

SceneGen: Learning to Generate Realistic Traffic Scenes

Shuhan Tan^{1,2*} Kelvin Wong^{1,3*} Shenlong Wang^{1,3}

Sivabalan Manivasagam^{1,3} Mengye Ren^{1,3} Raquel Urtasun^{1,3}

¹Uber Advanced Technologies Group ²Sun Yat-Sen University ³University of Toronto

tanshh@mail2.sysu.edu.cn {kelvinwong, slwang, manivasagam, mren, urtasun}@cs.toronto.edu

Abstract

We consider the problem of generating realistic traffic scenes automatically. Existing methods typically insert actors into the scene according to a set of hand-crafted heuristics and are limited in their ability to model the true complexity and diversity of real traffic scenes, thus inducing a content gap between synthesized traffic scenes versus real ones. As a result, existing simulators lack the fidelity necessary to train and test self-driving vehicles. To address this limitation, we present SceneGen—a neural autoregressive model of traffic scenes that eschews the need for rules and heuristics. In particular, given the ego-vehicle state and a high definition map of surrounding area, SceneGen inserts actors of various classes into the scene and synthesizes their sizes, orientations, and velocities. We demonstrate on two large-scale datasets SceneGen’s ability to faithfully model distributions of real traffic scenes. Moreover, we show that SceneGen coupled with sensor simulation can be used to train perception models that generalize to the real world.

1. Introduction

The ability to simulate realistic traffic scenarios is an important milestone on the path towards safe and scalable self-driving. It enables us to build rich virtual environments in which we can improve our self-driving vehicles (SDVs) and verify their safety and performance [9, 31, 32, 53]. This goal, however, is challenging to achieve. As a first step, most large-scale self-driving programs simulate pre-recorded scenarios captured in the real world [32] or employ teams of test engineers to design new scenarios [9, 31]. Although this approach can yield realistic simulations, it is ultimately not scalable. This motivates the search for a way to generate realistic traffic scenarios *automatically*.

More concretely, we are interested in generating the layout of actors in a traffic scene given the SDV’s current state and a high definition map (HD map) of the surround-

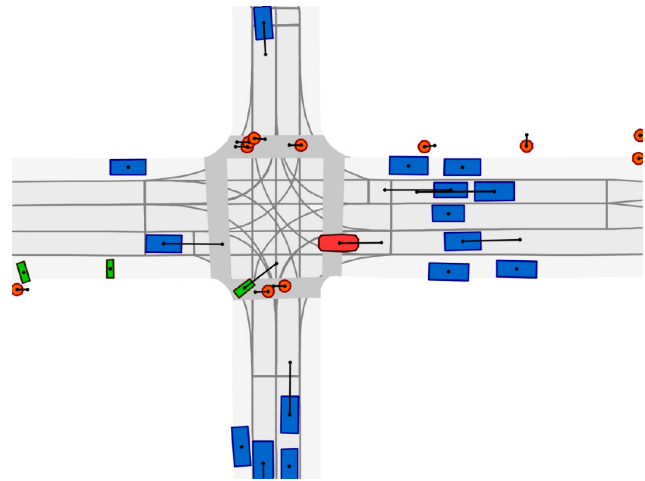


Figure 1: Given the SDV’s state and an HD map, SceneGen autoregressively inserts actors onto the map to compose a realistic traffic scene. The ego SDV is shown in red; vehicles in blue; pedestrians in orange; and bicyclists in green.

ing area. We call this task *traffic scene generation* (see Fig. 1). Here, each actor is parameterized by a class label, a bird’s eye view bounding box, and a velocity vector. Our lightweight scene parameterization is popular among existing self-driving simulation stacks and can be readily used in downstream modules; *e.g.*, to simulate LiDAR [9, 10, 32].

A popular approach to traffic scene generation is to use procedural models to insert actors into the scene according to a set of rules [55, 31, 9, 37]. These rules encode reasonable heuristics such as “pedestrians should stay on the sidewalk” or “vehicles should drive along lane centerlines”, and their parameters can be manually tuned to give reasonable results. Still, these simplistic heuristics cannot fully capture the complexity and diversity of real world traffic scenes, thus inducing a content gap between synthesized traffic scenes and real ones [26]. Moreover, this approach requires significant time and expertise to design good heuristics and tune their parameters.

To address these issues, recent methods use machine learning techniques to automatically tune model param-

*Indicates equal contribution. Work done at Uber ATG.

ters [52, 51, 24, 26, 8]. These methods improve the realism and scalability of traffic scene generation. However, they remain limited by their underlying hand-crafted heuristics and priors; *e.g.*, pre-defined scene grammars or assumptions about road topologies. As a result, they lack the capacity to model the true complexity and diversity of real traffic scenes and, by extension, the fidelity necessary to train and test SDVs in simulation. Alternatively, we can use a simple data-driven approach by sampling from map-specific empirical distributions [10]. But this cannot generalize to new maps and may yield scene-inconsistent samples.

In this paper, we propose SceneGen—a traffic scene generation model that eschews the need for hand-crafted rules and heuristics. Our approach is inspired by recent successes in deep generative modeling that have shown remarkable results in estimating distributions of a variety of data, without requiring complex rules and heuristics; *e.g.*, handwriting [18], images [49], text [39], *etc.* Specifically, SceneGen is a neural autoregressive model that, given the SDV’s current state and an HD map of the surrounding area, sequentially inserts actors into the scene—mimicking the process by which humans do this as well. As a result, we can sample realistic traffic scenes from SceneGen and compute the likelihood of existing ones as well.

We evaluate SceneGen on two large-scale self-driving datasets. The results show that SceneGen can better estimate the distribution over real traffic scenes than competing baselines and generate more realistic samples as well. Furthermore, we show that SceneGen coupled with sensor simulation can generate realistic labeled data to train perception models that generalize to the real world. With SceneGen, we take an important step towards developing SDVs safely and scalably through large-scale simulation. We hope our work here inspires more research along this direction so that one day this goal will become a reality.

2. Related Work

Traffic simulation: The study of traffic simulation can be traced back to at least the 1950s with Gerlough’s dissertation on simulating freeway traffic flow [16]. Since then, various traffic models have been used for simulation. Macroscopic models simulate entire populations of vehicles in the aggregate [30, 40] to study “macroscopic” properties of traffic flow, such as traffic density and average velocity. In contrast, microscopic models simulate the behavior of each individual vehicle over time by assuming a car-following model [36, 6, 34, 13, 17, 1, 44]. These models improve simulation fidelity considerably but at the cost of computational efficiency. Microscopic traffic models have been included in popular software packages such as SUMO [31], CORSIM [35], VISSIM [11], and MITSIM [55].

Recently, traffic simulation has found new applications in testing and training the autonomy stack of SDVs. How-

ever, existing simulators do not satisfy the level of realism necessary to properly test SDVs [52]. For example, the CARLA simulator [9] spawns actors at pre-determined locations and uses a lane-following controller to simulate the vehicle behaviors over time. This approach is too simplistic and so it induces a sim2real content gap [26]. Therefore, in this paper, we study how to generate snapshots of traffic scenes that mimic the realism and diversity of real ones.

Traffic scene generation: While much of the research into microscopic traffic simulation have focused on modeling actors’ behaviors, an equally important yet underexplored problem is how to generate realistic snapshots of traffic scenes. These snapshots have many applications; *e.g.*, to initialize traffic simulations [52] or to generate labeled data for training perception models [26]. A popular approach is to procedurally insert actors into the scene according to a set of rules [55, 31, 9, 37]. These rules encode reasonable heuristics such as “pedestrians should stay on the sidewalk” and “vehicles should drive along lane centerlines”, and their parameters can be manually tuned to give reasonable results. For example, SUMO [31] inserts vehicles into lanes based on minimum headway requirements and initializes their speeds according to a Gaussian distribution [52]. Unfortunately, it is difficult to scale this approach to new environments since tuning these heuristics require significant time and expertise.

An alternative approach is to learn a probabilistic distribution over traffic scenes from which we can sample new scenes [52, 51, 24, 10, 14, 15, 57]. For example, Wheeler *et al.* [52] propose a Bayesian network to model a joint distribution over traffic scenes in straight multi-lane highways. This approach was extended to model inter-lane dependencies [51] and generalized to handle a four-way intersection [24]. These models are trained to mimic a real distribution over traffic scenes. However, they consider a limited set of road topologies only and assume that actors follow reference paths in the map. As a result, they are difficult to generalize to real urban scenes, where road topologies and actor behaviors are considerably more complex; *e.g.*, pedestrians do not follow reference paths in general.

Recent advances in deep learning have enabled a more flexible approach to learn a distribution over traffic scenes. In particular, MetaSim [26] augments the probabilistic scene graph of Prakash *et al.* [37] with a graph neural network. By modifying the scene graph’s node attributes, MetaSim reduces the content gap between synthesized images versus real ones, without manual tuning. MetaSim2 [8] extends this idea by learning to sample the scene graph as well. Unfortunately, these approaches are still limited by their hand-crafted scene grammar which, for example, constrains vehicles to lane centerlines. We aim to develop a more general method that avoids requiring these heuristics.

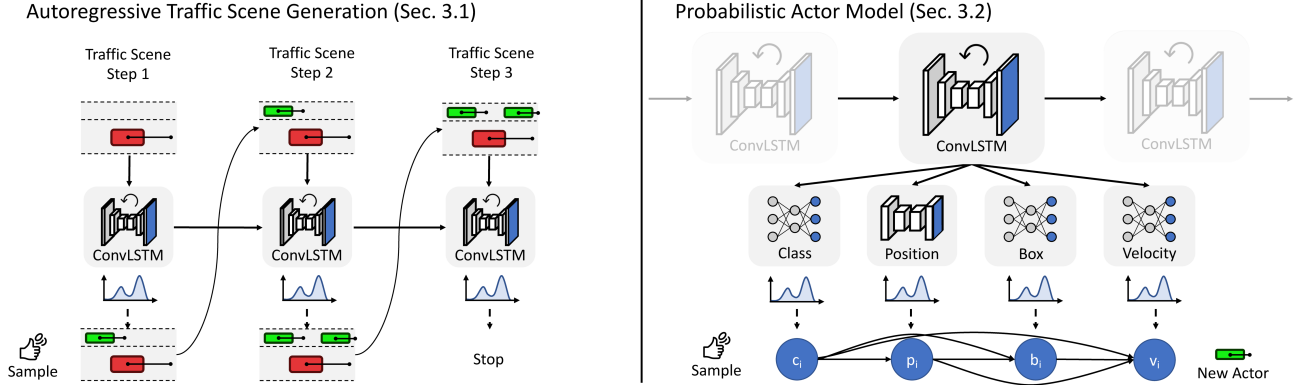


Figure 2: Overview of our approach. Given the ego SDV’s state and an HD map of the surrounding area, SceneGen generates a traffic scene by inserting actors one at a time (Sec. 3.1). We model each actor $\mathbf{a}_i \in \mathcal{A}$ probabilistically, as a product over distributions of its class $c_i \in \mathcal{C}$, position $\mathbf{p}_i \in \mathbb{R}^2$, bounding box $\mathbf{b}_i \in \mathbb{B}$, and velocity $\mathbf{v}_i \in \mathbb{R}^2$ (Sec. 3.2).

Autoregressive models: Autoregressive models factorize a joint distribution over n -dimensions into a product of conditional distributions $p(x) = \prod_{i=1}^n p(x_i|x_{<i})$. Each conditional distribution is then approximated with a parameterized function [12, 2, 45, 46, 47]. Recently, neural autoregressive models have found tremendous success in modeling a variety of data, including handwriting [18], images [49], audio [48], text [39], sketches [20], graphs [29], 3D meshes [33], indoor scenes [50] and image scene layouts [25]. These models are particularly popular since they can factorize a complex joint distribution into a product of much simpler conditional distributions. Moreover, they generally admit a tractable likelihood, which can be used for likelihood-based training, uncovering interesting/outlier examples, *etc.* Inspired by these advances, we exploit autoregressive models for traffic scene generation as well.

3. Traffic Scene Generation

Our goal is to learn a distribution over traffic scenes from which we can sample new examples and evaluate the likelihood of existing ones. In particular, given the SDV $\mathbf{a}_0 \in \mathcal{A}$ and an HD map $\mathbf{m} \in \mathcal{M}$, we aim to estimate the joint distribution over other actors in the scene $\{\mathbf{a}_1, \dots, \mathbf{a}_n\} \subset \mathcal{A}$,

$$p(\mathbf{a}_1, \dots, \mathbf{a}_n | \mathbf{m}, \mathbf{a}_0) \quad (1)$$

The HD map $\mathbf{m} \in \mathcal{M}$ is a collection of polygons and polylines that provide semantic priors for a region of interest around the SDV; *e.g.*, lane boundaries, drivable areas, traffic light states. These priors provide important contextual information about the scene and allow us to generate actors that are consistent with the underlying road topology.

We parameterize each actor $\mathbf{a}_i \in \mathcal{A}$ with an eight-dimensional random variable containing its class label $c_i \in \mathcal{C}$, its bird’s eye view location $(x_i, y_i) \in \mathbb{R}^2$, its bounding box $\mathbf{b}_i \in \mathbb{B}^1$, and its velocity $\mathbf{v}_i \in \mathbb{R}^2$. Each bounding box

¹Pedestrians are not represented by bounding boxes. They are repre-

$\mathbf{b}_i \in \mathbb{B}$ is a 3-tuple consisting of the bounding box’s size $(w_i, l_i) \in \mathbb{R}_{>0}^2$ and heading angle $\theta_i \in [0, 2\pi)$. In our experiments, \mathcal{C} consists of three classes: vehicles, pedestrians, and bicyclists. See Fig. 1 for an example.

Modeling Eq. 1 is a challenging task since the actors in a given scene are highly correlated among themselves and with the map, and the number of actors in the scene is random as well. We aim to model Eq. 1 such that our model is easy to sample from and the resulting samples reflect the complexity and diversity of real traffic scenes. Our approach is to autoregressively factorize Eq. 1 into a product of conditional distributions. This yields a natural generation process that sequentially inserts actors into the scene one at a time. See Fig. 2 for an overview of our approach.

In the following, we first describe our autoregressive factorization of Eq. 1 and how we model this with a recurrent neural network (Sec. 3.1). Then, in Sec. 3.2, we describe how SceneGen generates a new actor at each step of the generation process. Finally, in Sec. 3.3, we discuss how we train and sample from SceneGen.

3.1. The Autoregressive Generation Process

Given the SDV $\mathbf{a}_0 \in \mathcal{A}$ and an HD map $\mathbf{m} \in \mathcal{M}$, our goal is to estimate a conditional distribution over the actors in the scene $\{\mathbf{a}_1, \dots, \mathbf{a}_n\} \subset \mathcal{A}$. As we alluded to earlier, modeling this conditional distribution is challenging since the actors in a given scene are highly correlated among themselves and with the map, and the number of actors in the scene is random. Inspired by the recent successes of neural autoregressive models [18, 49, 39], we propose to autoregressively factorize $p(\mathbf{a}_1, \dots, \mathbf{a}_n | \mathbf{m}, \mathbf{a}_0)$ into a product of simpler conditional distributions. This factorization simplifies the task of modeling the complex joint distribution $p(\mathbf{a}_1, \dots, \mathbf{a}_n | \mathbf{m}, \mathbf{a}_0)$ and results in a model with a tractable likelihood. Moreover, it yields a natural generation process

represented by a single point indicating their center of gravity.

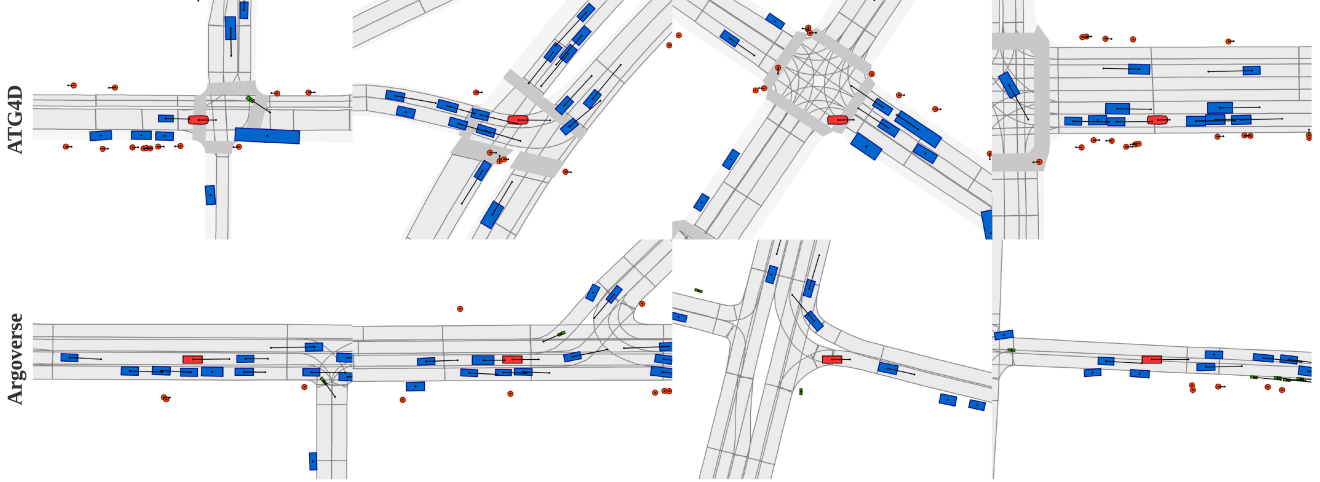


Figure 3: Traffic scenes generated by SceneGen conditioned on HD maps from ATG4D (top) and Argoverse (bottom).

that mimics how a human might perform this task as well.

In order to perform this factorization, we assume a fixed canonical ordering over the sequence of actors $\mathbf{a}_1, \dots, \mathbf{a}_n$,

$$p(\mathbf{a}_1, \dots, \mathbf{a}_n | \mathbf{m}, \mathbf{a}_0) = p(\mathbf{a}_1 | \mathbf{m}, \mathbf{a}_0) \prod_{i=1}^n p(\mathbf{a}_i | \mathbf{a}_{<i}, \mathbf{m}, \mathbf{a}_0) \quad (2)$$

where $\mathbf{a}_{<i} = \{\mathbf{a}_1, \dots, \mathbf{a}_{i-1}\}$ is the set of actors up to and including the $i-1$ -th actor in canonical order. In our experiments, we choose a left-to-right, top-to-bottom order based on each actor’s position in bird’s eye view coordinates. We found that this intuitive ordering works well in practice.

Since the number of actors per scene is random, we introduce a stopping token \perp to indicate the end of our sequential generation process. In practice, we treat \perp as an auxiliary actor that, when generated, ends the generation process. Therefore, for simplicity of notation, we assume that the last actor \mathbf{a}_n is always the stopping token \perp .

Model architecture: Our model uses a recurrent neural network to capture the long-range dependencies across our autoregressive generation process. The basis of our model is the ConvLSTM architecture [42]—an extension of the classic LSTM architecture [22] to spatial data—and the input to our model at the i -th generation step is a bird’s eye view multi-channel image encoding the SDV \mathbf{a}_0 , the HD map \mathbf{m} , and the actors generated so far $\{\mathbf{a}_1, \dots, \mathbf{a}_{i-1}\}$.

For the i -th step of the generation process: Let $\mathbf{x}^{(i)} \in \mathbb{R}^{C \times H \times W}$ denote the multi-channel image, where C is the number of feature channels and $H \times W$ is the size of the image grid. Given the previous hidden and cell states $\mathbf{h}^{(i-1)}$ and $\mathbf{c}^{(i-1)}$, the new hidden and cell states are given by:

$$\mathbf{h}^{(i)}, \mathbf{c}^{(i)} = \text{ConvLSTM}(\mathbf{x}^{(i)}, \mathbf{h}^{(i-1)}, \mathbf{c}^{(i-1)}; \mathbf{w}) \quad (3)$$

$$\mathbf{f}^{(i)} = \text{CNN}_b(\mathbf{h}^{(i)}; \mathbf{w}) \quad (4)$$

where ConvLSTM is a two-layer ConvLSTM, CNN_b is a five-layer convolutional neural network (CNN) that extract backbone features, and \mathbf{w} are the neural network parameters. The features $\mathbf{f}^{(i)}$ summarize the generated scene so far $\mathbf{a}_{<i}$, \mathbf{a}_0 , and \mathbf{m} , and we use $\mathbf{f}^{(i)}$ to predict the conditional distribution $p(\mathbf{a}_i | \mathbf{a}_{<i}, \mathbf{m}, \mathbf{a}_0)$, which we describe next. See our appendix for details.

3.2. A Probabilistic Model of Actors

Having specified the generation process, we now turn our attention to modeling each actor probabilistically. As discussed earlier, each actor $\mathbf{a}_i \in \mathcal{A}$ is parameterized by its class label $c_i \in \mathbb{C}$, location $(x_i, y_i) \in \mathbb{R}^2$, oriented bounding box $\mathbf{b}_i \in \mathbb{B}$ and velocity $\mathbf{v}_i \in \mathbb{R}^2$. To capture the dependencies between these attributes, we factorize $p(\mathbf{a}_i | \mathbf{a}_{<i}, \mathbf{m}, \mathbf{a}_0)$ as follows:

$$p(\mathbf{a}_i) = p(c_i)p(x_i, y_i | c_i)p(\mathbf{b}_i | c_i, x_i, y_i)p(\mathbf{v}_i | c_i, x_i, y_i, \mathbf{b}_i) \quad (5)$$

where we dropped the condition on $\mathbf{a}_{<i}$, \mathbf{m} , and \mathbf{a}_0 to simplify notation. Thus, the distribution over an actor’s location is conditional on its class; its bounding box is conditional on its class and location; and its velocity is conditional on its class, location, and bounding box. Note that if \mathbf{a}_i is the stopping token \perp , we do not model its location, bounding box, and velocity. Instead, we have $p(\mathbf{a}_i) = p(c_i)$, where c_i is the auxiliary class c_\perp .

Class: To model a distribution over an actor’s class, we use a categorical distribution whose support is the set of classes $\mathbb{C} \cup \{c_\perp\}$ and whose parameters π_c are predicted by a neural network:

$$\pi_c = \text{MLP}_c(\text{avg-pool}(\mathbf{f}^{(i)}); \mathbf{w}) \quad (6)$$

$$c_i \sim \text{Categorical}(\pi_c) \quad (7)$$

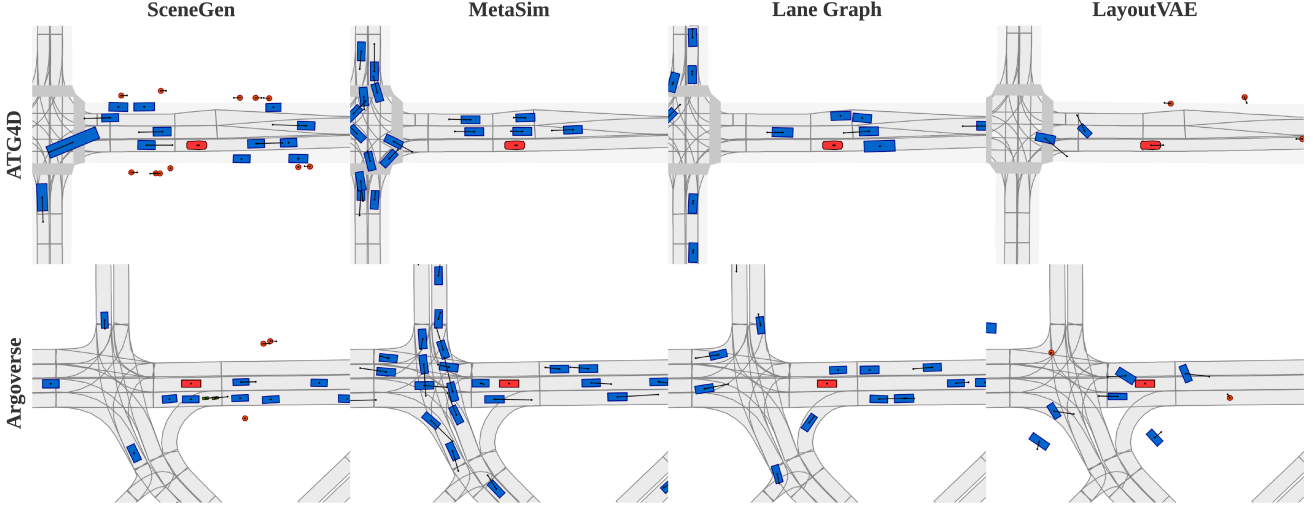


Figure 4: Qualitative comparison of traffic scenes generated by SceneGen and various baselines.

where $\text{avg-pool}: \mathbb{R}^{C \times H \times W} \rightarrow \mathbb{R}^C$ is average pooling over the spatial dimensions and MLP_c is a three-layer multi-layer perceptron (MLP) with softmax activations.

Location: We apply uniform quantization to the actor’s position and model the quantized values using a categorical distribution. The support of this distribution is the set of $H \times W$ quantized bins within our region of interest and its parameters π_{loc} are predicted by a class-specific CNN. This approach allows the model to express highly multi-modal distributions without making assumptions about the distribution’s shape [49]. To recover continuous values, we assume a uniform distribution within each quantization bin.

Let k denote an index into one of the $H \times W$ quantized bins, and suppose $\lfloor \mathbf{p}_k \rfloor \in \mathbb{R}^2$ (*resp.*, $\lceil \mathbf{p}_k \rceil \in \mathbb{R}^2$) is the minimum (*resp.*, maximum) continuous coordinates in the k -th bin. We model $p(x_i, y_i | c_i)$ as follows:

$$\pi_{\text{loc}} = \text{CNN}_{\text{loc}}(\mathbf{f}^{(i)}; c_i, \mathbf{w}) \quad (8)$$

$$k \sim \text{Categorical}(\pi_{\text{loc}}) \quad (9)$$

$$(x_i, y_i) \sim \text{Uniform}(\lfloor \mathbf{p}_k \rfloor, \lceil \mathbf{p}_k \rceil) \quad (10)$$

where $\text{CNN}_{\text{loc}}(\cdot; c_i, \mathbf{w})$ is a CNN with softmax activations for the class c_i . During inference, we mask and re-normalize π_{loc} such that quantized bins with invalid positions according to our canonical ordering have zero probability mass. Note that we do not mask during training since this resulted in worse performance.

After sampling the actor’s location $(x_i, y_i) \in \mathbb{R}^2$, we extract a feature vector $\mathbf{f}_{x_i, y_i}^{(i)} \in \mathbb{R}^C$ by spatially indexing into the k -th bin of $\mathbf{f}^{(i)}$. This feature vector captures local information at (x_i, y_i) and is used to subsequently predict the actor’s bounding box and velocity.

Bounding box: An actor’s bounding box $\mathbf{b}_i \in \mathbb{B}$ consists of its width and height $(w_i, l_i) \in \mathbb{R}_{>0}^2$ and its heading $\theta_i \in [0, 2\pi)$. We model the distributions over each of these independently. For an actor’s bounding box size, we use a mixture of K bivariate log-normal distributions:

$$[\pi_{\text{box}}, \mu_{\text{box}}, \Sigma_{\text{box}}] = \text{MLP}_{\text{box}}(\mathbf{f}_{x_i, y_i}^{(i)}; c_i, \mathbf{w}) \quad (11)$$

$$k \sim \text{Categorical}(\pi_{\text{box}}) \quad (12)$$

$$(w_i, l_i) \sim \text{LogNormal}(\mu_{\text{box}, k}, \Sigma_{\text{box}, k}) \quad (13)$$

where π_{box} are mixture weights, each $\mu_{\text{box}, k} \in \mathbb{R}^2$ and $\Sigma_{\text{box}, k} \in \mathbb{S}_+^2$ parameterize a component log-normal distribution, and $\text{MLP}_{\text{box}}(\cdot; c_i, \mathbf{w})$ is a three-layer MLP for the class c_i . This parameterization allows our model to naturally capture the multi-modality of actor sizes in real world data; *e.g.*, the size of sedans versus trucks.

Similarly, we model an actor’s heading angle with a mixture of K Von-Mises distributions:

$$[\pi_{\theta}, \mu_{\theta}, \kappa_{\theta}] = \text{MLP}_{\theta}(\mathbf{f}_{x_i, y_i}^{(i)}; c_i, \mathbf{w}) \quad (14)$$

$$k \sim \text{Categorical}(\pi_{\theta}) \quad (15)$$

$$\theta_i \sim \text{VonMises}(\mu_{\theta, k}, \kappa_{\theta, k}) \quad (16)$$

where π_{θ} are mixture weights, each $\mu_{\theta, k} \in [0, 2\pi)$ and $\kappa_{\theta, k} > 0$ parameterize a component Von-Mises distribution, and $\text{MLP}_{\theta}(\cdot; c_i, \mathbf{w})$ is a three-layer MLP for the class c_i . The Von-Mises distribution is a close approximation of a normal distribution wrapped around the unit circle [38] and has the probability density function

$$p(\theta | \mu, \kappa) = \frac{e^{\kappa \cos(\theta - \mu)}}{2\pi I_0(\kappa)} \quad (17)$$

where I_0 is the modified Bessel function of order 0. We use a mixture of Von-Mises distributions to capture the multi-modality of headings in real world data; *e.g.*, a vehicle can

Method	ATG4D							Argoverse						
	NLL	Feat.	Loc.	Class	Size	Speed	Head	NLL	Feat.	Loc.	Class	Size	Speed	Head
Prob. Gram.	-	0.20	0.10	0.24	0.46	0.34	0.31	-	0.38	0.14	0.26	0.41	0.57	0.38
MetaSim	-	0.12	0.10	0.24	0.45	0.35	0.15	-	0.18	0.14	0.26	0.50	0.52	0.18
Procedural	-	0.38	0.10	0.24	0.17	0.34	0.07	-	0.58	0.16	0.26	0.23	0.59	0.17
Lane Graph	-	0.17	0.11	0.24	0.30	0.32	0.16	-	0.11	0.16	0.26	0.31	0.32	0.28
LayoutVAE	210.80	0.15	0.09	0.12	0.18	0.33	0.29	200.78	0.25	0.13	0.11	0.21	0.41	0.29
SceneGen	59.86	0.11	0.10	0.20	0.06	0.27	0.08	67.11	0.14	0.13	0.21	0.17	0.17	0.21

Table 1: Negative log-likelihood (NLL) and maximum mean discrepancy (MMD) results on ATG4D and Argoverse. NLL is reported in *nats*, averaged across all scenes in the test set. MMD is computed between distributions of features extracted by a motion forecasting model and various scene statistics (see main text for description). For all metrics, lower is better.

go straight or turn at an intersection. To sample from a mixture of Von-Mises distributions, we sample a component k from a categorical distribution and then sample θ from the Von-Mises distribution of the k -th component [3].

Velocity: We parameterize the actor’s velocity $\mathbf{v}_i \in \mathbb{R}^2$ as $\mathbf{v}_i = (s_i \cos \omega_i, s_i \sin \omega_i)$, where $s_i \in \mathbb{R}_{\geq 0}$ is its speed and $\omega_i \in [0, 2\pi)$ is its direction. Note that this parameterization is not unique since ω_i can take any value in $[0, 2\pi)$ when $\mathbf{v}_i = 0$. Therefore, we model the actor’s velocity as a mixture model where one of the $K \geq 2$ components corresponds to $\mathbf{v}_i = 0$. More concretely, we have

$$\boldsymbol{\pi}_v = \text{MLP}_v(\mathbf{f}_{x_i, y_i}^{(i)}; c_i, \mathbf{w}) \quad (18)$$

$$k \sim \text{Categorical}(\boldsymbol{\pi}_v) \quad (19)$$

where for $k > 0$, we have $\mathbf{v}_i = (s_i \cos \omega_i, s_i \sin \omega_i)$, with

$$[\mu_s, \sigma_s] = \text{MLP}_s(\mathbf{f}_{x_i, y_i}^{(i)}; c_i, \mathbf{w}) \quad (20)$$

$$[\mu_\omega, \kappa_\omega] = \text{MLP}_\omega(\mathbf{f}_{x_i, y_i}^{(i)}; c_i, \mathbf{w}) \quad (21)$$

$$s_i \sim \text{LogNormal}(\mu_{s,k}, \sigma_{s,k}) \quad (22)$$

$$\omega_i \sim \text{VonMises}(\mu_{\omega,k}, \kappa_{\omega,k}) \quad (23)$$

and for $k = 0$, we have $\mathbf{v}_i = 0$. As before, we use three-layer MLPs to predict the parameters of each distribution.

For vehicles and bicyclists, we parameterize $\omega_i \in [0, 2\pi)$ as an offset relative to the actor’s heading $\theta_i \in [0, 2\pi)$. This is equivalent to parameterizing their velocities with a bicycle model [43], which we found improves sample quality.

3.3. Learning and Inference

Sampling: Pure sampling from deep autoregressive models can lead to degenerate examples due to their “unrealistic long tails” [23]. Therefore, we adopt a sampling strategy inspired by *nucleus sampling* [23]. Specifically, at each generation step, we sample from each of SceneGen’s output distributions M times and keep the most likely sample. We found this to help avoid degenerate traffic scenes while

maintaining sample diversity. Furthermore, we reject vehicles and bicyclists whose bounding boxes collide with those of the actors sampled so far.

Training: We train our model to maximize the log-likelihood of real traffic scenes in our training dataset:

$$\mathbf{w}^* = \arg \max_{\mathbf{w}} \sum_{i=1}^N \log p(\mathbf{a}_{i,1}, \dots, \mathbf{a}_{i,n} | \mathbf{m}_i, \mathbf{a}_{i,0}; \mathbf{w}) \quad (24)$$

where \mathbf{w} are the neural network parameters and N is the number of samples in our training set. In practice, we use the Adam optimizer [27] to minimize the average negative log-likelihood over mini-batches. We use teacher forcing and backpropagation-through-time to train through the generation process, up to a fixed window as memory allows.

4. Experiments

We evaluate SceneGen on two self-driving datasets: Argoverse [7] and ATG4D [54]. Our results show that SceneGen can generate more realistic traffic scenes than the competing methods (Sec. 4.3). We also demonstrate how SceneGen with sensor simulation can be used to train perception models that generalize to the real world (Sec. 4.4).

4.1. Datasets

ATG4D: ATG4D [54] is a large-scale dataset collected by a fleet of SDVs in cities across North America. It consists of 5500 25-second logs which we split into a training set of 5000 and an evaluation set of 500. Each log is subsampled at 10Hz to yield 250 traffic scenes, and each scene is annotated with bounding boxes for vehicles, pedestrians, and bicyclists. Each log also provides HD maps that encode lane boundaries, drivable areas, and crosswalks as polygons, and lane centerlines as polylines. Each lane segment is annotated with attributes such as its type (car vs. bike), turn direction, boundary colors, and traffic light state.

# Mixtures	Scene	Vehicle	Pedestrian	Bicyclist
1	125.97	7.26	10.36	9.16
3	68.22	2.64	8.52	7.34
5	64.05	2.35	8.27	7.22
10	59.86	1.94	8.32	6.90

Table 2: Ablation of the number of mixture components in ATG4D. Scene NLL is averaged across scenes and NLL per class is the average NLL per actor of that class.

L	DA	C	TL	Scene	Veh.	Ped.	Bic.
				93.73	4.90	8.85	7.17
✓				63.33	2.12	8.69	7.10
✓	✓			57.66	1.73	8.40	6.84
✓	✓	✓		57.96	1.77	8.32	6.61
✓	✓	✓	✓	59.86	1.94	8.32	6.90

Table 3: Ablation of map in ATG4D (in NLL). **L** is lanes; **DA** drivable areas; **C** crosswalks; and **TL** traffic lights.

In our experiments, we subdivide the training set into two splits of 4000 and 1000 logs respectively. We use the first split to train the traffic scene generation models and the second split to train the perception models in Sec. 4.4.

Argoverse: Argoverse [7] consists of two datasets collected by a fleet of SDVs in Pittsburgh and Miami. We use the Argoverse 3D Tracking dataset which contains track annotations for 65 training logs and 24 validation logs. Each log is subsampled at 10Hz to yield 13,122 training scenes and 5015 validation scenes. As in ATG4D, Argoverse provides HD maps annotated with drivable areas and lane segment centerlines and their attributes; *e.g.*, turn direction. However, Argoverse does not provide crosswalk polygons, lane types, lane boundary colors, and traffic lights.

4.2. Experiment Setup

Baselines: Our first set of baselines is inspired by recent work on probabilistic scene grammars and graphs [37, 26, 8]. In particular, we design a probabilistic grammar of traffic scenes (**Prob. Grammar**) such that actors are randomly placed onto lane segments using a hand-crafted prior [37]. Sampling from this grammar yields a *scene graph*, and our next baseline (**MetaSim**) uses a graph neural network to transform the attributes of each actor in the scene graph. Our implementation follows Kar *et al.* [26], except we use a training algorithm that is supervised with heuristically generated ground truth scene graphs.²

Our next set of baselines is inspired by methods that reason directly about the road topology of the scene [52, 51, 24, 32]. Given a *lane graph* of the scene, **Procedural** uses

²We were unable to train MetaSim using their unsupervised losses.

a set of rules to place actors such that they follow lane centerlines, maintain a minimum clearance to leading actors, *etc.* Each actor’s bounding box is sampled from a Gaussian KDE fitted to the training dataset [5] and velocities are set to satisfy speed limits and a time gap between successive actors on the lane graph. Similar to MetaSim, we also consider a learning-based version of Procedural that uses a lane graph neural network [28] to transform each actor’s position, bounding box, and velocity (**Lane Graph**).

Since the HD maps in ATG4D and Argoverse do not provide reference paths for pedestrians, the aforementioned baselines cannot generate pedestrians.³ Therefore, we also compare against **LayoutVAE** [25]—a variational autoencoder for image layouts that we adapt for traffic scene generation. We modify LayoutVAE to condition on HD maps and output oriented bounding boxes and velocities for actors of every class. Please see our appendix for details.

Metrics: Our first metric measures the negative log-likelihood (**NLL**) of real traffic scenes from the evaluation distribution, measured in *nats*. NLL is a standard metric to compare generative models with tractable likelihoods. However, as many of our baselines do not have likelihoods, we compute a sample-based metric as well: maximum mean discrepancy (**MMD**) [19]. For two distributions p and q , MMD measures a distance between p and q as

$$\text{MMD}^2(p, q) = \mathbb{E}_{x, x' \sim p}[k(x, x')] + \mathbb{E}_{y, y' \sim q}[k(y, y')] - 2\mathbb{E}_{x \sim p, y \sim q}[k(x, y)] \quad (25)$$

for some kernel k . Following [56, 29], we compute MMD using Gaussian kernels with the total variation distance to compare distributions of scene statistics between generated and real traffic scenes. Our scene statistics measure the distribution of locations, classes, bounding box sizes, speeds, and heading angles (relative to the SDV) in each scene. To peer into the global properties of the traffic scene, we also compute MMD in the feature space of a pre-trained motion forecasting model that takes a rasterized image of the scene as input [53]. This is akin to the popular IS [41], FID [21], and KID [4] metrics for evaluating generative models, except we use a feature extractor trained on traffic scenes. Please see our appendix for details.

Additional details: Each traffic scene is a $80m \times 80m$ region of interest centered on the ego SDV. By default, SceneGen uses $K = 10$ mixture components and conditions on all available map elements for each dataset. We train SceneGen using the Adam optimizer [27] with a learning rate of $1e-4$ and a batch size of 16, until convergence. When sampling each actor’s position, heading, and velocity, we sample $M = 10$ times and keep the most likely sample.

³In Argoverse, these baselines generate vehicles only since bike lanes are not given. This highlights the challenge of designing good heuristics.



Figure 5: ATG4D scene with a traffic violation.

4.3. Results

Quantitative results: Tab. 1 summarizes the NLL and MMD results for ATG4D and Argoverse. Overall, SceneGen achieves the best results across both datasets, demonstrating that it can better model real traffic scenes and synthesize realistic examples as well. Interestingly, all learning-based methods outperform the hand-tuned baselines with respect to MMD on deep features—a testament to the difficulty of designing good heuristics.

Qualitative results: Fig. 3 visualizes samples generated by SceneGen on ATG4D and Argoverse. Fig. 4 compares traffic scenes generated by SceneGen and various baselines. Although MetaSim and Lane Graph generate reasonable scenes, they are limited by their underlying heuristics; *e.g.*, actors follow lane centerlines. LayoutVAE generates a greater variety of actors; however, the model does not position actors on the map accurately, rendering the overall scene unrealistic. In contrast, SceneGen’s samples reflect the complexity of real traffic scenes much better. That said, SceneGen occasionally generates near-collision scenes that are plausible but unlikely; *e.g.*, Fig. 3 top-right.

Ablation studies: In Tab. 2, we sweep over the number of mixture components used to parameterize distributions of bounding boxes and velocities. We see that increasing the number of components consistently lowers NLL, reflecting the need to model the multi-modality of real traffic scenes. We also ablate the input map to SceneGen: starting from an unconditional model, we progressively add lanes, drivable areas, crosswalks, and traffic light states. From Tab. 3, we see that using more map elements generally improves NLL. Surprisingly, incorporating traffic lights slightly degrades performance, which we conjecture is due to infrequent traffic light observations in ATG4D.

Discovering interesting scenes: We use SceneGen to find unlikely scenes in ATG4D by searching for scenes with the highest NLL, normalized by the number of actors. Fig. 5 shows an example of a traffic violation found via this procedure; the violating actor has an NLL of 21.28.

4.4. Sim2Real Evaluation

Our next experiment demonstrates that SceneGen coupled with sensor simulation can generate realistic labeled

Method	Vehicle		Pedestrian		Bicyclist	
	0.5	0.7	0.3	0.5	0.3	0.5
Prob. Gram.	81.1	66.6	-	-	11.2	10.6
MetaSim	76.3	63.3	-	-	8.2	7.5
Procedural	80.2	63.0	-	-	6.5	3.8
Lane Graph	82.9	71.7	-	-	7.6	6.9
LayoutVAE	85.9	76.3	49.3	41.8	18.4	16.4
SceneGen	90.4	82.4	58.1	48.7	19.6	17.9
Real Scenes	93.7	86.7	69.3	61.6	29.2	25.9

Table 4: Detection AP on real ATG4D scenes.



Figure 6: Outputs of detector trained with SceneGen scenes.

data for training perception models. For each method under evaluation, we generate 250,000 traffic scenes conditioned on the SDV and HD map in each frame of the 1000 held-out logs in ATG4D. Next, we use LiDARsim [32] to simulate the LiDAR point cloud corresponding to each scene. Finally, we train a 3D object detector [54] using the simulated LiDAR and evaluate its performance on real scenes and LiDAR in ATG4D.

From Tab. 4, we see that SceneGen’s traffic scenes exhibit the lowest sim2real gap. Here, Real Scenes is simulated LiDAR from ground truth placements. This reaffirms our claim that the underlying rules and priors used in MetaSim and Lane Graph induce a content gap. By eschewing these heuristics altogether, SceneGen learns to generate significantly more realistic traffic scenes. Intriguingly, LayoutVAE performs competitively despite struggling to position actors on the map. We conjecture that this is because LayoutVAE captures the diversity of actor classes, sizes, headings, *etc.* well. However, by accurately modeling actor positions as well, SceneGen further reduces the sim2real gap, as compared to ground truth traffic scenes.

5. Conclusion

We have presented SceneGen—a neural autoregressive model of traffic scenes from which we can sample new examples as well as evaluate the likelihood of existing ones. Unlike prior methods, SceneGen eschews the need for rules or heuristics, making it a more flexible and scalable approach for modeling the complexity and diversity of real world traffic scenes. As a result, SceneGen is able to generate realistic traffic scenes, thus taking an important step towards safe and scalable self-driving.

References

- [1] M. Bando, K. Hasebe, A. Nakayama, A. Shibata, and Y. Sugiyama. Dynamical model of traffic congestion and numerical simulation. *Physical Review E*, 1995.
- [2] Yoshua Bengio and Samy Bengio. Modeling high-dimensional discrete data with multi-layer neural networks. In *NeurIPS*, 1999.
- [3] Donald Best and Nicholas Fisher. Efficient simulation of the von mises distribution. *Journal of the Royal Statistical Society. Series C. Applied Statistics*, 1979.
- [4] Mikolaj Binkowski, Dougal J. Sutherland, Michael Arbel, and Arthur Gretton. Demystifying MMD gans. In *ICLR*, 2018.
- [5] Christopher M. Bishop. *Pattern Recognition and Machine Learning*. 2006.
- [6] Robert E. Chandler, Robert Herman, and Elliott W. Montroll. Traffic dynamics: Studies in car following. In *Operations Research*, 1958.
- [7] Ming-Fang Chang, John Lambert, Patsorn Sangkloy, Jagjeet Singh, Slawomir Bak, Andrew Hartnett, De Wang, Peter Carr, Simon Lucey, Deva Ramanan, and James Hays. Argoverse: 3d tracking and forecasting with rich maps. In *CVPR*, 2019.
- [8] Jeevan Devaranjan, Amlan Kar, and Sanja Fidler. Meta-sim2: Unsupervised learning of scene structure for synthetic data generation. 2020.
- [9] Alexey Dosovitskiy, Germán Ros, Felipe Codevilla, Antonio M. López, and Vladlen Koltun. CARLA: an open urban driving simulator. In *CoRL*, 2017.
- [10] Jin Fang, Dingfu Zhou, Feilong Yan, Tongtong Zhao, Feihu Zhang, Yu Ma, Liang Wang, and Ruigang Yang. Augmented lidar simulator for autonomous driving, 2019.
- [11] Martin Fellendorf. Vissim: A microscopic simulation tool to evaluate actuated signal control including bus priority. 1994.
- [12] Brendan J. Frey, Geoffrey E. Hinton, and Peter Dayan. Does the wake-sleep algorithm produce good density estimators? In *NeurIPS*, 1995.
- [13] Denos C. Gazis, Robert Herman, and Richard W. Rothery. Nonlinear follow-the-leader models of traffic flow. In *Operations Research*, 1961.
- [14] Andreas Geiger, Martin Lauer, and Raquel Urtasun. A generative model for 3d urban scene understanding from movable platforms. In *CVPR*, 2011.
- [15] Andreas Geiger, Christian Wojek, and Raquel Urtasun. Joint 3d estimation of objects and scene layout. In *NeurIPS*, 2011.
- [16] Daniel Gerlough. *Simulation of Freeway Traffic on a General-purpose Discrete Variable Computer*. 1955.
- [17] Peter Gipps. Computer program multsim for simulating output from vehicle detectors on a multi-lane signal-controlled road. 1976.
- [18] Alex Graves. Generating sequences with recurrent neural networks. *CoRR*, 2013.
- [19] Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Schölkopf, and Alexander J. Smola. A kernel two-sample test. *JMLR*, 2012.
- [20] David Ha and Douglas Eck. A neural representation of sketch drawings. In *ICLR*, 2018.
- [21] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *NeurIPS*, 2017.
- [22] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 1997.
- [23] Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. The curious case of neural text degeneration. In *ICLR*, 2020.
- [24] Stefan Jesenski, Jan Erik Stellet, Florian A. Schiegg, and J. Marius Zöllner. Generation of scenes in intersections for the validation of highly automated driving functions. In *IV*, 2019.
- [25] Akash Abdu Jyothi, Thibaut Durand, Jiawei He, Leonid Sigal, and Greg Mori. Layoutvae: Stochastic scene layout generation from a label set. In *ICCV*, 2019.
- [26] Amlan Kar, Aayush Prakash, Ming-Yu Liu, Eric Cameracci, Justin Yuan, Matt Rusiniak, David Acuna, Antonio Torralba, and Sanja Fidler. Meta-sim: Learning to generate synthetic datasets. In *ICCV*, 2019.
- [27] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.
- [28] Ming Liang, Bin Yang, Rui Hu, Yun Chen, Renjie Liao, Song Feng, and Raquel Urtasun. Learning lane graph representations for motion forecasting. In *ECCV*, 2020.
- [29] Renjie Liao, Yujia Li, Yang Song, Shenlong Wang, William L. Hamilton, David Duvenaud, Raquel Urtasun, and Richard S. Zemel. Efficient graph generation with graph recurrent attention networks. In *NeurIPS*, 2019.
- [30] Michael James Lighthill and Gerald Beresford Whitham. On kinematic waves. II. A theory of traffic flow on long crowded roads. In *Royal Society of London. Series A, Mathematical and Physical Sciences*, 1955.
- [31] Pablo Álvarez López, Michael Behrisch, Laura Bieker-Walz, Jakob Erdmann, Yun-Pang Flötteröd, Robert Hilbrich, Leonhard Lüken, Johannes Rummel, Peter Wagner, and Evamarie WieBner. Microscopic traffic simulation using SUMO. In *ITSC*, 2018.
- [32] Sivabalan Manivasagam, Shenlong Wang, Kelvin Wong, Wenyuan Zeng, Mikita Sazanovich, Shuhan Tan, Bin Yang, Wei-Chiu Ma, and Raquel Urtasun. Lidarsim: Realistic lidar simulation by leveraging the real world. In *CVPR*, 2020.
- [33] Charlie Nash, Yaroslav Ganin, S. M. Ali Eslami, and Peter W. Battaglia. Polygen: An autoregressive generative model of 3d meshes. *ICML*, 2020.
- [34] G. F. Newell. Nonlinear effects in the dynamics of car following. *Operations Research*, 1961.
- [35] Larry Owen, Yunlong Zhang, Lei Rao, and Gene Mchale. Traffic flow simulation using corsim. *Winter Simulation Conference*, 2001.
- [36] Louis A. Pipes. An operational analysis of traffic dynamics. In *Journal of Applied Physics*, 1953.
- [37] Aayush Prakash, Shaad Boochoon, Mark Brophy, David Acuna, Eric Cameracci, Gavriel State, Omer Shapira, and Stan Birchfield. Structured domain randomization: Bridging the reality gap by context-aware synthetic data. In *ICRA*, 2019.

- [38] Sergey Prokudin, Peter V. Gehler, and Sebastian Nowozin. Deep directional statistics: Pose estimation with uncertainty quantification. In *ECCV*, 2018.
- [39] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. 2018.
- [40] Paul Richards. Shock waves on the highway. In *Operations Research*, 1956.
- [41] Tim Salimans, Ian J. Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. In *NeurIPS*, 2016.
- [42] Xingjian Shi, Zhourong Chen, Hao Wang, Dit-Yan Yeung, Wai-Kin Wong, and Wang-chun Woo. Convolutional LSTM network: A machine learning approach for precipitation nowcasting. In *NeurIPS*, 2015.
- [43] Saied Taheri. *An investigation and design of slip control braking systems integrated with four wheel steering*. 1990.
- [44] Martin Treiber, Ansgar Hennecke, and Dirk Helbing. Congested traffic states in empirical observations and microscopic simulations. In *Automatisierungstechnik*, 2000.
- [45] Benigno Uria, Iain Murray, and Hugo Larochelle. NADE: the real-valued neural autoregressive density-estimator. *CoRR*, 2013.
- [46] Benigno Uria, Iain Murray, and Hugo Larochelle. RNADE: the real-valued neural autoregressive density-estimator. In *NeurIPS*, 2013.
- [47] Benigno Uria, Iain Murray, and Hugo Larochelle. A deep and tractable density estimator. In *ICML*, 2014.
- [48] Aäron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew W. Senior, and Koray Kavukcuoglu. Wavenet: A generative model for raw audio. In *ISCA*, 2016.
- [49] Aäron van den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. In *ICML*, 2016.
- [50] Kai Wang, Manolis Savva, Angel X. Chang, and Daniel Ritchie. Deep convolutional priors for indoor scene synthesis. *TOG*, 2018.
- [51] Tim Allan Wheeler and Mykel J. Kochenderfer. Factor graph scene distributions for automotive safety analysis. In *ITSC*, 2016.
- [52] Tim Allan Wheeler, Mykel J. Kochenderfer, and Philipp Robbel. Initial scene configurations for highway traffic propagation. In *ITSC*, 2015.
- [53] Kelvin Wong, Qiang Zhang, Ming Liang, Bin Yang, Renjie Liao, Abbas Sadat, and Raquel Urtasun. Testing the safety of self-driving vehicles by simulating perception and prediction. *ECCV*, 2020.
- [54] Bin Yang, Ming Liang, and Raquel Urtasun. HDNET: exploiting HD maps for 3d object detection. In *CoRL*, 2018.
- [55] Qi Yang and Haris N. Koutsopoulos. A microscopic traffic simulator for evaluation of dynamic traffic management systems. *Transportation Research Part C: Emerging Technologies*, 1996.
- [56] Jiaxuan You, Rex Ying, Xiang Ren, William L. Hamilton, and Jure Leskovec. Graphrnn: Generating realistic graphs with deep auto-regressive models. In *ICML*, 2018.
- [57] Hongyi Zhang, Andreas Geiger, and Raquel Urtasun. Understanding high-level semantics by modeling traffic patterns. In *ICCV*, 2013.