Testing the Safety of Self-driving Vehicles by Simulating Perception and Prediction

Kelvin Wong^{1,2*}, Qiang Zhang^{1,3*}, Ming Liang¹, Bin Yang^{1,2}, Renjie Liao^{1,2}, Abbas Sadat¹, and Raquel Urtasun^{1,2}

¹ Uber Advanced Technologies Group, Toronto, Canada ² University of Toronto, Toronto, Canada ³ Shanghai Jiao Tong University, Shanghai, China {kelvin.wong, ming.liang, byang10, rjliao, asadat, urtasun}@uber.com zhangqiang2016@sjtu.edu.cn

Abstract. We present a novel method for testing the safety of selfdriving vehicles in simulation. We propose an alternative to sensor simulation, as sensor simulation is expensive and has large domain gaps. Instead, we directly simulate the outputs of the self-driving vehicle's perception and prediction system, enabling realistic motion planning testing. Specifically, we use paired data in the form of ground truth labels and real perception and prediction outputs to train a model that predicts what the online system will produce. Importantly, the inputs to our system consists of high definition maps, bounding boxes, and trajectories, which can be easily sketched by a test engineer in a matter of minutes. This makes our approach a much more scalable solution. Quantitative results on two large-scale datasets demonstrate that we can realistically test motion planning using our simulations.

Keywords: Simulation · Perception & Prediction · Self-Driving Vehicles

1 Introduction

Self-driving vehicles (SDVs) have the potential to become a safer, cheaper, and more scalable form of transportation. But while great progress has been achieved in the last few decades, there still remain many open challenges that impede the deployment of these vehicles at scale. One such challenge concerns how to test the safety of these vehicles and, in particular, their motion planners [13, 44]. Most large-scale self-driving programs in industry use simulation for this purpose, especially in the case of testing safety-critical scenarios, which can be costly—even *unethical*—to perform in the real world. To this end, test engineers first create a large bank of test scenarios, each comprised of a high definition (HD) map and a set of actors represented by bounding boxes and trajectories. These mocked objects are then given as input to the motion planner. Finally, metrics computed on the simulation results are used to assess progress.

^{*} Indicates equal contribution. Work done during Qiang's internship at Uber ATG.

2



Fig. 1. Perception and prediction simulation. Our goal is to simulate the outputs of the SDV's perception and prediction system in order to realistically test its motion planner. For each timestep, our system ingests an HD map and a set of actors (bounding boxes and trajectories) and produces noisy outputs similar to those from the real system. To test the motion planner, we mock real outputs with our simulated ones.

However, in order to provide realistic testing, the mocked objects need to reflect the noise of real perception and prediction⁴ systems [34, 7, 33, 62, 50, 31]. Unfortunately, existing approaches typically assume perfect perception or use simple heuristics to generate noise [18]. As a result, they yield unrealistic assessments of the motion planner's safety. For example, under this testing regime, we will never see the SDV slamming its brakes due to a false positive detection.

An alternative approach is to use sensor simulation to test the SDV's full autonomy stack, end-to-end. Sensor simulation is a popular area of research, particularly in the case of images [42, 15, 1, 53, 25, 32, 60]. However, most existing sensor simulators are costly and difficult to scale since they are based on virtual worlds created by teams of artists; e.g., TORCS [54], CARLA [12], Air-Sim [46]. Rendering these virtual worlds also results in observations that have very different statistics from real sensor data. As a result, there are large domain gaps between these virtual worlds and our physical one. Recently, LiDARSim [35] leveraged real-world data to produce realistic LiDAR simulations at scale, narrowing the fidelity gap significantly. However, current autonomy stacks use a host of different sensors, including LiDAR [63, 59, 30], radar [8, 57], cameras [10, 51, 33, and ultrasonics, and thus all of these sensors must be simulated consistently for this approach to be useful in testing the full autonomy stack. These challenges make sensor simulation a very exciting area of research, but also one that is potentially far from deployment in real-world systems that must meet requirements developed by safety, systems engineering, and testing teams.

In this paper, we propose to simulate the SDV's perception and prediction system instead; see Fig. 1. To this end, we provide a comprehensive study of a variety of noise models with increasing levels of sophistication. Our best model is

⁴ We use the terms *prediction* and *motion forecasting* interchangeably.

a convolutional neural network that, given a simple representation of the scene, produces realistic perception and prediction simulations. Importantly, this input representation can be sketched by a test engineer in a matter of minutes, making our approach cheap and easy to scale. We validate our model on two self-driving datasets and show that our simulations closely match the outputs of a real perception and prediction system. We also demonstrate that they can be used to realistically test motion planning. We hope to inspire work in this important field so that one day we can certify the safety of SDVs and deploy them at scale.

2 Related Work

Sensor simulation: The use of sensor simulation in self-driving dates back to at least the seminal work of Pomerleau [40] who used both simulated and real road images to train a neural network to drive. Since then, researchers and engineers have developed increasingly realistic sensor simulators for self-driving across various modalities. For example, [42, 15, 1, 53, 25] use photo-realistic rendering techniques to synthesize images to train neural networks and [32, 60] leverage real sensor data to generate novel views. Likewise, [17, 61, 14, 12] use physics-based ray-casting to simulate LiDAR while [35] enhances its realism with learning. And in radar, [19] propose a ray-tracing based simulator and [52] use a fully-learned approach. However, despite much progress in recent years, there remain sizeable domain gaps between simulated sensor data and real ones [35]. Moreover, developing a realistic sensor simulator requires significant effort from domain experts [25], which limits the scalability of doing so across an entire sensor suite. In this paper, we sidestep these challenges by instead simulating a much simpler scene representation: the SDV's perception and prediction outputs.

Virtual environments: Training and testing robots in the phyiscal world can be a slow, costly, and even dangerous affair; virtual environments are often used to circumvent these difficulties. For example, in machine learning and robotics, popular benchmarks include computer games [4, 24, 47, 26, 3], indoor environments [29, 45, 56, 55], robotics simulators [11, 49, 28], and self-driving simulators [54, 9, 12, 46]. These virtual worlds have motivated a wealth of research in fields ranging from embodied vision to self-driving. However, they also require significant effort to construct, and this has unfortunately limited the diversity of their content. For example, CARLA [12] originally had just two artist-generated towns consisting of 4.3km of drivable roads. In this paper, we use a lightweight scene representation that simplifies the task of generating new scenarios.

Knowledge distillation: Knowledge distillation was first popularized by Hinton *et al.* [22] as a way to compress neural networks by training one network with the (soft) outputs of another. Since then, researchers have found successful applications of distillation in subfields across machine learning [21, 16, 38, 43, 20]. In this paper, we also train our simulation model using outputs from an

SDV's perception and prediction system. In this sense, our work is closely related with distillation. However, unlike prior work in distillation, we assume no direct knowledge of the target perception and prediction system; *i.e.*, we treat these modules as black boxes. Moreover, the inputs to our simulation model differ from the inputs to the target system. This setting is more suitable for self-driving, where perception and prediction systems can be arbitrarily complex pipelines.

3 Perception and Prediction Simulation

Our goal is to develop a framework for testing the SDV's motion planner as it will behave in the real world. One approach is to use sensor simulation to test the SDV's full autonomy stack, end-to-end. However, this can be a complex and costly endeavor that requires constructing realistic virtual worlds and developing high-fidelity sensor simulators. Moreover, there remains a large domain gap between the sensor data produced by existing simulators and our physical world.

In this work, we study an alternative approach. We observe that the autonomy stack of today's SDVs employ a cascade of interpretable modules: perception, prediction, and motion planning. Therefore, rather than simulate the raw sensor data, we simulate the SDV's intermediate perception and prediction outputs instead, thus leveraging the compositionally of its autonomy stack to bypass the challenges of sensor simulation. Testing the SDV's motion planner can then proceed by simply mocking real perception and prediction outputs with our simulated ones. We call this task *perception and prediction simulation*.

Our approach is predicated on the hypothesis that there exists systemic noise in modern perception and prediction systems that we could simulate. Indeed, our experiments show that this is the case in practice. Therefore, we study a variety of noise models with increasing levels of sophistication. Our best model is a convolutional neural network that, given a simple representation of the scene, learns to produce realistic perception and prediction simulations. This enables us to realistically test motion planning in simulation. See Fig. 1 for an overview.

In this section, we first formulate the task of perception and prediction simulation and define some useful notation. Next, we describe a number of noise models in order of increasing sophistication and highlight several key modeling choices that informs the design of our best model. Finally, we describe our best model for this task and discuss how to train it in an end-to-end fashion.

3.1 Problem Formulation

Given a sensor reading at timestep t, the SDV's perception and prediction system ingests an HD map and sensor data and produces a class label \hat{c}_i , a bird's eye view (BEV) bounding box \hat{b}_i , and a set of future states $\hat{s}_i = \{\hat{s}_{i,t+\delta}\}_{\delta=1}^H$ for each actor i that it detects in the scene, where H is the prediction horizon. Each state $\hat{s}_{i,t+\delta} \in \mathbb{R}^3$ consists of the actor's 2D BEV position and orientation at some timestep $t + \delta$ in the future.⁵ Note that this is the typical output parame-

4

⁵ Actors' future orientations are approximated from their predicted waypoints using finite differences, and their bounding box sizes remain constant over time.



Fig. 2. Perturbation models for perception and prediction simulation. NoNoise assumes perfect perception and prediction. GaussianNoise and Multimodal-Noise use marginal noise distributions to perturb each actor's shape, position, and whether it is misdetected. ActorNoise accounts for inter-actor variability by predicting perturbations conditioned on each actor's bounding box and positions over time.

terization for an SDV's perception and prediction system [34, 7, 33, 62, 50, 31], as it is lightweight, interpretable, and easily ingested by existing motion planners.

For each timestep in a test scenario, our goal is to simulate the outputs of the SDV's perception and prediction system without using sensor data—neither real nor simulated. Instead, we use a much simpler representation of the world such that we can: (i) bypass the complexity of developing realistic virtual worlds and sensor simulators; and (ii) simplify the task of constructing new test scenarios.

Our scenario representation consists of an HD map \mathcal{M} , a set of actors \mathcal{A} , and additional meta-data for motion planning, such as the SDV's starting state and desired route. The HD map \mathcal{M} contains semantic information about the static scene, including lane boundaries and drivable surfaces. Each actor $a_i \in \mathcal{A}$ is represented by a class label c_i , a bounding box \mathbf{b}_i , and a set of states $\mathbf{s}_i = \{\mathbf{s}_{i,t}\}_{t=0}^T$, where T is the scenario duration. Note that \mathcal{A} is a *perfect* perception and prediction of the world, not the (noisy) outputs of a real online system.

This simple representation can be easily sketched by a test engineer in a matter of seconds or minutes, depending on the complexity and duration of the scenario. The test engineer can start from scratch or from existing logs collected in real traffic or in structured tests at a test track by adding or removing actors, varying their speeds, changing the underlying map, *etc*.

3.2 Perturbation Models for Perception and Prediction Simulation

One family of perception and prediction simulation methods builds on the idea of perturbing the actors \mathcal{A} of the input test scenario with noise approximating that found in real systems. In this section, we describe a number of such methods in order of increasing sophistication; see Fig. 2. Along the way, we highlight several key modeling considerations that will motivate the design of our best model.

NoNoise: For each timestep t of the test scenario, we can readily simulate perfect perception and prediction by outputting the class label c_i , the bounding box b_i , and the future states $\{s_{i,t+\delta}\}_{\delta=1}^H$ for each actor $a_i \in \mathcal{A}$. Indeed, most existing methods to test motion planning similarly use perfect perception [18]. This approach gives an important signal as an upper bound on the motion planner's performance in the real world. However, it is also unrealistic as it yields

6

an overly optimistic evaluation of the motion planner's safety. For example, this approach cannot simulate false negative detections; thus, the motion planner will never be tested for its ability to exercise caution in areas of high occlusion.

GaussianNoise: Due to its assumption of perfect perception and prediction, the previous method does not account for the noise present in real perception and prediction systems. As such, it suffers a sim-to-real domain gap. In domain randomization, researchers have successfully used random noise to bridge this gap during training [39, 48, 41, 36, 37]. This next approach investigates whether random noise can be similarly used to bridge the sim-to-real gap *during testing*. Specifically, we model the noise present in real perception and prediction systems with a marginal distribution p_{noise} over all actors. For each timestep t in the test scenario, we perturb each actor's bounding box \mathbf{b}_i and future states $\{\mathbf{s}_{i,t+\delta}\}_{\delta=1}^H$ with noise drawn from p_{noise} . In our experiments, we use noise drawn from a Gaussian distribution $\mathcal{N}(0, 0.1)$ to perturb each component in $\mathbf{b}_i = (x, y, \log w, \log h, \sin \theta, \cos \theta)$, where (x, y) is the box's center, (w, h) is its width and height, and θ is its orientation. We similarly perturb each state in $\{\mathbf{s}_{i,t+\delta}\}_{\delta=1}^H$. To simulate misdetections, we randomly drop boxes with probability equal to the observed rate of false negative detections in our data.⁶

MultimodalNoise: Simple noise distributions such as the one used in GaussianNoise do not adequately reflect the complexity of the noise in perception and prediction systems. For example, prediction noise is highly multi-modal since vehicles can go straight or turn at intersections. Therefore, in this next approach, we instead use a Gaussian Mixture Model, which we fit to the empirical distribution of noise in our data via expectation-maximization [5]. As before, we simulate misdetections by dropping boxes with probability equal to the observed rate of false negative detections in our data.

ActorNoise: In MultimodalNoise, we use a marginal noise distribution over all actors to model the noise present in perception and prediction systems. This, however, does not account for inter-actor variability. For example, prediction systems are usually more accurate for stationary vehicles than for ones with irregular motion. In our next approach, we relax this assumption by conditioning the noise for each actor on its bounding box \boldsymbol{b}_i and past, present, and future states \boldsymbol{s}_i . We implement ActorNoise as a multi-layer perceptron that learns to predict perturbations to each component of an actor's bounding box \boldsymbol{b}_i and future states $\{\boldsymbol{s}_{i,t+\delta}\}_{\delta=1}^H$. We also predict each actor's probability of misdetection. To train ActorNoise, we use a combination of a binary cross entropy loss for misdetection classification and a smooth ℓ_1 loss for box and waypoint regression.

⁶ True positive, false positive, and false negative detections are determined by IoU following the detection AP metric. In our experiments, we use a 0.5 IoU threshold for cars and vehicles and 0.3 IoU for pedestrians and bicyclists.



Fig. 3. ContextNoise for perception and prediction simulation. Given BEV rasterized images of the scene (drawn from bounding boxes and HD maps), our model simulates outputs similar to those from the real perception and prediction system. It consists of: (i) a shared backbone feature extractor; (ii) a perception head for simulating bounding box outputs; and (iii) a prediction head for simulating future states outputs.

3.3 A Contextual Model for Perception and Prediction Simulation

So far, we have discussed several perturbation-based models for perception and prediction simulation. However, these methods have two limitations. First, they cannot simulate false positive misdetections. More importantly, they do not use contextual information about the scene, which intuitively should correlate with the success of a perception and prediction system. For example, HD maps provide valuable contextual information to determine what actor behaviors are possible.

To address these limitations, we propose to use a convolutional neural network that takes as input BEV rasterized images of the scene (drawn from bounding boxes and HD maps) and learns to simulate dense bounding boxes and future state outputs similar to those from the real perception and prediction system. This is the native parameterization of the perception and prediction system used in our experiments. Our model architecture is composed of three components: (i) a shared backbone feature extractor; (ii) a perception head for simulating bounding box outputs; and (iii) a prediction head for simulating future states outputs. We call this model *ContextNoise*. See Fig. 3 for an overview.

Input representation: For each timestep t of the input scenario, our model takes as input BEV raster images of the scene in ego-centric coordinates. In particular, for each class of interest, we render the actors of that class as bounding boxes in a sequence of occupancy masks [2, 23] indicating their past, present, and future positions. Following [7, 58], we rasterize the HD map \mathcal{M} into multiple binary images. We represent lane boundaries as polylines and drivable surfaces as filled polygons. Occlusion is an important source of systemic errors for perception and prediction systems. For example, a heavily occluded pedestrian is more likely to be misdetected. To model this, we render a temporal sequence of 2D occlusion

masks using a constant-horizon ray-casting algorithm [18]. By stacking these binary images along the feature channel, we obtain our final input representation.

Backbone network: We use the backbone architecture of [33] as our shared feature extractor. Specifically, it is a convolutional neural network that computes a feature hierarchy at three scales of input resolution: 1/4, 1/8, and 1/16. These multi-scale features are then upscaled to 1/4 resolution and fused using residual connections. This yields a $C \times H/4 \times W/4$ feature map, where C is the number of output channels and H and W is the height and width of the input raster image. Note that we use this backbone to extract features from BEV raster images (drawn from bounding boxes and HD maps), not voxelized LiDAR point clouds as it was originally designed for. We denote the resulting feature map by:

$$\mathcal{F}_{\text{bev}} = \text{CNN}_{\text{bev}} \left(\mathcal{A}, \mathcal{M} \right) \tag{1}$$

Perception head: Here, our goal is to simulate the bounding box outputs of the real perception and prediction system. To this end, we use a lightweight header to predict dense bounding box outputs for every class. Our dense output parameterization allows us to naturally handle false positive and false negative misdetections. In detail, for each class of interest, we use one convolution layer with 1×1 kernels to predict a bounding box $\tilde{\boldsymbol{b}}_i$ and detection score $\tilde{\alpha}_i$ at every BEV pixel *i* in \mathcal{F}_{bev} . We parameterize $\tilde{\boldsymbol{b}}_i$ as $(\Delta x, \Delta y, \log w, \log h, \sin \theta, \cos \theta)$, where $(\Delta x, \Delta y)$ are the position offsets to the box center, (w, h) are its width and height, and θ is its orientation [59]. We use non-maximum suppression to remove duplicates. This yields a set of simulated bounding boxes $\mathcal{B}_{\text{sim}} = {\{\tilde{\boldsymbol{b}}_i\}_{i=1}^N}$.

Prediction head: Our goal now is to simulate a set of future states for each bounding box $\tilde{b}_i \in \mathcal{B}_{sim}$. To this end, for each $\tilde{b}_i \in \mathcal{B}_{sim}$, we first extract a feature vector f_i by bilinearly interpolating \mathcal{F}_{bev} around its box center. We then use a multi-layer perceptron to simulate its future positions:

$$\tilde{\boldsymbol{x}}_i = \mathrm{MLP}_{\mathrm{pred}}\left(\boldsymbol{f}_i\right)$$
(2)

where $\tilde{\boldsymbol{x}}_i \in \mathbb{R}^{H \times 2}$ is a set of 2D BEV waypoints over the prediction horizon H. We also simulate its future orientation $\tilde{\boldsymbol{\theta}}_i$ using finite differences. Together, $\{\tilde{\boldsymbol{x}}_i\}_{i=1}^N$ and $\{\tilde{\boldsymbol{\theta}}_i\}_{i=1}^N$ yield a set of simulated future states $\mathcal{S}_{\text{sim}} = \{\tilde{\boldsymbol{s}}_i\}_{i=1}^N$. Combining \mathcal{S}_{sim} with \mathcal{B}_{sim} , we have our final perception and prediction simulation.

Learning: We train our model with a multi-task loss function:

$$\mathcal{L} = \ell_{\rm perc} + \ell_{\rm pred} \tag{3}$$

where ℓ_{perc} is the perception loss and ℓ_{pred} is the prediction loss. Note that these losses are computed between our simulations and the outputs of the real perception and prediction system. Thus, we train our model using datasets that provide

both real sensor data (to generate real perception and prediction outputs) and our input scenario representations (to give as input to our model).⁷

Our perception loss is a multi-task detection loss. For object classification, we use a binary cross-entropy loss with online negative hard-mining, where positive and negative BEV pixels are determined according to their distances to an object's center [59]. For box regression at positive pixels, we use a smooth ℓ_1 loss for box orientation and an axis-aligned IoU loss for box location and size.

Our prediction loss is a sum of smooth ℓ_1 losses over future waypoints for each true positive bounding box, where a simulated box is positive if its IoU with a box from the real system exceeds a certain threshold. In our experiments, we use a threshold of 0.5 for cars and vehicles and 0.3 for pedestrians and bicyclists.

4 Experimental Evaluation

In this section, we benchmark a variety of noise models for perception and prediction simulation on two large-scale self-driving datasets (Section 4.3). Our best model achieves significantly higher simulation fidelity than existing approaches that assume perfect perception and prediction. We also conduct downstream experiments with two motion planners (Section 4.4). Our results show that there is a strong correlation between our ability to realistically simulate perception and prediction and our ability to realistically test motion planning.

4.1 Datasets

nuScenes: nuScenes [6] consists of 1000 traffic scenarios collected in Boston and Singapore, each containing 20 seconds of video captured by a 32-beam LiDAR sensor at 20Hz. In this dataset, keyframes sampled at 2Hz are annotated with object labels within a 50m radius. We generate additional labels at unannotated frames by linearly interpolating labels from adjacent keyframes [33]. We use the official training and validation splits and perform evaluation on the *car* class. To prevent our simulation model from overfitting to the training split, we partition the training split into two halves: one to train the perception and prediction model and the other our simulation model. Note that we do not use HD maps in our nuScenes experiments due to localization issues in some maps.⁸

ATG4D: ATG4D [59] consists of 6500 challenging traffic scenarios collected by a fleet of self-driving vehicles in cities across North America. Each scenario contains 25 seconds of video captured by a Velodyne HDL-64E at 10Hz, resulting in 250 LiDAR sweeps per video. Each sweep is annotated with bounding boxes and trajectories for the vehicle, pedestrian, and bicyclist classes within a 100m radius and comes with localized HD maps. We split ATG4D into two training splits of 2500 scenarios each, a validation split of 500, and a test split of 1000.

 $^{^7}$ Our representation uses bounding boxes and trajectories. Most self-driving datasets provide this as ground truth labels for the standard perception and prediction task.

For perception and prediction simulation, we use these labels as *inputs* instead.

⁸ As of nuScenes map v1.0.



Fig. 4. Simulation results on ATG4D. We visualize PLT [44] motion planning results when given real perception and prediction (top) versus simulations from NoNoise (middle) and ContextNoise (bottom). ContextNoise faithfully simulates a misprediction due to multi-modality and induces a lane-change behavior from the motion planner.

4.2 Experiment Setup

Autonomy stack: We simulate the outputs of PnPNet [33]—a state-of-the-art joint perception and prediction model. PnPNet takes as input an HD map and the past 0.5s of LiDAR sweeps and outputs BEV bounding boxes and 3.0s of future waypoints (in 0.5s increments) for each actor that it detects. Since our focus is on simulating perception and prediction, we use the variant of PnPNet without tracking. We configure PnPNet to use a common detection score threshold of 0. In the ATG4D validation split, this corresponds to a recall rate of 94% for vehicles, 78% for pedestrians, and 62% for bicyclists.

To gauge the usefulness of using our simulations to test motion planning, we conduct downstream experiments with two motion planners. Our first motion planner is adaptive cruise control (ACC), which implements a car-following algorithm. Our second motion planner is PLT [44]—a jointly learnable behavior and trajectory planner. PLT is pretrained on the ManualDrive dataset [44], which consists of 12,000 logs in which the drivers were instructed to drive smoothly.

Experiment details: In nuScenes, we use a $100m \times 100m$ region of interest centered on the SDV for training and evaluation. In ATG4D, we use one encompassing 70m in front of the SDV and 40m to its left and right. Our rasters have a resolution of 0.15625m per pixel, resulting in 640×640 input images for nuScenes and 448×512 for ATG4D. All of our noise models ingest 0.5s of actor states in the past and 3.0s into the future (in 0.5s increments). We train ActorNoise and ContextNoise using the Adam optimizer [27] with a batch size of 32 and an initial learning rate of 4e-4, which we decay by 0.1 after every five epochs for a total of 15 epochs. We re-train PnPNet for our experiments following [33].

4.3 Perception and Prediction Simulation Results

In this section, we benchmark a variety of noise models for perception and prediction simulation. Our best model, ContextNoise, produces simulations that closely match the outputs of the real perception and prediction system.

Metrics: We use two families of metrics to evaluate the similarity between our simulated outputs and those from the real perception and prediction system. This is possible since our datasets provide both real sensor data and our input scenario representations. Our first family of metrics measures the similarity between simulated bounding boxes and real ones. To this end, we report detection average precision (AP) and maximum recall at various IoU thresholds depending on the class and dataset. Our second family of metrics measures the similarity between simulated future states and real ones. We use average displacement error (ADE) over 3.0s and final displacement error (FDE) at 3.0s for this purpose. These metrics are computed on true positive bounding boxes at 0.5 IoU for cars and vehicles and 0.3 IoU for pedestrians and bicyclists. In order to fairly compare models with different maximum recall rates, we report ADE and FDE for all methods at a common recall point, if it is attained. All metrics for Gaussian-Noise and MultimodalNoise are averaged over 25 sample runs. Note that we use random ordering to compute AP, ADE, and FDE for the methods that do not produce ranking scores: NoNoise, GaussianNoise, and MultimodalNoise.

Quantitative results: Tables 1 and 2 show the results of our experiments on nuScenes and ATG4D respectively. Overall, ContextNoise attains the best performance. In contrast, simple marginal noise models such as GaussianNoise and MultimodalNoise perform worse than the method that uses no noise at all. This attests to the importance of using contextual information for simulating the noise in real perception and prediction systems. In addition, we highlight the fact that only ContextNoise improves maximum recall over NoNoise. This is at least in part due to its dense output parameterization, which can naturally model misdetections due to mislocalization, misclassification, *etc.* Finally, we note that ContextNoise's improvements in prediction metrics are most evident for the car and vehicle classes; for rarer classes, such as pedestrians and bicyclists, ContextNoise and ActorNoise perform similarly well.

	Perception Metrics AP (%) Max Rec			$s \uparrow$ ecall (%)	P: ADE	rediction (cm)	ion Metrics \downarrow FDE (cm)		
Car	$0.5 { m IoU}$	$0.7 { m IoU}$	$0.5 \mathrm{IoU}$	$0.7 \ \mathrm{IoU}$	$50\%~{\rm R}$	70% R	$50\%~{\rm R}$	70% R	
GaussianNoise	4.9	0.9	13.0	1.7	-	-	-	-	
MultimodalNoise	12.8	4.9	21.1	13.1	-	-	-	-	
NoNoise	51.5	39.0	72.0	62.7	85	84	147	146	
ActorNoise	65.7	55.0	72.1	63.5	64	66	97	100	
ContextNoise	72.2	59.1	80.3	68.9	54	61	81	90	

 Table 1. Perception and prediction simulation metrics on nuScene validation. R denotes the common recall point at which prediction metrics are computed.

	Perception Metrics \uparrow				Prediction Metrics \downarrow			
	AP	(%)	Max Re	ecall $(\%)$	ADE	(cm)	FDE	(cm)
Vehicle	$0.5 { m IoU}$	$0.7 { m IoU}$	$0.5 {\rm IoU}$	$0.7~{\rm IoU}$	$70\%~{\rm R}$	$90\%~{\rm R}$	$70\%~{\rm R}$	$90\%~{\rm R}$
GaussianNoise	16.5	0.4	34.4	5.2	-	-	-	-
MultimodalNoise	30.7	12.1	46.8	29.4	-	-	-	-
NoNoise	71.7	56.9	93.1	82.9	70	70	127	128
ActorNoise	86.6	70.4	93.2	82.9	65	57	109	93
ContextNoise	91.8	82.3	95.7	87.8	46	51	72	78
Pedestrian	$0.3 { m IoU}$	0.5 IoU	$0.3 \mathrm{IoU}$	$0.5 \ IoU$	$60\%~{\rm R}$	$80\%~{\rm R}$	$60\%~\mathrm{R}$	$80\%~{\rm R}$
GaussianNoise	13.8	3.0	30.0	13.8	-	-	-	-
MultimodalNoise	30.3	21.7	44.2	37.4	-	-	-	-
NoNoise	57.4	52.3	84.0	80.2	41	41	70	70
ActorNoise	67.1	61.6	84.0	80.0	36	35	55	54
ContextNoise	75.1	66.6	88.2	80.3	34	34	51	52
Bicyclist	$0.3 { m IoU}$	0.5 IoU	$0.3 \mathrm{IoU}$	$0.5 \ IoU$	$50\%~{\rm R}$	70% R	$50\%~{\rm R}$	$70\%~{\rm R}$
GaussianNoise	4.7	0.4	17.8	5.4	-	-	-	-
MultimodalNoise	8.4	3.2	24.0	14.7	-	-	-	-
NoNoise	30.6	21.6	79.7	66.8	54	55	95	95
ActorNoise	60.4	44.1	79.7	67.8	54	49	88	78
ContextNoise	66.8	52.8	89.8	76.5	52	50	80	75

Table 2. Perception and prediction simulation metrics on ATG4D test.

4.4 Motion Planning Evaluation Results

Our ultimate goal is to use perception and prediction simulation to test motion planning. Therefore, we conduct downstream experiments in ATG4D to quantify the efficacy of doing so for two motion planners: ACC and PLT.

Metrics: Our goal is to evaluate how similarly a motion planner will behave in simulation versus the physical world. To quantify this, we compute the ℓ_2 distance between a motion planner's trajectory given simulated perception and prediction outputs versus its trajectory when given real outputs instead. We

	ℓ_2 Dis	Z_2 Distance (cm) ↓ Collision Sim. (%) ↑			Driving Diff. (%) \downarrow			
	1.0s	2.0s	3.0s	IoU	Recall	Beh.	Jerk	Acc.
PLT								
GaussianNoise	2.6	8.4	15.9	34.5	92.7	0.30	0.10	1.05
MultimodalNoise	2.7	9.4	18.0	25.2	93.6	0.33	1.22	1.25
NoNoise	1.4	4.8	9.5	52.9	58.2	0.18	0.44	0.03
ActorNoise	1.0	3.6	7.0	57.6	63.6	0.12	0.27	0.13
ContextNoise	0.8	2.9	5.6	65.1	74.3	0.10	0.05	0.06
ACC								
GaussianNoise	6.4	32.5	79.9	36.5	96.7	-	5.14	0.03
MultimodalNoise	5.1	26.2	64.9	36.5	96.7	-	3.84	0.11
NoNoise	1.9	10.0	25.2	52.9	32.4	-	0.20	0.17
ActorNoise	1.6	8.1	20.0	58.6	66.3	-	0.40	0.13
ContextNoise	1.4	7.2	17.6	61.3	74.1	-	0.14	0.03

Table 3. Motion planning evaluation metrics on ATG4D test.

report this metric for $\{1.0, 2.0, 3.0\}$ seconds into the future. In addition, we also measure their differences in terms of passenger comfort metrics; *i.e.*, jerk and lateral acceleration. Finally, we report the proportion of scenarios in which PLT chooses a different behavior when given simulated outputs instead of real ones.⁹

An especially important metric to evaluate the safety of a motion planner measures the proportion of scenarios in which the SDV will collide with an obstacle. To quantify our ability to reliably measure this in simulation, we report the intersection-over-union of collision scenarios and its recall-based variant:

$$IoU_{col} = \frac{|R_{+} \cap S_{+}|}{|R_{+} \cap S_{+}| + |R_{+} \cap S_{-}| + |R_{-} \cap S_{+}|} \quad \text{Recall}_{col} = \frac{|R_{+} \cap S_{+}|}{|R_{+}|} \quad (4)$$

where R_+ and S_+ are the sets of scenarios in which the SDV collides with an obstacle after 3.0s given real and simulated perception and prediction respectively, and R_- and S_- are similarly defined for scenarios with no collisions.

Quantitative results: Table 3 shows our experiment results on ATG4D. They show that by realistically simulating the noise in real perception and prediction systems, we can induce similar motion planning behaviors in simulation as in the real world, thus making our simulation tests more realistic. For example, ContextNoise yields a 41.1% and 30.2% relative reduction in ℓ_2 distance at 3.0s over NoNoise for PLT and ACC respectively. Importantly, we can also more reliably measure a motion planner's collision rate using ContextNoise versus NoNoise. This is an important finding since existing methods to test motion planning in simulation typically assume perfect perception or use simple heuristics to generate noise. Our results show that more sophisticated noise modeling is necessary.

⁹ Note that ACC always uses the same driving behavior.

Inputs AP $(\%)$ \uparrow				F	DE (cm)	$\ell_2 @ 3.0s (cm) \downarrow$			
Varian	t А О М	Veh.	Ped.	Bic.	Veh.	Ped.	Bic.	PLT	ACC
1	\checkmark	85.0	64.0	59.9	87	56	70	4.7	15.0
2	\checkmark \checkmark	85.5	63.8	61.7	86	55	72	4.7	14.4
3	\checkmark \checkmark \checkmark	86.9	68.6	64.1	76	52	70	4.2	14.2

Table 4. Ablation of ContextNoise input features on ATG4D validation. We progressively add each input feature described in Section 3.3. A denotes actor occupancy images; O denotes occlusion masks; and M denotes HD maps. AP is computed using 0.7 IoU for vehicles and 0.5 IoU for pedestrians and bicyclists. FDE at 3.0s is computed at 90% recall for vehicles, 80% for pedestrians, and 70% for bicyclists.

4.5 Ablation Study

To understand the usefulness of contextual information for simulation, we ablate the inputs to ContextNoise by progressively augmenting it with actor occupancy images, occlusion masks, and HD maps. From Table 4, we see that adding contextual information consistently improves simulation performance. These gains also directly translate to more realistic evaluations of motion planning.

4.6 Qualitative Results

We also visualize results from the PLT motion planner when given real perception and prediction versus simulations from NoNoise and ContextNoise. As shown in Fig. 4, ContextNoise faithfully simulates a misprediction due to multimodality and induces a lane-change behavior from the motion planner—the same behavior as if the motion planner was given real perception and prediction. In contrast, NoNoise induces an unrealistic keep-lane behavior instead.

5 Conclusion

In this paper, we introduced the problem of perception and prediction simulation in order to realistically test motion planning. To this end, we have studied a variety of noise models. Our best model has proven to be a convolutional neural network that, given a simple representation of the scene, learns to produce realistic perception and prediction simulations. Importantly, this representation can be easily sketched by a test engineer in a matter of minutes. We have validated our model on two large-scale self-driving datasets and showed that our simulations closely match the outputs of real perception and prediction systems. We have only begun to scratch the surface of this task. We hope our findings here will inspire advances in this important field so that one day we can certify the safety of self-driving vehicles and deploy them at scale.

15

References

- Alhaija, H.A., Mustikovela, S.K., Mescheder, L.M., Geiger, A., Rother, C.: Augmented reality meets computer vision: Efficient data generation for urban driving scenes. Int. J. Comput. Vis. (2018)
- Bansal, M., Krizhevsky, A., Ogale, A.S.: ChauffeurNet: Learning to drive by imitating the best and synthesizing the worst. In: Robotics: Science and Systems XV, University of Freiburg, Freiburg im Breisgau, Germany, June 22-26, 2019 (2019)
- Beattie, C., Leibo, J.Z., Teplyashin, D., Ward, T., Wainwright, M., Küttler, H., Lefrancq, A., Green, S., Valdés, V., Sadik, A., Schrittwieser, J., Anderson, K., York, S., Cant, M., Cain, A., Bolton, A., Gaffney, S., King, H., Hassabis, D., Legg, S., Petersen, S.: Deepmind lab. CoRR (2016)
- 4. Bellemare, M.G., Naddaf, Y., Veness, J., Bowling, M.: The arcade learning environment: An evaluation platform for general agents. J. Artif. Intell. Res. (2013)
- 5. Bishop, C.M.: Pattern recognition and machine learning, 5th Edition. Springer (2007)
- Caesar, H., Bankiti, V., Lang, A.H., Vora, S., Liong, V.E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., Beijbom, O.: nuScenes: A multimodal dataset for autonomous driving. CoRR (2019)
- Casas, S., Luo, W., Urtasun, R.: IntentNet: Learning to predict intention from raw sensor data. In: 2nd Annual Conference on Robot Learning, CoRL 2018, Zürich, Switzerland, 29-31 October 2018, Proceedings (2018)
- Chadwick, S., Maddern, W., Newman, P.: Distant vehicle detection using radar and vision. In: International Conference on Robotics and Automation, ICRA 2019, Montreal, QC, Canada, May 20-24, 2019 (2019)
- Chen, C., Seff, A., Kornhauser, A.L., Xiao, J.: DeepDriving: Learning affordance for direct perception in autonomous driving. In: 2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015 (2015)
- Chen, X., Kundu, K., Zhang, Z., Ma, H., Fidler, S., Urtasun, R.: Monocular 3d object detection for autonomous driving. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016 (2016)
- 11. Coumans, E., Bai, Y.: PyBullet, a python module for physics simulation for games, robotics and machine learning. http://pybullet.org (2016–2019)
- Dosovitskiy, A., Ros, G., Codevilla, F., López, A., Koltun, V.: CARLA: an open urban driving simulator. In: 1st Annual Conference on Robot Learning, CoRL 2017, Mountain View, California, USA, November 13-15, 2017, Proceedings (2017)
- Fan, H., Zhu, F., Liu, C., Zhang, L., Zhuang, L., Li, D., Zhu, W., Hu, J., Li, H., Kong, Q.: Baidu apollo EM motion planner. CoRR (2018)
- Fang, J., Yan, F., Zhao, T., Zhang, F., Zhou, D., Yang, R., Ma, Y., Wang, L.: Simulating LIDAR point cloud for autonomous driving using real-world scenes and traffic flows. CoRR (2018)
- Gaidon, A., Wang, Q., Cabon, Y., Vig, E.: Virtual worlds as proxy for multi-object tracking analysis. CoRR (2016)
- Geras, K.J., Mohamed, A., Caruana, R., Urban, G., Wang, S., Aslan, Ö., Philipose, M., Richardson, M., Sutton, C.A.: Compressing lstms into cnns. CoRR (2015)
- Gschwandtner, M., Kwitt, R., Uhl, A., Pree, W.: BlenSor: Blender sensor simulation toolbox. In: Advances in Visual Computing - 7th International Symposium, ISVC 2011, Las Vegas, NV, USA, September 26-28, 2011. Proceedings, Part II (2011)

- 16 K. Wong, Q. Zhang, M. Liang, B. Yang, R. Liao, A. Sadat, and R. Urtasun
- Gu, T., Dolan, J.M.: A lightweight simulator for autonomous driving motion planning development. In: ICIS 2015 (2015)
- 19. Gubelli, D., Krasnov, O.A., Yarovyi, O.: Ray-tracing simulator for radar signals propagation in radar networks. In: 2013 European Radar Conference (2013)
- Guo, X., Li, H., Yi, S., Ren, J.S.J., Wang, X.: Learning monocular depth by distilling cross-domain stereo networks. In: Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part XI (2018)
- Gupta, S., Hoffman, J., Malik, J.: Cross modal distillation for supervision transfer. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016 (2016)
- Hinton, G.E., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network. CoRR (2015)
- Jain, A., Casas, S., Liao, R., Xiong, Y., Feng, S., Segal, S., Urtasun, R.: Discrete residual flow for probabilistic pedestrian behavior prediction. In: 3rd Annual Conference on Robot Learning, CoRL 2019, Osaka, Japan, October 30 - November 1, 2019, Proceedings (2019)
- Johnson, M., Hofmann, K., Hutton, T., Bignell, D.: The malmo platform for artificial intelligence experimentation. In: Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016 (2016)
- Kar, A., Prakash, A., Liu, M., Cameracci, E., Yuan, J., Rusiniak, M., Acuna, D., Torralba, A., Fidler, S.: Meta-sim: Learning to generate synthetic datasets. In: 2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019 (2019)
- Kempka, M., Wydmuch, M., Runc, G., Toczek, J., Jaskowski, W.: ViZDoom: A doom-based AI research platform for visual reinforcement learning. CoRR (2016)
- Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings (2015)
- Koenig, N.P., Howard, A.: Design and use paradigms for Gazebo, an open-source multi-robot simulator. In: 2004 IEEE/RSJ International Conference on Intelligent Robots and Systems, Sendai, Japan, September 28 - October 2, 2004 (2004)
- 29. Kolve, E., Mottaghi, R., Gordon, D., Zhu, Y., Gupta, A., Farhadi, A.: AI2-THOR: an interactive 3d environment for visual AI. CoRR (2017)
- Lang, A.H., Vora, S., Caesar, H., Zhou, L., Yang, J., Beijbom, O.: Pointpillars: Fast encoders for object detection from point clouds. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019 (2019)
- Li, L., Yang, B., Liang, M., Zeng, W., Ren, M., Segal, S., Urtasun, R.: End-toend contextual perception and prediction with interaction transformer. In: 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2020, October 25-29, 2020 (2020)
- 32. Li, W., Pan, C., Zhang, R., Ren, J., Ma, Y., Fang, J., Yan, F., Geng, Q., Huang, X., Gong, H., Xu, W., Wang, G.P., Manocha, D., Yang, R.: AADS: augmented autonomous driving simulation using data-driven algorithms. Sci. Robotics (2019)
- 33. Liang, M., Yang, B., Zeng, W., Chen, Y., Hu, R., Casas, S., Urtasun, R.: PnPNet: End-to-end perception and prediction with tracking in the loop. In: 2020 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 16-18, 2020 (2020)

- 34. Luo, W., Yang, B., Urtasun, R.: Fast and furious: Real time end-to-end 3d detection, tracking and motion forecasting with a single convolutional net. In: 2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018. pp. 3569–3577 (2018)
- Manivasagam, S., Wang, S., Wong, K., Zeng, W., Sazanovich, M., Tan, S., Yang, B., Ma, W., Urtasun, R.: LiDARsim: Realistic lidar simulation by leveraging the real world. In: 2020 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 16-18, 2020 (2020)
- Mehta, B., Diaz, M., Golemo, F., Pal, C.J., Paull, L.: Active domain randomization. In: 3rd Annual Conference on Robot Learning, CoRL 2019, Osaka, Japan, October 30 - November 1, 2019, Proceedings (2019)
- 37. OpenAI, Akkaya, I., Andrychowicz, M., Chociej, M., Litwin, M., McGrew, B., Petron, A., Paino, A., Plappert, M., Powell, G., Ribas, R., Schneider, J., Tezak, N., Tworek, J., Welinder, P., Weng, L., Yuan, Q., Zaremba, W., Zhang, L.: Solving rubik's cube with a robot hand. CoRR (2019)
- Papernot, N., McDaniel, P.D., Wu, X., Jha, S., Swami, A.: Distillation as a defense to adversarial perturbations against deep neural networks. In: IEEE Symposium on Security and Privacy, SP 2016, San Jose, CA, USA, May 22-26, 2016 (2016)
- Peng, X.B., Andrychowicz, M., Zaremba, W., Abbeel, P.: Sim-to-real transfer of robotic control with dynamics randomization. In: 2018 IEEE International Conference on Robotics and Automation, ICRA 2018, Brisbane, Australia, May 21-25, 2018 (2018)
- Pomerleau, D.: ALVINN: an autonomous land vehicle in a neural network. In: Touretzky, D.S. (ed.) Advances in Neural Information Processing Systems 1, [NIPS Conference, Denver, Colorado, USA, 1988] (1988)
- Pouyanfar, S., Saleem, M., George, N., Chen, S.: ROADS: randomization for obstacle avoidance and driving in simulation. In: IEEE Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops 2019, Long Beach, CA, USA, June 16-20, 2019 (2019)
- 42. Ros, G., Sellart, L., Materzynska, J., Vázquez, D., López, A.M.: The SYNTHIA dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016 (2016)
- Rusu, A.A., Colmenarejo, S.G., Gülçehre, Ç., Desjardins, G., Kirkpatrick, J., Pascanu, R., Mnih, V., Kavukcuoglu, K., Hadsell, R.: Policy distillation. In: 4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings (2016)
- 44. Sadat, A., Ren, M., Pokrovsky, A., Lin, Y., Yumer, E., Urtasun, R.: Jointly learnable behavior and trajectory planning for self-driving vehicles. In: 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2019, Macau, SAR, China, November 3-8, 2019 (2019)
- 45. Savva, M., Chang, A.X., Dosovitskiy, A., Funkhouser, T.A., Koltun, V.: MINOS: multimodal indoor simulator for navigation in complex environments. CoRR (2017)
- Shah, S., Dey, D., Lovett, C., Kapoor, A.: AirSim: High-fidelity visual and physical simulation for autonomous vehicles. CoRR (2017)
- Tessler, C., Givony, S., Zahavy, T., Mankowitz, D.J., Mannor, S.: A deep hierarchical approach to lifelong learning in minecraft. In: Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA (2017)

- 18 K. Wong, Q. Zhang, M. Liang, B. Yang, R. Liao, A. Sadat, and R. Urtasun
- 48. Tobin, J., Fong, R., Ray, A., Schneider, J., Zaremba, W., Abbeel, P.: Domain randomization for transferring deep neural networks from simulation to the real world. In: 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2017, Vancouver, BC, Canada, September 24-28, 2017 (2017)
- Todorov, E., Erez, T., Tassa, Y.: MuJoCo: A physics engine for model-based control. In: 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2012, Vilamoura, Algarve, Portugal, October 7-12, 2012 (2012)
- 50. Wang, T.H., Manivasagam, S., Liang, M., Yang, B., Zeng, W., Urtasun, R.: V2VNet: Vehicle-to-vehicle communication for joint perception and prediction. In: Computer Vision - ECCV 2020 - 16th European Conference, August 23-28, 2020, Proceedings (2020)
- 51. Wang, Y., Chao, W., Garg, D., Hariharan, B., Campbell, M.E., Weinberger, K.Q.: Pseudo-lidar from visual depth estimation: Bridging the gap in 3d object detection for autonomous driving. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019 (2019)
- 52. Wheeler, T.A., Holder, M., Winner, H., Kochenderfer, M.J.: Deep stochastic radar models. CoRR (2017)
- 53. Wrenninge, M., Unger, J.: Synscapes: A photorealistic synthetic dataset for street scene parsing. CoRR (2018)
- 54. Wymann, B., Dimitrakakisy, C., Sumnery, A., Guionneauz, C.: TORCS: The open racing car simulator (2015)
- 55. Xia, F., Shen, W.B., Li, C., Kasimbeg, P., Tchapmi, M., Toshev, A., Martín-Martín, R., Savarese, S.: Interactive Gibson benchmark: A benchmark for interactive navigation in cluttered environments. IEEE Robotics Autom. Lett. (2020)
- Xia, F., Zamir, A.R., He, Z., Sax, A., Malik, J., Savarese, S.: Gibson Env: Realworld perception for embodied agents. In: 2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018 (2018)
- Yang, B., Guo, R., Liang, M., Casas, S., Urtasun, R.: Exploiting radar for robust perception of dynamic objects. In: Computer Vision - ECCV 2020 - 16th European Conference, August 23-28, 2020, Proceedings (2020)
- Yang, B., Liang, M., Urtasun, R.: HDNET: exploiting HD maps for 3d object detection. In: 2nd Annual Conference on Robot Learning, CoRL 2018, Zürich, Switzerland, 29-31 October 2018, Proceedings (2018)
- Yang, B., Luo, W., Urtasun, R.: PIXOR: real-time 3d object detection from point clouds. In: 2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018 (2018)
- Yang, Z., Chai, Y., Anguelov, D., Zhou, Y., Sun, P., Erhan, D., Rafferty, S., Kretzschmar, H.: SurfelGAN: Synthesizing realistic sensor data for autonomous driving. In: 2020 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 16-18, 2020 (2020)
- Yue, X., Wu, B., Seshia, S.A., Keutzer, K., Sangiovanni-Vincentelli, A.L.: A lidar point cloud generator: from a virtual world to autonomous driving. In: Proceedings of the 2018 ACM on International Conference on Multimedia Retrieval, ICMR 2018, Yokohama, Japan, June 11-14, 2018 (2018)
- Zhang, Z., Gao, J., Mao, J., Liu, Y., Anguelov, D., Li, C.: STINet: Spatio-temporalinteractive network for pedestrian detection and trajectory prediction. In: 2020 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 16-18, 2020 (2020)

 Zhou, Y., Tuzel, O.: Voxelnet: End-to-end learning for point cloud based 3d object detection. In: 2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018 (2018)