

# Guided Conditional Diffusion for Controllable Traffic Simulation

Ziyuan Zhong, Davis Rempe, Danfei Xu, Yuxiao Chen,  
 Sushant Veer, Tong Che, Baishakhi Ray, and Marco Pavone

**Abstract**—Controllable and realistic traffic simulation is critical for developing and verifying autonomous vehicles. Typical heuristic-based traffic models offer flexible control to make vehicles follow specific trajectories and traffic rules. On the other hand, data-driven approaches generate realistic and human-like behaviors, improving transfer from simulated to real-world traffic. However, to the best of our knowledge, no traffic model offers both controllability and realism. In this work, we develop a conditional diffusion model for controllable traffic generation (CTG) that allows users to control desired properties of trajectories at test time (e.g., reach a goal or follow a speed limit) while maintaining realism and physical feasibility through enforced dynamics. The key technical idea is to leverage recent advances from diffusion modeling and differentiable logic to guide generated trajectories to meet rules defined using signal temporal logic (STL). We further extend guidance to multi-agent settings and enable interaction-based rules like collision avoidance. CTG is extensively evaluated on the nuScenes dataset for diverse and composite rules, demonstrating improvement over strong baselines in terms of the controllability-realism tradeoff. Demo videos can be found at <https://aiasd.github.io/ctg.github.io>

## I. INTRODUCTION

Simulation is crucial to comprehensively evaluate modern autonomous vehicles (AVs). Due to the difficulty and danger of running large-scale real-world tests [1], AV developers rely on extensive testing in simulation to produce reliable systems [2]. To be most useful, simulators must embody both *realism* and *controllability*, especially for the models of traffic. *Realistic* traffic allows developments made in simulation to faithfully transfer to the real world, while *controllability* enables constructing fine-grained traffic scenarios to analyze specific AV behavior. Yet, developing realistic traffic models is still an open challenge [3], [4], and little attention has been devoted to making such models easily controllable.

Existing driving simulators commonly synthesize agent behaviors by either replaying recorded driving logs or using heuristic-based controllers [5]–[7]. Such behaviors can often be controlled through a user-friendly high-level programming

Ziyuan Zhong is with Columbia University, this work was carried out while at NVIDIA, ziyuan.zhong@columbia.com. Davis Rempe is with Stanford University, this work was carried out while at NVIDIA, drempe@stanford.edu. Danfei Xu is with Georgia Tech, and with NVIDIA Research, danfei@gatech.edu, danfeix@nvidia.com. Yuxiao Chen, Sushant Veer, Tong Che are with NVIDIA Research, yuxiaoc, sveer, tongc@nvidia.com. Baishakhi Ray is with Columbia University, rayb@cs.columbia.edu. Marco Pavone is with Stanford University, and with NVIDIA Research, pavone@stanford.edu, mpavone@nvidia.com

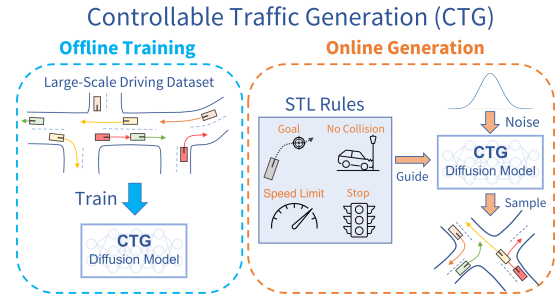


Fig. 1: **Overview.** CTG uses two key stages to enable controllable simulation. (Left) a conditional diffusion model is trained to generate realistic trajectories. (Right) Guided model sampling uses STL rules to meet desired objectives.

API [5], [6] or scenario editor UI [7]. However, these methods lack *realism* as they either cannot react to other traffic agents or are not expressive enough to appear human-like. To address these issues, recent works propose to learn generative models of traffic behavior from large-scale driving datasets [3], [4]. While generated behaviors from these models appear reasonable, they merely reflect the distribution of the training data: they lack any mechanism to *control* the generated traffic flow, making these methods less useful in practice. For example, to test an AV against cut-in from a neighboring lane, a user needs an interface to steer the vehicle to perform the maneuver. Unfortunately, the opaque nature of neural network models makes this difficult for current approaches.

We look to bridge the gap between realism and controllability by studying the problem of *controllable traffic behavior generation*. We seek a model that can generate trajectories to meet specific desirable objectives from a user at inference time, which may differ from objectives used in training. This is different from traffic models that aim to learn maximum-likelihood behaviors, and therefore are not flexible to new objectives. Moreover, controllable generation is especially challenging for learned approaches, which typically struggle to produce samples outside of the training distribution.

To enable such flexible generation, we leverage recent advances in diffusion modeling, which have achieved state-of-the-art performance in several domains including images [8], [9], audio [10]–[12], and pedestrian trajectories [13]. Importantly, diffusion models allow a notion of control at generation time through so-called *guidance*, which has benefited several tasks including conditional image generation [9], language generation [14], and offline reinforcement learning [15]. Inspired by these works, we propose a conditional agent-centric diffusion model (Fig. 1) that allows flexible

traffic behavior generation. Unlike prior diffusion models for trajectories [13], [15], our model is conditioned on a holistic context surrounding the vehicle including the local roadmap and neighboring agents. Since traffic has an inherent ground-truth transition function, we also enforce realistic vehicle dynamics into the design of the diffusion model state space to guarantee that generated trajectories are physically feasible.

To achieve *controllable* generation, Diffuser [15] proposed to *guide* each step of the denoising process in a diffusion model by perturbing network outputs with the gradient of some differentiable objective to encourage desired properties. For traffic simulation, however, deriving and implementing objectives like collision avoidance, goal reaching, and road rules is complex due to its spatio-temporal and multi-agent nature. We propose to leverage the established syntax of STL [16]. As a formal language designed for specifying spatio-temporal constraints, STL allows one to easily and scalably define driving rules; it also incorporates a notion of *robustness* that measures how well rules are satisfied. Concretely, we use this measure of rule satisfaction as the objective function for guiding the diffusion model (Fig. 1, right) by leveraging differentiable frameworks [17], [18] to make STL compatible with guidance. Since our diffusion model generates trajectories independently for each agent in a scene, we further propose a joint guidance procedure for rules involving multi-agent interactions (*e.g.*, no collisions), which simultaneously denoises all agents in the scene to mitigate interaction rule violations.

We evaluate CTG on the nuScenes [19] dataset, demonstrating the ability to meet user constraints while maintaining realistic trajectory generation. In summary, we contribute (1) the problem formulation of controllable traffic generation, (2) a conditional diffusion-based method to generate realistic traffic satisfying physical feasibility and user-specified STL rules, and (3) extensive evaluation comparing CTG to several strong baselines, demonstrating its superiority in terms of the trade-off between controllability and realism.

## II. RELATED WORK AND BACKGROUND

**Traffic simulation.** Traffic simulation approaches can be categorized into rule-based and learning-based. Rule-based approaches rely on analytical models such as cellular automata and intelligent driver model [20]. These approaches typically have fixed routes for vehicles to follow and separate longitudinal and lateral motions of agents, thus having limited expressiveness for behavior simulation. Learning-based approaches mimic real-world driving behavior based on trajectory datasets [21]–[23]. For example, TrafficSim [3] uses a trajectory prediction model to perform scene-level traffic simulation. BITS [4] decouples the problem into a high-level latent inference and a low-level driving behavior imitation. However, none of these methods allow a user to specify customized properties of the generated traffic behaviors at inference time.

Recent work in adversarial or safety-critical scenario generation can be seen as one instance of controllable traffic simulation, namely generating trajectories that cause an AV

to misbehave in some way [24]. STRIVE [25] generates near-collision scenarios by searching in the latent space of a trained trajectory prediction model via a test-time optimization. Abeyirigoonawardena et al. [26] and Chen et al. [27] use Bayesian optimization and reinforcement learning, respectively, to generate adversarial trajectories for vehicles in a specific scenario like intersection crossing or lane changing. While these works focus specifically on adversarial objectives, our approach is general and can generate trajectories to meet several different objectives.

**Diffusion modeling.** Controllable diffusion models have been explored with classifier [9], classifier-free [28], and reconstruction [29] guidance for image and video generation. Li et al. [14] use a diffusion model with different pre-trained classifiers to guide language generation on different natural language tasks. Diffuser [15] uses a diffusion model to plan robot behavior (state-action trajectories). Gu et al. [13] model pedestrian trajectories for forecasting. Different from these works that use no or limited context, our diffusion model is conditioned on decision-relevant information such as the map and states of nearby agents. Additionally, we leverage known vehicle dynamics models to ensure physical feasibility of the generated trajectories. Diffuser [15] also introduces test-time guidance to generate trajectories that optimize a given reward function. We build on this formulation to generate controllable traffic trajectories, but rather than learning reward functions we use analytical loss functions based on STL rules that are easy and scalable for driving applications.

### A. Signal Temporal Logic (STL)

STL [16] formulas are interpreted over signals,  $\tau = s_t, \dots, s_{t+T}$ , an ordered finite sequence of states  $s_i \in \mathbb{R}^n$ . A signal represents a sequence of real-valued, discrete-time outputs from a system. STL formulas are recursively defined based on the following context-free grammar:

$$\phi ::= \top \mid \mu_c \mid \neg\phi \mid \phi \wedge \psi \mid \phi \mathcal{U}_{[a,b]} \psi. \quad (1)$$

The grammar defines a list of expressions, each separated by the pipe ( | ), that can be used to construct an STL formula. In particular, a formula  $\phi$  is generated by selecting expressions recursively. We briefly summarize Eq. (1) here, but refer readers to Leung et al. [17] (Section 2.2) for a pedagogical introduction to STL. The core of an STL formula are predicates  $\mu_c$  of the form  $\mu(z) > c$ , where  $c \in \mathbb{R}$  and  $\mu : \mathbb{R}^n \rightarrow \mathbb{R}$  is a differentiable function. These define the desired properties/constraints of a signal. Additionally,  $\top$  means true,  $\phi$  and  $\psi$  are STL formulas, and  $[a, b] \subseteq \mathbb{R}_{\geq 0}$  is a time interval (assumed to be  $[0, \infty)$  if omitted). The symbol  $\neg$  is negation, while other symbols describe how to combine multiple formulas:  $\wedge$  (conjunction/and) and  $\mathcal{U}$  (until), which is a temporal operator. Other common logical connectives ( $\vee$  disjunction/or,  $\Rightarrow$  implies), and temporal operators ( $\diamond$  eventually,  $\square$  always) can be derived from Eq. (1).

Importantly, STL includes a formal notion of robustness, which measures how much a signal satisfies or violates a formula [16]. We use *robustness formulas* as guidance functions for the proposed conditional diffusion model. In contrast

to Leung and Pavone [30] that focus on a single-controller setting and rules fixed at training time, we consider a more flexible inference-time guidance with multiple agents.

### III. CONTROLLABLE TRAFFIC GENERATION

Next, we detail our approach CTG. Sec. III-A formulates the problem of controllable traffic generation. We then describe the two stages of CTG. The offline stage trains a dynamics-enforced conditional diffusion model to capture diverse behaviors from real-world driving data (Sec. III-B). Then during online inference, CTG generates rule-compliant behaviors by sampling the model using a novel iterative joint STL guidance process (Sec. III-C). Taken together, the conditional diffusion model and STL-based guidance enable realistic and controllable generation of traffic trajectories.

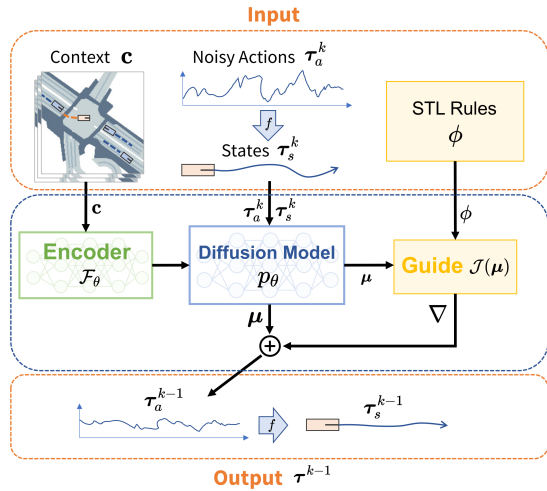


Fig. 2: **Single test-time denoising step.** Given current noisy actions and states along with encoded context, the diffusion model predicts the mean of the next-step distribution which is then perturbed according to the desired guidance function.

#### A. Problem Formulation

For some target vehicle (*target*) that we would like to simulate, let the state at a timestep  $t$  be  $s_t^{\text{target}} = (x_t^{\text{target}}, y_t^{\text{target}}, v_t^{\text{target}}, \theta_t^{\text{target}})$ , including 2D location, speed, and yaw. Similarly, let the action (*i.e.*, control) be  $a_t^{\text{target}} = (v_t^{\text{target}}, \dot{\theta}_t^{\text{target}})$  with acceleration and yaw rate. We denote  $\mathbf{c} = (I, S)$  to be decision-relevant context for the target agent. This consists of a local (agent-centric) semantic map  $I$  and the  $H$  previous states of both the target agent and its  $M$  neighbors  $S_{t-H:t}^{\text{target}} = \{s_{t-H:t}^{\text{target}}, s_{t-H:t}^1, \dots, s_{t-H:t}^M\}$ . To obtain state  $s_{t+1}^m$  for vehicle  $m$  at time  $t+1$ , we assume a transition function  $f$  that computes  $s_{t+1}^m = f(s_t^m, a_t^m)$  given the previous state  $s_t^m$  and control  $a_t^m$ . We use a unicycle dynamics model for  $f$ .

Our goal is to generate realistic and rule-satisfying traffic behavior for the target agent given (1) the decision context  $\mathbf{c}$  and (2) a function  $r : \mathbb{R}^{4T} \times \mathbb{R}^{2T} \rightarrow \mathbb{R}$  to measure rule satisfaction of a state and action trajectory. A model should generate a future trajectory for the target agent  $s_{t:t+T}^{\text{target}}$  over the next  $T$  time steps. Ideally, this trajectory maximizes satisfaction  $r(s_{t:t+T}^{\text{target}}, a_{t:t+T}^{\text{target}})$  to avoid violating the given

rule. However, in many cases there is an inherent tradeoff between rule satisfaction and trajectory realism: *e.g.*, if a user seeks a simulated vehicle with a speed much slower/faster than the speed limit, this is naturally “unrealistic” in the context of city streets. Therefore we must strike a balance between meeting user-specified constraints while still maintaining realistic behavior. As described next, our method does this by first training a rule-agnostic traffic generation model on real-world data to capture *realism*, which is then guided for rule-specific compliance only during inference.

#### B. Conditional Diffusion for Traffic Modeling

Diffusion models [8], [31] pose data generation as an iterative denoising process by learning to reverse a forward diffusion process. As shown in Fig. 2, our diffusion model operates primarily on a (future) trajectory of states and actions [15], but is conditional as it receives the context as input at each step of denoising. Starting from Gaussian noise, the diffusion model is applied iteratively to predict a clean, denoised trajectory of states and actions.

**Trajectory representation.** In this section, we denote the (future) trajectory that the model operates on as:

$$\tau := \begin{bmatrix} \tau_a \\ \tau_s \end{bmatrix}, \quad \tau_a := [a_0 \dots a_{T-1}], \quad \tau_s := [s_1 \dots s_T].$$

Unlike [15] which directly predicts states and actions jointly, our model only predicts actions  $\tau_a$  and we leverage the known dynamics  $f$  to infer states  $\tau_s$  via rollout starting at the initial state  $s_0$  (included in the past context). In other words, in the following formulation,  $\tau_s$  always refers to a state trajectory resulting from actions, or, more formally:  $\tau_s = f(s_0, \tau_a)$ . This ensures physical feasibility of the state trajectory throughout the denoising process.

**Formulation.** Let  $\tau_a^k$  be the action trajectory at the  $k$ th diffusion step where  $k=0$  is at the original clean trajectory. The forward diffusion process acting on  $\tau_a^0$  is defined as:

$$q(\tau_a^{1:K} | \tau_a^0) := \prod_{k=1}^K q(\tau_a^k | \tau_a^{k-1}) \quad (2)$$

$$q(\tau_a^k | \tau_a^{k-1}) := \mathcal{N}(\tau_a^k; \sqrt{1 - \beta_k} \tau_a^{k-1}, \beta_k \mathbf{I}),$$

where  $\beta_1, \beta_2, \dots, \beta_K$  are a fixed variance schedule that controls the scale of the injected noise at each diffusion step. As noise is gradually added, the signal is corrupted into an isotropic Gaussian distribution  $\mathcal{N}(\mathbf{0}, \mathbf{I})$ .

For trajectory generation, we seek to reverse this diffusion process using a learned conditional denoising model (Fig. 2) that is iteratively applied starting from sampled noise. Given the context information  $\mathbf{c}$ , the reverse diffusion process is:

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

$$p_\theta(\tau_a^{k-1} | \tau_a^k, \mathbf{c}) := \mathcal{N}(\tau_a^{k-1}; \mu_\theta(\tau_a^k, k, \mathbf{c}), \Sigma_\theta(\tau_a^k, k, \mathbf{c})), \quad (3)$$

where  $p(\tau_a^K) = \mathcal{N}(\mathbf{0}, \mathbf{I})$  and  $\theta$  denotes the parameters of the diffusion model. Note that at each step, the model receives *both* actions  $\tau_a^k$  and the resulting states  $\tau_s^k = f(s_0, \tau_a^k)$  as input. Following [8], [13], the variance term of the transition is a fixed schedule as  $\Sigma^k = \Sigma_\theta(\tau_a^k, k, \mathbf{c}) = \sigma_k^2 \mathbf{I}$ .

**Training.** At each training iteration, the context  $\mathbf{c}$  and ground truth clean trajectory  $\tau^0$  are sampled from a real-world driving dataset and the denoising step  $k$  is uniformly sampled from  $\{1, \dots, K\}$ . We compute the noisy input  $\tau^k$  from  $\tau^0$  by first corrupting the action trajectory  $\tau_a^k = \sqrt{\bar{\alpha}_k} \tau_a^0 + \sqrt{1 - \bar{\alpha}_k} \epsilon$ ,  $\epsilon \sim \mathcal{N}(0, \mathbf{I})$  with  $\bar{\alpha}_k = \prod_{l=0}^k 1 - \beta_l$ , and then computing the corresponding state  $\tau_s^k = f(s_0, \tau_a^k)$ . The diffusion model indirectly parameterizes  $\mu_\theta$  in Eq. (3) by instead predicting the uncorrupted trajectory  $\hat{\tau}^0 = [\hat{\tau}_a^0; f(s_0, \hat{\tau}_a^0)]$  where  $\hat{\tau}_a^0 = \hat{\tau}_a^0(\tau^k, k, \mathbf{c})$  is the direct network output (see [14], [15], [32]). Finally, we use a simplified loss function to train the model:

$$L(\theta) = \mathbb{E}_{\epsilon, k, \tau^0, \mathbf{c}} [|\tau^0 - \hat{\tau}^0|^2]. \quad (4)$$

Note that both the action and state trajectories are supervised by this loss since we found that the additional state trajectory information improves generation quality.

**Implementation details.** The input context  $\mathbf{c}$  containing agent-centric map information and past trajectories is represented in a rasterized format. This context is processed by a ResNet [33] encoder  $\mathcal{F}_\theta$  before being passed to the diffusion model. Similar to Diffuser [15], we use a diffusion model architecture like U-Net containing several blocks of temporal 1D convolutions over the input trajectory. We incorporate conditioning information by first concatenating with the diffusion step input  $k$  and then adding this conditioning feature to the convolutional features at each block of the U-Net. The diffusion process uses a cosine variance schedule [15], [32] and  $K = 100$  diffusion steps for all experiments.

### C. Guided Generation with Signal Temporal Logic

To enforce desired rules on realistic samples from the trained diffusion model, we introduce an iterative guidance algorithm with rules specified as STL formulas.

**Conditional guidance formulation.** Diffuser [15] introduces the notion of *guidance* to sample trajectories from an *unconditional* diffusion model to meet some pre-defined objective. We can do the same for our conditional diffusion model by defining a binary random variable  $\mathcal{O}$  that indicates if a trajectory is optimal, with  $p(\mathcal{O} = 1) = \exp(r(\tau_s, \tau_a))$  based on the rule satisfaction “reward”  $r$ . To approximately sample from the distribution of optimal trajectories, the steps of the denoising process can be modified to [15], [31]:

$$p_\theta(\tau_a^{k-1} | \tau^k, \mathbf{c}, \mathcal{O}) \approx \mathcal{N}(\tau_a^{k-1}; \mu + \Sigma g, \Sigma), \quad (5)$$

where  $\mu = \mu_\theta$ ,  $\Sigma = \Sigma_\theta$  as in Eq. (3), and the added gradient is computed from a *guide*  $\mathcal{J}$  based on satisfaction:

$$g = \nabla \mathcal{J}(\mu) = \nabla_\mu r(f(s_0, \mu), \mu). \quad (6)$$

The process of perturbing the predicted means from the diffusion model using gradients of a specified objective is summarized in Alg. 1. Different from Diffuser [15], an iterative inner gradient descent with clipping using the Adam optimizer is incorporated rather than using a single-step gradient update. This gives flexibility to trade off rule compliance and realism by adjusting learning rate and the number of optimization steps. When generating a future trajectory in

---

### Algorithm 1 Guided Sampling

---

```

1: Require encoder  $\mathcal{F}_\theta$ , conditional diffusion model  $\mu_\theta$ , transition function  $f$ , guide  $\mathcal{J}$ , scale  $\alpha$ , covariances  $\Sigma^k$ , diffusion steps  $K$ , inner gradient descent steps  $M$ , number of actions to take before re-planning  $l$ .
2: while not done do
3:   Observe state  $s_0$  and context  $\mathbf{c}$ 
4:   Initialize trajectory  $\tau_a^K \sim \mathcal{N}(0, \mathbf{I})$ ;  $\tau_s^K = f(s_0, \tau_a^K)$ ;  $\tau^K = [\tau_a^K; \tau_s^K]$ 
5:   for  $k = K, \dots, 1$  do
6:      $\mu = \mu_\theta(\tau^k, k, \mathcal{F}_\theta(\mathbf{c}))$ 
7:      $\mu^{(0)} = \mu$ 
8:     for  $j = 1, \dots, M$  do
9:        $\mu^{(j)} = \mu^{(j-1)} + \alpha \nabla \mathcal{J}(\mu^{(j-1)})$ 
10:       $\Delta \mu = |\mu^{(j)} - \mu^{(0)}|$ ;  $\Delta \mu \leftarrow \text{clip}(\Delta \mu, -\beta_k, \beta_k)$ 
11:       $\mu^{(j)} \leftarrow \mu^{(0)} + \Delta \mu$ 
12:       $\tau_a^{k-1} \sim \mathcal{N}(\mu^{(M)}, \Sigma^k)$ ;  $\tau_s^{k-1} = f(s_0, \tau_a^{k-1})$ ;
13:       $\tau^{k-1} = [\tau_a^{k-1}; \tau_s^{k-1}]$ 
13:   Execute first  $l$  actions of trajectory  $\tau_a^0$ 

```

---

practice, we guide several samples from the diffusion model and choose the one with the best rule satisfaction according to  $\mathcal{J}$  at the end of denoising. We refer to this as *filtration*.

**STL as guidance.** Instead of training a classifier or reward function for  $\mathcal{J}$  as in prior works [14], [15], our guidance functions are implemented analytically based on STL. For each rule we wish to apply through guidance, the grammar in Eq. (1) is used to generate a corresponding STL formula describing how the trajectory should be constrained. Examples of various STL rules used are shown in Tab. I; note the formulas are relatively simple despite the complex behavior they describe. By construction, every STL formula admits a *robustness formula* measuring the degree of rule satisfaction, which is used as the guide  $\mathcal{J}$ . Since it is necessary to compute a gradient through this function, STL formulas are implemented using differentiable frameworks [17], [18]. When multiple rules are used, we take the sum of their corresponding robustness formulas to form  $\mathcal{J}$ .

**Multi-agent guidance.** A particular challenge is applying *scene-level* rules that involve multiple agents (*e.g.*, no collisions). For this purpose, guided sampling is performed in a batched fashion over all agents in the same scene simultaneously. This way guidance can be computed across all trajectories, and the corresponding gradients are collected and propagated back to agents as needed.

**Simulating traffic.** To perform closed-loop traffic simulation of a scene with many agents, the same model is used for each agent in a standard control loop: for each agent at each step of simulation, a guided sample is generated from the model and the first few actions are taken before re-planning at a prescribed frequency. In all experiments (Sec. IV), each scene is rolled out for 20 seconds starting from a ground truth driving log, and the re-plan rate is 2 Hz as in [4].

## IV. EXPERIMENTS

We conduct experiments to validate that: (1) CTG can generate controllable traffic behaviors that satisfy user-specified rules, and (2) compared to strong baselines, CTG achieves better rule satisfaction while maintaining realism. As discussed in Sec. III-A, there is often a trade-off between realism and rule compliance; ideally a method will strike a reasonable balance and achieve good performance for both. After describing the experimental design (Sec. IV-A), we

TABLE I: STL formulas and evaluation metrics for rules described in Sec. IV-B and Sec. IV-C. STL notation is detailed in Sec. II-A.

Rule	STL Formula	Evaluation Metric
speed limit	$\phi_{\text{speedlimit}} := \bigwedge_i \square (v_{i,t} - v_{\text{limit}_i}) < \epsilon$	$h_{\text{speedlimit}} := \sum_i \sum_t \max(0, v_{i,t} - v_{\text{limit}_i})$
target speed	$\phi_{\text{targetspeed}} := \bigwedge_{i,t}  v_{i,t} - v_{\text{target}_{i,t}}  < \epsilon$	$h_{\text{targetspeed}} := \sum_i \sum_t  v_{i,t} - v_{\text{target}_{i,t}} $
no collision	$\phi_{\text{collision}} := \square \bigwedge_{i \neq j} \ (x_{i,t}, y_{i,t}) - (x_{j,t}, y_{j,t})\  > \epsilon$	$h_{\text{collision}} := \sum_{i \neq j} \mathbb{I}[\ (x_{i,t}, y_{i,t}) - (x_{j,t}, y_{j,t})\  \leq \epsilon]$
no off-road	$\phi_{\text{offroad}} := \bigwedge_{i,w} \square \ (x_{i,t}, y_{i,t}) - (x_{\text{off-road}_w}, y_{\text{off-road}_w})\  > \epsilon$	$h_{\text{offroad}} := \sum_i \mathbb{I}[\min_w \ (x_{i,t}, y_{i,t}) - (x_{\text{off-road}_w}, y_{\text{off-road}_w})\  \leq \epsilon]$
goal waypoint	$\phi_{\text{waypoint}} := \bigwedge_i \diamond \ (x_{i,t}, y_{i,t}) - (x_{\text{goal}_i}, y_{\text{goal}_i})\  < \epsilon$	$h_{\text{waypoint}} := \sum_i \min_t \ (x_{i,t}, y_{i,t}) - (x_{\text{goal}_i}, y_{\text{goal}_i})\ $
stop sign	$\phi_{\text{stopsign}} := \bigwedge_i \square (\phi_i \text{ in stop region} \Rightarrow \diamond \square_{[0,m]} (\phi_i \text{ in stop region} \wedge \phi_{\text{stop}}))$	$h_{\text{stopsign}} := \sum_i \min_i \text{ in stop region }  v_{i,t} $
stop sign + no off-road	$\phi_{\text{stopsign}} \wedge \phi_{\text{offroad}}$	$h_{\text{stopsign}}, h_{\text{offroad}}$
goal waypoint + target speed	$\phi_{\text{waypoint}} \wedge \phi_{\text{targetspeed}}$	$h_{\text{waypoint}}, h_{\text{targetspeed}}$

TABLE II: Quantitative results (Single Rule). The top two methods for each metric are bolded.

	speed limit			target speed			no collision			no off-road			goal waypoint		
	rule	real	fail	rule	real	fail	rule	real	fail	rule	real	fail	rule	real	fail
SimNet	0.739	0.898	0.353	1.989	0.898	0.353	0.137	0.898	0.353	0.453	0.900	0.353	7.543	0.900	0.353
SimNet+opt	0.038	0.770	0.470	0.630	1.234	0.593	0.045	1.149	0.398	0.427	1.242	0.416	1.947	1.162	0.467
TrafficSim	0.737	1.362	0.443	1.922	1.346	0.440	0.140	1.542	0.416	0.485	1.564	0.401	7.733	1.574	0.458
TrafficSim+opt	0.042	1.444	0.325	<b>0.610</b>	1.634	0.404	0.075	1.063	<b>0.265</b>	0.423	1.836	0.413	2.414	1.766	0.532
BITS	0.188	1.068	<b>0.256</b>	1.054	0.968	<b>0.246</b>	<b>0.038</b>	1.220	0.314	0.432	1.131	<b>0.296</b>	4.493	1.152	<b>0.332</b>
BITS+opt	<b>0.033</b>	1.190	0.487	0.681	1.434	0.542	<b>0.028</b>	1.617	0.354	0.435	1.261	0.343	<b>1.533</b>	1.280	0.442
CTG (w/o f+g)	1.380	<b>0.396</b>	0.301	2.662	<b>0.396</b>	0.301	0.172	<b>0.396</b>	0.301	<b>0.369</b>	<b>0.396</b>	<b>0.301</b>	7.052	<b>0.396</b>	<b>0.301</b>
CTG	<b>0.019</b>	<b>0.359</b>	<b>0.165</b>	<b>0.150</b>	<b>0.855</b>	<b>0.179</b>	0.040	<b>0.569</b>	<b>0.271</b>	<b>0.341</b>	<b>0.501</b>	0.455	<b>1.943</b>	<b>0.564</b>	0.387

TABLE III: Quantitative results (Multiple Rules).

	stop sign + no off-road				goal waypoint + target speed			
	rule1	rule2	real	fail	rule1	rule2	real	fail
SimNet	2.282	0.454	0.898	0.353	3.803	1.610	0.898	0.353
SimNet+opt	1.527	0.480	0.796	0.443	3.238	0.980	1.189	0.659
TrafficSim	2.670	0.484	1.605	0.409	<b>1.583</b>	4.544	1.369	0.450
TrafficSim+opt	<b>0.849</b>	0.405	1.577	<b>0.281</b>	2.817	<b>0.995</b>	1.398	0.529
BITS	2.023	0.434	1.032	<b>0.254</b>	2.677	1.077	0.919	<b>0.286</b>
BITS+opt	1.299	0.471	1.083	0.399	3.171	1.019	1.455	0.601
CTG (w/o f+g)	2.573	<b>0.369</b>	<b>0.396</b>	0.301	3.868	2.202	<b>0.396</b>	<b>0.301</b>
CTG	<b>0.528</b>	<b>0.338</b>	<b>0.288</b>	0.326	<b>1.205</b>	<b>0.231</b>	<b>0.738</b>	0.329

compare to baselines in single-rule (Sec. IV-B) and multi-rule (Sec. IV-C) settings both quantitatively and qualitatively, and finally conduct an ablation study (Sec. IV-D).

### A. Experimental Setup

**Datasets.** nuScenes [19] is a large-scale real-world driving dataset, which consists of 5.5 hours of accurate trajectories across two cities with diverse scenarios and dense traffic. We train all models on scenes from the train split and evaluate on 100 scenes randomly sampled from the validation split. In the current work, we focus only on vehicle simulation and defer other types (*e.g.*, pedestrians, cyclists) to future works. **Metrics.** Our evaluation focuses on controllability, realism, and stability (*i.e.*, avoiding collisions and off-road driving). We consider rule-specific violation metrics (**rule**), detailed in Tab. I, to evaluate *controllability*. Metrics are computed for each scene ( $i$  denotes each vehicle in the scene), then averaged across all testing scenes. To evaluate *realism*, we follow [4] and compare data statistics between generated traffic simulations and ground truth trajectories in the dataset. This comparison is computed via the Wasserstein distance between the normalized histograms of the driving profiles for the simulated and recorded trajectories. We define an aggregated metric – realism deviation (**real**) – as the mean of the realism for three properties from [4]: longitudinal acceleration magnitude, latitudinal acceleration magnitude, and jerk. We further provide failure rate (**fail**) to evaluate the *stability* of generated trajectories. This is measured as the average fraction of agents experiencing a critical failure, *i.e.* collision or road departure, in a scene.

**Baselines.** Since there are no comparable works on rule-compliant traffic generation, we augment state-of-the-art traffic simulation models by adding a test-time optimization to meet specified rules. For a fair comparison, this optimization uses the same loss function as used for guidance in CTG. **SimNet** [34] is a deterministic behavior-cloning model. We apply an optimization on its output action trajectory (**SimNet+opt**). **TrafficSim** [3] is a CVAE-based trajectory generation method. We consider a variant with filtration (**TrafficSim**) using our loss function, and another with both filtration and latent space optimization (**TrafficSim+opt**). **BITS** [4] is a bi-level imitation learning model and we adapt its sampling ranking function to use our loss function (**BITS**). We also use a variant that employs optimization on the output action trajectory (**BITS+opt**). Finally, we compare to CTG *without* filtration and guidance (**CTG w/o f+g**), *i.e.* random samples from the diffusion model.

### B. Single Rule Evaluation

We first evaluate how well methods satisfy a single specified rule, which is crucial for applications such as traffic scene editing. We apply five STL rules formulated in Tab. I: speed limit, target speed, goal waypoint, no collision, and no off-road. For rules that require specific parameters to be set (*e.g.*, the goal waypoint location), we select reasonable values based on the ground truth log in the dataset to avoid setting out-of-distribution values (*e.g.* off-road waypoints). **Speed Limit.** Vehicles should not exceed a speed limit threshold. Since the speed limit of each road is not available in the dataset, we set the limit per scene to be the speed at the 75% quantile of all moving vehicles in that scene. **Target Speed.** Vehicles should follow a specified speed at each time step. For each vehicle, the speed is set to 50% of its speed in the ground truth scene, similar to a traffic jam. **Goal Waypoint.** Vehicles should reach a specified waypoint at any time in the future. We set waypoints to be the position at 15s along the ground truth data trajectory for each vehicle.




Fig. 3: Target speed (a)-(c) and goal waypoint (d)-(f) rule results. Rule violations are in parentheses ( $m/s$  and  $m$ , respectively).

Hitting a waypoint from ground truth data is not trivial since sampled trajectories often greatly deviate from the dataset.

**No Collision.** Vehicles should not collide with each other.

**No Off-road.** Vehicles should not leave the drivable area.

Quantitative results are shown in Table II. In general, CTG achieves lower values for rule violation, realism deviation, and failure rate than the baselines. Among all five settings, CTG has the lowest rule violation in three and is competitive in the others. For realism deviation and failure rate, CTG is usually top two. Fig. 3 shows qualitative results for target speed and waypoint rules. Compared to TrafficSim+opt and BC+opt, the strongest baselines for target speed, CTG has the lowest rule violation using more realistic trajectories. For the waypoint example, although BC+opt provides better rule satisfaction than CTG, both BC+opt and BITS+opt predict curvy, unrealistic trajectories resulting in multiple collisions.

### C. Multiple Rules Evaluation

**Stop Sign and No Off-road.** Vehicles should stop if they enter a stop sign region and not go off-road. This rule (see Tab. I) is relatively involved using “Implies” ( $\Rightarrow$ ) and “Eventually” ( $\diamond$ ) operators, making it a good test for our STL-based approach. Stop regions are  $20 \times 20$  m boxes with centers set to be 5s along the ground truth data trajectories. **Goal Waypoint and Target Speed.** Vehicles should reach their goal following the specified target speeds. Waypoints are set 10s along the ground truth trajectories, and the target speeds at each time step are the same as the ground truth scene. This setting is similar to a “reactive replay” use case in AV testing, where a user reconstructs a driving log using the traffic model, which allows agents to realistically react to any subsequent changes to AV behavior or environment.

Results are shown in Table III. For both settings, CTG variants achieve the top two lowest rule violation and realism deviation, with only slightly higher failure rates.

### D. Ablation Study

To analyze design choices, we conduct an ablation study under the speed limit rule setting. Results are shown in

Table IV where the top row is our proposed version of CTG. The first section compares different combinations of guidance (g) and filtration (f). Without guidance, rule violation increases greatly. Filtration is more effective paired with guidance than by itself, as it can choose the best from several already-guided samples that may satisfy rules to differing degrees. In the next part of the table, variants using an additional output action optimization (a) are evaluated. Replacing guidance with optimization is worse on all metrics, while combining guidance with optimization reduces rule violation at the cost of higher realism deviation and failure rate. Next, more inner optimization steps (op) are used in guidance, improving rule violation but giving worse realism and failure rate. Finally, in the bottom section, we evaluate a variant that only supervises the action trajectory, and variants where unicycle dynamics (dyn) are not enforced. Supervising only actions (instead of states and actions) gives more faithful accelerations resulting in lower realism deviation, but failure is more frequent without state supervision to regularize. Additionally, we find that enforcing dynamics using the unicycle model is key: all metrics degrade without this.

TABLE IV: Ablation study of CTG features: dyn (dynamics enforced), f (filtration), g (guidance), a (action trajectory post-optimization), and op (number of inner optimization steps).

Features					Loss	Metrics		
dyn	f	g	a	op	action only?	rule	real	fail
✓	✓	✓		1		0.0189	0.3588	<b>0.1647</b>
✓				1		1.3802	0.3964	0.3011
✓		✓		1		0.0387	0.4336	0.2227
✓	✓			1		1.3268	0.3869	0.2942
✓		✓	✓	1		0.019	0.447	0.1879
✓	✓	✓	✓	1		<b>0.0132</b>	0.4431	0.1845
✓	✓		✓	1		0.0329	0.4591	0.1881
✓			✓	1		0.0819	0.4339	0.2117
✓	✓	✓		3		0.0158	0.4591	0.2017
✓	✓	✓		5		0.0156	0.456	0.1966
✓	✓	✓		1	✓	0.018	<b>0.2057</b>	0.2824
		✓		1		0.0525	3.4204	0.612
		✓		1		1.2797	3.3251	0.7523

## V. CONCLUSION

We proposed CTG, a conditional diffusion model that enables simulating realistic and physically-feasible vehicle trajectories that are also user-controllable through STL rules. CTG opens several exciting future research directions. Currently, it only models vehicles, but cyclists and pedestrians are also important agents to simulate. An improved model must consider interactions between all these agent types, and needs to capture different behavior distributions and guidance rules for pedestrians and cyclists. Additionally, using collision and off-road guidance to enable very long-term, robust traffic simulation is an important application. Outside of AV, the proposed guidance framework may benefit many tasks where learned models in-the-loop must be reactive and follow novel objectives online.

## VI. ACKNOWLEDGMENTS

The authors thank Or Litany, Sanja Fidler, and Karen Leung for valuable discussions and feedback.

## REFERENCES

- [1] N. Kalra and S. M. Paddock, *Driving to Safety: How Many Miles of Driving Would It Take to Demonstrate Autonomous Vehicle Reliability?* RAND Corporation, 2016. [Online]. Available: <http://www.jstor.org/stable/10.7249/j.ctt1btc0xw>
- [2] Waymo, "Waymo safety report," <https://storage.googleapis.com/waymo-uploads/files/documents/safety/2021-03-waymo-safety-report.pdf>, February 2021.
- [3] S. Suo, S. Regalado, S. Casas, and R. Urtasun, "TrafficSim: Learning to simulate realistic multi-agent behaviors," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021, pp. 10400–10409.
- [4] D. Xu, Y. Chen, B. Ivanovic, and M. Pavone, "Bits: Bi-level imitation for traffic simulation," 2022.
- [5] P. A. Lopez, M. Behrisch, L. Bieker-Walz, J. Erdmann, Y.-P. Flötteröd, R. Hilbrich, L. Lücken, J. Rummel, P. Wagner, and E. Wiessner, "Microscopic traffic simulation using sumo," in *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, 2018, pp. 2575–2582.
- [6] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "CARLA: An open urban driving simulator," ser. Proceedings of Machine Learning Research, S. Levine, V. Vanhoucke, and K. Goldberg, Eds., vol. 78, 13–15 Nov 2017, pp. 1–16. [Online]. Available: <http://proceedings.mlr.press/v78/dosovitskiy17a.html>
- [7] LG Electronics, "SVL Simulator: An Autonomous Vehicle Simulator, A ROS/ROS2 Multi-robot Simulator for Autonomous Vehicles," <https://github.com/lgsvl/simulator>, 2021.
- [8] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds., vol. 33. Curran Associates, Inc., 2020, pp. 6840–6851.
- [9] P. Dhariwal and A. Nichol, "Diffusion models beat gans on image synthesis," *Advances in Neural Information Processing Systems*, vol. 34, pp. 8780–8794, 2021.
- [10] N. Chen, Y. Zhang, H. Zen, R. J. Weiss, M. Norouzi, and W. Chan, "Wavegrad: Estimating gradients for waveform generation," in *International Conference on Learning Representations*, 2021.
- [11] Z. Kong, W. Ping, J. Huang, K. Zhao, and B. Catanzaro, "Diffwave: A versatile diffusion model for audio synthesis," in *International Conference on Learning Representations*, 2021.
- [12] V. Popov, I. Vovk, V. Gogoryan, T. Sadekova, and M. Kudinov, "GradTts: A diffusion probabilistic model for text-to-speech," in *Proceedings of the 38th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, vol. 139. PMLR, 18–24 Jul 2021, pp. 8599–8608.
- [13] T. Gu, G. Chen, J. Li, C. Lin, Y. Rao, J. Zhou, and J. Lu, "Stochastic trajectory prediction via motion indeterminacy diffusion," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2022, pp. 17113–17122.
- [14] X. L. Li, J. Thickstun, I. Gulrajani, P. Liang, and T. B. Hashimoto, "Diffusion-lm improves controllable text generation," 2022.
- [15] M. Janner, Y. Du, J. B. Tenenbaum, and S. Levine, "Planning with diffusion for flexible behavior synthesis," in *International Conference on Machine Learning*, 2022.
- [16] O. Maler and D. Nickovic, "Monitoring temporal properties of continuous signals," in *Formal Techniques, Modelling and Analysis of Timed and Fault-Tolerant Systems*, Y. Lakhnech and S. Yovine, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2004, pp. 152–166.
- [17] K. Leung, N. Aréchiga, and M. Pavone, "Backpropagation through signal temporal logic specifications: Infusing logical structure into gradient-based methods," *The International Journal of Robotics Research*.
- [18] A. Paszke et. al, "Pytorch: An imperative style, high-performance deep learning library," in *Advances in Neural Information Processing Systems* 32, 2019, pp. 8024–8035.
- [19] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "nusenes: A multimodal dataset for autonomous driving," in *CVPR*, 2020.
- [20] E. Brockfeld, R. D. Kühne, A. Skabardonis, and P. Wagner, "Toward benchmarking of microscopic traffic flow models," *Transportation Research Record*, vol. 1852, no. 1, pp. 124–129, 2003.
- [21] Y. Chen, B. Ivanovic, and M. Pavone, "Scept: Scene-consistent, policy-based trajectory predictions for planning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2022, pp. 17103–17112.
- [22] T. Salzmann, B. Ivanovic, P. Chakravarthy, and M. Pavone, "Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data," in *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVIII 16*. Springer, 2020, pp. 683–700.
- [23] J. Wang, A. Pun, J. Tu, S. Manivasagam, A. Sadat, S. Casas, M. Ren, and R. Urtasun, "AdvSim: Generating safety-critical scenarios for self-driving vehicles," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 9909–9918.
- [24] D. Rempe, J. Phillion, L. J. Guibas, S. Fidler, and O. Litany, "Generating useful accident-prone driving scenarios via a learned traffic prior," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [25] Y. Abeyirigoonawardena, F. Shkurti, and G. Dudek, "Generating adversarial driving scenarios in high-fidelity simulators," in *2019 International Conference on Robotics and Automation (ICRA)*, 2019, pp. 8271–8277.
- [26] B. Chen, X. Chen, Q. Wu, and L. Li, "Adversarial evaluation of autonomous vehicles in lane-change scenarios," *IEEE Transactions on Intelligent Transportation Systems*, pp. 1–10, 2021.
- [27] J. Ho and T. Salimans, "Classifier-free diffusion guidance," *arXiv preprint arXiv:2207.12598*, 2022.
- [28] J. Ho, T. Salimans, A. Gritsenko, W. Chan, M. Norouzi, and D. J. Fleet, "Video diffusion models," *arXiv preprint arXiv:2204.03458*, 2022.
- [29] K. Leung and M. Pavone, "Semi-supervised trajectory-feedback controller synthesis for signal temporal logic specifications," 2022.
- [30] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli, "Deep unsupervised learning using nonequilibrium thermodynamics," in *Proceedings of the 32nd International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, F. Bach and D. Blei, Eds., vol. 37. Lille, France: PMLR, 07–09 Jul 2015, pp. 2256–2265.
- [31] A. Q. Nichol and P. Dhariwal, "Improved denoising diffusion probabilistic models," in *International Conference on Machine Learning*. PMLR, 2021, pp. 8162–8171.
- [32] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [33] L. Bergamini, Y. Ye, O. Scheel, L. Chen, C. Hu, L. Del Pero, B. Osiński, H. Grimmet, and P. Ondruska, "Simnet: Learning reactive self-driving simulations from real-world observations," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021.