mechanism function used in this paper is a soft attention mechanism function similar to [49]. Assume there are M social vehicles that are near to the ego vehicle, the attention values which are output by the social attention mechanism function can be defined as:

$$Attention(Q^A, K^A, V^A) = softmax(\frac{Q^A K^{A^T}}{\sqrt{d_k}})V^A \quad (11)$$

where $Q^A = [q_e]$ is a single query calculated by processing the ego vehicle's features with a nonlinear projection P_q, $K^A = [k_e, k_1, ..., k_M]$ are keys calculated by processing the ego and nearby vehicles' features with another nonlinear projection P_k. Through the dot product between Q^A and K^A, the similarities between the query and the keys are generated, which are then scaled by $1/\sqrt{d_k}$. d_k is the dimension of q_e and $k_j, j = e, 1, ..., M$. The output of this softmax function is defined as *attention weights*, which indicates ego's attention to nearby social vehicles including itself. Finally, the dot product between the *attention weights* and values $V^A = [v_e, v_1, ..., v_M]$, which are calculated by processing the ego and nearby vehicles' features with the nonlinear projection P_v, is the result of the social attention mechanism function. Note that the nonlinear projections P_q, P_k and P_v are all realized by multi-layer perceptron in this paper.

2) RL settings:

Observation: A 48-dimension vector is used to represent RL controlled ego vehicle's current observation, which can be written as $o = [o_r, o_e, o_1, ...o_5]$. As shown in Table II, the observation vector o can be divided into three major parts: o_r contains ego vehicle's target position, target speed value, and route trend. Specifically, we obtain the ego vehicle's relative position and relative speed according to the predicted trajectory as its target position and target speed. The route trend consists of heading errors between the ego vehicle and 5 closest forward prediction waypoints. o_e refers to ego vehicle's shape and its movement trend, specifically, the length and width of the ego vehicle, ego vehicle's current speed value and its position in the next second if it remains current speed and heading. $o_1, ...o_5$ describe the observation of 5 social vehicles, each of them contains the length, width, coordinate, current speed value, and current heading's cosine and sine value of a social vehicle.

Note that all position and heading observations mentioned above are relative to the ego's coordinate system. We design a simple filter to select social vehicles, whose rules can be described below:

$$\begin{cases} d_{ei} < 30m \\ y_{e_i} > -12m \end{cases}$$

where d_{ei} indicates the distance between ego vehicle and a nearby vehicle i, y_{e_i} indicates a nearby vehicle i's Y-axis coordinate under the ego's coordinate system. Only the closest 5 vehicles are considered as social vehicles if there are more than 5 vehicles that are kept through the filter.

TABLE II: Observation Representations

Component	Description	Length
o_r	ego's target posture and route trend	1*8
o_e	ego's shape and movement trend	1*5
$o_i(i = 1, 2, ..., 5)$	social vehicles' shape and posture	5*7
o	complete observation vector	48

Action: Since most collisions are caused by unreasonable longitudinal behavior, we consider the ego vehicle's longitudinal target speed as the action space. Specifically, the output

of the policy network is a normalized scalar, and it is linearly mapped to $(0, 10)m/s$ as the longitudinal target speed. Then, a longitudinal PID controller $\mathrm{PID_{acc}}$ is applied to control the acceleration based on the output value. For the lateral movement, another PID controller $\mathrm{PID_{angle}}$ is designed to control the front wheel angle to track the trajectories given by stage 1.

Reward: The main purpose of using RL is to avoid collision while track the predicted trajectories. Thus, the reward function r is designed as $r = r_{collision} + r_{tracking} + r_{step}$ with

$$\begin{cases} r_{collision} = -500 * (1 + v\_norm) \\ r_{tracking} = 1 - 0.2 * d_{ep} \\ r_{step} = -0.5 \end{cases}$$

where v_norm is the ego vehicle's normalized speed value. It is easier for the ego vehicle to learn to avoid active collisions through dynamic adjustment $r_{collision}$ with v_norm. We use $r_{tracking}$ to reduce the distance between ego vehicle and its tracking trajectory, where d_{ep} refers to ego vehicle's distance to its current predicted position. r_{step} is used to prevent RL from learning a negative strategy such as staying at the startpoint.

Combining the trajectory prediction in stage 1 and trajectory modification in stage 2, the pseudo-code of the proposed TrajGen is given in Algorithm 1.

D. Simulator: I-Sim

To train the RL agent, we develop I-Sim, a data-driven simulator with an OpenAI Gym environment based on the INTERACTION dataset. The overview of I-Sim with a three-layer architecture is shown in Fig. 4, which can support effective RL and inverse RL training in parallel.

1) The bottom layer can load centimeter-accurate HD Map files and recorded_trackfiles for agents from the INTERACTION dataset. Based on this layer, we can obtain the replayed scenarios for the INTERACTION dataset. In addition, the bottom layer is responsible for providing the agent's kinematic model and geometry calculation functions such as the distance and angle calculation function between two rectangles. The kinematic model can be considered as the transition modeling for the traffic vehicles in MDP. In I-Sim, we use a bicycle model [50] for agents given by

$$\begin{cases} x_{t+1} = x_t + v_t \cos(\psi_t + \omega) \times \Delta t \\ y_{t+1} = y_t + v_t \sin(\psi_t + \omega) \times \Delta t \\ \psi_{t+1} = \psi_t + \frac{v_t}{l_r} \sin(\omega) \times \Delta t \\ v_{t+1} = v_t + acc \times \Delta t \end{cases} \quad (12)$$

with $\omega = \tan^{-1}(\frac{l_r}{l_r + l_f} \tan(\delta_f))$, where t and Δt indicate current timestamp and delta time separately, x and y represent coordinates, ψ represents vehicle's yaw angle, and v represents vehicle's speed. The vechcle is controlled through the acceleration acc and the front wheel angle δ_f. Note that l_f and l_r are distances from the front and rear wheels to the vehicle center of gravity. To get l_f and l_r of vehicles with different shapes for simplicity, we assume that the wheelbase length is a fixed ratio of 0.6 of the vehicle's length and the gravity center of the vehicle is in its middle. To guarantee the

---

Algorithm 1: Algorithm for TrajGen

input : INTERACTION dataset with trajectory data and HD Map.

output: Multi-modal trajectories for N agents.

1 **Stage 1: Train trajectory prediction model**
2 Randomly initialize LaneGCN network, and the hyper-parameters K, α, β.
3 Search each agent's routes using depth limited DFS over the map graph.
4 **for** *iterations* $iter = 0$ **to** $ITER$ **do**
5   **for** *modes* $k = 1$ **to** K **do**
6   **for** *agent* $n = 1$ **to** N **do**
7   output the confidence score $c^{n,k}$ and the predicted trajectory $z^{n,k}$.
8   **end**
9   **end**
10   Compute the total loss described in Eq.(4).
11   Update all parameters via ADAM optimizer.
12 **end**

13 **Stage 2: Train trajectory modification RL model**
14 Initialize critic's networks Q_{ϕ_1}, Q_{ϕ_2} and social attentive actor's network μ_θ with random parameters ϕ_1, ϕ_2, θ for TD3 algorithm.
15 Initialize target networks $\phi_1' \leftarrow \phi_1$, $\phi_2' \leftarrow \phi_2$, $\theta' \leftarrow \theta$.
16 Initialize replay buffer $\mathcal{B}$.
17 **for** *iterations* $iter = 1$ **to** $ITER$ **do**
18   Select action with exploration noise $a = \mu(s) + \delta, \delta \sim \mathcal{N}(0, \sigma)$.
19   Observe reward r and next state s'.
20   Store transition tuple (s, a, r, s') in $\mathcal{B}$.
21   Sample a mini-batch of transitions (s, a, r, s') from $\mathcal{B}$.
22   Update critic's parameters ϕ_1 and ϕ_2 by minimizing $L(\phi_i)$ described in Eq.(6).
23   **if** $t \bmod d$ **then**
24   Update actor's parameters θ with Eq.(8).
25   Update target networks' parameters ϕ_1', ϕ_2' and θ' with Eq.(9) and Eq.(10).
26   **end**
27 **end**

---

feasibility, the acceleration range is set to $[-3, 3]m/s^2$, and the max front wheel angle is set to ± 30 degrees according to the INTERACTION dataset.

2) The middle layer is designed to give observation modeling and visualization tools. In I-Sim, we provide two kinds of common observation representations, which are vector observation and BEV observation. Note that the vector observation is used in the proposed TrajGen. The visualization tools can be used to analyze the agent behaviors and trajectory quality qualitatively.

3) The top layer is used for communication with the server-client pattern, which can carry out parallel simulation for efficient RL training. In detail, the client can obtain required data from multiple servers in parallel. After the RL policy

is updated, the corresponding output action for each server is transmitted to its controlled agent. Based on the bicycle model in Eq.(12), the next states can be obtained. Meanwhile, it also supports the scenario generation ability by providing multiple control APIs, which allows us to obtain observations and generate trajectories for multiple vehicles based on the proposed method simultaneously in one single server thread.
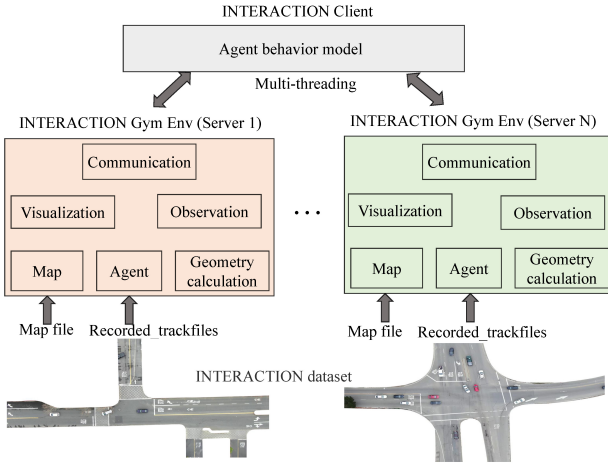

Fig. 4: An overview of I-Sim that simulates massive mixed traffic based on INTERACTION dataset.

IV. Experiments

A. Dataset

We use the INTERACTION dataset [18], which contains different highly interactive and complex urban scenarios from different countries. To validate the trajectory prediction performance with the proposed auxiliary RouteLoss in stage 1, we first compare the performances on the INTERACTION dataset with intersection scenarios for $T = 3s$ prediction horizon, which is a common setting for the trajectory prediction task. Note that the trajectories are sampled at 10Hz with $(0, 2]s$ for observation and $(2, 5]s$ for future prediction. Since the unrolling horizon is $8s$ for our trajectory generation task, we additionally extract 11938 snippets with $10s$ length, where 8437 are set for training, 2365 for validation, and 1136 for testing. The prediction model performances are also validated for the $T = 8s$ prediction horizon.

To evaluate the trajectory quality generated by TrajGen in highly interactive scenarios, we extract 425 snippets containing safety-critical situations from the USA_Intersection_EP0 in INTERACTION dataset as the validation dataset. Firstly, we deploy the trained LaneGCN with RouteLoss predictor on this safety-critical validation set with $10s$ length for RL training in stage 2. During the RL training process, we use the HD Map and $(0,2]s$ trajectories as the scenario initialization, and the subsequent $T = 8s$ for forward simulation in the developed I-Sim. After training, the converged RL model is used to modify those trajectories generated by the predictor, which can reduce collisions and satisfy the kinematic constraints.

B. Metrics

Evaluating behavior simulation is still challenging since there is no singular quantitative metric that can fully capture the quality of the generated simulation scenarios. In this paper, we design the metrics in four aspects as shown in Table I, which are all important to evaluate the trajectory quality for the simulation scenarios.

1) Fidelity: We use distance-based scenario reconstruction metrics to evaluate the fidelity between generated trajectories with the ground truth. The fidelity of generated trajectories should be a primary concern.

For the prediction task in stage 1, since the time horizons of generated trajectories are the same, we adopt the widely used minADE/minFDE by selecting the best matching prediction modal, and meanADE/meanFDE by averaging over the set of $K = 6$ predicted trajectories. In addition, we also provide the off-road rate (OR) as follows.

$$\text{OR} = \frac{Num(\text{off-road trajectories})}{Num(\text{all generated trajectories})}$$

where $Num(\cdot)$ means the number.

For simulation scenarios evaluation, since the time horizons of trajectories generated by TrajGen may be different for different agents, we use the root mean square error (RMSE) of position [12] to evaluate the mismatch between the generated trajectories with the ground truth.

$$\text{RMSE} = \frac{1}{N} \sum_{n=1}^{N} \sqrt{\frac{1}{T_{pred}} \sum_{t=1}^{T_{pred}} (z_t^n - \hat{z}_t^n)^2} \qquad (13)$$

2) Reactivity: Motivated by [14], we measure the agent's reactivity based on the collision rate in synthetic scenarios, where a static car is placed in front of a moving car controlled by the behavior model. This requires the trailing car to react by stopping. 233 scenarios are chosen from the safety-critical validation set for reactivity testing. We report the synthetic collision rate (SCR) as follows.

$$\text{SCR} = \frac{Num(\text{collision scenarios})}{Num(\text{all synthetic scenarios})}$$

3) Feasibility: We consider the feasibility of generated trajectories from the following two aspects.

Kinematic feasibility: The acceleration and angular velocity is considered for kinematic feasibility. The distribution of maximum acceleration for all the ground truth trajectories is shown in Fig 5. We monitor the absolute value of max acceleration for the generated trajectories and raise an Acceleration Failure (AF) when it exceeds the range of $4m/s^2$. In addition, KL (Kullback-Leibler) divergence of the angular velocity distribution between the generated trajectory data with the log trajectory data is given. The smaller the KL divergence is, the closer it is to the angle velocity distribution of real-world data.

Common sense feasibility: we measure the trajectory collision rate (TCR) as follows.

$$\text{TCR} = \frac{Num(\text{collision trajectories})}{Num(\text{all generated trajectories})}$$
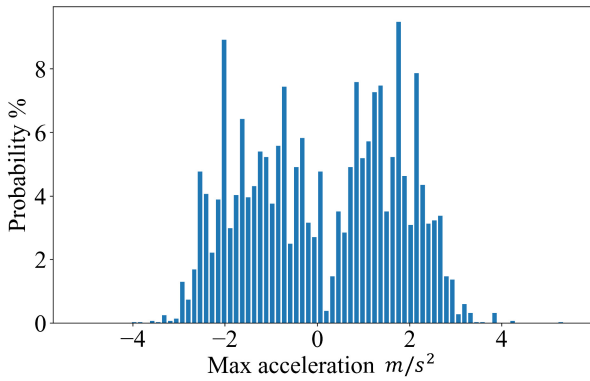
Fig. 5: The distribution of ground truth max acceleration.

4) Diversity: Following [13], we use a map-aware average self distance (MASD) metric to measure the diversity of the sampled scenarios. In particular, we measure the average distance between the two most distinct sampled trajectories that do not violate traffic rules for each agent in each scenario.

$$\text{MASD} = \max_{k,k' \in 1,...,l} \frac{1}{NT_{pred}} \sum_{n=1}^{N} \sum_{t=1}^{T_{pred}} \|z_t^{n,k} - z_t^{n,k'}\|^2 \quad (14)$$

C. Model Details

LaneGCN model: The input of LaneGCN model contains the HD Map, and the $2s$ history trajectories of agents. The output are the $K = 6$ multi-modal trajectories. To calculate RouteLoss, we first sample $T_{seq} = 10$ waypoints with the equal time interval including startpoint and endpoint from the predicted trajectory. The waypoints and the projected route points from DFS-searched nearest route are used to compute L_{route}. The hyper-parameters $\alpha = 1$ and $\beta = 0.05$ in Eq.(4). The learning rate is set to 1e-3 and the batch size is 16.

TD3 model with attention: The policy network consists of three modules of encoder, attention and decoder. The encoder module has three encoders for three kinds of observation o_r, o_e and $o_j, j = 1, ..., 5$, each of them contains two dense layers with the same shape of 64, which takes observation as input and outputs features. These features are then fed into the attention module, which contains three nonlinear projections P_q, P_k, P_v with the same shape of 64. The *attention values* are computed by Eq.(11), and the decoder module processes the *attention values* through two dense layers with the same shape of 256, obtaining the action value. The Q-value approximation networks hold the same structure with each other. Each of them contains two modules of encoder and decoder, which has a lighter structure of layers and units compared with the policy network. The observation of the ego and social vehicles are concatenated together to be encoded as a whole rather than encoded separately by different encoders, and the action value as a part of the network's input is concatenated with the encoded state features. All layers mentioned above are activated by the tanh function. The learning rate of actor's network is set to 1e-4, and the learning rate of critic's networks is set to 5e-5. The update frequency d is 2 and the batch size is 256.

For the longitudinal and lateral PID controllers, the parameters are designed as follows:

$$\begin{cases} \text{PID}_{acc} : k_p = 1.0, k_i = 0, k_d = 0.05 \\ \text{PID}_{angle} : k_p = 1.4, k_i = 0.05, k_d = 0.25 \end{cases}$$
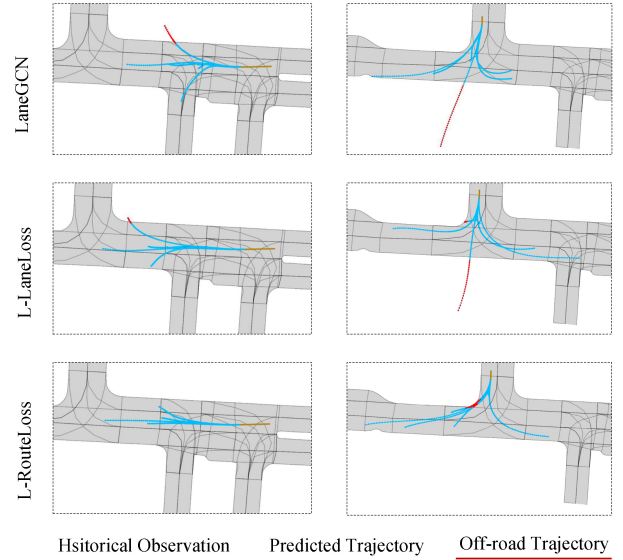

Fig. 6: Qualitative results for $T = 8s$ on typical intersection cases. Lane boundary is shown in black, agent's past trajectory in yellow, predicted multi-modal trajectories in light blue, and the off-road part in red.

D. Comparison Results

Here we provide a quantitative evaluation of our proposed trajectory generation method. We generally aim to answer the following three questions.

Q1: Trajectory prediction. Is the auxiliary RouteLoss effective at reducing the off-road trajectories in stage 1?

Q2: Trajectory modification. Can collision behaviors be reduced by the RL model in stage 2 and the kinematic constraints be satisfied with I-Sim?

Q3: Multi-agent behaviors. Is the proposed TrajGen suitable for multi-agent behaviors?

To answer Q1, we implement the LaneGCN [26], LaneGCN with LaneLoss [45], and LaneGCN with RouteLoss proposed in this paper on INTERACTION validation set for the prediction horizon $T = 3s$ and $T = 8s$. For simplicity, we use L-LaneLoss and L-RouteLoss represent the LaneGCN with LouteLoss and the LaneGCN with RouteLoss, respectively. As shown in Table III, the L-RouteLoss achieves better performances for $T = 3s$ in most metrics, especially OR is significantly decreased, which validates the effectiveness of the designed auxiliary RouteLoss. For $T = 8s$, its performances are still better in most metrics. Qualitative results for typical intersection testing cases are shown in Fig. 6. It can be seen that the off-road trajectories are reduced with the auxiliary Routeloss. However, OR of L-LaneLoss and L-RouteLoss are close. It may be the off-road trajectory with an endpoint in the lane shown in Fig. 3(c) is rare for the long horizon prediction.

TABLE III: Comparison results of the proposed model on INTERACTION validation set

T	Model	meanADE$_6$	meanFDE$_6$	minADE$_6$	minFDE$_6$	OR(%)
3s	LaneGCN [26]	1.42	3.92	0.75	1.87	2.05
	L-LaneLoss[45]	1.10	**3.07**	0.32	0.70	2.13
	L-RouteLoss(ours)	**1.08**	3.09	**0.30**	**0.69**	**1.64**
8s	LaneGCN [26]	4.65	14.03	1.96	4.20	10.82
	L-LaneLoss[45]	4.50	13.40	1.43	2.85	**7.69**
	L-RouteLoss(ours)	**4.38**	**13.27**	**1.41**	**2.83**	7.88

To answer Q2, a comparative experiment is conducted between scenarios generated by L-RouteLoss and by TrajGen. The RL model of TrajGen is trained on 2849 trajectories from 425 scenarios generated by L-RouteLoss. We randomly select one trajectory at each training episode to be revised using RL model. The episode's result is "collision" if RL controlled vehicle collides with other traffic participants. If RL controlled vehicle reaches trajectory's endpoint within the 10s time horizon, the episode's result is "success", otherwise, the episode's result is "timeout". The RL model's training curve of episode results can be seen in Fig. 7, which shows that the RL model achieves a relatively high success rate and low collision and timeout rate within 5000 episodes of training. Then, the trained RL model will be used to modify those infeasible trajectories generated by the predictor.
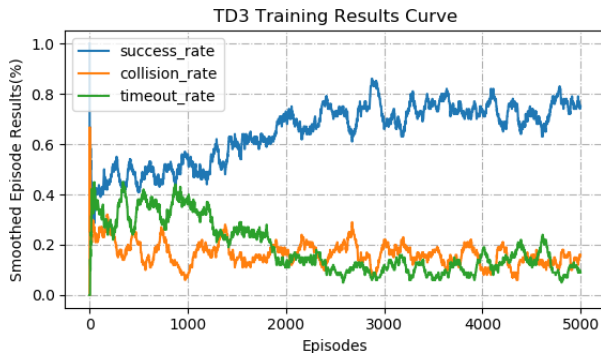

Fig. 7: RL model's training curve of episode results.

Among the 2849 trajectories generated by L-RouteLoss, there are 828 collision trajectories and 19 kinematic unsatisfied trajectories. Excluding 4 repeated trajectories, there are 843 testing trajectories that need to be modified. Then, the trained RL model is used to control one collision or kinematic unsatisfied agent to figure whether the collision or kinematic unsatisfied can be avoided. As shown in Table IV, TrajGen achieves lower RMSE compared with the trajectory prediction model L-RouteLoss, which means the higher matching degree and fidelity with ground truth trajectory. For the reactivity, 233 synthetic scenarios' test shows that TrajGen's collision rate decreases 15.8% compared with L-RouteLoss. For the feasibility, with the help of RL model, collision trajectories are reduced from 828 to 435, thus TCR is also reduced from 28.9% to 15.3%, which means nearly half of the collision trajectories are reduced as shown in Scenario 1 of Fig. 8. Meanwhile, it shows that the AF and KL divergence of TrajGen, which considers

the bicycle model during the training process, is significantly reduced compared with the L-RouteLoss prediction model. In addition, the MASD of L-RouteLoss and TrajGen are very close, which means TrajGen can maintain the diversity from the first multi-modal prediction stage.

For comparison, we also train a GAIL model [29] as the trajectory generator. In general, we should first focus on fidelity for simulation scenarios, where the RMSE of GAIL is too poor to be acceptable. The reason is that GAIL cannot keep the vehicle driving in the lane as shown in Fig. 8. As the vehicle controlled by GAIL can not drive along the lane well, it may have moved out of the drivable area before colliding with a static car in synthetic scenarios. This is the reason why the SCR of GAIL is lower than TrajGen. In addition, the MASD of GAIL is zero since it can not generate multi-modal trajectories. In other words, GAIL is almost impossible to generate high-quality trajectories with a long horizon. Overall, the proposed TrajGen outperforms the trajectory prediction method and GAIL in terms of comprehensive metrics.

TABLE IV: Comparison quality of generated simulation scenarios

Model	Fidelity	Reactivity	Feasibility			Diversity
	RMSE↓	SCR↓	TCR↓	AF↓	KL↓	MASD↑
Log replay	/	100%	/	/	/	/
GAIL [32]	20.96	**19.9%**	45.8%	0	3.58	0
L-RouteLoss(ours)	4.21	41.6%	28.9%	19	3.44	**6.71**
TrajGen(ours)	**3.58**	25.8%	**15.3%**	0	**1.05**	6.56

To answer Q3, we deploy the TrajGen model on the two collision agents using the parameters sharing technique. A typical scenario can be seen in Scenario 2 of Fig. 8, where the predicted collision trajectories can not be reduced if we choose either agent to deploy TrajGen. The results show that the collision can be avoided by depolying TrajGen on those two agents, and the generated scenarios are naturalistic. However, for the GAIL model with parameter sharing, the multiple controlled vehicles are out of drivable area which generates an invalid scenario. Therefore, the proposed TrajGen is also suitable for multi-agent behaviors based on parameters sharing.

V. CONCLUSION

In this paper, we have proposed a novel framework to generate realistic and diverse trajectories with a reactive and feasible agent behavior model based on naturalistic driving data. TrajGen is a two-stage trajectory generation model for background vehicles in simulation scenarios, which is trained directly from real-world INTERACTION dataset. In addition, we develop a data-driven simulator I-Sim that can be used for RL training based on real-world traffic scenarios. The experimental results show that TrajGen can reduce nearly half of the collision trajectories, and the quality of generated trajectories is improved obviously in terms of fidelity, reactivity, feasibility, and diversity.

Scenario generation with data-driven behavior model is still a challenging and important task for autonomous driving now.
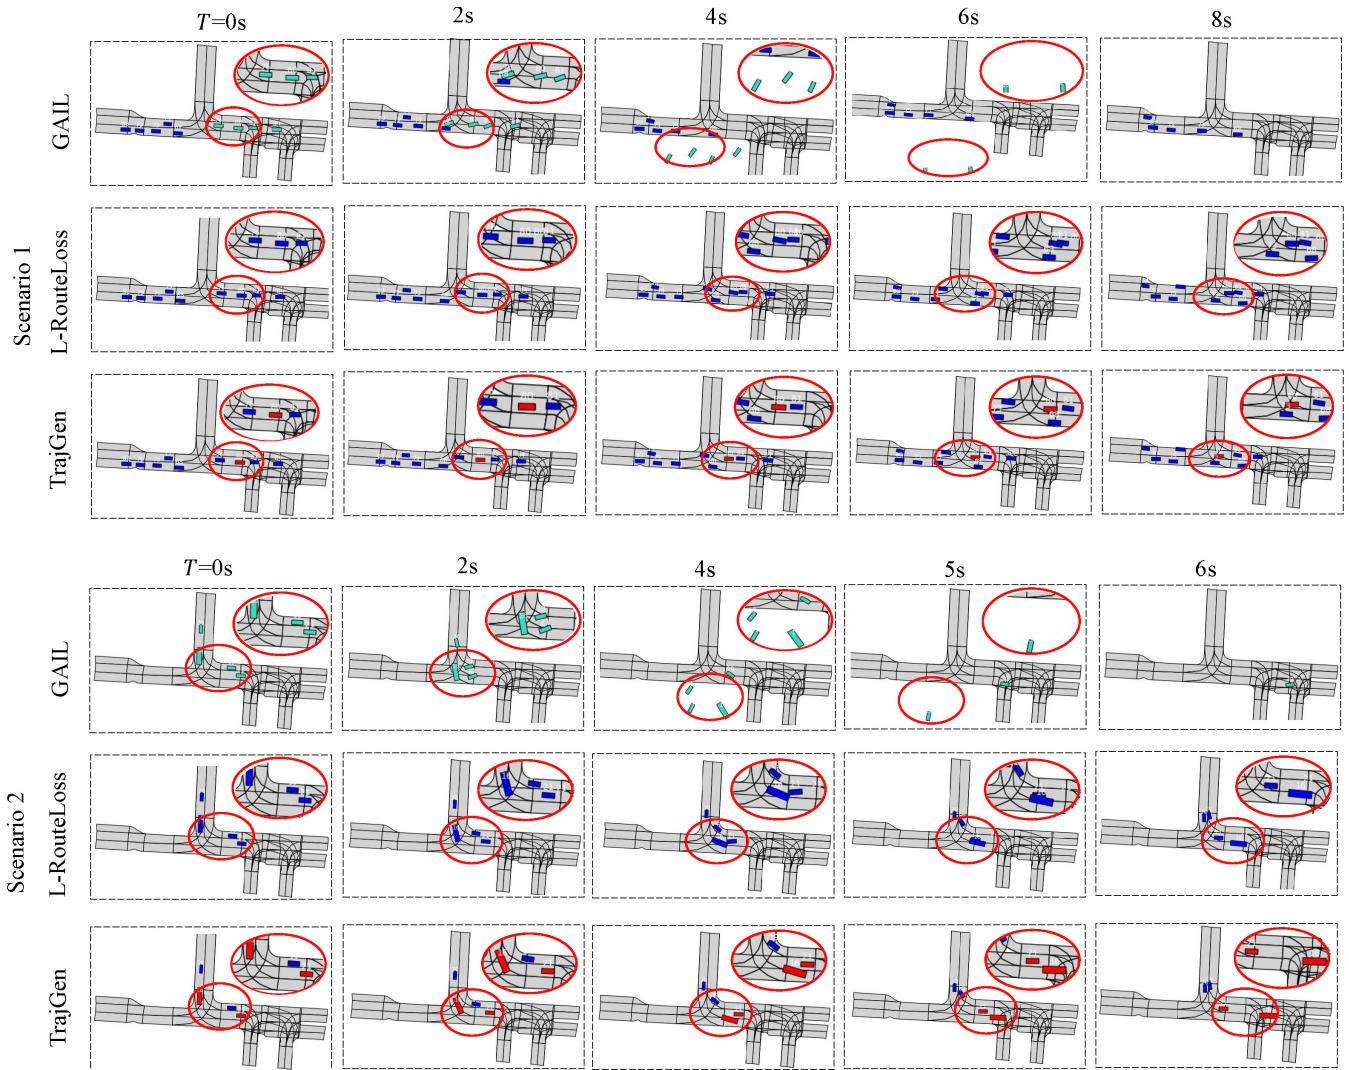
Fig. 8: Simulated traffic scenarios with sampled key timestamps from trajectory generator GAIL[29], L-RouteLoss and TrajGen. The cyan, blue and red vehicles are controlled by GAIL, L-RouteLoss, and TrajGen, respectively. Since the interactive times of focused vehicles are not the same for scenario 1 and 2, the sampled key timestamps are also different.

In future work, we will further exploit the generalization ability of TrajGen in different kinds of complicated scenarios. In addition, evaluating the existing planning models and finding their corner case scenarios based on the scenario generation approach is another interesting and important work.

REFERENCES

[1] H. Li, Q. Zhang, and D. Zhao, "Deep reinforcement learning-based automatic exploration for navigation in unknown environment," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 31, no. 6, pp. 2064–2076, 2019.

[2] Y. Guo, Q. Zhang, J. Wang, and S. Liu, "Hierarchical reinforcement learning-based policy switching towards multi-scenarios autonomous driving," in *2021 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2021, pp. 1–8.

[3] B. Osiński, P. Mio, A. Jakubowski, P. Zicina, M. Martyniak, C. Galias, A. Breuer, S. Homoceanu, and H. Michalewski, "Carla real traffic scenarios–novel training ground and benchmark for autonomous driving," *arXiv preprint arXiv:2012.11329*, 2020.

[4] J. M. Scanlon, K. D. Kusano, T. Daniel, C. Alderson, A. Ogle, and T. Victor, "Waymo simulated driving behavior in reconstructed fatal crashes within an autonomous vehicle operating domain," 2021.

[5] A. Kesting, M. Treiber, and D. Helbing, "Enhanced intelligent driver model to access the impact of driving strategies on traffic capacity," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 368, no. 1928, pp. 4585–4605, 2010.

[6] ——, "General lane-changing model mobil for car-following models," *Transportation Research Record*, vol. 1999, no. 1, pp. 86–94, 2007.

[7] S. Feng, X. Yan, H. Sun, Y. Feng, and H. X. Liu, "Intelligent driving intelligence test for autonomous vehicles with naturalistic and adversarial environment," *Nature Communications*, vol. 12, no. 1, pp. 1–14, 2021.

[8] F.-Y. Wang, "Parallel control and management for intelligent transportation systems: Concepts, architectures, and applications," *IEEE Transactions on Intelligent Transportation Systems*, vol. 11, no. 3, pp. 630–638, 2010.

[9] S. Dai, Z. Li, L. Li, N. Zheng, and S. Wang, "A flexible and explainable vehicle motion prediction and inference framework combining semi-supervised aog and st-lstm," *IEEE Transactions on Intelligent Transportation Systems*, vol. PP, no. 99, pp. 1–21, 2020.

[10] X. Zhao, Y. Chen, J. Guo, and D. Zhao, "A spatial-temporal attention model for human trajectory prediction." *IEEE CAA J. Autom. Sinica*, vol. 7, no. 4, pp. 965–974, 2020.

[11] R. P. Bhattacharyya, D. J. Phillips, C. Liu, J. K. Gupta, K. Driggs-Campbell, and M. J. Kochenderfer, "Simulating emergent properties of human driving behavior using multi-agent reward augmented imitation learning," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 789–795.

[12] G. Zheng, H. Liu, K. Xu, and Z. Li, "Objective-aware traffic simulation via inverse reinforcement learning," in *2021 International Joint Conference on Artificial Intelligence (IJCAI)*, 2021.

[13] S. Suo, S. Regalado, S. Casas, and R. Urtasun, "Trafficsim: Learning to simulate realistic multi-agent behaviors," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021, pp. 10 400–10 409.

[14] L. Bergamini, Y. Ye, O. Scheel, L. Chen, C. Hu, L. Del Pero, B. Osinski, H. Grimmett, and P. Ondruska, "Simnet: Learning reactive self-driving simulations from real-world observations," in *2021 International Conference on Robotics and Automation (ICRA)*, 2021.

[15] J. Chen, S. E. Li, and M. Tomizuka, "Interpretable end-to-end urban autonomous driving with latent deep reinforcement learning," *IEEE Transactions on Intelligent Transportation Systems*, pp. 1–11, 2021.

[16] G. Wang, J. Hu, Z. Li, and L. Li, "Harmonious lane changing via deep reinforcement learning," *IEEE Transactions on Intelligent Transportation Systems*, vol. PP, no. 99, pp. 1–9, 2021.

[17] P. Hang, C. Lv, Y. Xing, C. Huang, and Z. Hu, "Human-like decision making for autonomous driving: A noncooperative game theoretic approach," *IEEE Transactions on Intelligent Transportation Systems*, vol. 22, no. 4, pp. 2076–2087, 2020.

[18] W. Zhan, L. Sun, D. Wang, H. Shi, A. Clausse, M. Naumann, J. Kummerle, H. Konigshof, C. Stiller, A. de La Fortelle, *et al.*, "Interaction dataset: An international, adversarial and cooperative motion dataset in interactive driving scenarios with semantic maps," *arXiv preprint arXiv:1910.03088*, 2019.

[19] Y. Chai, B. Sapp, M. Bansal, and D. Anguelov, "Multipath: Multiple probabilistic anchor trajectory hypotheses for behavior prediction," in *Conference on Robot Learning (CoRL)*. PMLR, 2020, pp. 86–99.

[20] H. Song, D. Luan, W. Ding, M. Y. Wang, and Q. Chen, "Learning to predict vehicle trajectories with model-based planning," *arXiv preprint arXiv:2103.04027*, 2021.

[21] H. Zhao, J. Gao, T. Lan, C. Sun, B. Sapp, B. Varadarajan, Y. Shen, Y. Shen, Y. Chai, C. Schmid, *et al.*, "Tnt: Target-driven trajectory prediction," in *Conference on Robot Learning (CoRL)*, 2020, pp. 525–533.

[22] N. Rhinehart, R. McAllister, K. Kitani, and S. Levine, "Precog: Prediction conditioned on goals in visual multi-agent settings," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019, pp. 2821–2830.

[23] S. Narayanan, R. Moslemi, F. Pittaluga, B. Liu, and M. Chandraker, "Divide-and-conquer for lane-aware diverse trajectory prediction," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 15 799–15 808.

[24] C. Tang and R. R. Salakhutdinov, "Multiple futures prediction," *Advances in Neural Information Processing Systems*, vol. 32, pp. 15 424–15 434, 2019.

[25] J. Gao, C. Sun, H. Zhao, Y. Shen, D. Anguelov, C. Li, and C. Schmid, "Vectornet: Encoding hd maps and agent dynamics from vectorized representation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020, pp. 11 525–11 533.

[26] M. Liang, B. Yang, R. Hu, Y. Chen, R. Liao, S. Feng, and R. Urtasun, "Learning lane graph representations for motion forecasting," in *European Conference on Computer Vision (ECCV)*. Springer, 2020, pp. 541–556.

[27] Z.-H. Yin, L. Sun, L. Sun, M. Tomizuka, and W. Zhan, "Diverse critical interaction generation for planning and planner evaluation," *arXiv preprint arXiv:2103.00906*, 2021.

[28] Z. Huang, J. Wu, and C. Lv, "Driving behavior modeling using naturalistic human driving data with inverse reinforcement learning," *IEEE Transactions on Intelligent Transportation Systems*, vol. PP, no. 99, pp. 1–13, 2021.

[29] J. Ho and S. Ermon, "Generative adversarial imitation learning," *Advances in Neural Information Processing Systems*, vol. 29, pp. 4565–4573, 2016.

[30] R. P. Bhattacharyya, D. J. Phillips, B. Wulfe, J. Morton, A. Kuefler, and M. J. Kochenderfer, "Multi-agent imitation learning for driving simulation," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 1534–1539.

[31] M. O'Kelly, A. Sinha, H. Namkoong, R. Tedrake, and J. C. Duchi, "Scalable end-to-end autonomous vehicle testing via rare-event simulation," *Advances in Neural Information Processing Systems*, vol. 31, pp. 9827–9838, 2018.

[32] D. Zhao, Z. Xia, and Q. Zhang, "Model-free optimal control based intelligent cruise control with hardware-in-the-loop demonstration [research frontier]," *IEEE Computational Intelligence Magazine*, vol. 12, no. 2, pp. 56–69, 2017.

[33] B. Chen, X. Chen, Q. Wu, and L. Li, "Adversarial evaluation of autonomous vehicles in lane-change scenarios," *IEEE Transactions on Intelligent Transportation Systems*, pp. 1–10, 2021.

[34] S. Shiroshita, S. Maruyama, D. Nishiyama, M. Y. Castro, K. Hamzaoui, G. Rosman, J. DeCastro, K.-H. Lee, and A. Gaidon, "Behaviorally diverse traffic simulation via reinforcement learning," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 2103–2110.

[35] D. Li, D. Zhao, Q. Zhang, and Y. Chen, "Reinforcement learning and deep learning based lateral control for autonomous driving [application notes]," *IEEE Computational Intelligence Magazine*, vol. 14, no. 2, pp. 83–98, 2019.

[36] P. A. Lopez, M. Behrisch, L. Bieker-Walz, J. Erdmann, Y.-P. Flötteröd, R. Hilbrich, L. Lücken, J. Rummel, P. Wagner, and E. Wießner, "Microscopic traffic simulation using sumo," in *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2018, pp. 2575–2582.

[37] Z. Tang, M. Naphade, M.-Y. Liu, X. Yang, S. Birchfield, S. Wang, R. Kumar, D. Anastasiu, and J.-N. Hwang, "Cityflow: A city-scale benchmark for multi-target multi-camera vehicle tracking and re-identification," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 8797–8806.

[38] E. Leurent, "An environment for autonomous driving decision-making," https://github.com/eleurent/highway-env, 2018.

[39] Y. Luo, P. Cai, Y. Lee, and D. Hsu, "Simulating autonomous driving in massive mixed urban traffic," *arXiv preprint arXiv:2011.05767*, 2020.

[40] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "Carla: An open urban driving simulator," in *Conference on robot learning (CoRL)*. PMLR, 2017, pp. 1–16.

[41] J. Bernhard, K. Esterle, P. Hart, and T. Kessler, "Bark: Open behavior benchmarking in multi-agent environments," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 6201–6208.

[42] M. Zhou, J. Luo, J. Villella, Y. Yang, D. Rusu, J. Miao, W. Zhang, M. Alban, I. Fadakar, Z. Chen, *et al.*, "Smarts: Scalable multi-agent reinforcement learning training school for autonomous driving," *arXiv preprint arXiv:2010.09776*, 2020.

[43] Q. Li, Z. Peng, Z. Xue, Q. Zhang, and B. Zhou, "Metadrive: Composing diverse driving scenarios for generalizable reinforcement learning," *arXiv preprint arXiv:2109.12674*, 2021.

[44] M. Niedoba, H. Cui, K. Luo, D. Hegde, F.-C. Chou, and N. Djuric, "Improving movement prediction of traffic actors using off-road loss and bias mitigation," in *Workshop on 'Machine Learning for Autonomous Driving' at Conference on Neural Information Processing Systems*, 2019.

[45] Anonymous, "Improving diversity of multiple trajectory prediction based on trajectory proposal with lane loss," in *Submitted to 5th Annual Conference on Robot Learning (CoRL)*, 2021, under review. [Online]. Available: https://openreview.net/forum?id=PKJdn4h3kpP

[46] O. Scheel, L. Bergamini, M. Wolczyk, B. Osinski, and P. Ondruska, "Urban driver: Learning to drive from real-world demonstrations using policy gradients," in *5th Annual Conference on Robot Learning (CoRL)*, 2021, accepted. [Online]. Available: https://openreview.net/forum?id=ibktAcINCaj

[47] S. Fujimoto, H. Van Hoof, and D. Meger, "Addressing function approximation error in actor-critic methods," *arXiv preprint arXiv:1802.09477*, 2018.

[48] E. Leurent and J. Mercat, "Social attention for autonomous decision-making in dense traffic," *arXiv preprint arXiv:1911.12250*, 2019.

[49] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," *arXiv*, 2017.

[50] P. Polack, F. Altché, B. d'Andréa Novel, and A. de La Fortelle, "The kinematic bicycle model: A consistent model for planning feasible trajectories for autonomous vehicles?" in *2017 IEEE intelligent vehicles symposium (IV)*. IEEE, 2017, pp. 812–818.