

SAFE-SIM: Safety-Critical Closed-Loop Traffic Simulation with Diffusion-Controllable Adversaries

Wei-Jer Chang¹, Francesco Pittaluga², Masayoshi Tomizuka¹, Wei Zhan¹, and Manmohan Chandraker^{2,3}

¹ UC Berkeley ² NEC Labs America ³ UC San Diego

Abstract. Evaluating the performance of autonomous vehicle planning algorithms necessitates simulating long-tail safety-critical traffic scenarios. However, traditional methods for generating such scenarios often fall short in terms of controllability and realism; they also neglect the dynamics of agent interactions. To address these limitations, we introduce SAFE-SIM, a novel diffusion-based controllable closed-loop safety-critical simulation framework. Our approach yields two distinct advantages: 1) generating realistic long-tail safety-critical scenarios that closely reflect real-world conditions, and 2) providing controllable adversarial behavior for more comprehensive and interactive evaluations. We develop a novel approach to simulate safety-critical scenarios through an adversarial term in the denoising process of diffusion models, which allows an adversarial agent to challenge a planner with plausible maneuvers while all agents in the scene exhibit reactive and realistic behaviors. Furthermore, we propose novel guidance objectives and a partial diffusion process that enables users to control key aspects of the scenarios, such as the collision type and aggressiveness of the adversarial agent, while maintaining the realism of the behavior. We validate our framework empirically using the nuScenes and nuPlan datasets across multiple planners, demonstrating improvements in both realism and controllability. These findings affirm that diffusion models provide a robust and versatile foundation for safety-critical, interactive traffic simulation, extending their utility across the broader autonomous driving landscape. Project website: <https://safe-sim.github.io/>.

1 Introduction

A key safety feature of autonomous vehicles (AVs) is their ability to navigate near-collision events in real-world scenarios. However, these events rarely occur on roads and testing AVs in such high-risk situations on public roads is unsafe. Therefore, simulation is indispensable in the development and assessment of AVs, providing a safe and reliable means to study their safety and dependability. A critical aspect of simulation is modeling the behavior of other road users, since AVs must learn to interact with them safely.

A common method of safety-critical testing of AVs involves manually designing scenarios that could potentially lead to failures, such as collisions. While this approach allows for targeted testing, it is inherently limited in scalability and

Method	Safety-Critical	Controllable	Controllable Adversary	Evaluate Planner	Closed-Loop	Real-World
CTG [39]	×	✓	×	×	✓	✓
CTG++ [38]	×	✓	×	×	✓	✓
STRIVE [26]	✓	×	×	✓	✓	✓
DiffScene [34]	✓	✓	×	✓	×	×
SAFE-SIM (Ours)	✓	✓	✓	✓	✓	✓

Table 1: Comparison of methods. Our contribution is the development of a framework for (a) safety-critical (b) closed-loop (c) controllable adversarial simulations. These aspects are not concurrently present in previous frameworks. We formulate a novel partial diffusion with novel guidance functions for stable long-term simulation and are the first to enable an ego planner to be tested against controllable adversaries with varied behavior patterns.

lacks the comprehensiveness required for thorough evaluation [7, 8, 33]. Some recent works focus on automatically generating challenging scenarios that cause planners to fail, but their emphasis has been mostly on static scenario generation rather than dynamic, *closed-loop* simulations. This results in a critical gap: the behavior of other agents often does not adapt or respond to the planner’s actions, which is essential for a comprehensive safety evaluation. Furthermore, the results from these simulations often lack *controllability*, typically producing only a single adversarial outcome per scenario without the flexibility to explore a range of conditions and responses.

In this work, we introduce SAFE-SIM, a *closed-loop* simulation framework for generating safety-critical scenarios, with a particular emphasis on *controllability* and *realism* for the behavior of agents, which allows simulations over a long-horizon as needed to evaluate AV planning algorithms (Fig. 1). Different from prior works [26, 34, 38, 39] that primarily adhere to rule-constraint satisfaction, our approach enhances *controllability* by modulating adversarial vehicle behaviors within identical scenarios, thereby facilitating a broader exploration of potential outcomes. See Tab. 1 for a comprehensive comparison of these approaches.

Our approach builds upon recent developments in controllable diffusion models [18, 25, 39]. Specifically, we adopt a test-time guidance to direct the denoising phase of the diffusion process, using the gradients from differentiable objectives to enhance scenario generation, enabling generation of adversarial scenarios in which and adversarial agent collides with the ego agent behaving according to specific planning policy. Additionally, we develop an novel approach, which we refer to as Partial Diffusion that introduces trajectory proposals into the diffusion process to provide a high degree of controllability over the type of collision scenario. Overall, our balanced integration of adversarial objectives with regularization during the guidance phase combined with Partial Diffusion allows for refined control over the conditions of the generated scenarios, ensuring both their realism and relevance to safety-critical testing.

In our study, we use the nuScenes [2] and nuPlan [3] datasets to evaluate the efficacy of our method in generating safety-critical closed-loop simulations. Our results demonstrate a marked improvement in the controllability and real-

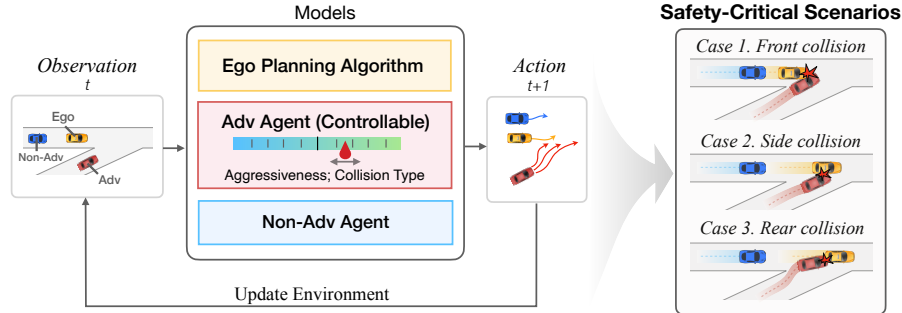


Fig. 1: Overview of SAFE-SIM Framework for Controllable Safety-Critical Closed-Loop Simulation. This framework evaluates a planner within scenarios featuring multiple controllable reactive agents. These agents have two distinct roles: adversarial agents, which actively challenge the planner by exhibiting controllable adversarial behaviors such as specific collision types and levels of aggressiveness, and non-adversarial agents, which follow normal driving behavior to maintain the realism of the entire scene. Such a setup facilitates the generation of various realistic, interactive, and safety-critical scenarios, providing a thorough evaluation of the planner’s capabilities.

ism of scenarios compared to previous adversarial scenario generation methods. Furthermore, we showcase the advantage of our proposed framework in varying the safety-criticality and collision types of scenarios. These attributes make our approach particularly well-suited for the closed-loop simulation of AVs, providing a more reliable and comprehensive framework for safety evaluation.

2 Related Work

2.1 Traffic Simulation

Traffic simulation can broadly be categorized into two main groups: heuristic-based and learning-based methods. In heuristic-based methods, agents are controlled by human-specified rules, such as the Intelligent Driver Model (IDM) [30] to follow a leading vehicle while maintaining a safe following distance. However, these methods have modeling capacity issues and may not reflect the real traffic distribution; a large domain gap limits its usage for planning evaluation. To close the gap, data-driven approaches learn from real driving datasets to imitate real-world behavior [23, 28, 29, 35]. TrafficSim [28] utilizes a trained variational autoencoder for scene-level traffic simulation, while BITS [35] combines high-level goal inference with low-level driving behavior imitation to enhance the realism of simulated driving behavior. Recently, the Waymo SimAgents challenge has focused on whether simulators can accurately represent real-world driving distributions [20]. However, little work has specifically focused on behavior simulation for generating safety-critical, long-tail scenarios.

Diffusion models [27, 31, 32] have shown significant promise for synthetic image generation. One of the key advantages of diffusion models is controllability,

which can take the form of classifier [6], classifier-free [15], and reconstruction [13] guidance. Recently, controllable diffusion models have been employed for planning and traffic simulation [25, 38, 39] via guidance. We adopt trajectory diffusion models to develop a novel approach for generate safety-critical realistic traffic simulations in which an adversarial agent collides with an ego planner agent. Other works [19, 24] use diffusion with guidance or conditioning to achieve controllable scene initialization. In contrast, we emphasize closed-loop, controllable adversarial behavior simulation based on real-world data initialization rather than scene initialization.

2.2 Safety-Critical Traffic Simulation

Safety-critical traffic generation plays a crucial role in training and evaluating AV systems, enhancing their capability to navigate diverse real-world scenarios and enhancing robustness. Gradient-based methods that leverage back-propagation to create safety-critical scenarios have been proposed to evaluate AV prediction and planning models [4, 11]. Hanselmann et al. use kinematic gradients to modify vehicle trajectories, with the goal of improving the robustness of imitation learning planners [11]. Cao et al. have developed a model with differentiable dynamics, enabling the generation of realistic adversarial trajectories for trajectory prediction models through backpropagation techniques [4]. Black-box optimization approaches include perturbing actions based on kinematic bicycle models [33] and using Bayesian Optimization to create adversarial self-driving scenarios that escalate collision risks with simulated entities [1]. Zhang et al. target trajectory prediction models via white- and black-box attacks that adversarially perturb real driving trajectories [37]. For an extensive review of this topic, we direct readers to Ding et al. [9].

The field has recently seen advancements in data-driven methods for safety-critical scenario generation [26, 34, 36]. For instance, Xu et al. introduce a diffusion-based approach in CARLA, applying various adversarial optimization objectives to guide the diffusion process for safety-critical scenario generation [34]. Rempe et al. proposed STRIVE to utilize gradient-based adversarial optimization on the latent space, constrained by a graph-based CVAE traffic motion model, to generate realistic safety-critical scenarios for rule-based planners [26]. However, a common limitation in these approaches is the absence of closed-loop interaction, essential for accurately simulating interactive real-world driving.

3 Problem Formulation

We consider a simulated interactive traffic scenario consisting of N agents; one is the ego vehicle controlled by the *planner* π , and the remaining $N - 1$ are reactive agents modeled by a function g . Our objective is to create a safety-critical *closed-loop* collision simulation, where reactive agents demonstrate *realistic, controllable* behavior. Of the $N - 1$ reactive agents, one or a subset is considered the adversarial agents (denoted as agent a), meant to collide with the ego vehicle.

The adversarial agent, formulated within the reactive agent model g , is governed by an adversarial term designed (detailed in Sec. 5) to be both controllable and adversarial to the planner π . This setup allows the adversarial agent to pose direct challenges to π , testing its resilience in complex scenarios. Concurrently, the other non-adversarial agents, also controlled by g with varying parameters, emulate authentic, reactive behaviors, thus enriching the simulation scenario with realistic and diverse traffic conditions. The dual role of the adversarial agent and non-adversarial agents ensures that while it challenges π , the overall simulation environment plausibly represents real-world driving conditions.

At any given timestep t , the states of the N vehicles are represented as $\mathbf{s}_t = [\mathbf{s}_t^1, \dots, \mathbf{s}_t^N]$, where $\mathbf{s}_t^i = (x_t^i, y_t^i, v_t^i, \theta_t^i)$ indicates the 2D position, speed, and yaw of vehicle i . The corresponding actions for each vehicle are $\mathbf{a}_t = [\mathbf{a}_t^1, \dots, \mathbf{a}_t^N]$, with $\mathbf{a}_t^i = (\dot{v}_t^i, \dot{\theta}_t^i)$ representing the acceleration and yaw rate. To predict the state at the next timestep $t + 1$, a transition function f is used, which computes $\mathbf{s}_{t+1} = f(\mathbf{s}_t, \mathbf{a}_t)$ based on current state and action. We adopt unicycle dynamics as the transition function.

Each agent’s decision context is \mathbf{c}_t^i , which includes the agent-centric map I^i and the T_{hist} historical states of neighboring vehicles from time $t - T_{\text{hist}}$ to t , defined as $\mathbf{s}_{t-T_{\text{hist}}:t} = \{\mathbf{s}_{t-T_{\text{hist}}}, \dots, \mathbf{s}_t\}$. In closed-loop traffic simulation, each agent continuously generates and updates its trajectory based on the current decision context \mathbf{c}_t^i . After generating a trajectory, the simulation executes the first few steps of the planned actions before updating \mathbf{c}_t^i and re-planning. See Sec. 6.2 for more implementation details.

Planner π The planner π determines the ego vehicle’s future trajectory over a time horizon t to $t + T$. The planned state sequence is denoted by $s_{t:t+T}^1 = \pi(\mathbf{c}_t^1)$, where $\pi(\mathbf{c}_t^1)$ processes the historical states and map data within \mathbf{c}_t^1 to plan future states based on the current scene context.

Reactive Agents g The reactive agent model g , parameterized by θ , is designed to simulate the behavior of the $N - 1$ non-ego vehicles, represented by the set $\{s_{t:t+T}^i\}_{i=2}^N$. Each vehicle’s state sequence, $s_{t:t+T}^i$, is generated by $g_\theta(\mathbf{c}_t^i, \psi_i)$, which incorporates the decision context \mathbf{c}_t^i and a set of control parameters ψ_i unique to each agent. These parameters ψ_i enable the fine-tuning of individual behaviors within the simulation. In our approach, we train the model g on real-world driving data to ensure the trajectories it produces are not only controllable, supporting the generation of various safety-critical scenarios, but also realistic.

4 Diffusion Models for Traffic Simulation

For closed-loop safety-critical traffic simulation, the reactive agents, especially the adversarial agent, should be 1) controllable, and 2) realistic. With recent advances in controllable diffusion models [18, 25, 39], we adopt trajectory diffusion models to generate realistic simulations.

We define the model’s operational trajectory as τ , which comprises both action and state sequences: $\tau := [\tau_a, \tau_s]$. Specifically, $\tau_a := [a_0, \dots, a_{T-1}]$ represents the

sequence of actions, while $\tau_s := [s_1, \dots, s_T]$ denotes the corresponding sequence of states. Following the approach described in [39], our model predicts the action sequence τ_a , and the state sequence τ_s can be derived starting from the initial state s_0 and dynamic model f .

A diffusion model generates a trajectory by reversing a process that incrementally adds noise. Starting with an actual trajectory τ_0 sampled from the data distribution $q(\tau_0)$, a sequence of increasingly noisy trajectories $(\tau_1, \tau_2, \dots, \tau_K)$ is produced via a forward noising process. Each trajectory τ_k at step k is generated by adding Gaussian noise parameterized by a predefined variance schedule β_k [14]:

$$q(\tau_{1:K}|\tau_0) := \prod_{k=1}^K q(\tau_k|\tau_{k-1}), \quad (1)$$

$$q(\tau_k|\tau_{k-1}) := \mathcal{N}(\tau_k; \sqrt{1 - \beta_k}\tau_{k-1}, \beta_k\mathbf{I}). \quad (2)$$

The noising process gradually obscures the data, where the final noisy version $q(\tau_K)$ approaches $\mathcal{N}(\tau_K; \mathbf{0}, \mathbf{I})$. The trajectory generation process is then achieved by learning the reverse of this noising process. Given a noisy trajectory τ_K , the model learns to denoise it back to τ_0 through a sequence of reverse steps. Each reverse step is modeled as:

$$p_\theta(\tau_{k-1}|\tau_k, \mathbf{c}) := \mathcal{N}(\tau_{k-1}; \mu_\theta(\tau_k, k, \mathbf{c}), \Sigma_k), \quad (3)$$

where θ are learned functions that predict the mean μ of the reverse step, and Σ_k is a fixed schedule. By iteratively applying the reverse process, the model learns a trajectory distribution, effectively generating a plausible future trajectory from a noisy start.

During the trajectory prediction phase, the model estimates the final clean trajectory denoted by $\hat{\tau}_0$. This estimated trajectory is used to compute the mean μ as described in [21]. For more details, see supplementary material Sec. D.

5 Diffusion Models for Safety-Critical Traffic Simulation

The diffusion model, once trained on realistic trajectory data, inherently reflects the behavioral patterns present in its training distribution. However, to effectively simulate and analyze safety-critical scenarios, there is a crucial need for a mechanism that allows for the controlled manipulation of agent behaviors [18, 39]. This is particularly important for generating adversarial behaviors and ensuring scene consistency in simulations.

5.1 Guiding Reactive Agents

Our approach specifically introduces guidance to the sampled trajectories at each denoising step, aligning them with predefined objectives $J(\tau)$. The concept of guidance involves using the gradient of J to subtly perturb the predicted mean of the model at each denoising step. This process enables the generation

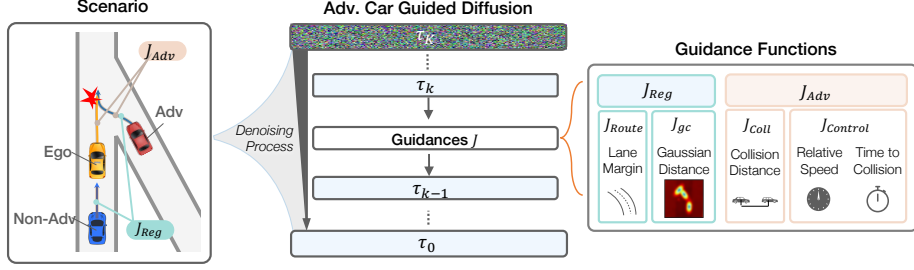


Fig. 2: Guided Diffusion Process for the Adversarial Agent. This process optimizes the adversarial agent’s trajectory using the adversarial cost function J_{adv} to the ego vehicle. In particular, we introduce $J_{control}$ to vary the adversarial behavior. Simultaneously, it applies regularization through J_{reg} for maintaining realism.

of trajectories that not only reflect realistic behavior but also cater to specific simulation needs, such as adversarial testing and maintaining scene consistency over extended periods. We adopt the reconstruction guidance (clean guidance) introduced from [16, 25]:

$$\tilde{\tau}_0 = \hat{\tau}_0 - \alpha \Sigma_k \nabla_{\tau_k} J(\hat{\tau}_0) \quad (4)$$

This strategy improves guidance robustness, yielding smoother, more stable trajectories without the usual numerical issues from noisy data.

In practice, diversifying the behavior of adversarial agents within the *same* scenarios is crucial for a thorough assessment of AVs. Despite the significance of this challenge, it remains largely unexplored in previous works [26, 39]

The loss function for the non-reactive agents, $J(\tau)$, consists of a collision term J_{coll} , which encourages collisions between the adversarial agent and the ego agent, two control terms J_v and J_{ttc} , which control the relative speed and time-to-collision between the ego and adversarial agent respectively, a regularization term J_{Gauss} , which discourages collisions between the reactive agents, and a route guidance term J_{route} , which discourages the reactive agents from going outside the road:

$$J(\tau) = \rho \underbrace{(J_{coll} + J_v + J_{ttc})}_{J_{adv}(\tau)} + \underbrace{J_{route} + J_{Gauss}}_{J_{reg}(\tau)}, \quad (5)$$

where ρ denotes a scalar weight that determines whether a reactive agent behaves adversarially towards the ego agent, i.e., whether it attempts to collide with the ego agent.

Collision with Planner We define J_{coll} to encourage the collision between the adversarial agent and the ego agent, given by:

$$J_{coll} = - \sum_{t=1}^T d(t), \quad (6)$$

where $d(t)$ represents the distance between the ego and the adversarial agent at each time step of the planning horizon T . The adversarial agent is either pre-selected based on lane proximity or dynamically selected based on the distance to the ego agent. Details are in the supplementary Sec. D.4.

Safety Criticality of Collisions We control the relative speed J_v between the ego and adversary at each time step (v_t^1 and v_t^a), and the time-to-collision (TTC) cost J_{ttc} [22] to control the safety criticality of potential collisions, with the latter given by:

$$J_{\text{ttc}} = \sum_{t=1}^T -\exp\left(-\frac{\tilde{t}_{\text{col}(t)}^2}{2\lambda_t} - \frac{\tilde{d}_{\text{col}(t)}^2}{2\lambda_d}\right), \quad (7)$$

where $\tilde{t}_{\text{col}(t)}$ is the time to collision at time t , $\tilde{d}_{\text{col}(t)}$ is the distance to collision and λ_t and λ_d are bandwidth parameters for time and distance. This formula uses a constant velocity assumption. Intuitively, the time-to-collision cost favors scenarios with high relative speeds and challenging collision angles for the ego vehicle to avoid. For details of J_v and J_{ttc} , see supplementary Sec. C.2.

Route Guidance Given an agent’s trajectory τ and the corresponding route r —the predefined path on a lane graph from its starting point to its destination—we compute the normal distance of each point τ_t on the trajectory to the route at each timestep. We then penalize deviations from the route that exceed a predefined margin d . This process is captured by the following route guidance cost function:

$$J_{\text{route}}(\tau, r) = \sum_{t=1}^T \max(0, |d_n(\tau_t, r) - d_m|), \quad (8)$$

where $d_n(\tau_t, r)$ denotes the normal distance from the point τ_t on the trajectory to the nearest point on the route r at timestep t , and d_m represents the acceptable deviation margin from the route. In contrast to the off-road loss in prior studies [39], our proposed route guidance system more effectively indicates each agent’s intended path, improving adherence to traffic rules as demonstrated in Sec. 6. The flexibility of route guidance supports diverse agent interactions, such as modifying routes to encourage lane changes among reactive agents [29].

Gaussian Collision Guidance Given the trajectories of agents, we calculate the Gaussian distance for each pair of agents (i, j) at each timestep t from 1 to T . The Gaussian distance between the agents takes into account both the tangential (d_t) and normal (d_n) components of the projected distances. The aggregated Gaussian distance is:

$$J_{\text{Gauss}} = \sum_{t=1}^T \sum_{i,j} \exp\left(-\frac{1}{2\sigma^2} \left(\lambda \cdot d_t^{ij}(t)^2 + d_n^{ij}(t)^2\right)\right) \quad (9)$$

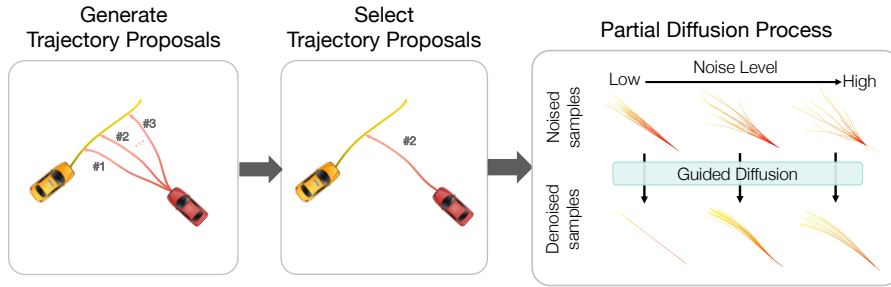


Fig. 3: Framework for Partial Diffusion. We generate proposals based on domain knowledge (e.g., collision types). Users can adjust noise levels to balance between user control and the model’s data distribution.

where $d_t^{ij}(t)$ and $d_n^{ij}(t)$ represent the tangential and normal distances from agent j ’s trajectory point at time t to agent i ’s heading axis, respectively, and σ is the standard deviation for these distances. In this formulation, λ is a scaling factor applied to the tangential distance $d_t^{ij}(t)$. This approach contrasts with the disk approximation method, which primarily penalizes the Euclidean distance between agents. By accounting for both tangential and normal components, the Gaussian collision distance method significantly reduces collision rate, which we discuss in Sec. 6.

5.2 Partial Diffusion: Controlling Collision Types

We introduce a novel approach through a partial diffusion process, utilizing trajectory proposals to initiate the diffusion process. This methodology enables the variation in collision types by the adversarial agent within the diffusion, tailoring the adversarial outcomes to specific evaluation needs, the results are discussed in Sec. 6.6.

Figure 3 illustrates our framework, which is divided into three main steps to generate trajectory proposals for various collision scenarios. First, we create initial trajectory proposals (τ_0) aimed at capturing different types of collisions. The next critical step involves setting the partial diffusion ratio γ , which defines the specific point in the process, $k_p = \gamma \cdot K$, at which we start modifying the trajectory. Starting from step k_p , we adjust the trajectory by adding a precise level of Gaussian noise $\epsilon \sim N(0, I)$: $\hat{\tau}_{k_p} = \sqrt{\alpha_{k_p}} \tau_0 + \sqrt{1 - \alpha_{k_p}} \epsilon$. The final stages include removing noise and using guided diffusion for the rest of the k_p steps to refine the trajectory into a realistic path that suits our collision scenario goals.

To generate the trajectory proposals, we develop a rule-based approach in which we first identify the centerlines of the ego and adversarial agent and then search for potential intersections of their respective centerlines. If such an intersection exists, we generate the proposals by selecting an acceleration value and lateral offset from the centerline that is likely to cause the desired collision type based on the projected plan of the ego agent. Note that the trajectory




Fig. 4: Partial Diffusion Results of Rule-Based Planner on NuScenes Dataset. The safety-critical scenarios show the framework’s ability to create realistic and challenging situations, varying collision types based on different trajectory proposals. The gradient lines reflects the planned trajectories in the next 3.2 seconds.

proposals are updated in a closed-loop manner to account for the interaction between the ego and the adversarial agent.

This method allows for precise control over the diffusion trajectory, enabling adversarial agents to create customized collision scenarios. Users can adjust γ to fine-tune the balance between explicit control and the model’s trained data distribution.

6 Experiments

We validate the efficacy of our proposed framework via experiments with real-world driving data. Our results demonstrate that the framework can generate *realistic* and *controllable* adversarial behavior to challenge the planner.

6.1 Dataset

We conduct our experiments on two large-scale real-world driving datasets: nuScenes [2], which consists of 5.5 hours of driving data from two cities, and nuPlan, which consists of 1500 hours of driving data from four cities. We train the model on scenes from the nuScenes train split and evaluate it on the scenes from the nuScenes validation splits and nuPlan mini validation splits. We focus on vehicle-to-vehicle interactions.

6.2 Implementation Details

Baselines. We compare our approach against STRIVE [26] using their open-source implementation and our re-implementation of DiffScene [34]. STRIVE is recognized for its proficiency in generating adversarial safety-critical scenarios using a learned traffic model and adversarial optimization in the latent space.

Planner. Our experiments utilize a range of different planners: 1) A rule-based planner, as implemented in STRIVE [26], which operates on the lane graph

and employs the constant velocity model to predict future trajectories of non-ego vehicles, generating multiple trajectory candidates and selecting the one least likely to result in a collision; 2) a hybrid planner BITS [35]; 3) PDM-Closed [5], the winner of the 2023 nuPlan Planning Challenge; 4) a deterministic a learning-based Behavior Cloning (BC) planner, which utilizes a ResNet Encoder, followed by an MLP decoder to generate future trajectories; and 5) an Intelligent Driver Model (IDM) planner [30].

Diffusion Model. We follow the architecture described in [39]. We represent the context using an agent-centric map and past trajectories on a rasterized map. The traffic scene is encoded using a ResNet structure [12], while the input trajectory is processed through a series of 1D temporal convolution blocks in an UNet-like architecture, as detailed in [18]. The model uses $K = 100$ diffusion steps. During the inference phase, we generate a sample of potential future trajectories for each reactive agent in a given scene. From these, we select the trajectory that yields the lowest guidance cost. This process is referred to as *filtering*.

Closed-loop Simulation. The simulation framework for our experiments is built upon an open-source traffic behavior framework [35]. Within this framework, both the planner and reactive agents update their plans at a frequency of 2Hz.

6.3 Evaluation Metrics

Our goal is to validate that the proposed method can generate safety-critical scenarios that are both *realistic* and *controllable*. For *realism* assessment, in accordance with [39], we compare statistical data between simulated trajectories and actual ground trajectories. This involves calculating the Wasserstein distance between their driving profiles’ normalized histograms, focusing on the mean of mean values for three properties: longitudinal acceleration, latitudinal acceleration, and jerk. To evaluate *controllability*, we measure metrics related to parameters we can control, specifically **relative speed** and **time-to-collision cost** between the ego and adversarial agents.

Our method aims to evaluate the extent of control over adversarial behaviors within a single scenario. To do this, we focus on measuring **collision diversity**—a metric that quantifies the range of differences in collision angles, relative speeds, and collision points. By calculating the variance of these parameters within the same scenario, we can determine how diverse and controllable the adversarial behaviors are, ensuring the scenario’s realism and controllability.

For a more nuanced understanding of *realism*, we analyze the **collision rate**—the average fractions of agents colliding, and the **offroad rate**—the percentage of reactive agents going off-road, both of which are considered failure rates [35]. All metrics are averaged across scenarios. For detailed definitions and metrics, see supplementary Sec. C.

Dataset	Method	Collision (%) \uparrow	Other Offroad (%) \downarrow	Adv Offroad (%) \downarrow	Collision Rel Speed (m/s) \downarrow	Realism \downarrow	Time (s) \downarrow
nuScenes	SAFE-SIM	43.2	1.8	11.4	-0.12	0.38	104.5 \pm 17.7
	STRIVE	36.4	2.2	11.4	5.52	0.85	427.2 \pm 169.8
	DiffScene	18.2	11.4	9.0	16.4	0.52	105.4 \pm 22.5
nuPlan	SAFE-SIM	80	9.4	11.7	6.75	0.27	173.4 \pm 73.3
	DiffScene	56.7	14.0	5.0	-2.81	0.42	176.7 \pm 77.5

Table 2: Safety-critical Traffic Simulation. We compare our approach against STRIVE [26] and DiffScene [34] for safety-critical traffic simulation with a rule-based planner. SAFE-SIM outperforms STRIVE on all metrics and demonstrates higher collision rates and better realism than DiffScene.

Planner	Ego-Adv Coll (%) \uparrow	Ego-Other Coll (%) \uparrow	Adv Offroad (%) \downarrow	Coll Speed (m/s) \downarrow	Ego Accel (m/s ²) \downarrow	Realism \downarrow
BC	38.8	37.3	9.0	2.47	0.54	0.79
IDM	49.3	58.2	3.0	-0.40	0.94	0.78
Lane-Graph	34.3	37.3	1.5	2.68	1.62	0.57
BITS	16.4	19.4	6.0	3.07	0.95	0.79
PDM-Closed	26.9	50.7	1.5	2.73	1.30	0.86

Table 3: Safety-Critical Simulation for Different Planners. SAFE-SIM can generate diverse safety-critical scenarios tailored to different planners, including rule-based, learning-based, and hybrid planners.

6.4 Evaluation of Safety-Critical Traffic Simulation with Baseline Methods

We compared our method with STRIVE [26] and DiffScene [34], utilizing the Lane-graph-based planner from STRIVE. The evaluation focused on collision rates between the ego and the adversarial agent (“Collision”), adversarial agent off-road rates (“Adv Offroad”), other agents’ off-road rates (“Other Offroad”), collision relative speed between the ego and adversary (“Collision Rel Speed”), realism of all non-ego agents (“Realism”), and simulation time.

The results, presented in Tab. 2, demonstrate that our method excels in all metrics compared to STRIVE, especially in collisions and realism. Compared to DiffScene, SAFE-SIM also exhibits a higher collision rate with better realism. Note that we trained models on nuScenes and tested them in nuPlan without additional fine-tuning. As illustrated in Figure 4, the qualitative examples from the NuScenes dataset demonstrate how our framework can challenge the rule-based planner with various driving situations. See supplementary Sec. A for more qualitative examples.

6.5 Safety-Critical Simulation with Different Planners

In Tab. 3, we demonstrate our framework’s ability to generate collisions across various planner types: rule-based, learning-based, and hybrid planners. Notably, the IDM planner exhibits the highest Ego-Adv collision rate. This heightened rate

TTC Cost Weight	TTC Cost	Coll Speed (m/s)	Coll Angle (deg)	Coll Rate (%) ↑	Realism ↓
0.0	0.18	2.45	-7.43	48.2	0.76
1.0	0.21	2.30	0.43	53.6	0.79
2.0	0.26	3.78	-17.0	60.7	0.81

Table 4: Controlling Time to Collision (TTC). The table shows the impact of different TTC Cost weights on collision scenarios. Increasing the TTC Cost weight results in an increase in collision rate, suggesting a heightened challenge for the ego vehicle in avoiding collisions.

J_{adv}	J_{reg}	Partial Diffusion	Collision Metrics			Diversity		
			Collision (%) ↑	Adv Offroad (%) ↓	Realism ↓	Collision Angle Var (rad) ↑	Collision Rel Speed Var (m/s) ↑	Collision Point Var (m) ↑
✓	✓	×	23.9	13.8	0.58	2.22	2.99	1.62
✓	✓	✓	29.0	14.6	0.57	3.10	1.96	5.44
✓	×	×	53.5	23.1	0.58	3.34	4.81	2.47

Table 5: Ablation Study on Controllability. This study examines each component of our proposed method. Partial diffusion significantly increases the variance of the collision point, resulting in a greater diversity of collision scenarios.

can be attributed to the IDM’s focus on vehicles near the same lane, potentially overlooking other vehicles in the scene. Consequently, in scenarios with the IDM planner, our framework can induce collisions at relatively lower speeds.

6.6 Evaluation: Controlling Safety-Criticality

A key feature of SAFE-SIM is its ability to generate controllable adversarial behaviors, offering variations not possible with previous methods.

Controlling Time-to-Collision. We control the orientation and relative speed together using the time-to-collision (TTC) cost, as described in Sec. 5. We manipulate the scenario’s safety-criticality by adjusting the relative weight of the TTC cost. To assess the impact of these adjustments, we measure the average TTC cost shortly before a collision occurs (0.5 seconds). Our observations, detailed in Tab. 4, show that increasing the TTC weight raises the TTC cost. Notably, while the relative collision speed remains fairly consistent, the collision angle shifts, indicating a greater difficulty in avoiding ego-adversary collisions. Additionally, our method can control other aspects, such as relative speed, as detailed in supplementary Sec. F.1.

6.7 Ablation Study on Controllability

We performed an ablation study to assess the influence of different guidance strategies in diffusion models on the quality of simulations. This study, detailed

in Table 5, aimed to quantify collision diversity. We compare against our baseline approach, consisting of regularized and adversarial objectives ($J_{reg} + J_{adv}$), by predetermining the selection of adversarial agents and conducting experiments with multiple seeds (three) in our framework to evaluate collision diversity. In partial diffusion, we manipulated the trajectory proposal selection mechanism across centerlines with normal offsets of -2.0, 0.0, and 2.0.

Effectiveness of Partial Diffusion. Table 5 reveals that partial diffusion significantly enhances both the collision point and angle diversity in comparison to the baseline approach. This underscores the capability of partial diffusion to generate a wider array of collision scenarios through various trajectory proposals, illustrating its potential to explore diverse collision dynamics effectively.

Impact of the Regularization Term (J_{reg}). Incorporating J_{reg} results in a notable decrease in the adversarial-collision rate, highlighting the importance of the regularization term in enhancing simulation realism. Without J_{reg} , the collision rate between the adversarial agent and ego vehicle increases, but the adversarial vehicle goes offroad more often, leading to scenarios that differ significantly from realistic behavior.

6.8 Limitation and Failure Cases

We identified areas for improvement in SAFE-SIM (see supp. Fig. A4). In certain cases, the adversarial agent unrealistically collides with non-adversarial agents before reaching the ego agent. Additionally, some scenarios result in collisions where the ego planner is not at fault. While understanding how the ego planner can avoid such cases is important, creating more scenarios where the ego is at fault would be beneficial.

7 Conclusion

In this work, we present a closed-loop simulation framework utilizing guided diffusion models for creating safety-critical scenarios to assess autonomous vehicle (AV) algorithms. Our research is in line with the goals of SO-TIF, focusing on how autonomous vehicles respond to dynamic scenarios like aggressive driving, underscoring our dedication to safety across diverse and unforeseeable conditions. Our framework introduces innovative guidance objectives tailored for controllable, stable, long-term safety-critical simulations. A key aspect of our method lies in its ability to vary the types of adversarial behavior within collision scenarios. By integrating adversarial objectives and partial diffusion, we enable fine-grained control over adversarial actions. This versatility enables our framework to produce a broader range of realistic and manageable scenarios, setting a new standard in adversarial scenario generation beyond the limitations of existing approaches.

Future directions for our research include: 1) exploring the application of our framework in closed-loop policy training, and 2) developing automated methods for adjusting controllable parameters. These methods aim to facilitate the generation of diverse, long-tail scenarios. We believe this framework holds significant promise for enhancing real-world AV safety.

Acknowledgements

This work was part of W.J. Chang’s summer internship at NEC Labs America, and he is also supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE 2146752. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation. The authors would like to thank Chih-Ling Chang for her insightful suggestions and assistance with figures and presentations.

References

1. Abeysirigoonawardena, Y., Shkurti, F., Dudek, G.: Generating adversarial driving scenarios in high-fidelity simulators. In: 2019 International Conference on Robotics and Automation (ICRA). pp. 8271–8277. IEEE (2019)
2. Caesar, H., Bankiti, V., Lang, A.H., Vora, S., Liong, V.E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., Beijbom, O.: nuscenes: A multimodal dataset for autonomous driving. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 11621–11631 (2020)
3. Caesar, H., Kabzan, J., Tan, K.S., Fong, W.K., Wolff, E., Lang, A., Fletcher, L., Beijbom, O., Omari, S.: nuplan: A closed-loop ml-based planning benchmark for autonomous vehicles. arXiv preprint arXiv:2106.11810 (2021)
4. Cao, Y., Xiao, C., Anandkumar, A., Xu, D., Pavone, M.: Advdo: Realistic adversarial attacks for trajectory prediction. In: European Conference on Computer Vision. pp. 36–52. Springer (2022)
5. Dauner, D., Hallgarten, M., Geiger, A., Chitta, K.: Parting with misconceptions about learning-based vehicle motion planning. In: CoRL (2023)
6. Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. *Advances in neural information processing systems* **34**, 8780–8794 (2021)
7. Ding, W., Chen, B., Li, B., Eun, K.J., Zhao, D.: Multimodal safety-critical scenarios generation for decision-making algorithms evaluation. *IEEE Robotics and Automation Letters* **6**(2), 1551–1558 (2021). <https://doi.org/10.1109/LRA.2021.3058873>
8. Ding, W., Chen, B., Xu, M., Zhao, D.: Learning to collide: An adaptive safety-critical scenarios generating method. In: 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 2243–2250. IEEE (2020)
9. Ding, W., Xu, C., Arief, M., Lin, H., Li, B., Zhao, D.: A survey on safety-critical driving scenario generation—a methodological perspective. *IEEE Transactions on Intelligent Transportation Systems* (2023)
10. Gu, T., Chen, G., Li, J., Lin, C., Rao, Y., Zhou, J., Lu, J.: Stochastic trajectory prediction via motion indeterminacy diffusion. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 17113–17122 (2022)
11. Hanselmann, N., Renz, K., Chitta, K., Bhattacharyya, A., Geiger, A.: King: Generating safety-critical driving scenarios for robust imitation via kinematics gradients. In: European Conference on Computer Vision. pp. 335–352. Springer (2022)
12. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)

13. Ho, J., Chan, W., Saharia, C., Whang, J., Gao, R., Gritsenko, A., Kingma, D.P., Poole, B., Norouzi, M., Fleet, D.J., et al.: Imagen video: High definition video generation with diffusion models. arXiv preprint arXiv:2210.02303 (2022)
14. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in neural information processing systems* **33**, 6840–6851 (2020)
15. Ho, J., Salimans, T.: Classifier-free diffusion guidance. arXiv preprint arXiv:2207.12598 (2022)
16. Ho, J., Salimans, T., Gritsenko, A., Chan, W., Norouzi, M., Fleet, D.J.: Video diffusion models. arXiv:2204.03458 (2022)
17. Ivanovic, B., Song, G., Gilitschenski, I., Pavone, M.: trajdata: A unified interface to multiple human trajectory datasets. arXiv preprint arXiv:2307.13924 (2023)
18. Janner, M., Du, Y., Tenenbaum, J.B., Levine, S.: Planning with diffusion for flexible behavior synthesis. arXiv preprint arXiv:2205.09991 (2022)
19. Lu, J., Wong, K., Zhang, C., Suo, S., Urtasun, R.: Scenecontrol: Diffusion for controllable traffic scene generation. In: *IEEE International Conference on Robotics and Automation (ICRA)* (2024)
20. Montali, N., Lambert, J., Mouglin, P., Kuefler, A., Rhinehart, N., Li, M., Gulino, C., Emrich, T., Yang, Z., Whiteson, S., et al.: The waymo open sim agents challenge. *Advances in Neural Information Processing Systems* **36** (2024)
21. Nichol, A.Q., Dhariwal, P.: Improved denoising diffusion probabilistic models. In: *International Conference on Machine Learning*. pp. 8162–8171. PMLR (2021)
22. Nishimura, H., Mercat, J., Wulfe, B., McAllister, R.T., Gaidon, A.: Rap: Risk-aware prediction for robust planning. In: *Conference on Robot Learning*. pp. 381–392. PMLR (2023)
23. Pillion, J., Peng, X.B., Fidler, S.: Trajenglish: Traffic modeling as next-token prediction. In: *The Twelfth International Conference on Learning Representations* (2024)
24. Pronovost, E., Ganesina, M.R., Hendy, N., Wang, Z., Morales, A., Wang, K., Roy, N.: Scenario diffusion: Controllable driving scenario generation with diffusion. *Advances in Neural Information Processing Systems* **36**, 68873–68894 (2023)
25. Rempe, D., Luo, Z., Bin Peng, X., Yuan, Y., Kitani, K., Kreis, K., Fidler, S., Litany, O.: Trace and pace: Controllable pedestrian animation via guided trajectory diffusion. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 13756–13766 (2023)
26. Rempe, D., Pillion, J., Guibas, L.J., Fidler, S., Litany, O.: Generating useful accident-prone driving scenarios via a learned traffic prior. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 17305–17315 (2022)
27. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 10684–10695 (2022)
28. Suo, S., Regalado, S., Casas, S., Urtasun, R.: Trafficsim: Learning to simulate realistic multi-agent behaviors. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 10400–10409 (2021)
29. Suo, S., Wong, K., Xu, J., Tu, J., Cui, A., Casas, S., Urtasun, R.: Mixsim: A hierarchical framework for mixed reality traffic simulation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 9622–9631 (2023)
30. Treiber, M., Hennecke, A., Helbing, D.: Congested traffic states in empirical observations and microscopic simulations. *Physical review E* **62**(2), 1805 (2000)

31. Unknown: Midjourney. <https://www.midjourney.com>, accessed: 2023-11-16
32. Unknown: Openai dall-e-2. <https://openai.com/product/dall-e-2>, accessed: 2023-11-16
33. Wang, J., Pun, A., Tu, J., Manivasagam, S., Sadat, A., Casas, S., Ren, M., Urtasun, R.: Advsim: Generating safety-critical scenarios for self-driving vehicles. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9909–9918 (2021)
34. Xu, C., Zhao, D., Sangiovanni-Vincentelli, A., Li, B.: Diffscene: Diffusion-based safety-critical scenario generation for autonomous vehicles. In: The Second Workshop on New Frontiers in Adversarial Machine Learning (2023)
35. Xu, D., Chen, Y., Ivanovic, B., Pavone, M.: Bits: Bi-level imitation for traffic simulation. In: 2023 IEEE International Conference on Robotics and Automation (ICRA). pp. 2929–2936. IEEE (2023)
36. Yin, Z.H., Sun, L., Sun, L., Tomizuka, M., Zhan, W.: Diverse critical interaction generation for planning and planner evaluation. In: 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 7036–7043. IEEE (2021)
37. Zhang, Q., Hu, S., Sun, J., Chen, Q.A., Mao, Z.M.: On adversarial robustness of trajectory prediction for autonomous vehicles. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15159–15168 (2022)
38. Zhong, Z., Rempe, D., Chen, Y., Ivanovic, B., Cao, Y., Xu, D., Pavone, M., Ray, B.: Language-guided traffic simulation via scene-level diffusion. arXiv preprint arXiv:2306.06344 (2023)
39. Zhong, Z., Rempe, D., Xu, D., Chen, Y., Veer, S., Che, T., Ray, B., Pavone, M.: Guided conditional diffusion for controllable traffic simulation. In: 2023 IEEE International Conference on Robotics and Automation (ICRA). pp. 3560–3566. IEEE (2023)