

Causal Composition Diffusion Model for Closed-loop Traffic Generation

Haohong Lin^{1*}, Cyrus Huang², Tung Phan², David Hayden²,
Huan Zhang³, Ding Zhao¹, Siddhartha Srinivasa², Eric Wolff², Hongge Chen²
¹CMU, ²Cruise LLC, ³UIUC

haohongl@cmu.edu, hongge.chen@getcruise.com

Abstract

Simulation is critical for safety evaluation in autonomous driving, particularly in capturing complex interactive behaviors. However, generating **realistic** and **controllable** traffic scenarios in long-tail situations remains a significant challenge. Existing generative models suffer from the conflicting objective between user-defined controllability and realism constraints, which is amplified in safety-critical contexts. In this work, we introduce the **Causal Compositional Diffusion Model (CCDiff)**, a structure-guided diffusion framework to address these challenges. We first formulate the learning of controllable and realistic closed-loop simulation as a constrained optimization problem. Then, CCDiff maximizes controllability while adhering to realism by automatically identifying and injecting causal structures directly into the diffusion process, providing structured guidance to enhance both realism and controllability. Through rigorous evaluations on benchmark datasets and in a closed-loop simulator, CCDiff demonstrates substantial gains over state-of-the-art approaches in generating realistic and user-preferred trajectories. Our results show CCDiff’s effectiveness in extracting and leveraging causal structures, showing improved closed-loop performance based on key metrics such as collision rate, off-road rate, FDE, and comfort. For more details, welcome to check our [project website](#).

1. Introduction

Reliable closed-loop traffic simulation is essential for assessing autonomous vehicle (AV) safety in diverse and complex scenarios [1–5]. Simulations must be both *realistic*, capturing the intricacies of real-world driving, and *controllable*, allowing customization aligned with user preferences. However, balancing realism with controllability remains a significant challenge. Previous works often prioritize one aspect, optimizing either realism or user-specified objectives [3, 5]. How to jointly achieve both objectives under safety-critical conditions remains fruitful yet unresolved.

*This work was done when Haohong was an intern at Cruise LLC.

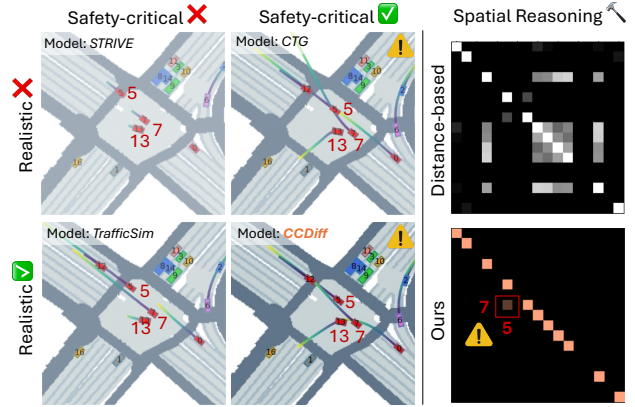


Figure 1. Comparison of safety-critical scenario generation methods, featuring CCDiff alongside existing methods (STRIVE, CTG, and TrafficSim). The illustrated scenario involves Car 13 executing an unprotected left turn, prompting Car 7 to change lanes and interfere with Car 5. Unlike other methods, CCDiff successfully achieves both realism and controllability in generating this safety-critical scenario. In the right column, CCDiff’s spatial reasoning method is compared to a distance-based baseline approach. CCDiff accurately captures the causal relationships between key agents, identifying crucial interactions with greater precision and spatial alignment than distance-based reasoning.

Traffic agent simulation often resorts to either (i) data-driven approaches that generate the most probable trajectories based on scene context or (ii) rule-based approaches that maximize alignment with a user’s control. However, both approaches face key limitations for effective scenario generation.

Data-driven scenario generation faces two primary challenges. First, the rarity of collision and near-miss events in public datasets limits the ability of data-driven methods to generate safety-critical scenarios. As shown in prior studies [2, 6], even a small domain mismatch, such as changes in road structure or the behavior of surrounding vehicles, can cause significant regressions. Second, closed-loop simulation requires that generated trajectories continuously interact with the simulated environment, so current predictions

influence future predictions. This feedback loop often creates compounding errors, leading to distributional shifts that challenge the generation of both controllable and realistic behaviors over long horizons.

On the other hand, rule-based approaches to simulation [7, 8] offer precise user control, but often fail to capture the nuanced, adaptive behaviors of real-world driving, especially in unpredictable scenarios. Their rigidity can make generated behaviors feel scripted and unrealistic, particularly in closed-loop simulations where each action influences future states. This lack of adaptability often leads to compounding errors and a drift from realistic behavior distributions, limiting their effectiveness in complex, long-horizon interactions.

Recent advances in deep generative models have enabled scalable traffic behavior simulation [9, 10], facilitating realistic scenario generation from massive offline datasets. Notably, prior works compose explicit rules into scenario generation, such as causal graphs (CG), signal temporal logic (STL), or large language models (LLM), which act as structured constraints to improve the controllability [3, 4, 11–13]. However, interactive driving scenarios cannot be fully encapsulated by explicit rules alone. Rule-based models struggle to generalize effectively in many safety-critical corner cases, where certain rules may need to be adapted.

Our key insight in driving scenarios is that interactions between agents follow an inherent causal structure: each agent’s actions depend primarily on the states of a subset of nearby agents. Following this observation, we frame the problem as a Constrained Factored Markov Decision Process (MDP), shaped by these causal dependencies to mirror real-world interactions. Unlike previous approaches that manage conflicts between controllability and realism through reweighting, we directly utilize the causal structure by selectively masking agents with conflicting behaviors. This structure enables our model to uphold both realism and controllability constraints simultaneously, even in safety-critical situations. To implement this approach, we introduce the **Causal Composition Diffusion model (CCDiff)**—a structure-enhanced diffusion model that combines structure-aware classifier-free guidance with compositional classifier-based guidance. By integrating these elements, our model achieves a balanced, flexible generation of realistic and controllable driving scenarios, as demonstrated in Figure 1. Our contributions can be summarized as follows:

- We formulate the learning of controllable and realistic closed-loop simulation as a constrained optimization problem, which aims to maximize the user’s control preferences while satisfying realism constraints.
- We propose *CCDiff*, a principled algorithm to solve the constrained optimization problem by identifying the causal structure and injecting it as a structured guidance to the diffusion model.
- We systematically evaluate the performance of *CCDiff*

with state-of-the-art in closed-loop scenario generation on the nuScenes dataset [14], showing benefits in the controllability and realism in generating safety-critical driving scenarios.

2. Related Work

Causal reasoning for behavior models Causal reasoning has seen extensive applications in trajectory modeling, with previous studies leveraging causal structures to enhance the robustness and generalizability of open-loop behavior prediction models. These efforts have included causal representation learning [15], backdoor adjustment [16], counterfactual analysis [17], and realistic causal interventions [11, 18]. Despite these advances, applying a causal approach to broader trajectory prediction tasks [19] often demands substantial human annotation efforts [20]. To address the challenges of automating spatiotemporal reasoning in traffic scenarios, state-of-the-art methods [21–24] employ factorized attention mechanisms. However, while previous work applies causal structured reasoning primarily in open-loop settings, the efficacy of causal behavior modeling in closed-loop scenarios for autonomous driving remains under-explored.

Generative models for traffic simulation Prior arts balance the trade-off between *realism* and *controllability* in safety-critical scenario generation by incorporating various constraints, such as inference-time sampling strategies [10], retrieval-augmented generation [5], low-rank fine-tuning [25], and language-conditioned generation [13]. In closed-loop simulation methods, compositional constraints in the training loss are often integrated into the simulation pipeline. For instance, TrafficSim [26] achieves a balance between realism and common sense using a time-adaptive multi-task loss design; SimNet [27] factorizes trajectory sequences using Markov processes; STRIVE [28] imposes structured priors to constrain samples, avoiding unrealistic outcomes; and BITS [29] optimizes closed-loop performance via bi-level imitation. Yet these prior methods struggle to resolve conflicts between *controllability* and *realism* objectives when these are at odds during inference.

Diffusion model for sequential decision making Diffusion models [30–32] have shown strong controllability in density estimation and generation tasks. Scenario Diffusion [33] adopts latent diffusion, utilizing multi-source conditioning to generate realistic scenarios. In closed-loop traffic simulation, several prior works incorporate compositional classifier-based guidance to steer the diffusion model’s sampling process [34, 35], including signal temporal logic (STL) guidance, language-based guidance [4], adversarial guidance [36–38], and game-theoretic guidance [39].

Appendix Table 4 systematically compares the key features among the prior works and our proposed ap-

proach. While related works often focus on generating rule-compliant normal scenarios or enhancing safety-critical scenarios purely through classifier guidance, a fundamental challenge of achieving a balance between *controllability* and *realism* under safety-critical conditions remains unresolved.

3. Problem Formulation

3.1. Constrained Factored Markov Decision Process

We formulate the closed-loop traffic simulation as an MDP problem, then utilize diffusion model for sequential modeling to learn a controllable simulation policy π . Since we would like to exploit the causal structure between the state, action, and reward space, we define the Constrained Factored MDP as follows:

Definition 1 (Constrained Factored MDP). A *Constrained Factored Markov Decision Process* (CFMDP) is a Markov Decision Process where the state space \mathcal{S} and reward function R are factorized to exploit the structure of the problem. A CFMDP is defined by the tuple: $\mathcal{M}_F = (\mathcal{S}, \mathcal{A}, P, R, C, s_0)$.

The factored state space, denoted as $\mathcal{S} = \mathcal{S}^{(1)} \times \mathcal{S}^{(2)} \times \dots \times \mathcal{S}^{(N)}$, represents the motion trajectory space at the current step for each agent i . The factored action space, $\mathcal{A} = \mathcal{A}^{(1)} \times \mathcal{A}^{(2)} \times \dots \times \mathcal{A}^{(N)}$, consists of interventions on the subsequent driving behaviors for each agent in the scenario. The joint transition dynamics $P(s_t | s_{t-1}, a_{t-1}) = \prod_{i=1}^N p_i(s_t^{(i)} | s_{t-1}^{(i)}, a_{t-1}^{(i)})$, are defined over the state $s \in \mathcal{S}$ and action $a \in \mathcal{A}$ pairs. In our case, P is the *deterministic* vehicle dynamics for each agent in our setting. The reward objective $R(s, a) = \sum_{j=1}^{d_r} R^{(j)}(s^{(I_j)}, a)$ include collision, off-road events, over-speed, or other objectives, where each subset I_j specifies the state factors impacting the j -th reward. For a learned policy π , the constraint function $C(s, a) = \mathbb{D}(\pi_\beta(\cdot | s_t) \| \pi(\cdot | s_t))$ indicates the realism level of generated trajectories with respect to the dataset policies π_β , where a lower constraint value implies greater realism. The initial state s_0 lies in the factored state space \mathcal{S} .

We then formulate the closed-loop scenario generation problem as a constrained optimization problem that aims to find an intervention policy π that maximizes the controllability $R(\tau)$ while maintaining an acceptable deviation in the realism $C(\tau)$:

$$\max_{\pi} \mathbb{E}_{\tau \sim (P, \pi)} [R(\tau)], \quad s.t. \quad \mathbb{E}_{\tau \sim (P, \pi)} [C(\tau)] \leq \kappa,$$

where the cumulative reward, $R(\tau)$, represents the total reward accumulated from individual reward factors along the trajectory τ : $R(\tau) = \sum_{j=1}^{d_r} R^{(j)}(\tau^{(I_j)}) = \mathbb{E} \left[\sum_{t=1}^T \sum_{j=1}^{d_r} R^{(j)}(s_t^{(I_j)}, a_t^{(I_j)}) \right]$. The cumulative cost $C(\tau)$ quantifies the realism constraints, measuring how

closely the generated trajectory τ resembles the ground-truth trajectory τ^* . Following the approach in [3, 29], we use the Total Variation (TV) distance between the estimated intervention policy $\pi(a_t | s_t)$ and the dataset policy $\pi_\beta(a_t | s_t)$: $C(\tau) \triangleq \sum_{t=1}^T \mathbb{D}(\pi_\beta(a_t | s_t) \| \pi(a_t | s_t))$.

To further incorporate the structure in this multi-agent decision-making problem [40], we define the Decision Causal Graph (DCG) below.

Definition 2 (Decision Causal Graph). For every timestep t , we define a causal graph $G \in \mathbb{R}^{N \times N}$, where $G_{ij} = 0$ if and only if the future action of agent j is conditionally independent with the i -th agent's history: $a_t^{(j)} \perp\!\!\!\perp s_t^{(i)} | s_t^{(-i)}$. And $G_{ij} = 1$ means there exists a causal edge $s_t^{(i)} \rightarrow a_t^{(j)}$.

Following the definition above, we can also define a set of policy $\pi(a_t^{(1)}, \dots, a_t^{(N)} | s_t) = \prod_{i=1}^N \pi^{(i)}(a_t^{(i)} | \mathbf{PA}_t^G(i))$, where $\mathbf{PA}_t^G(i) \in \{s^{(1)}, s^{(2)}, \dots, s^{(N)}\}$ is the causal parents to the i -th agents in graph G when making decisions. A diagram of Factored DCG is illustrated in Figure 2(a).

3.2. Diffusion Model for Sequence Modeling

We then solve the traffic simulation inspired by the recent advancement in diffusion-guided sequential data generation [3, 41, 42]. Denote $\tau(k) \triangleq \{(s_t(k), a_t(k))\}_{t=1}^T$ represent the joint state-action trajectory at the k -th diffusion step, $k \in \{0, 1, \dots, K\}$, where $\tau(0)$ denotes the original clean trajectory. The forward diffusion process, acting on $\tau(0)$, gradually corrupts it with Gaussian noise:

$$q(\tau(1:K) | \tau(0)) \triangleq \prod_{k=1}^K q(\tau(k) | \tau(k-1)),$$

$$q(\tau(k) | \tau(k-1)) \triangleq \mathcal{N}(\tau(k); \sqrt{1 - \beta_k} \tau(k-1), \beta_k \mathbf{I}),$$

where β_1, \dots, β_K are pre-defined variance schedules at each diffusion step. Over the forward process, the trajectory is transformed into a standard Gaussian distribution: $q(\tau(K)) \approx \mathcal{N}(\mathbf{0}, \mathbf{I})$. For scenario generation, the reverse diffusion process iteratively denoises from noise to recover the original trajectories. Given a context c (e.g., map features), the reverse process is:

$$p_{\phi, \psi}(\tau(0:K) | c) = p(\tau(K)) \prod_{k=1}^K p_{\theta}(\tau(k-1) | \tau(k), c),$$

$$p_{\phi, \psi}(\tau(k-1) | \tau(k), c) = \mathcal{N}(\tau(k-1); \pi_{\phi, \psi}(\tau(k), k, c), \sigma_k^2 \mathbf{I}),$$

where $p(\tau(K)) = \mathcal{N}(\mathbf{0}, \mathbf{I})$ is the Gaussian prior, and $\pi_{\phi, \psi}$ is the scene encoder parameterized by ϕ, ψ , which will be covered in later sections.

4. Methodology

4.1. Realism Constrained Score Matching

We denote a factored optimality of time-step t , a set of binary random variables as $\mathcal{O}_t = \{\mathcal{O}_t^{(j)}\}_{j=1}^{d_r}$ [41, 43]. The

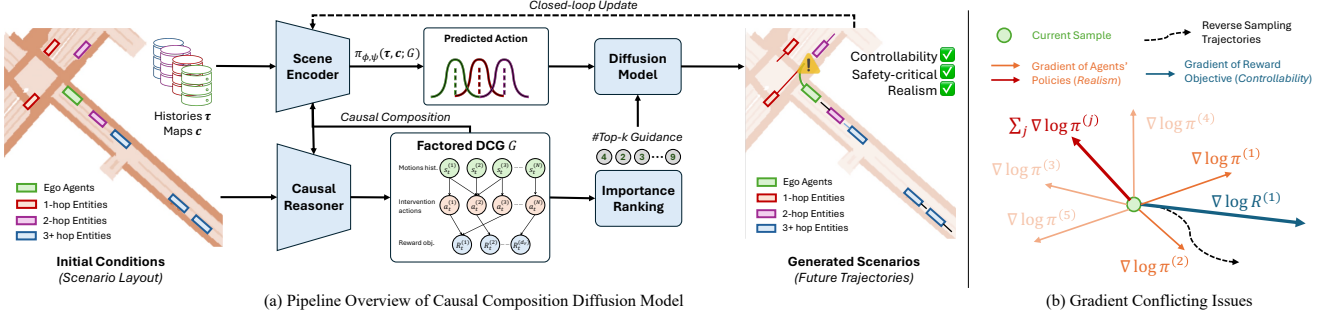


Figure 2. (a): Overview of Causal Composition Diffusion Model. The scene encoder encodes the history and then uses causal reasoning for a structured scene encoding and causal ranking. Finally, we exert guidance only to the top-K agents and eliminate the non-causal agents that would not contribute to the guidance objective to maintain better realism. (b): Summing up the score functions over all the agents achieves sub-optimal performance due to the conflict between the gradients of realism and controllability objectives.

joint optimality in all reward objectives can be written as $p(\mathcal{O}_t^{(j)} = 1 | \tau_t) \propto \exp(R^{(j)}(\tau_t^{(j)}))$. We slightly exploit the notations $\tau \triangleq \{(s_t, a_t)\}_{t=1}^T$ as the trajectories of state action pairs, where $s_t \in \mathbb{R}^{N \times d_s}$ is the state trajectories for all N agents. Given the CFMDP in Definition 1, with known transition (vehicle) dynamics, we can factorize the objective of the optimal closed-loop scenario generation as follows:

$$\begin{aligned} \max_{\pi} P(\mathcal{O}_t = 1, \tau_t | \tau_{t-1}) &\Leftrightarrow \max_{\pi} P(\mathcal{O}_t = 1 | \tau_t) P(\tau_t | \tau_{t-1}) \Leftrightarrow \\ \max_{\pi} P(\mathcal{O}_t = 1 | s_t, a_t) \pi(a_t | s_t) P(s_t | s_{t-1}, a_{t-1}) &\Leftrightarrow \\ \max_{\pi} \underbrace{\prod_{j=1}^{d_r} \exp(R^{(j)}(s_t^{(I_j)}, \pi(s_t)))}_{\text{Controllability}} \underbrace{\prod_{i=1}^N \pi^{(i)}(a_t^{(i)} | s_t)}_{\text{Realism}}, \end{aligned} \quad (1)$$

where the first term corresponds to *controllability*, i.e. the likelihood of optimality specified by some user-specified reward objective, and the second term corresponds to the *realism*, the likelihood of generated behaviors. We denote $\nabla \log P \triangleq \nabla \log P(\mathcal{O}_t = 1, \tau_t | \tau_{t-1})$. The score function of the maximum likelihood objective in equation (1) can be written as [30, 41]:

$$\nabla \log P = \sum_{j=1}^{d_r} \nabla_{\tau} R^{(j)}(s_t^{(I_j)}, \pi(s_t)) + \sum_{i=1}^N \nabla_{\tau} \log \pi^{(i)}(a_t^{(i)} | s_t) \quad (2)$$

Unlike the normal scenarios where optimizing the imitation basically adheres with the rule compliance reward, safety-critical guidance $R^{(j)}$ can suffer from gradient conflict [44–47]. Namely, for some $i \in [1, N], j \in [1, d_r]$, if

$$\langle \nabla_{\tau} \log \pi^{(i)}(a_t^{(i)} | s_t), \nabla_{\tau} R^{(j)}(\tau^{(I_j)}) \rangle < 0,$$

using a weighted sum of all the objectives as classifier-based guidance would achieve sub-optimal performance, as illustrated in Figure 2(b). In order to resolve this gradient conflicting issues, we need to prioritize to control the agents index $i \in [N]$ that could maximize the reward while maintaining a high likelihood of the learned policies, i.e., a lower

realism gap between the learned policies and behavior policies $\pi_{\beta}: \mathbb{D}(\pi_{\beta} || \hat{\pi})$. We denote the flag of controllable agents as $\rho \in [0, 1]^N$, the target simulation policies:

$$\pi^{(i)}(a_t^{(i)} | s_t) = \begin{cases} \pi_{\rho, \psi}^{(i)}(a_t^{(i)} | s_t), & \rho_i = 1 \\ \pi_{\beta}^{(i)}(a_t^{(i)} | s_t), & \rho_i = 0 \end{cases}$$

We can use the Lagrangian multiplier [48] and structured projected gradient descent [49] to solve the constrained optimization with the following maximum likelihood estimation problems:

$$\begin{aligned} \max_{\pi \in \Pi, \rho \in \{0,1\}^N} \prod_{j=1}^{d_r} \exp(R^{(j)}(\tau_t^{(I_j)}; \rho)) \prod_{i \in [N], \rho_i = 1} \pi^{(i)}(a_t^{(i)} | \mathbf{PA}_t^G(i)) \\ \text{s.t. } |G| \leq C_{\text{sparsity}}, \quad \sum_i \rho_i \leq N_c. \end{aligned} \quad (3)$$

We can then control the realism level by changing the constraint level of $N_c, C_{\text{sparsity}} \in \mathbb{Z}^+$.

4.2. Proposed Method: CCDiff

We hereby introduce *CCDiff* to optimize the simulation policy π of equation (3) in a scalable and efficient way by decomposing the constrained optimization problem into several small components. We illustrate the pipeline in Figure 2(a). To promote *realism*, *CCDiff* first encodes the motion histories of different agents based on the spatial attention, then discovers the decision causal graph G based on the factorized attention masks and kinematic factors. *CCDiff* then utilizes causal interactive patterns in G to extract the importance rank ρ . Finally, *CCDiff* optimizes its *controllability* by masking out those unimportant agents based on ρ to guide the diffusion reverse sampling process in a structured way. We zoom into the details below.

Causal Composition Scene Encoder Inspired by [50, 51], the goal of causal composition scene encoder is to generate

histories $\phi(s_{\text{ego}})$, and $\pi_{\phi,\psi}(a_t^{(i)}|\mathbf{PA}_t^G(i))$ is the intervened encoded results given some parental agents in causal graph G . This formulation implies that the guided distribution corresponds to a geometric mixture:

$$\pi(a_t^{(i)}) \propto \underbrace{\pi(a_t^{(i)}|\mathbf{PA}_t^I(i))^{1-w}}_{\text{Original}} \cdot \underbrace{\pi(a_t^{(i)}|\mathbf{PA}_t^G(i))^w}_{\text{Intervened}}.$$

Thus, classifier-free guidance in diffusion models can be viewed as a do-intervention [52] by specifying the causal parents of i -th agents in DCG during the generative process. The guidance scale $w \in [1, 2)$ acts analogously to the strength of intervention, extrapolating the original and intervened distributions. With the causal ranking, we mask out the agents with conflicted gradients as a reweighted classifier-based guidance:

$$\sum_{j=1}^{d_r} \nabla_{\mathbf{a}(t_j)} R^{(j)}(\tau) \approx \sum_{j=1}^{d_r} \rho_j(\tau) \odot \nabla_{\mathbf{a}}[R^{(j)}(\tau)] \quad (5)$$

In practice, we use the distance-based guidance objective over the trajectories, including the map collision guidance and the agent collision guidance [3]. We also use the same classifier function for all the baseline methods, see detailed description in the appendix C.3.

Training and Inference We train the model with using the classical DDPM [31] diffusion with classifier-free guidance [32]. Specifically, the loss we solve is $\min_{\phi,\psi} \mathbb{E} \|\pi_{\phi,\psi}(\tau(k), \mathbf{c}, k; G(\tau)) - \mathbf{a}(0)\|^2$. We introduce counterfactual conditions by randomly dropping the DCG G of the scene transformer by replacing the decision causal graph G as a diagonal matrix, so all agents' actions are only conditioned on the ego history. At inference time, we combine both classifier-based and classifier-free guidance to facilitate better controllability, see algorithm 1.

Algorithm 1 *CCDiff* for Scenario Generation

Require: Dropout p_{uncond} , threshold C_ρ, C_{tc}
Require: Guidance loss $\{\mathcal{J}_i\}_{i=1}^N$, trajectories τ , map \mathbf{c} .
while $k = K, \dots, 1$ **do** ▷ *Inference Sampling*
 $\pi_{\text{uncond}}, \alpha_{\text{attn}} \leftarrow \pi_{\phi,\psi}(\tau(k), \mathbf{c}, k; \emptyset)$
 $G(\tau) \leftarrow M(\tau) \cdot \alpha_{\text{attn}}$ ▷ Causal masking
 $\rho(\tau) \leftarrow \text{ranking}(G)$ ▷ Importance ranking
 $\hat{\pi} \leftarrow (1 - \omega)\pi_{\phi,\psi}(\tau(k), \mathbf{c}, k; G) + \omega\pi_{\text{uncond}}$
 $\mathbf{a}(0) \sim \hat{\pi}(\cdot|\tau(k), \mathbf{c}; G)$
 $\mathbf{a}(k-1) \leftarrow \mathbf{a}(k-1) + \sum_{i=1}^{d_r} \rho_i(\tau) \cdot \nabla R(\mathbf{s}, \mathbf{a}(k-1))$
 $\tau(k-1) \leftarrow f_{\text{dyn}}(\mathbf{s}, \mathbf{a}(k-1))$ ▷ Vehicle dynamics
end while
return Generated trajectory $\tau(0)$.

5. Experiment

In the following parts of the experiments, we aim to answer the following three research questions (RQs): **RQ1:** Under different sizes of total controllable agents, how are the

realism and controllability of the safety-critical scenarios generated by *CCDiff* compared to the baselines? **RQ2:** With a longer generation horizon and lower frequency, how are the realism and controllability of the safety-critical scenarios generated by *CCDiff* compared to the baselines? **RQ3:** How much does the causal reasoning module in *CCDiff* contribute to the overall performance?

The remaining parts of the experiment section first introduce our experiment settings, then compare our methods with baselines in the *controllability* and *realism* to answer the research questions. Finally, we conduct ablation studies to show the effect of individual modules in *CCDiff*.

5.1. Experiment Settings

Datasets We use the nuScenes dataset [14] and traffic behavior simulation (tbsim) [29] for model training and evaluation. We train all models on scenes from the train split and evaluate on 100 scenes randomly sampled from the validation split. During evaluation phase, we initialize all the models with the same set of initial layouts and initial history trajectories of 3 seconds, the model is responsible of generating the future 10 seconds of trajectories for the driving agents in a closed-loop manner.

Baselines We implement the following baselines in the above platform settings. To systematically illustrate the effectiveness of *CCDiff*, we include the following SOTAs for comparison: **SimNet** [27], **TrafficSim** [26], **BITS** [29], **Strive** [28], and **CTG** [3]. We compare all the baselines with *CCDiff* in the publicly available nuScenes dataset [14] and baseline implementations¹. For a fair comparison with all the baselines, we use the rasterized map used in the previous works and encode them with ResNet-18 for the map conditioning \mathbf{c} for all the methods.

Metrics We compare the performance of *CCDiff* and all the baselines with the following categories of metrics:

- **Controllability Score (CS):** we use the scenario-wise collision rate (SCR) used in [13, 26] as the controllability metrics. Among all the testing scenarios, we calculate the proportion of the scenarios where at least one collision event occurred between different agents. We then standardize SCR among all the methods to get the CS, a higher-the-better score between 0 and 1.
- **Realism Score (RS):** How to quantify realism is an open problem in evaluating traffic scenarios. In order to get a more interpretable and direct way to quantify realism, we adopt three widely-used quantitative metrics to evaluate the realism of the scenarios: (i) scenario off-road rate (ORR) used in [5, 13], (ii) final displacement error (FDE, m) and (iii) comfort distance (CFD), which is used in [3, 29] to quantify the realism of the similarity in the smoothness of agents' trajectories in the generated

¹<https://github.com/NVlabs/CTG>

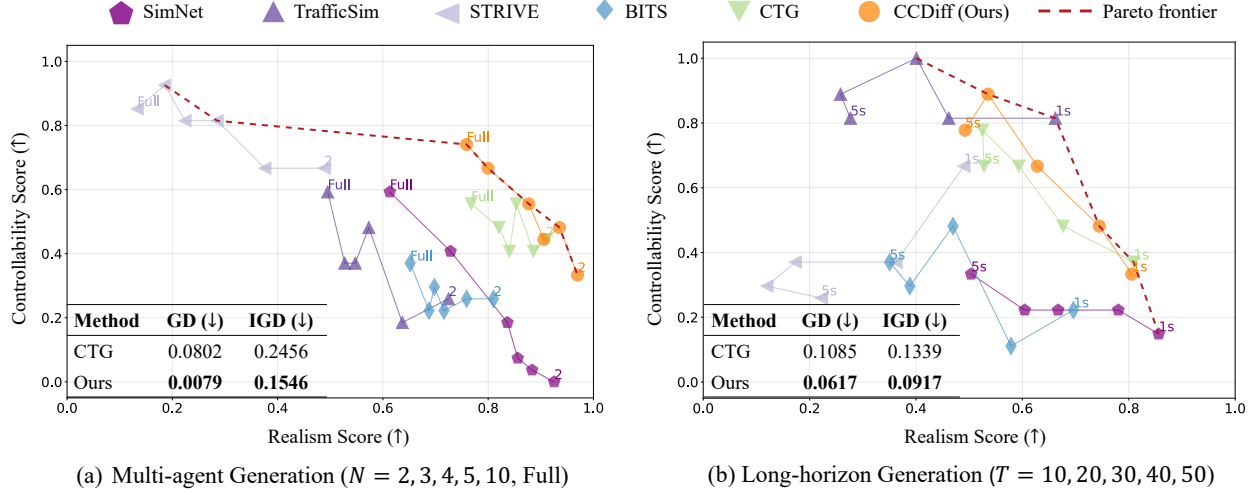


Figure 4. Plot of the controllability v.s. realism in the multi-agent and long-horizon generation settings. *CCDiff* outperforms baselines in both the Generational Distance (GD) and Inverted Generational Distance (IGD), with better proximity to the Pareto frontier, and better coverage of the optimal solution along the frontier in this multi-objective optimization. Our method is more realistic and controllable compared to other approaches consistently in both multi-agent scenario generation and long-horizon scenario generation.

scenarios. We standardize all the metrics among all the methods respectively and average them to get the **RS**, a higher-the-better score between 0 and 1.

- **Multi-objective optimization metrics:** with the **RS** and **CS**, we further quantify the optimality of the solution based on generational distance (**GD**) and inverted generational distance (**IGD**), the average minimum distance between the methods and Pareto frontier [44, 53].

5.2. Multi-agent Scenario Generation (RQ1)

To address **RQ1**, we train our model on the training split of the nuScenes dataset and vary the number of controllable agents from 2 agents to the full sets of agents for all baselines at inference time by running a closed-loop generation at 2Hz (0.5s). We then evaluate the CS and RS of generated scenarios under different numbers of controllable agents and report the comparison between *CCDiff* with SOTAs in Table 1 and Figure 4(a).

From the table, we can see that *CCDiff* outperforms SimNet, TrafficSim, and BITS in both controllability and realism for almost all the cases, whereas TrafficSim only outperforms *CCDiff* in one controllability score. STRIVE shows some good controllability in generating safety-critical scenarios, yet its realism score is the poorest among all. The closest competitor, CTG, shows comparable performance in realism, yet *CCDiff* outperforms CTG in the controllability metrics, especially when the size of controllable agents goes larger as we gradually scale up the number of controllable agents from 2 to 5 and eventually to the full size of agents at the scene. From Figure 4(a), we can also see that *CCDiff* enjoys significantly better realism and controllability score on the most upper right side, Pareto front. More detailed results

for the SCR, ORR, FDE, and CFD are illustrated in the appendix Table 5. The consistent benefits of *CCDiff* in RS and CS compared to other baselines confirm that composing causal structure in the diffusion model facilitates the algorithm to generate reasonable safety-critical scenarios. We also show qualitative studies for long-horizon generation in the Appendix Figure 26.

Table 1. Comparison in Controllability (CS) and Realism (RS) among all the baselines in K -agent scenario generation. Red means *CCDiff* outperforms the baseline in the corresponding metrics, green means the baseline is better. *CCDiff* has **best** or **second best** performance in 10 out of 12 metrics.

Method	Metric	$K=2$	3	4	5	10	Full
SimNet	CS (↑)	0.00	0.04	0.07	0.19	0.41	0.59
	RS (↑)	0.93	0.88	0.86	0.84	0.73	0.61
TrafficSim	CS (↑)	0.26	0.19	0.48	0.37	0.37	0.59
	RS (↑)	0.72	0.64	0.57	0.55	0.53	0.50
STRIVE	CS (↑)	0.67	0.67	0.81	0.81	0.93	0.85
	RS (↑)	0.49	0.38	0.29	0.22	0.19	0.13
BITS	CS (↑)	0.26	0.26	0.22	0.30	0.22	0.37
	RS (↑)	0.81	0.76	0.72	0.70	0.69	0.65
CTG	CS (↑)	0.44	0.41	0.56	0.41	0.48	0.56
	RS (↑)	0.91	0.89	0.85	0.84	0.82	0.77
Ours	CS (↑)	0.33	0.48	0.44	0.56	0.67	0.74
	RS (↑)	0.97	0.94	0.91	0.88	0.80	0.76

5.3. Long-horizon Closed-loop Generation (RQ2)

To address **RQ2**, we evaluate *CCDiff*'s performance in long-horizon safety-critical scenario generation by changing the

Table 2. Comparison in **RS** and **CS** in long-horizon scenario generation of 5 controllable agents over T seconds. **Red** means *CCDiff* outperforms the baseline in the corresponding metrics, **green** means the baseline is better, and **yellow** means a tie. *CCDiff* has **best** and **second best** performance in 10 out of 12 metrics.

Method	Metric	$T=1s$	2s	3s	4s	5s
SimNet	CS (\uparrow)	0.15	0.22	0.22	0.22	0.33
	RS (\uparrow)	0.86	0.78	0.67	0.60	0.50
TrafficSim	CS (\uparrow)	0.81	0.81	1.00	0.89	0.81
	RS (\uparrow)	0.66	0.46	0.40	0.26	0.28
STRIVE	CS (\uparrow)	0.67	0.37	0.37	0.30	0.26
	RS (\uparrow)	0.49	0.36	0.17	0.12	0.22
BITS	CS (\uparrow)	0.22	0.11	0.48	0.30	0.37
	RS (\uparrow)	0.70	0.58	0.47	0.39	0.35
CTG	CS (\uparrow)	0.37	0.48	0.67	0.78	0.67
	RS (\uparrow)	0.81	0.68	0.59	0.53	0.53
CCDiff	CS (\uparrow)	0.33	0.48	0.67	0.89	0.78
	RS (\uparrow)	0.81	0.74	0.63	0.54	0.49

simulation frequency in T seconds. We test the generation results with $T \in \{0.5s, 1s, 2s, 3s, 4s, 5s\}$, which corresponds to a closed-loop simulation frequency between 0.2Hz to 2Hz. We demonstrate the comparison of realism and controllability results in Table 2 and Figure 4(b). *CCDiff* consistently outperform BITS and STRIVE in both **CS** and **RS** as the planning horizon enlarges. TrafficSim outperforms all the other baselines with the best controllability, while its realism in long-horizon generation is second-worst and only better than BITS. SimNet marginally outperforms *CCDiff* in realism, yet its controllability is the worst among all. *CCDiff* outperforms CTG with a comparable realism score and higher controllability at longer horizons. Figure 4(b) confirms our approach has the best proximity to the multi-objective Pareto frontier and has best coverage in the realistic zone. We also show qualitative studies for long-horizon generation in Appendix Figure 9-27.

5.4. Ablation study (RQ3)

To answer **RQ3**, we evaluate our methods with different ablation variants related to the causal composition, including (i) *CCDiff* w/o encoder, which removes sparsity constraints $\lambda_{\text{sparsity}}$ of the causal composition scene encoder, (ii) *CCDiff* w/o factored guide, which replaces the factorized guidance with the whole state space guidance, (iii) *CCDiff* w/ human and (iv) *CCDiff* w/ distance which replace the causal ranking algorithms with distance-based ranking and human ranking. We demonstrate the quantitative results in Table 6. The non-causal guidance variants and non-causal encoder variants show a performance drop in the collision rate (controllabil-

Table 3. Ablation study on *CCDiff*'s variants. Evaluation of Controllability (CO, OR) and Realism (FDE and CFD) over different agent scales. For each metric we highlight the **best** and the **second best** results.

Enc.	Guide	Rank	Metrics	$K=1$	2	3	4	5
	\checkmark	\checkmark	SCR (\uparrow)	0.32	0.43	0.44	0.43	0.42
			ORR (\downarrow)	0.53	1.10	0.98	0.91	0.91
			FDE (\downarrow)	2.18	4.00	5.41	5.87	5.79
			CFD (\downarrow)	1.09	1.00	1.14	1.22	1.22
\checkmark		\checkmark	SCR (\uparrow)	0.31	0.38	0.45	0.40	0.40
			ORR (\downarrow)	0.33	0.81	0.76	1.00	1.06
			FDE (\downarrow)	2.21	4.33	5.28	6.13	6.82
			CFD (\downarrow)	1.51	1.81	1.60	1.84	1.94
\checkmark	\checkmark	Dist	SCR (\uparrow)	0.32	0.33	0.34	0.36	0.37
			ORR (\downarrow)	0.63	1.38	1.50	1.59	1.49
			FDE (\downarrow)	2.80	4.15	5.15	5.79	5.96
			CFD (\downarrow)	1.24	1.79	2.44	2.03	2.34
\checkmark	\checkmark	Human	SCR (\uparrow)	0.28	0.34	0.35	0.33	0.31
			ORR (\downarrow)	0.67	1.66	1.65	1.73	1.93
			FDE (\downarrow)	3.13	5.80	6.74	7.40	7.84
			CFD (\downarrow)	1.31	2.21	2.51	2.83	3.14
\checkmark	\checkmark	\checkmark	SCR (\uparrow)	0.29	0.40	0.44	0.43	0.46
			ORR (\downarrow)	0.37	0.61	0.72	0.99	1.02
			FDE (\downarrow)	2.16	4.17	5.22	5.99	6.59
			CFD (\downarrow)	1.70	1.88	1.92	1.93	2.25

ity), yet the w/o encoder variants outperform in kinematic comfort. Among all the ablation variants, different ranking strategies result in the *largest* performance drop for the multi-agent controllable generation settings. Compared to our causal ranking, human ranking and distance-based ranking strategy suffers from a performance drop with 5 to 10% in the collision rate and, more than 0.5% in the off-road rate, 1m for FDE, and also larger CFD. This signifies the importance of applying the proper guidance to the correct agents at the traffic scene when doing safety-critical generation.

6. Conclusion

In this paper, we propose *CCDiff*, a causal composition diffusion model that aims to improve the *controllability* and *realism* in closed-loop safety-critical scenario generation for autonomous driving. Based on the formulation of constrained factored MDP, *CCDiff* promotes realism by first identifying the underlying causal structure between agents, then incorporating it in the scene encoder and ranking the importance of agents based on causal knowledge. *CCDiff* uses both interventional classifier-free guidance and masked classifier guidance to improve controllability in safety-critical scenario generation. In multi-agent generation and long-horizon generation settings, *CCDiff* outperforms SOTA methods over nuScenes data in closed-loop evaluation. One limitation of the current work is that the design of the causal reasoning pipeline relies on hyperparameter tuning and it is hard to directly evaluate. It would be interesting to construct a traffic reasoning benchmark and incorporate a foundation model to further scale up the traffic reasoning and generation process.

References

- [1] Shuo Feng, Haowei Sun, Xintao Yan, Haojie Zhu, Zhengxia Zou, Shengyin Shen, and Henry X Liu. Dense reinforcement learning for safety validation of autonomous vehicles. *Nature*, 615(7953):620–627, 2023. 1
- [2] Wenhao Ding, Chejian Xu, Mansur Arief, Haohong Lin, Bo Li, and Ding Zhao. A survey on safety-critical driving scenario generation—a methodological perspective. *IEEE Transactions on Intelligent Transportation Systems*, 2023. 1
- [3] Ziyuan Zhong, Davis Rempe, Danfei Xu, Yuxiao Chen, Sushant Veer, Tong Che, Baishakhi Ray, and Marco Pavone. Guided conditional diffusion for controllable traffic simulation. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3560–3566. IEEE, 2023. 1, 2, 3, 6, 5
- [4] Ziyuan Zhong, Davis Rempe, Yuxiao Chen, Boris Ivanovic, Yulong Cao, Danfei Xu, Marco Pavone, and Baishakhi Ray. Language-guided traffic simulation via scene-level diffusion. In *Conference on Robot Learning*, pages 144–177. PMLR, 2023. 2, 5, 1
- [5] Wenhao Ding, Yulong Cao, Ding Zhao, Chaowei Xiao, and Marco Pavone. Realgen: Retrieval augmented generation for controllable traffic scenarios. *arXiv preprint arXiv:2312.13303*, 2023. 1, 2, 6
- [6] Xuemin Hu, Shen Li, Tingyu Huang, Bo Tang, Rouxing Huai, and Long Chen. How simulation helps autonomous driving: A survey of sim2real, digital twins, and parallel intelligence. *IEEE Transactions on Intelligent Vehicles*, 2023. 1
- [7] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. Carla: An open urban driving simulator. In *Conference on robot learning*, pages 1–16. PMLR, 2017. 2
- [8] Pablo Alvarez Lopez, Michael Behrisch, Laura Bieker-Walz, Jakob Erdmann, Yun-Pang Flötteröd, Robert Hilbrich, Leonhard Lücken, Johannes Rummel, Peter Wagner, and Evamarie Wießner. Microscopic traffic simulation using sumo. In *2018 21st international conference on intelligent transportation systems (ITSC)*, pages 2575–2582. IEEE, 2018. 2
- [9] Jingkan Wang, Ava Pun, James Tu, Sivabalan Manivasagam, Abbas Sadat, Sergio Casas, Mengye Ren, and Raquel Urtasun. Advsim: Generating safety-critical scenarios for self-driving vehicles. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9909–9918, 2021. 2
- [10] Shuhan Tan, Kelvin Wong, Shenlong Wang, Sivabalan Manivasagam, Mengye Ren, and Raquel Urtasun. Scenegen: Learning to generate realistic traffic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 892–901, 2021. 2, 1
- [11] Wenhao Ding, Haohong Lin, Bo Li, and Ding Zhao. Causalaf: Causal autoregressive flow for safety-critical driving scenario generation. In *Conference on robot learning*, pages 812–823. PMLR, 2023. 2, 1
- [12] Jay Patrikar, Sushant Veer, Apoorva Sharma, Marco Pavone, and Sebastian Scherer. Rulefuser: Injecting rules in evidential networks for robust out-of-distribution trajectory prediction. *arXiv preprint arXiv:2405.11139*, 2024.
- [13] Shuhan Tan, Boris Ivanovic, Xinshuo Weng, Marco Pavone, and Philipp Kraehenbuehl. Language conditioned traffic generation. In *Conference on Robot Learning*, pages 2714–2752. PMLR, 2023. 2, 6, 1
- [14] H Caesar, V Bankiti, AH Lang, S Vora, VE Liong, Q Xu, A Krishnan, Y Pan, G Baldan, and O Beijbom. nuscenes: A multimodal dataset for autonomous driving. *arxiv*. 2019. 2, 6
- [15] Yuejiang Liu, Riccardo Cadei, Jonas Schweizer, Sherwin Bahmani, and Alexandre Alahi. Towards robust and adaptive motion forecasting: A causal representation perspective. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17081–17092, 2022. 2, 5
- [16] Chunjiang Ge, Shiji Song, and Gao Huang. Causal intervention for human trajectory prediction with cross attention mechanism. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 658–666, 2023. 2
- [17] Guangyi Chen, Junlong Li, Jiwen Lu, and Jie Zhou. Human trajectory prediction via counterfactual analysis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9824–9833, 2021. 2, 5
- [18] Mozghan Pourkeshavarz, Junrui Zhang, and Amir Rasouli. Cadet: a causal disentanglement approach for robust trajectory prediction in autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14874–14884, 2024. 2, 5
- [19] Scott Ettinger, Shuyang Cheng, Benjamin Caine, Chenxi Liu, Hang Zhao, Sabeek Pradhan, Yuning Chai, Ben Sapp, Charles R Qi, Yin Zhou, et al. Large scale interactive motion forecasting for autonomous driving: The waymo open motion dataset. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9710–9719, 2021. 2
- [20] Rebecca Roelofs, Liting Sun, Ben Caine, Khaled S Refaat, Ben Sapp, Scott Ettinger, and Wei Chai. Causalagents: A robustness benchmark for motion forecasting using causal relationships. *arXiv preprint arXiv:2207.03586*, 2022. 2

- [21] Jiquan Ngiam, Benjamin Caine, Vijay Vasudevan, Zhengdong Zhang, Hao-Tien Lewis Chiang, Jeffrey Ling, Rebecca Roelofs, Alex Bewley, Chenxi Liu, Ashish Venugopal, et al. Scene transformer: A unified architecture for predicting multiple agent trajectories. *arXiv preprint arXiv:2106.08417*, 2021. 2
- [22] Zikang Zhou, Luyao Ye, Jianping Wang, Kui Wu, and Kejie Lu. Hivt: Hierarchical vector transformer for multi-agent motion prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8823–8833, 2022. 5
- [23] Zikang Zhou, Jianping Wang, Yung-Hui Li, and Yu-Kai Huang. Query-centric trajectory prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17863–17873, 2023.
- [24] Nigamaa Nayakanti, Rami Al-Rfou, Aurick Zhou, Kratarth Goel, Khaled S Refaat, and Benjamin Sapp. Wayformer: Motion forecasting via simple & efficient attention networks. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2980–2987. IEEE, 2023. 2
- [25] Robert Dyro, Matthew Foutter, Ruolin Li, Luigi Di Lillo, Edward Schmerling, Xilin Zhou, and Marco Pavone. Realistic extreme behavior generation for improved av testing. *arXiv preprint arXiv:2409.10669*, 2024. 2
- [26] Simon Suo, Sebastian Regalado, Sergio Casas, and Raquel Urtasun. Trafficsim: Learning to simulate realistic multi-agent behaviors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10400–10409, 2021. 2, 6, 1, 5
- [27] Luca Bergamini, Yawei Ye, Oliver Scheel, Long Chen, Chih Hu, Luca Del Pero, Błażej Osiński, Hugo Grimmett, and Peter Ondruska. Simnet: Learning reactive self-driving simulations from real-world observations. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5119–5125. IEEE, 2021. 2, 6, 1, 5
- [28] Davis Rempe, Jonah Philion, Leonidas J Guibas, Sanja Fidler, and Or Litany. Generating useful accident-prone driving scenarios via a learned traffic prior. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17305–17315, 2022. 2, 6, 1, 5
- [29] Danfei Xu, Yuxiao Chen, Boris Ivanovic, and Marco Pavone. Bits: Bi-level imitation for traffic simulation. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2929–2936. IEEE, 2023. 2, 3, 6, 1, 5
- [30] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32, 2019. 2, 4
- [31] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 6
- [32] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022. 2, 5, 6
- [33] Ethan Pronovost, Meghana Reddy Ganesina, Noureldin Hendy, Zeyu Wang, Andres Morales, Kai Wang, and Nick Roy. Scenario diffusion: Controllable driving scenario generation with diffusion. *Advances in Neural Information Processing Systems*, 36:68873–68894, 2023. 2
- [34] Chiyu Jiang, Andre Cornman, Cheolho Park, Benjamin Sapp, Yin Zhou, Dragomir Anguelov, et al. Motion-diffuser: Controllable multi-agent motion prediction using diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9644–9653, 2023. 2
- [35] Chiyu Max Jiang, Yijing Bai, Andre Cornman, Christopher Davis, Xiukun Huang, Hong Jeon, Sakshum Kulshrestha, John Wheatley Lambert, Shuangyu Li, Xuanyu Zhou, et al. Scenediffuser: Efficient and controllable driving simulation initialization and rollout. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*. 2
- [36] Chejian Xu, Ding Zhao, Alberto Sangiovanni-Vincentelli, and Bo Li. Diffscene: Diffusion-based safety-critical scenario generation for autonomous vehicles. In *The Second Workshop on New Frontiers in Adversarial Machine Learning*, 2023. 2
- [37] Wei-Jer Chang, Francesco Pittaluga, Masayoshi Tomizuka, Wei Zhan, and Manmohan Chandraker. Controllable safety-critical closed-loop traffic simulation via guided diffusion. *arXiv preprint arXiv:2401.00391*, 2023.
- [38] Yuting Xie, Xianda Guo, Cong Wang, Kunhua Liu, and Long Chen. Advdiffuser: Generating adversarial safety-critical driving scenarios via guided diffusion. *arXiv preprint arXiv:2410.08453*, 2024. 2
- [39] Zhiyu Huang, Zixu Zhang, Ameya Vaidya, Yuxiao Chen, Chen Lv, and Jaime Fernández Fisac. Versatile scene-consistent traffic scenario generation as optimization with diffusion. *arXiv preprint arXiv:2404.02524*, 2024. 2
- [40] St John Grimbly, Jonathan Shock, and Arnu Pretorius. Causal multi-agent reinforcement learning: Review and open problems. *arXiv preprint arXiv:2111.06721*, 2021. 3
- [41] Michael Janner, Yilun Du, Joshua B Tenenbaum, and Sergey Levine. Planning with diffusion for flexible

- behavior synthesis. *arXiv preprint arXiv:2205.09991*, 2022. 3, 4, 5
- [42] Anurag Ajay, Yilun Du, Abhi Gupta, Joshua Tenenbaum, Tommi Jaakkola, and Pulkit Agrawal. Is conditional generative modeling all you need for decision-making? *arXiv preprint arXiv:2211.15657*, 2022. 3
- [43] Sergey Levine. Reinforcement learning and control as probabilistic inference: Tutorial and review. *arXiv preprint arXiv:1805.00909*, 2018. 3
- [44] Tianhe Yu, Saurabh Kumar, Abhishek Gupta, Sergey Levine, Karol Hausman, and Chelsea Finn. Gradient surgery for multi-task learning. *Advances in Neural Information Processing Systems*, 33:5824–5836, 2020. 4, 7, 6
- [45] Anh-Dung Dinh, Daochang Liu, and Chang Xu. Pixelasparam: A gradient view on diffusion sampling with guidance. In *International Conference on Machine Learning*, pages 8120–8137. PMLR, 2023.
- [46] Yihang Yao, Zuxin Liu, Zhepeng Cen, Peide Huang, Tingnan Zhang, Wenhao Yu, and Ding Zhao. Gradient shaping for multi-constraint safe reinforcement learning. In *6th Annual Learning for Dynamics & Control Conference*, pages 25–39. PMLR, 2024.
- [47] Qianli Ma, Xuefei Ning, Dongrui Liu, Li Niu, and Linfeng Zhang. Decouple-then-merge: Towards better training for diffusion models. *arXiv preprint arXiv:2410.06664*, 2024. 4
- [48] Joshua Achiam, David Held, Aviv Tamar, and Pieter Abbeel. Constrained policy optimization. In *International Conference on Machine Learning*, pages 22–31. PMLR, 2017. 4
- [49] Sohail Bahmani, Petros T Boufounos, and Bhiksha Raj. Learning model-based sparsity via projected gradient descent. *IEEE Transactions on Information Theory*, 62(4):2092–2099, 2016. 4
- [50] Philippe Brouillard, Sébastien Lachapelle, Alexandre Lacoste, Simon Lacoste-Julien, and Alexandre Drouin. Differentiable causal discovery from interventional data. *Advances in Neural Information Processing Systems*, 33:21865–21877, 2020. 4
- [51] Nino Scherrer, Olexa Bilaniuk, Yashas Annadani, Anirudh Goyal, Patrick Schwab, Bernhard Schölkopf, Michael C Mozer, Yoshua Bengio, Stefan Bauer, and Nan Rosemary Ke. Learning neural causal models with active interventions. *arXiv preprint arXiv:2109.02429*, 2021. 4
- [52] Judea Pearl. *Causality*. Cambridge university press, 2009. 6
- [53] Carlos A Coello Coello. *Evolutionary algorithms for solving multi-objective problems*. Springer, 2007. 7, 6

CCDiff: Causal Composition Diffusion Model for Closed-loop Traffic Generation (Supplementary Materials)

A. Additional Related Works

Table 4. Key features of related works in scenario generation for autonomous vehicles.

Paper	Controllability	Realism	Closed-loop	Safety-Critical	Compositionality
TrafficSim [26]	✓	✓	✓	✗	✗
BITS [29]	✓	✓	✓	✗	✗
SimNet [27]	✓	✓	✓	✗	✗
STRIVE [28]	✓	✓	✓	✓	✗
CTG [3]	✓	✓	✓	✗	STL
SceneGen [10]	✓	✓	✗	✗	✗
RealGen [5]	✓	✓	✗	✓	✗
CTG++ [4]	✓	✓	✓	✓	LLM
LCTGen [13]	✓	✓	✗	✗	LLM
CausalAF [11]	✓	✓	✗	✓	CG
Ours	✓	✓	✓	✓	CG

B. Additional Algorithm Details

Algorithm 2 presents the training of *CCDiff* similar to DDPM and outputs denoising scene encoder $\pi_{\phi,\psi}(\cdot|\mathbf{s}, \mathbf{c}; G)$. Algorithm 3 presents the causal discovery and ranking algorithm

Algorithm 2 Training of *CCDiff*

Require: Dropout p_{uncond} , threshold C_ρ, C_{tte}
Require: Guidance loss $\{\mathcal{J}_i\}_{i=1}^N$, trajectories τ , map \mathbf{c} .
while $M \leq M_{\text{max}}$ **do**
 $M \leftarrow M + 1$
 $(\tau(0), \mathbf{c}) \sim \mathcal{D}$
 $\mathbf{G} \leftarrow \mathbf{G}(\tau(0))$ with probability $1 - p_{\text{uncond}}$
 $\mathbf{G} \leftarrow \mathbf{I}_N$ with probability p_{uncond}
 $k \sim \text{Unif}[K]$
 $\tau(k) = \sqrt{\bar{\alpha}_k} \tau(0) + \sqrt{1 - \bar{\alpha}_k} \epsilon$
 Update $\pi_{\phi,\psi}$ with $\nabla_{\phi,\psi} \|\pi_{\phi,\psi}(\tau(k), \mathbf{c}, k; \mathbf{G}) - \mathbf{a}(0)\|^2$
end while
return Denoising scene encoder $\pi_{\phi,\psi}(\cdot|\mathbf{s}, \mathbf{c}; G)$

Algorithm 3 Causal discovery and Ranking for *CCDiff*

Require: History trajectories τ , TTC Graph M , attention matrix α , Top-K agents k
 $G = M \cdot \alpha \triangleq (V, E, w)$
for all $v_i \in G$ **do**
 $C_i \leftarrow \{v_i\}, w_i = 0$
 for all $v_j \in V \setminus C_i$ **do**
 if $(v_j, v_i) \in E, \forall v \in C_i$ **then**
 $w_i \leftarrow w_i + \sum_{v \in C_i} w(v_j, v)$
 $C_i \leftarrow C_i \cup \{v_j\}$
 end if
 end for
end for
 $\rho \leftarrow \text{argsort}(C, w)[1:k]$
return Importance ranking ρ

C. Additional Experiment Details

C.1. Additional Quantitative Results

Table 5. Evaluation of Controllability and Realism across different scales of editable agents (N) and planning horizons (T). For each metric, we report the **best** and **second best** performance among all the methods. CCDiff has the best overall performance presented in the main text.

Methods	Metrics	$K=2$	3	4	5	10	Full	$T=1s$	2s	3s	4s	5s
SimNet	SCR (\uparrow)	0.31	0.32	0.33	0.36	0.42	0.47	0.35	0.37	0.37	0.37	0.40
	ORR (\downarrow)	1.76	2.19	2.62	2.67	2.90	3.17	2.09	3.87	6.16	8.36	9.93
	FDE (\downarrow)	3.76	4.34	4.98	5.26	6.63	8.03	4.11	3.78	4.90	4.83	3.83
	CFD (\downarrow)	2.56	2.95	2.86	3.16	5.00	7.00	4.02	5.03	5.51	5.57	8.04
TrafficSim	SCR (\uparrow)	0.38	0.36	0.44	0.41	0.41	0.47	0.53	0.53	0.58	0.55	0.53
	ORR (\downarrow)	2.09	2.25	2.45	2.48	2.66	2.73	3.56	6.36	8.98	10.96	12.21
	FDE (\downarrow)	4.25	5.06	5.77	6.23	6.79	7.13	8.32	6.48	8.61	8.66	7.32
	CFD (\downarrow)	7.76	9.53	10.64	10.99	10.96	11.57	5.00	10.06	7.89	10.38	9.90
STRIVE	SCR (\uparrow)	0.49	0.49	0.53	0.53	0.56	0.54	0.49	0.41	0.41	0.39	0.38
	ORR (\downarrow)	5.70	6.45	7.13	7.50	8.04	8.53	5.75	4.98	6.64	8.40	10.02
	FDE (\downarrow)	9.01	10.79	12.13	13.00	13.76	14.52	11.48	11.20	14.56	15.00	12.41
	CFD (\downarrow)	7.72	8.93	9.91	10.67	10.72	11.21	5.60	10.21	11.59	11.32	9.11
BITS	SCR (\uparrow)	0.38	0.38	0.37	0.39	0.37	0.41	0.37	0.34	0.44	0.39	0.41
	ORR (\downarrow)	0.53	0.51	0.56	0.63	0.56	0.60	1.44	3.68	5.63	7.56	9.39
	FDE (\downarrow)	3.20	3.95	4.42	4.67	5.05	5.35	4.68	4.69	6.10	6.36	5.44
	CFD (\downarrow)	7.43	8.32	9.15	9.42	9.46	10.23	8.79	10.35	10.75	11.30	11.65
CTG	SCR (\uparrow)	0.43	0.42	0.46	0.42	0.44	0.46	0.41	0.44	0.49	0.52	0.49
	ORR (\downarrow)	1.00	1.04	1.10	1.09	1.12	1.23	1.91	4.58	7.13	9.04	10.71
	FDE (\downarrow)	5.32	6.18	6.83	7.40	8.10	9.19	7.58	7.91	10.26	10.30	8.27
	CFD (\downarrow)	2.37	2.31	2.68	2.59	2.57	3.13	2.68	4.06	2.43	2.80	3.00
Ours	SCR (\uparrow)	0.40	0.44	0.43	0.46	0.49	0.51	0.40	0.44	0.49	0.55	0.52
	ORR (\downarrow)	0.61	0.72	0.99	1.02	1.80	2.05	2.92	4.52	7.10	9.35	10.51
	FDE (\downarrow)	4.17	5.22	5.99	6.59	7.84	8.26	7.06	5.54	6.86	7.00	5.71
	CFD (\downarrow)	1.88	1.92	1.93	2.25	2.83	3.47	2.37	4.08	4.25	4.97	6.33

Table 6. Ablation study on CCDiff’s variants. Evaluation of Controllability (CO, OR) and Realism (FDE and CFD) over different agent scales. For each metric, we highlight the **best** and the **second best** results. Causal ranking has the greatest impact to the final performance.

Enc.	Guide	Rank	Metrics	$K=2$	3	4	5	10	Full	$T=1s$	2s	3s	4s	5s
	✓	✓	SCR (\uparrow)	0.43	0.44	0.43	0.42	0.42	0.48	0.41	0.48	0.46	0.50	0.44
			ORR (\downarrow)	1.10	0.98	0.91	0.91	1.39	1.43	2.45	4.54	6.85	9.46	10.38
			FDE (\downarrow)	4.00	5.41	5.87	5.79	7.65	8.22	6.33	5.96	7.17	7.01	5.73
			CFD (\downarrow)	1.00	1.14	1.22	1.22	1.78	1.73	2.47	3.86	4.11	4.77	5.41
✓		✓	SCR (\uparrow)	0.38	0.45	0.40	0.40	0.39	0.40	0.41	0.44	0.46	0.48	0.48
			ORR (\downarrow)	0.81	0.76	1.00	1.06	1.47	1.60	2.78	4.83	7.39	9.40	10.44
			FDE (\downarrow)	4.33	5.28	6.13	6.82	8.65	9.20	7.03	6.14	8.02	6.99	5.55
			CFD (\downarrow)	1.81	1.60	1.84	1.94	2.92	2.62	2.87	3.64	4.27	5.37	6.46
✓	✓	Dist	SCR (\uparrow)	0.33	0.34	0.36	0.37	0.39	0.36	0.34	0.35	0.41	0.41	0.39
			ORR (\downarrow)	1.38	1.50	1.59	1.49	1.56	1.74	3.06	5.21	7.41	10.14	10.43
			FDE (\downarrow)	4.15	5.15	5.79	5.96	8.01	9.69	6.51	5.73	6.82	7.01	5.38
			CFD (\downarrow)	1.79	2.44	2.03	2.34	3.09	3.30	1.94	2.92	3.88	4.44	5.95
✓	✓	Human	SCR (\uparrow)	0.34	0.35	0.33	0.31	0.33	0.33	0.33	0.34	0.40	0.37	0.40
			ORR (\downarrow)	1.66	1.65	1.73	1.93	1.66	1.75	3.10	5.25	7.44	10.37	10.51
			FDE (\downarrow)	5.80	6.74	7.40	7.84	8.63	8.99	8.12	7.25	8.70	9.16	7.01
			CFD (\downarrow)	2.21	2.51	2.83	3.14	2.60	2.96	3.39	5.20	6.17	6.65	8.43
✓	✓	✓	SCR (\uparrow)	0.40	0.44	0.43	0.46	0.49	0.51	0.40	0.44	0.49	0.55	0.52
			ORR (\downarrow)	0.61	0.72	0.99	1.02	1.80	2.05	2.92	4.52	7.10	9.35	10.51
			FDE (\downarrow)	4.17	5.22	5.99	6.59	7.84	8.26	7.06	5.54	6.86	7.00	5.71
			CFD (\downarrow)	1.88	1.92	1.93	2.25	2.83	3.47	2.37	4.08	4.25	4.97	6.33

C.2. Additional Qualitative Results

C.2.1 Long-horizon Generation

We evaluate the long-horizon generation with different planning cycle for the scenarios with same length between *CCDiff* and all the baselines. We illustrate the qualitative examples below. The results demonstrate that *CCDiff* can consistently generate realistic cross-traffic violation scenarios for $1s \leq T \leq 5s$. In contrast, CTG baseline can only generate an opposite-lane collision when $T = 1s$.




Figure 5. Comparison of *CCDiff* and CTG on the controllability and realism under different sizes of controllable agents. We can see that *CCDiff* can consistently generate realistic cross-traffic violation scenarios, yet CTG can only generate one with shorter planning cycle in 1s.

C.2.2 Multi-agent Generation

We evaluate the multi-agent generation with different sizes of controllable agents K . We illustrate the qualitative examples of unprotected left turn scenarios below. The results demonstrate that with abundant controllable access to the agents at the scene ($K \geq 2$ in this case), *CCDiff* can consistently generate realistic unprotected left-turn scenarios compared to the CTG baseline.




Figure 6. Comparison of *CCDiff* and CTG on the controllability and realism under different sizes of controllable agents. We can see that when the number of controllable agents is greater than 1, *CCDiff* can consistently generate realistic unprotected left-turn violations, yet CTG can only generate one unrealistic right turn collision with 5 controllable agents.

C.3. Detailed description of baselines

SimNet [27]: SimNet frames the problem as a Markov Process, and models state distributions and transitions directly from raw observational data, eliminating the need for handcrafted models. Trained on 1,000 hours of driving logs, it dynamically generates novel and adaptive scenarios that enable closed-loop evaluations. The system reveals subtle issues, such as causal confusion, in state-of-the-art planning models that traditional non-reactive simulations fail to detect.

TrafficSim [26]: TrafficSim uses an implicit latent variable model like conditional variational autoencoder (CVAE). The system parameterizes a joint actor policy that simultaneously generates plans for the agents in a scene. The model is jointly trained with (i) ELBO objective inspired by CVAE and (ii) common-sense following with agents’ pair-wise collision loss. TrafficSim generates diverse, realistic traffic scenarios and can serve as effective data augmentation for improving autonomous motion planners.

STRIVE [28]: STRIVE employs a graph-based conditional variational autoencoder (VAE) to model realistic traffic motions and formulates scenario generation as an optimization problem in the latent space of this model. By perturbing real-world traffic data, STRIVE generates scenarios that stress-test planners. A subsequent optimization step ensures that the scenarios are useful for improving planner performance by being solvable and challenging. STRIVE has been successfully applied to attack two planners, showing its ability to produce diverse, accident-prone scenarios and improve planner robustness through hyperparameter tuning.

BITS [29]: BITS (Bi-level Imitation for Traffic Simulation) framework leverages the hierarchical structure of driving behaviors by decoupling the simulation into two levels: high-level intent inference and low-level driving behavior imitation. This structure enhances sample efficiency, behavior diversity, and long-horizon stability. BITS also integrates a planning module to ensure consistency over extended scenarios.

CTG [3]: CTG is a novel framework combining controllability and realism in traffic simulation by leveraging conditional diffusion models and Signal Temporal Logic (STL). The approach allows fine-grained control over trajectory properties, such as speed and goal-reaching, while maintaining realism and physical feasibility through enforced dynamics. Extending to multi-agent settings, the model incorporates interaction-based rules, such as collision avoidance, to simulate realistic agent interactions in traffic.

We list implementation details of all the methods are listed below with important hyperparameters and model structures information in Table 7.

Table 7. Hyper-parameters of models used in experiments of CCDiff and baselines

Parameter Name	Value	Parameter Name	Value
Step length	0.1s	Map Encoder	ResNet-18
History steps	31	Map feature dim.	256
Generation steps	52	Trajectory Encoder	MLP
Learning rate	0.0001	Trajectory feature dim.	128
Optimizer	Adam	Transformer decoder head	16
Batch size	100	Transformer decoder layers	2
Trajectory prediction loss weight	1.0	Guidance gradient Steps	30
Yaw regularization weight	0.1	Guidance constraint norm	100
EMA step	1	Guidance learning rate	0.001
EMA decay	0.995	Guidance optimizer	Adam
Denoising Steps	100	Guidance weight: off-road	1.0
Guidance discount factor	0.99	Guidance weight: collision	-50.0
Planning steps	10, 20, 30, 40, 50	TTC threshold	3.0 s
Controllable Agents	1, 2, 3, 4, 5, 10, Full	Distance threshold	50 m

Training and Inference Resources We conduct training and inference of all the models on 4x NVIDIA Tesla V100 with 16GB GPU memory each, and 48-core CPU Intel(R) Xeon(R) CPU @ 2.30GHz. The training of one model takes 3 hours per epoch on nuScenes training split, and we train 10 epochs for each baseline model and CCDiff. At inference time, the parallel evaluation takes an average of 3 minutes on each closed-loop testing scenario for all the methods under the same configuration (controllable agents and generation frequencies).

C.4. Detailed description of evaluation metrics

- **Controllability Score (CS)**: The computation of CS standardizes the scenario-wise collision rate (SCR) used in [13, 26]:

$$CS = \frac{SCR - \min(SCR)}{\max(SCR) - \min(SCR)}$$

We then standardize SCR among all the methods to get the **CS**, a higher-the-better score between 0 and 1.

- **Realism Score (RS)**: We average over three widely-used quantitative metrics to evaluate the realism of the scenarios: (i) scenario off-road rate (ORR) used in [5, 13], (ii) final displacement error (FDE, m) and (iii) comfort distance (CFD) in [3, 29] to quantify the realism of the similarity in the smoothness of agents' trajectories in the generated scenarios. We standardize all the metrics among all the methods respectively and average them to get the **RS**, a higher-the-better score between 0 and 1:

$$RS = 1.0 - \frac{1}{3} \left(\frac{ORR - \min(ORR)}{\max(ORR) - \min(ORR)} + \frac{FDE - \min(FDE)}{\max(FDE) - \min(FDE)} + \frac{CFD - \min(CFD)}{\max(CFD) - \min(CFD)} \right)$$

Specifically, FDE describes the trajectory closeness between the synthetic one and the original one, ORR describes how frequently the generated trajectories go off-road, while CFD measures the **smoothness** of the generated trajectories with their acceleration and jerk. All these raw metrics are lower the better, so after we revert it above, the resulting RS is a higher-the-better metric.

- **Multi-objective optimization metrics**: with the **RS** and **CS**, we further quantify the optimality of the solution based on generational distance (**GD**) and inverted generational distance (**IGD**), the average minimum distance between the methods and Pareto frontier [44, 53]:

$$GD = \left(\frac{1}{|\mathcal{D}|} \sum_{\mathbf{d} \in \mathcal{D}} \min_{\mathbf{p} \in \mathcal{P}} \|\mathbf{a} - \mathbf{p}\|^q \right)^{\frac{1}{q}},$$

where $\|\cdot\|$ denotes the Euclidean distance, and q is typically set to 2. Conversely, IGD measures the average distance from each solution in the Pareto frontier \mathcal{P} to its nearest solution in the obtained set \mathcal{D} , and is defined as

$$IGD = \left(\frac{1}{|\mathcal{P}|} \sum_{\mathbf{p} \in \mathcal{P}} \min_{\mathbf{d} \in \mathcal{D}} \|\mathbf{p} - \mathbf{d}\|^q \right)^{\frac{1}{q}}.$$

Both metrics provide insights into the convergence and diversity of the obtained solution set: lower values of GD indicate better convergence to the Pareto frontier. On the other hand, lower values of IGD suggest better coverage over the Pareto frontier. We visualize an example for GD and IGD in Figure 7.




Figure 7. Examples of GD and IGD used to evaluate the multi-objective optimization. Two axes f_1, f_2 represent two objectives.

Quantitative Analysis on the design Causal Masking Design We also analyze the importance of different features w.r.t. the collision samples in the generated scenarios. The results show that TTC feature has the highest statistical correlation with the controllability score (i.e. the collision rate) in our setting.




Figure 8. The number of cliques in the TTC graph is more informative causal features of safety-critical incidents (higher Pearson correlation) compared to Relative Distance and number of agents.

Table 8. Correlation analysis between the collision accidents and different causal structure features: standardized clique score for TTC graph, standardized clique score for distance graph, and the standardized number of agents at the scene. We list the Pearson correlation R^2 between the standardized controllability score for each scenario, as well as the significance level of each feature (p-value)

Causal Structure Feature	$R^2(\uparrow)$	p-value (\downarrow)
#Cliques in Dist. graph	0.01	0.89
#Agents	0.13	0.20
#Cliques in TTC graph (Ours)	0.49	2.2×10^{-7}

C.5. Additional Qualitative Analysis over Scenarios

In the following subsection, we present seven representative interactive scenarios that are safety-critical in urban traffic. We begin by analyzing the comparisons with baseline methods and highlighting the differences between distance-based graphs and TTC-based graphs. The results demonstrate that TTC-based graphs are generally sparser yet more informative, particularly for capturing safety-critical maneuvers.

Additionally, we provide examples of multi-agent, long-horizon trajectory generation for individual scenarios, showcasing the model’s ability to handle complex interactions over extended time frames.

C.5.1 Unprotected Left Turn

Baseline Comparison Below, we present the unprotected left-turn scenarios. The relational reasoning of the distance-based graph fails to capture the interaction between the two involved vehicles (11 and 14). We omit the multi-agent and long-horizon generation examples for this scenario, as these have already been analyzed in previous comparisons.

Among all the baselines, CTG, SimNet, and BITS closely follow the ground-truth trajectories, successfully generating a left-lane right turn without producing collision samples. In contrast, STRIVE generates unrealistic collisions with parked vehicles in the side lane. Notably, only CCDiff manages to produce realistic unprotected left-turn behaviors. Only the TTC mask captures the interaction between agents 11 and 14.




Figure 9. Qualitative of *CCDiff* and baselines in unprotected left turn scenarios.

C.5.2 Cross Traffic Violation

Baseline Comparison A cross-traffic violation occurs when a vehicle at a T-intersection fails to yield the right of way to a vehicle approaching from a perpendicular direction. Such violations often result in side-impact collisions, particularly when the violating driver misjudges the speed or distance of the cross-traffic vehicle. In *CCDiff*, agent 0 collides with agent 6, illustrating this scenario.

Among the baselines, BITS, TrafficSim, and CTG successfully avoid generating collision samples. However, SimNet also generates a collision between agent 0 and agent 6, failing to model the scenario accurately. Both TTC and distance mask manage to capture the interaction between agents 0 and 6.




Figure 10. Qualitative of *CCDiff* and baselines in cross traffic violation scenarios.

Multi-agent Generation We compare the multi-agent generation results of *CCDiff* with CTG. *CCDiff* can consistently generate the cross traffic violation when the controllable agents $K \geq 2$.




Figure 11. Qualitative comparison of *CCDiff* and CTG under cross traffic violation generation under different sizes of controllable agents.

Long-horizon Generation We compare the long-horizon generation results of *CCDiff* with CTG. *CCDiff* can consistently generate the cross traffic violation even with a generation horizon $T > 2s$, yet CTG generated scenarios are more conservative.




Figure 12. Qualitative comparison of *CCDiff* and CTG under cross traffic violation generation under different generation horizons.

C.5.3 Lane Cut-in

Baseline Comparison A lane cut-in at an intersection occurs when a vehicle abruptly changes lanes or merges into another lane while navigating through or approaching an intersection, often without sufficient clearance or signaling. This maneuver typically forces other vehicles in the affected lane to brake suddenly or adjust their trajectory, increasing the risk of collisions or near-misses. In our case, agent 3 will suddenly cut in from the left lane to the right lane and collide with agent 0.

Among all the baselines, CTG and SimNet generate some irregular behaviors and drive some of the controllable agents off-road. STRIVE generates relatively unrealistic right turn collision, and TrafficSim generates a wild unprotected left turn that is more unrealistic under this context. The TTC mask manages to capture the interaction between agents 0 and 3, while the distance mask misses it.

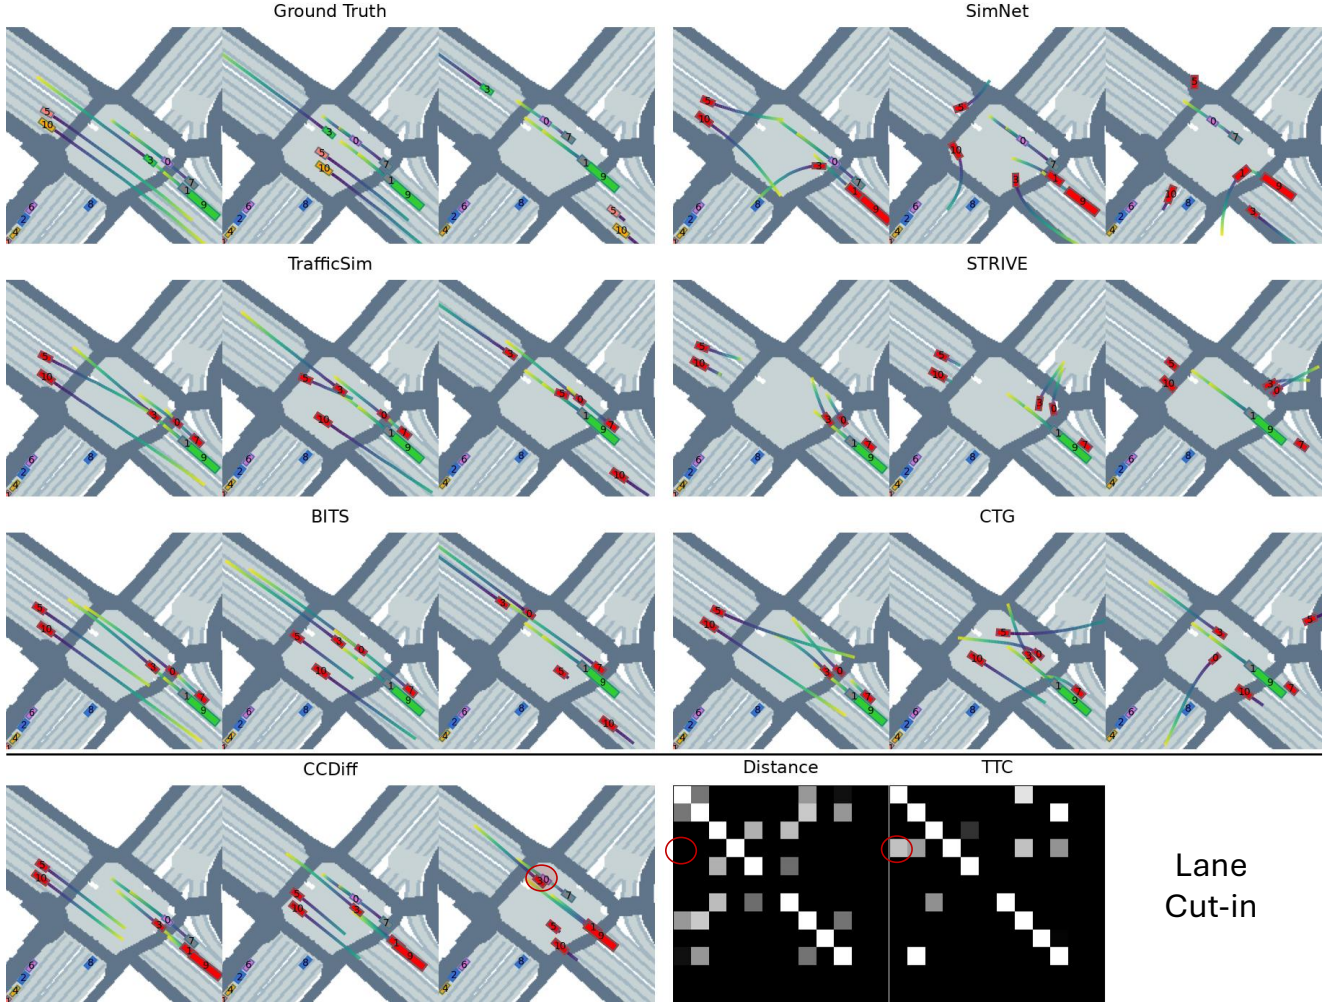


Figure 13. Qualitative of *CCDiff* and baselines in lane cut-in scenarios.

Multi-agent Generation We compare the multi-agent generation results of *CCDiff* with CTG. *CCDiff* can generate collision samples when $K = 5$, yet the CTG generates very wild behaviors that are unrealistic from the ground-truth trajectories.

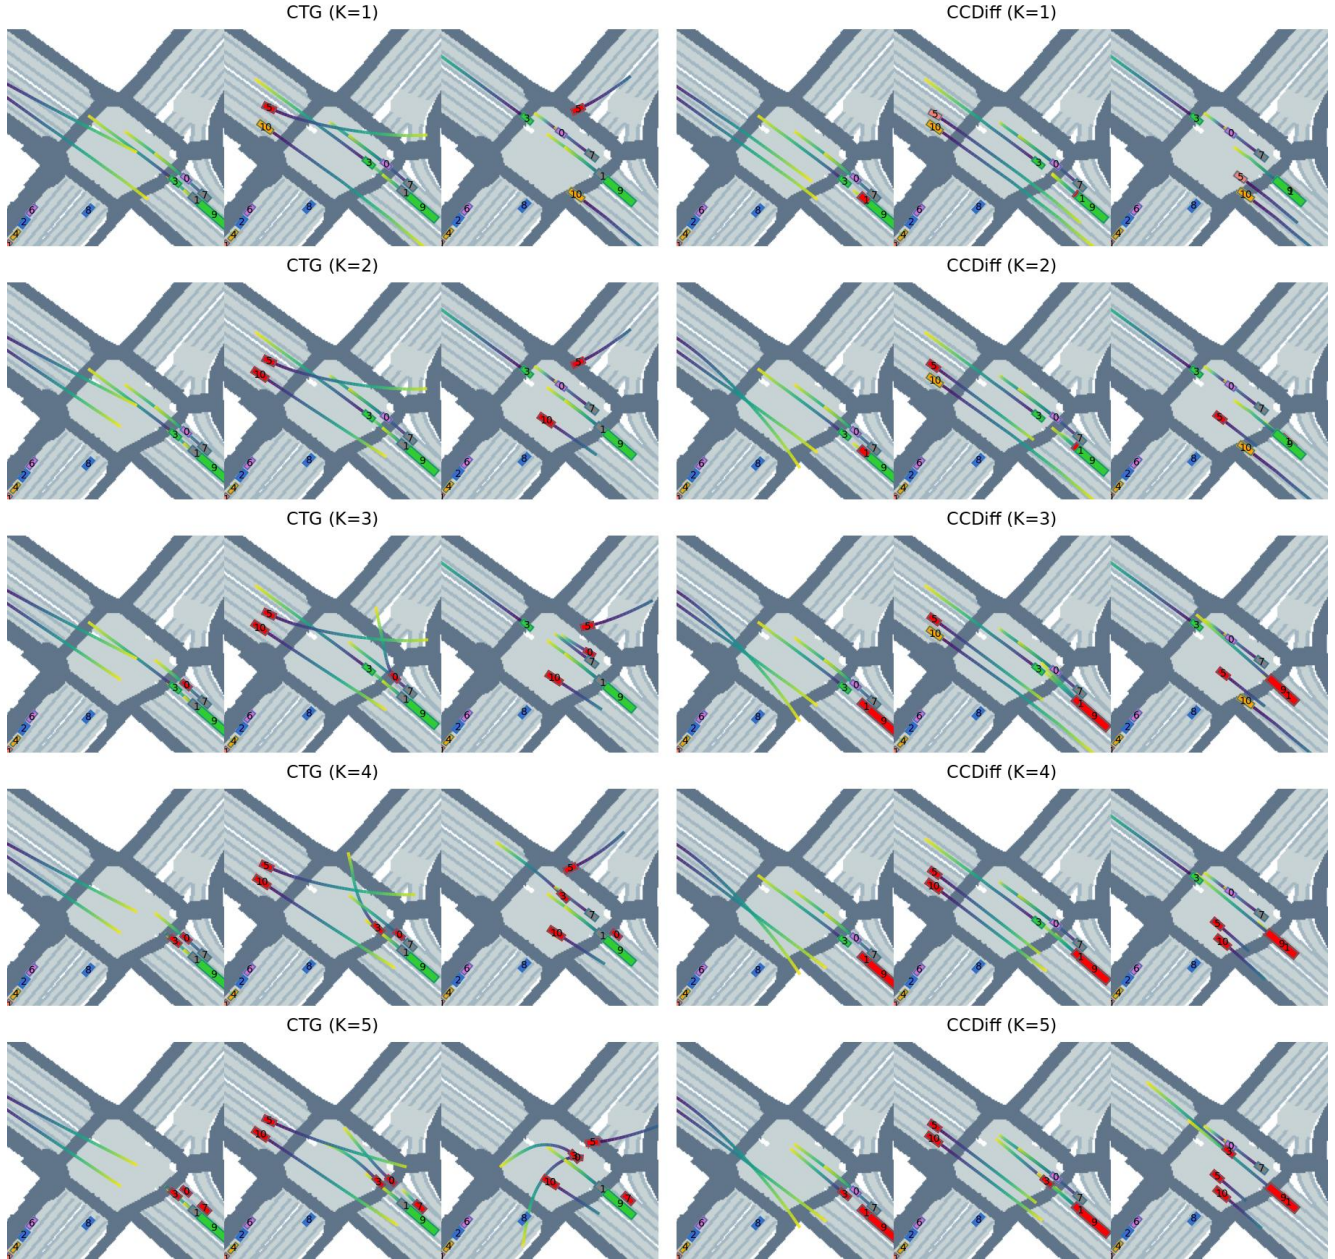


Figure 14. Qualitative comparison of *CCDiff* and CTG under cross traffic violation generation under different sizes of controllable agents.

Long-horizon Generation We compare the long-horizon generation results of *CCDiff* with CTG. *CCDiff* can consistently generate the cut-in violation scenarios with the generation horizon $1s \leq T \leq 4s$. In contrast, CTG attempts to generate some unprotected left turn in this context but fails.




Figure 15. Qualitative comparison of *CCDiff* and CTG under cross traffic violation generation under different generation horizons.

C.5.4 Emergency Break

Baseline Comparison The emergency break occurs when the middle vehicle (agent 0) brakes to keep distance from the forward vehicle (agent 9) suddenly, causing the trailing vehicle (agent 8) to collide with it due to insufficient stopping distance.

Among all the baselines, STRIVE generates some irregular behaviors, which drive some of the controllable agents off-road. TrafficSim, BITS, and SimNet fail to generate safety-critical samples. Notably, although CTG also generates some collision samples, it accelerates the trailing vehicle 8 to collide with the middle vehicle 0, which does not break yet. In comparison, in our case, the middle vehicle 0 breaks and causes a collision with trailing vehicle 8 at normal speed, which is more realistic. Both the TTC mask and distance mask capture the interaction among agents 0, 8, and 9 in this scenario.




Figure 16. Qualitative of *CCDiff* and baselines in the emergency break scenarios.

Multi-agent Generation We compare the multi-agent generation results of *CCDiff* with CTG. *CCDiff* can consistently generate safety-critical emergency breaking samples when $K \geq 2$, with a control of the most important vehicle 8 in this context. In contrast, CTG keeps accelerating the rear vehicle 8 instead of slowing down the middle vehicle 0.

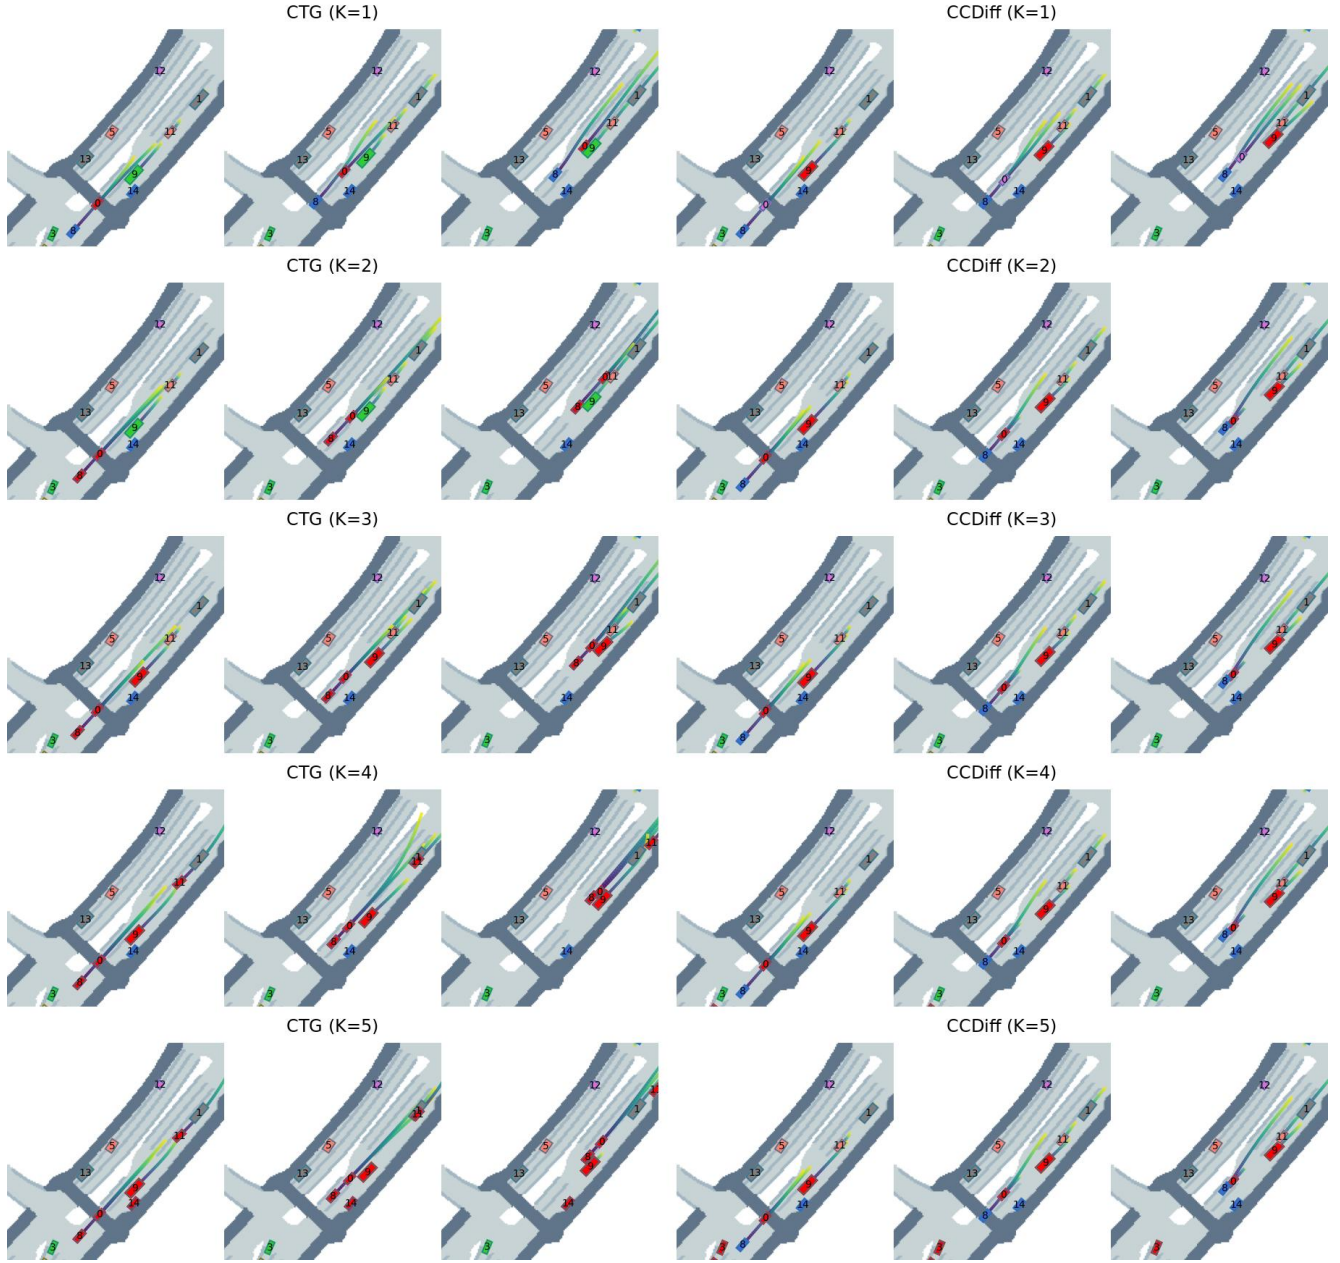


Figure 17. Qualitative comparison of *CCDiff* and CTG under cross traffic violation generation under different sizes of controllable agents.

Long-horizon Generation We compare the long-horizon generation results of *CCDiff* with CTG. *CCDiff* can consistently generate the cut-in violation scenarios with all different lengths of the generation horizon $1s \leq T \leq 5s$. In contrast, CTG attempts to accelerate the vehicle in the middle and cannot generate any near-miss samples with longer generation horizons.




Figure 18. Qualitative comparison of *CCDiff* and CTG under cross traffic violation generation under different generation horizons.

C.5.5 Chain-reaction Crash

Baseline Comparison A chain-reaction crash involving five vehicles (agents 1, 2, 5, 7, 8) occurs when a sudden stop or collision causes a cascade of impacts among closely spaced vehicles in the same lane. This happens before an intersection when vehicles fail to maintain a safe following distance, leading to multiple rear-end collisions.

Among all the baselines, SimNet and BITS fail to generate safety-critical scenarios. TrafficSim, STRIVE, and CTG generate collisions between agent 0 on the side lane with agent 2 with a very unrealistic cut-in behavior. In comparison, CCDiff generates realistic collisions where the trailing vehicles 1, 7, and 8 fail to break timely and collide with static front vehicle 5, waiting for the right turn of 2. Both TTC graph and distance graph captures the interaction of 5 and 7, 8. Yet distance-based graphs fail to capture the indirect interaction between 2 and 7, 8.

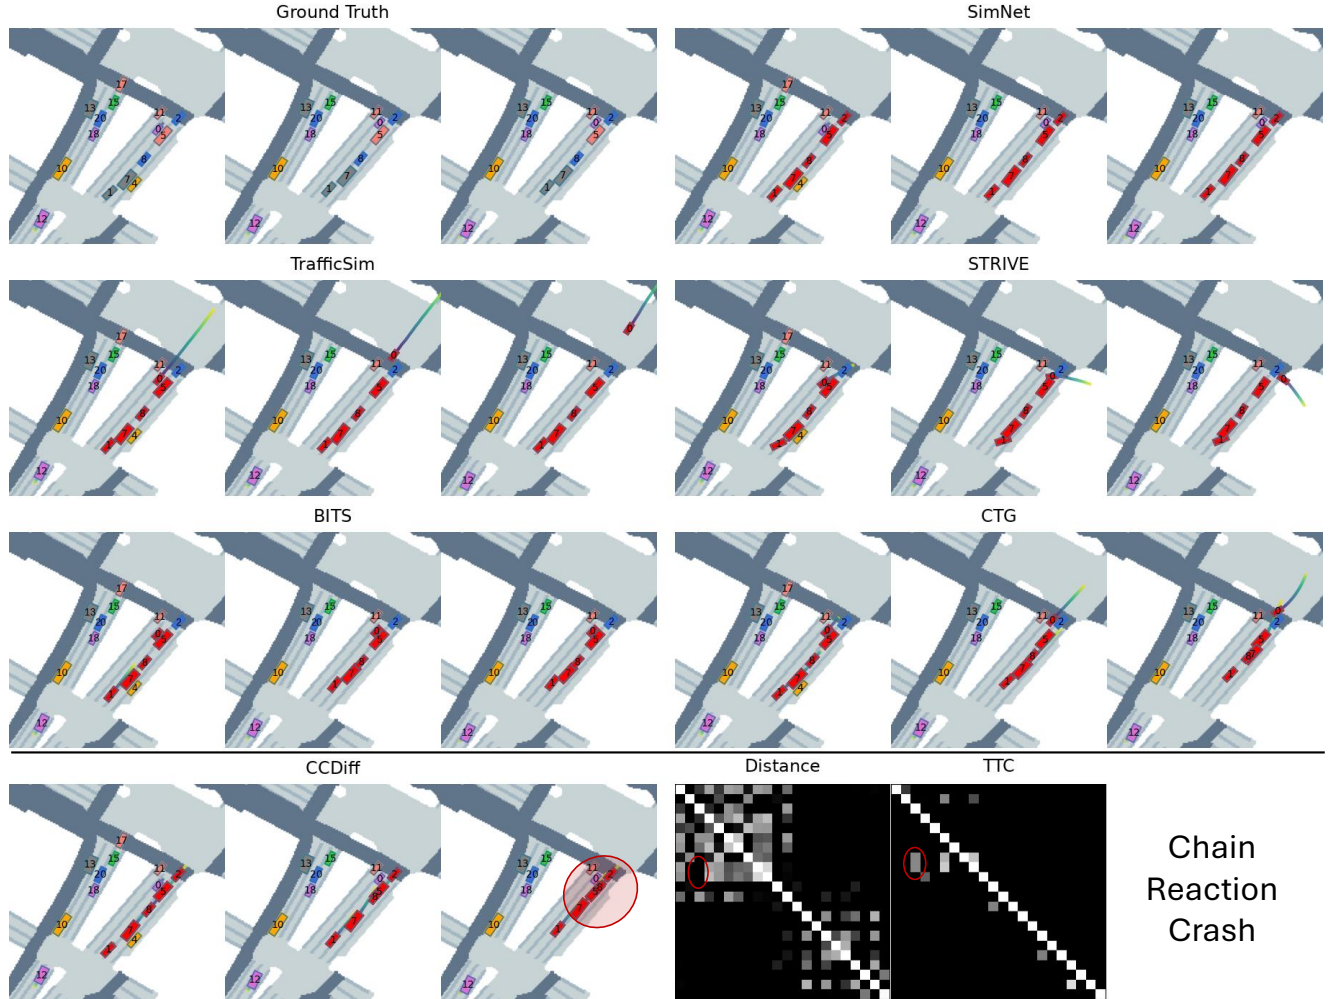


Figure 19. Qualitative of *CCDiff* and baselines in the chain-reaction crash scenarios.

Multi-agent Generation We compare the multi-agent generation results of *CCDiff* with CTG. *CCDiff* can consistently generate safety-critical emergency breaking samples when $K \geq 3$, with a control of the most important vehicle 7, 8 in this context. In contrast, CTG keep accelerating the side-lane vehicle 0 or rear vehicle 1 in a very unrealistic way.




Figure 20. Qualitative comparison of *CCDiff* and CTG under cross traffic violation generation under different sizes of controllable agents.

Long-horizon Generation We compare the long-horizon generation results of *CCDiff* with CTG. *CCDiff* can consistently generate the cut-in violation scenarios with all different lengths of the generation horizon $1s \leq T \leq 5s$. In contrast, the trajectories generated by CTG seem to diverge by a great deal when $T \geq 2s$.




Figure 21. Qualitative comparison of *CCDiff* and CTG under cross traffic violation generation under different generation horizons.

C.5.6 Adjacent Left-turn Side-wipe

Baseline Comparison An adjacent left-turn sideswipe occurs when two vehicles (agent 1, 11) in neighboring left-turn lanes collide as Agent 1 veers into Agent 11’s path.

Among all the baselines, STRIVE and CTG generate the motions of 1 and 11 to the straight lane reverse lane. TrafficSim generates the motions of 1 and 11 to the straight lane. BITS generally follows the original history scenarios with a rear collision between agents 18 and 11. CCDiff drifts 1 a little bit and let it veer into the agent 11’s path.

Both the Distance graph and the TTC graph could detect the close interaction between agents 1 and 11 in this case.




Figure 22. Qualitative of *CCDiff* and baselines in the adjacent left-turn side-wipe scenario.

Multi-agent Generation We compare the multi-agent generation results of *CCDiff* with CTG. *CCDiff* consistently generates safety-critical emergency braking scenarios when $K \geq 3$, effectively controlling the behavior of the most critical vehicle, agent 1, in this context. In contrast, CTG fails to accurately model the scenario, allowing agent 11 to continue in the wrong direction and being unable to generate collision samples, even when more agents are controllable.




Figure 23. Qualitative comparison of *CCDiff* and CTG under cross traffic violation generation under different sizes of controllable agents.

Long-horizon Generation We compare the long-horizon generation results of *CCDiff* with CTG. *CCDiff* can consistently generate the left-turn side-wipe scenarios, while CTG diverges and fails to generate collision samples at $T = 3s, 4s$.




Figure 24. Qualitative comparison of *CCDiff* and CTG under cross traffic violation generation under different generation horizons.

C.5.7 Multi-vehicle Merge-in

Baseline Comparison Multi-vehicle merge-in occurs when a vehicle from a side lane (agent 13) attempts to merge into a single-lane traffic flow (agents 6, 2, 29), causing disruptions or collisions involving three vehicles 2 and 29.

Among all the baselines, SimNet does not generate collision samples, TrafficSim and CTG generate collision between 13 and 2 and manipulates the trajectory of 13 in an abrupt way. Our scenario just slows down agents 6 and 2 with an expectation of merge-in from agent 13, which causes the trailing agent 29 collides to agent 2. The generated final scenario of CCDiff have the closest layout with the ground-truth trajectories compared to other baselines.

TTC mask in this case is more sparse with necessary information (agent 2 and 29) compared to the distance mask.




Figure 25. Qualitative of *CCDiff* and baselines in the multi-vehicle lane merge-in scenarios.

Multi-agent Generation We compare the multi-agent generation results of *CCDiff* with CTG. *CCDiff* can consistently generate safety-critical emergency breaking samples when $K \geq 4$, with a control of the most important vehicle 2, 6 in this context. In contrast, CTG keeps accelerating the side-lane vehicle 13 without generating any meaningful near-miss samples.

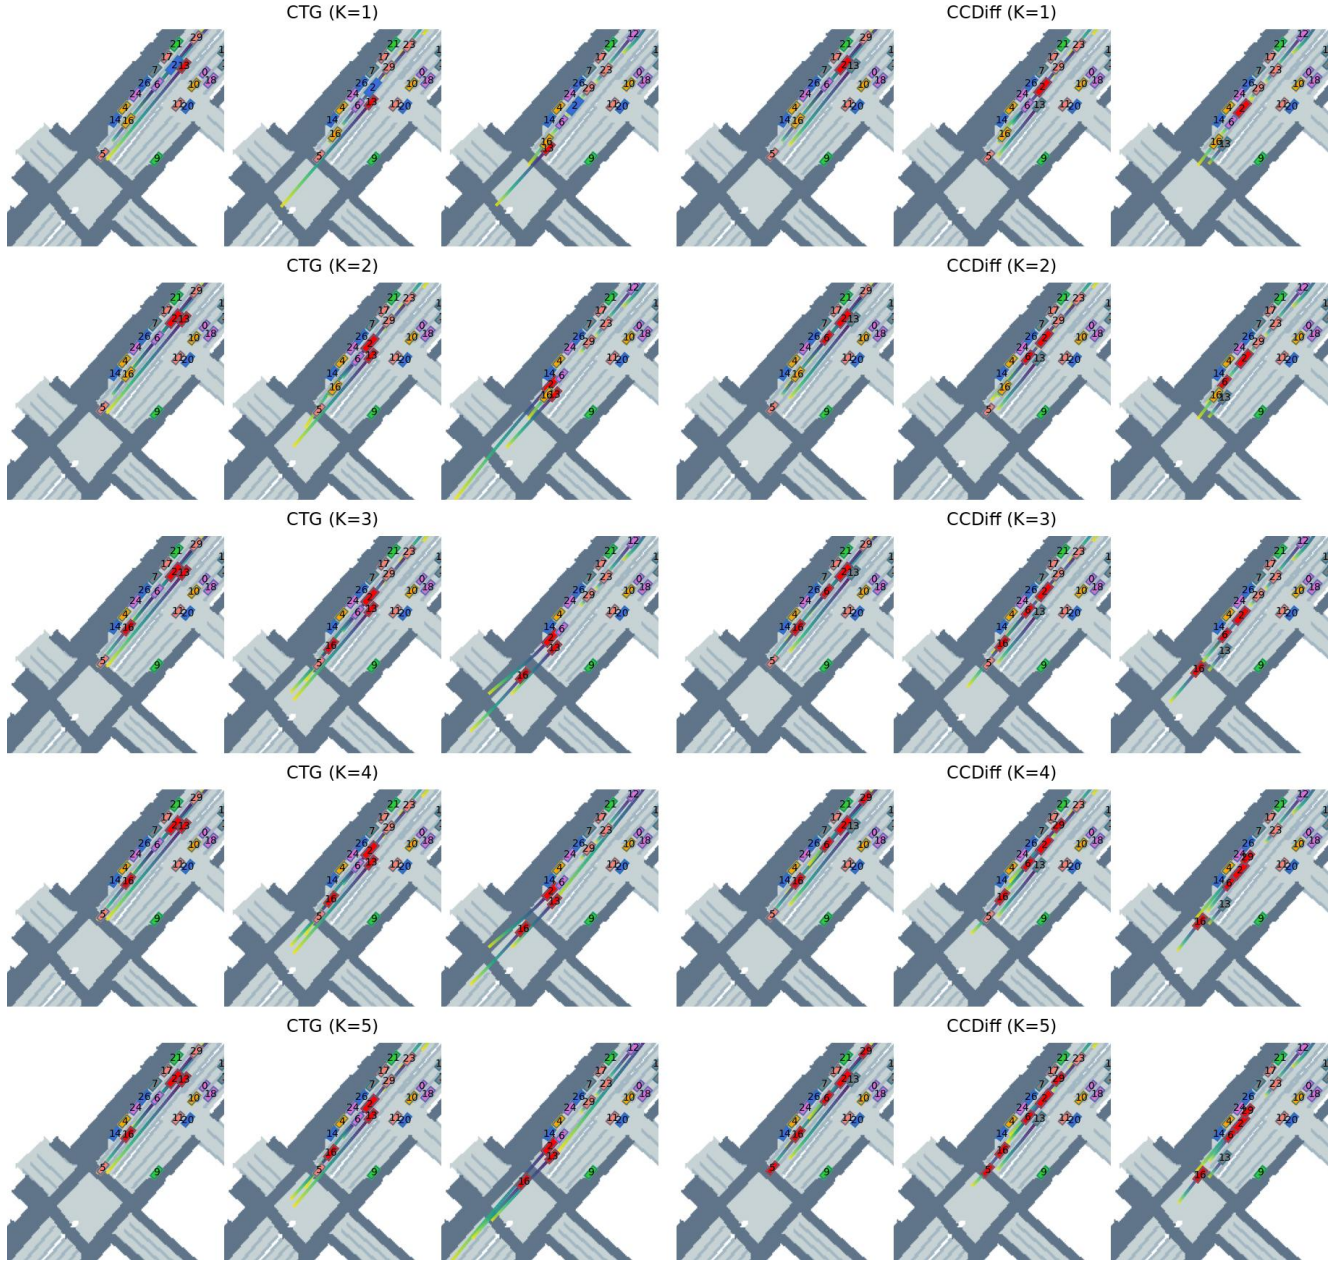


Figure 26. Qualitative comparison of *CCDiff* and CTG under cross traffic violation generation under different sizes of controllable agents.

Long-horizon Generation We compare the long-horizon generation results of *CCDiff* with CTG. *CCDiff* can consistently generate the multi-vehicle merge-in collision scenarios with all different lengths of the generation horizon $1s \leq T \leq 5s$. In contrast, CTG generates some cut-in collisions between 13 and 6 when $T \geq 2$, which is more unrealistic given the ground-truth layouts.




Figure 27. Qualitative comparison of *CCDiff* and CTG under cross traffic violation generation under different generation horizons.