

AdvSim: Generating Safety-Critical Scenarios for Self-Driving Vehicles

Jingkang Wang^{1,2} Ava Pun³ James Tu^{1,2} Sivabalan Manivasagam^{1,2} Abbas Sadat²
Sergio Casas^{1,2} Mengye Ren^{1,2} Raquel Urtasun^{1,2}
University of Toronto¹, Uber ATG², University of Waterloo³

{wangjk, manivasagam, sergio, mren, urtasun}@cs.toronto.edu

a5pun@uwaterloo.ca james.tu@mail.utoronto.ca abbas.sadat@gmail.com

Abstract

As self-driving systems become better, simulating scenarios where the autonomy stack may fail becomes more important. Traditionally, those scenarios are generated for a few scenes with respect to the planning module that takes ground-truth actor states as input. This does not scale and cannot identify all possible autonomy failures, such as perception failures due to occlusion. In this paper, we propose AdvSim, an adversarial framework to generate safety-critical scenarios for any LiDAR-based autonomy system. Given an initial traffic scenario, AdvSim modifies the actors' trajectories in a physically plausible manner and updates the LiDAR sensor data to match the perturbed world. Importantly, by simulating directly from sensor data, we obtain adversarial scenarios that are safety-critical for the full autonomy stack. Our experiments show that our approach is general and can identify thousands of semantically meaningful safety-critical scenarios for a wide range of modern self-driving systems. Furthermore, we show that the robustness and safety of these systems can be further improved by training them with scenarios generated by AdvSim.

1. Introduction

Self-driving vehicles (SDV) are safety critical applications in which the comprehensive testing is necessary before real-world deployment. As the performance of self-driving systems becomes better on natural and well-behaved scenarios, it becomes of key importance to find scenarios where the system is likely to fail. However, exhaustively searching over all possible scenarios to identify safety critical ones is computationally unfeasible, as there are exponentially many scenario variations due to the combinatorial number of possible lane topologies, actor configurations, trajectories, velocity profiles, appearance of actors and background, etc.

Conventional practice in industry for comprehensive testing is a semi-autonomic process that relies on human expertise to create an initial scenario set, where each sce-




Figure 1: AdvSim modifies traffic scenarios in a physically plausible manner to produce autonomy system failure. Here the SDV collides with a crossing pedestrian after the behavior of a nearby bus is modified by AdvSim. The collision is indicated by the red arrow.

nario contains at most 1 or 2 “actors of interest” (e.g., vehicles that interact with the SDV’s planned path) with specified initial locations and trajectories [34, 26]. Scenario variations are then programmatically created by varying the actors’ locations and velocity profiles. While such scenarios are valuable, they only evaluate simple interactions with the SDV and do not test complex multi-actor interactions, such as lane-merging and unprotected left-turns in dense traffic scenes. They do not test the autonomy system on the wide variety of scenarios that the SDV may encounter. Moreover, human involvement makes this process time-consuming and difficult to scale. Manual design may also result in missing testing configurations that identify unexpected failure modes, as it is difficult to assess coverage.

To address this problem, recent works aim to automate the scenario generation procedure by searching over the set of possible scenarios and identifying high-risk ones according to a specified cost function. Most previous works [14, 43, 23, 11, 13] only consider a motion planning module that has access to the ground-truth state of the actors in the scene. This overlooks the fact that many adversarial scenarios often involve actors that are hard to identify due to occlusion, or that have trajectory plans that can be difficult to localize and forecast. Such issues in the perception and motion forecasting modules of the autonomy system can




Figure 2: **Overview of our proposed adversarial scenario generation pipeline.** Our goal is to perturb the maneuvers of interactive actors in an existing scenario with adversarial behaviors that cause realistic autonomy system failures. Given an existing scenario and its original sensor data, we perturb the scenario and update accordingly how the SDV would observe the LiDAR sensor data based on the new scene configuration. We then evaluate the autonomy system on the modified scenario, compute an adversarial objective, and update the proposed perturbation using a search algorithm.

generate compounding errors that ultimately cause planning failures. While [30, 29, 1] test end-to-end image-based self-driving systems, the adversarial scenarios they generate are either generated at a small scale [1] or with respect to simplified imitation learning models that do not reflect the autonomy system of modern self-driving vehicles [30, 29]. Additionally, most past works either modify the scenario by changing high-level actor behavior [14, 13, 22, 30, 29, 1], or create physically unrealizable trajectories [43]. This does not allow for physically-plausible fine-grained control of the actor trajectory, such as creating a nudging actor.

In contrast, we want to find complex and realistic safety-critical scenarios at scale for the full self-driving system. Towards this goal, we frame the generation of worst-case scenarios as a black box adversarial attack that can test any LiDAR-based autonomy system. We explore adversarial perturbations with respect to physically feasible changes in actor behavior, since such perturbations provide insight into the different types of driving situations that are challenging. This contrasts previous works on black box attacks for perception systems [46, 24, 50, 42, 47] that perturb appearance and texture, but do not perturb actor behavior.

In this paper, we leverage real world traffic scenarios available in standard self-driving datasets and optimize the actors’ trajectories jointly to increase the risk of an autonomy system failure. As our perturbation modifies the actors’ trajectories, we need to adjust the sensor data to accurately reflect the actors’ new locations. We therefore adopt a high-fidelity LiDAR simulator [27] that modifies the sensor data accordingly taking into account occlusions. After running the black-box autonomy system with modified sensor data as input, we obtain the planned trajectory and evaluate how adversarial the scenario was. Our adversarial objective captures multiple safety factors such as collisions, violations in traffic rules, and uncomfortable driving behaviors. We demonstrate the flexibility and scalability of our approach by generating over 4000 adversarial scenarios for a wide range of modern autonomy systems. Finally, we leverage

AdvSim-generated safety-critical scenarios in training and further improve the safety of autonomy systems.

2. Related Work

We first give an overview of the development of modern autonomy stack. We then discuss the existing works on safety-critical scenario generation for the planning module. We also review works producing physically realizable adversarial examples for self-driving perception systems.

Self-Driving System: Industry typically decomposes autonomy systems into three sequential subtasks: object detection (perception), motion forecasting (prediction), and planning. However, these components are developed separately, and cannot correct compounding errors. Another approach is end-to-end self-driving, which traces back to the seminal work ALVINN [33]. Such direct control-based methods have advanced significantly in recent years thanks to deeper network architectures, more informative sensor inputs, and scalable learning methods [6, 21, 12, 28, 18, 38, 5, 31]. Recently, interpretable neural motion planners [48, 49, 36] provide an alternative that inherits the advantages of traditional pipelines and end-to-end approaches, by maintaining modularity and interpretability while enabling end-to-end learning. This first began with joint perception and prediction [25, 10], which neural planners extended to include planning. Specifically, NMP [48] shared feature representations between multiple subtasks and predicted a cost volume to represent the quality of possible locations in planning. DSDNet [49] proposed an energy-based model to parameterize the joint distribution of the actors’ future trajectories. P3 [36] developed a semantic occupancy representation and generated consistent ego-vehicle plans. Our work evaluates a wide range of autonomy systems, including modular and end-to-end interpretable ones.

Safety-Critical Scenario Generation: There are three main components for generating safety-critical scenarios: a scenario parameterization space to optimize over, a search

Method	LiDAR	Autonomy	Traj.-based	Real data	Scenarios	# Actors
BOAdv [1]		✓			2	2
Chen et al. [11]			✓		1	3
Klischat et al. [22]				✓	2	13
O’Kelly/Norden et al. [29, 30]		✓	✓		1	5
AdvSim (ours)	✓	✓	✓	✓	4,000+	5

Table 1: **Comparison with prior works.** AdvSim produces safety-critical scenarios that are physically plausible and adversarial to full LiDAR-based autonomy systems in scale.

algorithm that identifies critical scenario parameters, and an evaluation setting to evaluate the system under test. Previous works represent the action space of other agents either as a Frenet frame [22], initial position and velocity [13, 14], a high-level graph-based route [1], or steering and acceleration [43]. We choose to represent the behavior of actors as kinematic bicycle-model trajectories, allowing for physical feasibility and fine-grained behavior control. There are many potential choices of the search algorithm used to identify scenarios that cause autonomy failure, such as policy gradient [11, 43, 14], Bayesian optimization [1], evolutionary algorithms [22], and variants of monte-carlo sampling [13, 29, 39]. We build a general scenario generation algorithm and benchmark a wide variety of black-box search algorithms, providing insight into which search algorithms are effective. [11, 43, 14, 22] evaluate planners assuming groundtruth perception. [1, 29] use CARLA [15] to evaluate an image-based SDV planning algorithm. These works consider simplified planning modules or image-based systems that do not reflect modern autonomy systems. Moreover, they generate scenarios only for a handful of scene configurations. In contrast, we present an end-to-end adversarial scenario generation system that takes into account failures of the full autonomy stack. Our approach scales to datasets with diverse traffic patterns and map configurations. We summarize the differences of our paradigm with prior works in Table 1.

Physically Realizable Adversarial Examples: Physical adversarial perturbations expose real world threats for perception. Most previous works deal with threat models in the image space [7, 3, 47, 16] by imposing physical constraints such as different view angles and distances [7, 3] and color distortion [16, 47]. Recently, other works perturb meshes [46] or photo-realistic properties such as surface normals and lighting conditions [24, 50]. In self-driving, recent works focus on the robustness of LiDAR-based perception. Specifically, [9, 42] generate objects that are either invisible or detrimental to object detectors and [8, 41] directly spoof the LiDAR points using laser devices. [51, 45] craft adversarial vehicle textures in CARLA. We instead generate realistic adversarial maneuvers by other agents such that

Algorithm 1 Generating Adversarial Scenarios

Require: Sensory input \mathbf{x} , initial state \mathbf{s}_0 of the perturbed actor, adversarial objective \mathcal{L}_{adv} , number of queries N .

- 1: Pick the perturbed actor \mathbf{v}_{adv} heuristically
 - 2: Generate physically plausible trajectories set \mathcal{T}_{adv}
 - 3: Initialize observation set $\mathcal{H} = \emptyset$
 - 4: **for** $k = 1, \dots, N$ **do**
 - 5: Select $\delta^{(k)}$ based on black-box algorithms and historical observations \mathcal{H} .
 - 6: $\tau_{\text{adv}}^{(k)} = \Pi_{\tau \in \mathcal{T}_{\text{adv}}} [\text{BICYCLE}(\mathbf{s}_0, \delta^{(k)})] \triangleright (\text{Sec 3.2})$
 - 7: $\mathbf{x}_{\text{adv}}^{(k)} = f(\mathbf{x}, \tau_{\text{adv}}^{(k)}, \tau_{\text{sdv}})$ $\triangleright (\text{Sec 3.3})$
 - 8: Run the autonomy system and obtain the optimal SDV plan $\tau_0^{(k)} = \tau_0^*(\mathbf{x}_{\text{adv}}^{(k)})$
 - 9: Calculate the adversarial loss of the optimal plan: $\mathcal{L}_{\text{adv}}^{(k)} = \mathcal{L}_{\text{adv}}(\tau_0^{(k)}, \mathbf{x}_{\text{adv}}^{(k)}) \triangleright (\text{Sec 3.4})$
 - 10: Update observation set $\mathcal{H} = \mathcal{H} \cup \{(\tau_{\text{adv}}^{(k)}, \mathcal{L}_{\text{adv}}^{(k)})\}$
 - 11: **end for**
 - 12: $\tau_{\text{adv}}^* = \arg \max_{\tau_{\text{adv}}^{(k)}, k \in [N]} \mathcal{L}_{\text{adv}}^{(k)}$
-

the scenario is safety critical for autonomy.

3. Generating Safety-Critical Scenarios

Our objective is to generate realistic challenging scenarios that cause autonomy system failure. We frame our objective as a black box adversarial attack that exercises every component of the autonomy system, including object detection, motion forecasting and motion planning. As we search over the space of realistic perturbations in actor motions of an existing scenario, we must update the sensor data that the SDV observes and then evaluate the autonomy system. Our approach, AdvSim, works as follows: we first perturb the actors’ motion trajectories in an existing scenario, and generate the sequence of LiDAR point clouds that reflect the change in actor locations. With the adjusted sensor data, we run the autonomy stack and get the planned SDV motion path. Finally, we evaluate the output path with a proposed adversarial objective and adjust the scenario perturbation to be more challenging. An overview is shown in Fig. 2.

In what follows, we first define the autonomy system and our attack formulation in Sec 3.1. We then describe how we parameterize the adversarial actors’ behaviors (Sec 3.2) and conduct realistic LiDAR simulation to generate new LiDAR sweeps (Sec 3.3). Finally, we describe our adversarial objective and the suite of black-box optimization algorithms we benchmark to generate worst-case behaviors in Sec 3.4.

3.1. Problem Setup

Let $\mathcal{V} = \{\mathbf{v}_0, \mathbf{v}_1, \dots, \mathbf{v}_M\}$ be the set of vehicles that compose the scene, where \mathbf{v}_0 denotes the SDV, M is the number of other vehicles. The objective of a self-driving

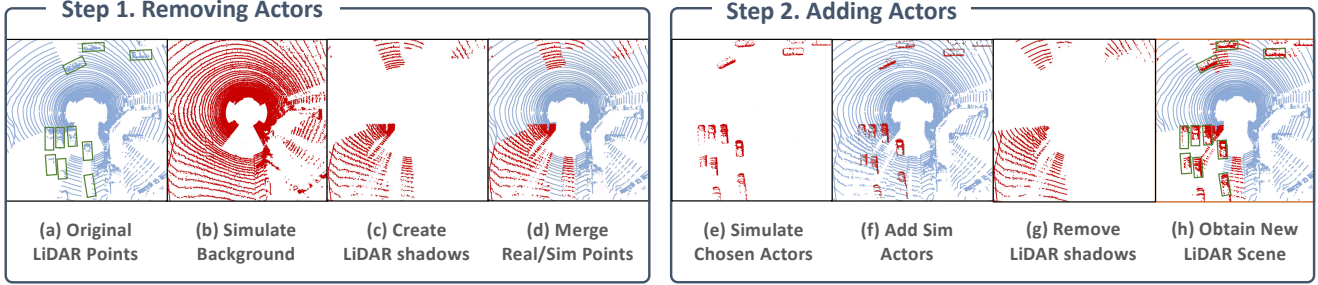


Figure 3: **Realistic LiDAR simulation for scenario perturbations.** Given a scenario perturbation on the actors’ motions, the previously recorded LiDAR data is modified to accurately reflect the updated scene configuration. We remove the original actor LiDAR observations and replace with simulated actor LiDAR observations at the perturbed locations, while ensuring sensor realism. The above example perturbs all actors left by 5 meters.

system is to find the best planned trajectory τ_0^* according to a cost function \mathcal{C} that comfortably and safely maneuvers around the scene, given the available sensor data inputs \mathbf{x} :

$$\tau_0^*(\mathbf{x}) = \arg \min_{\tau_0} \mathcal{C}(\tau_0, \mathbf{x}) \quad (1)$$

where τ_0 is the SDV’s planned trajectory. As \mathbf{x} consists of raw sensor data (i.e., LiDAR point clouds), High-Definition maps, and other relevant information (e.g., previous SDV states, traffic light states), this minimization represents the full autonomy system, not just the planning module.

Our goal is to increase the risk of the self-driving car by perturbing the behaviors of other actors in a physically plausible manner for an existing traffic scenario. Without loss of generality, we consider perturbing a single actor in the following discussion for brevity, but we apply AdvSim for multi-actor perturbations in experiments.

We characterize the behavior of an adversary by the trajectory τ_{adv} it will take in the future. As the perturbed actor’s trajectory τ_{adv} differs from its original behavior in the sensor data, the vehicle position and the occlusions it generates will change (see Fig. 3). Therefore, we must simulate the new LiDAR data given the adversary trajectory τ_{adv} and SDV trajectory τ_{sdv} to evaluate the system (Eq. 1). The generation of point clouds in the perturbed traffic scene is given as follows:

$$\mathbf{x}_{\text{adv}} = f(\mathbf{x}, \tau_{\text{adv}}, \tau_{\text{sdv}}) \quad (2)$$

where $f(\cdot)$ denotes the realistic LiDAR simulation (Sec 3.3) for perturbed input \mathbf{x}_{adv} given the adversary’s trajectory and original sensor data sequence \mathbf{x} .

We then define an adversarial objective \mathcal{L}_{adv} which we maximize to generate scenarios as follows

$$\tau_{\text{adv}}^* = \arg \max_{\tau_{\text{adv}}} \mathcal{L}_{\text{adv}}(\tau_0^*, \mathbf{x}_{\text{adv}}), \quad (3)$$

where $\tau_0^* = \tau_0^*(\mathbf{x}_{\text{adv}})$ is the optimal SDV’s planned trajectory under simulated scene \mathbf{x}_{adv} . The design of the adversarial loss \mathcal{L}_{adv} is deferred to Sec 3.4.

3.2. Modeling Adversarial Behaviors

To produce physically feasible actor behaviors, we parameterize the trajectory $\tau_{\text{adv}} = \{\mathbf{s}_t\}_{t=0}^T$ as a sequence of kinematic bicycle model states $\mathbf{s}_t = \{x_t, y_t, \theta_t, v_t, \kappa_t, a_t\}$ in the next T timesteps. Here (x, y) is the center position of the perturbed actor, θ is the heading, v and a are the forward velocity and acceleration, and κ is the vehicle path’s curvature. Candidate adversary trajectories can be generated by perturbing the change of curvature $\dot{\kappa}_t$ and acceleration values a_t within set bounds at different timesteps, and using the kinematic bicycle model to compute the other states [32].

Moreover, to enlarge the space of sampled adversarial behaviors, we also allow the perturbation of initial states $(x_0, y_0, \theta_0, v_0)$ within set bounds. In summary, the perturbation space can be depicted as

$$\delta = \{\Delta \mathbf{s}_0, (a_0, \dot{\kappa}_t|_{t=0}), \dots, (a_{T-1}, \dot{\kappa}_t|_{t=T-1})\}.$$

To increase the perturbed trajectory’s plausibility, we ensure it does not collide with other actors or the original expert trajectory of the SDV. In practice, we do this by first performing rejection sampling to create a set of physically feasible trajectories \mathcal{T}_{adv} and then projecting the trajectory generated by δ on to the physically feasible set, measured by L_2 distance. Our search space is low-dimensional and conducive to query-based black box optimization, while still allowing for fine-grained actor motion control.

3.3. Realistic LiDAR Simulation

Given an initial traffic scenario and the corresponding adversarial perturbation to the actors’ behaviors, we discuss how we modify the existing real LiDAR sweeps to reflect the perturbation. We adopt the high-fidelity LiDARsim [27] simulator, which leverages real world data to generate realistic background meshes and dynamic object assets, and then applies physics-based raycasting and machine learning to generate realistic LiDAR point clouds. Given a modified scene configuration, we use LiDARsim to render a simulated point cloud, and then update the real LiDAR sweep

with the modified regions. We choose to update the sensor data for modified regions only rather than generating the full sweep to speed up the query function f in Eq. 2. Specifically, we cache the simulated background LiDAR (Fig. 3 Step 1b) as the SDV trajectory is fixed during the actor perturbation. The sensor perturbation is illustrated in Fig. 3.

Modifying the LiDAR sensor data to reflect the scenario perturbation is non-trivial, as the LiDAR’s sensing characteristics cause specific visibility artifacts that should exist in the generated scene to be realistic and physically accurate. We perform two main steps for sensor simulation for modified scenarios: *actor removal* (removing the existing benign actors’ LiDAR point cloud and filling the LiDAR shadow created) and *actor addition* (inserting the adversarial actors’ LiDAR point cloud, while accounting for occlusion).

Removing Actors: Given an original LiDAR point cloud (Fig. 3a), we first remove the points within the bounding boxes of perturbed actors and simulate background points (Fig. 3b) using LiDARsim’s background mesh. We then convert the simulated and real LiDAR sweeps into a range image, allowing us to identify the specific rays missing in the real LiDAR sweep (Fig. 3c) that exist in the simulated LiDAR. By taking the element-wise minimum ray distance between the range images, we can merge the LiDAR point clouds. Fig. 3d shows the synthetic point clouds (red: simulated points; blue: original real points) after *actor removal*.

Adding Actors: Once we have removed the selected actors from the LiDAR sweep, we update the LiDAR with the actors at their new locations. We first render the simulated LiDAR for the actors at their new locations using LiDARsim’s vehicle asset bank (Fig. 3e). Fig. 3f shows the real LiDAR point cloud with the added actors. However, when a LiDAR ray hits an object, the remaining path of the ray becomes occluded, creating a LiDAR shadow. Similar to the actor removal process, we create range images of the simulated and real LiDAR, and merge the LiDAR point clouds, thereby removing the LiDAR points of the now-occluded regions (Fig. 3g) and obtaining the final modified LiDAR sweep (Fig. 3h). The generated scenes are realistic and match the desired perturbation in actors’ motions (Fig. 3).

3.4. Adversarial Scenario Search

Since we aim for a general adversarial scenario generation framework, we consider the autonomy system as a black box, where we access the evaluation scores through limited queries. Our goal is to find the perturbation that maximizes the SDV’s planned trajectory cost. In this section, we introduce the adversarial objective we optimize to produce worst-case scenarios and detail the search algorithms applied. We then summarize the AdvSim algorithm.

Adversarial Objective: To induce autonomy system failures, we propose a combination of three costs as our adversarial loss function. These costs are similar to those autonomy systems [37, 52] attempt to minimize over in Eq. 1. We first include l_{IL} , an imitation-learning based cost that encourages the SDV’s output plan to deviate from the recorded human trajectory in the original scenario. We compute this as a smooth ℓ_1 distance between output trajectory τ_0^* and the ground-truth for the entire scenario horizon. We also compute a cumulative collision cost l_{col}^t that encourages the perturbation to cause the SDV to collide with other actors in the scene. Finally, we add a safety cost $c_s^t(\mathbf{x}_{adv}, \tau_0^*)$ that encourages the output plan τ_0^* to have lane violations and be dangerous (i.e. high accelerations and jerk) at each timestep t . The full adversarial loss is defined as:

$$\mathcal{L}_{adv} = \min_{\tau_0} \left[l_{IL} + \sum_t l_{col}^t + \sum_t c_s^t(\mathbf{x}_{adv}, \tau_0^*) \right]$$

Our use of multiple different costs allows us to identify different types of autonomy system failures, such as unnatural trajectories, collisions, and hard braking.

Search Algorithms: AdvSim is a framework that can use any black-box search algorithm to identify autonomy system failures. The search algorithm attempts to find the safety critical scenarios by maximizing the adversarial objective \mathcal{L}_{adv} in Eq. 2. The search algorithm queries the autonomy system with a candidate perturbation τ_{adv} to obtain a query pair $(\tau_{adv}, \mathcal{L}_{adv})$ and maintains a history \mathcal{H} of past query pairs to generate the next candidate perturbation. We study a wide variety of black-box search algorithms including (1) Bayesian optimization [40, 35] (BO), (2) genetic algorithms [2] (GA), (3) random search [17] (RS) and (4) gradient estimation methods (NES [19] and Bandit-TD [20]). Specifically, BO maintains a surrogate model and select the next candidate based on the acquisition function and current model states. For GA algorithms, a group of candidate trajectories are evolved to maximize the objective and the best candidate is preserved at each iteration. For RS, the perturbations sampled from a pre-defined orthonormal basis are added or subtracted to original input iteratively. Another branch of query-based black-box search algorithms estimate the gradient through the target model. Specifically, NES maximizes the expectation of the objective under one search distribution and Bandit-TD further leverages temporal information to improve the query efficiency.

Overall Adversarial Scenario Generation Algorithm: We summarize our proposed AdvSim framework in Algorithm 1. Given an initial traffic scene, we pick the actors to be perturbed using heuristics, such as the closest reachable actors, and then sample physically plausible trajectories \mathcal{T}_{adv} to ensure that our perturbations remain in this set.

	Collision (%)		Planning Metrics ↓				Perception ↑	Prediction Metrics	
	up to 3s	up to 5s	L2 human		Comfortable		AP / F1 occ. (%)	L2 center / F1 occ.	
			@3s	@5s	Lat. (m/s^2)	Jerk (m/s^3)		@3s	@5s
IL: End-to-end Imitation Learning [6]									
Original	5.7	20.4	4.25	8.74	–	–	–	–	–
Adv. scenes	8.9 (+3.2)	25.2 (+4.8)	5.48 (+1.30)	10.77 (+2.03)	–	–	–	–	–
NMP: Neural Motion Planner [48]									
Original	2.6	14.5	3.30	9.31	3.46	6.08	AP (IoU = 0.7)	L2 center (↓)	
Adv. scenes	14.2 (+11.6)	42.2 (+27.7)	4.82 (+1.52)	13.24 (+3.93)	3.59 (+0.13)	7.28 (+0.97)	72.7 (-9.0)	1.63 (+0.20)	3.21 (+0.41)
PLT: Jointly Learnable Behavior and Trajectory Planning [37]									
Original	1.3	12.1	1.74	5.10	3.22	1.91	AP (IoU = 0.7)	L2 center (↓)	
Adv. scenes	5.5 (+4.2)	31.6 (+19.5)	2.55 (+0.81)	7.32 (+2.22)	3.51 (+0.29)	2.38 (+0.47)	75.2 (-7.7)	1.96 (+0.21)	3.75 (+0.43)
P3: Perceive, Predict, and Plan [36]									
Original	0.8	8.5	1.58	4.74	3.09	1.79	F1 occ. (@0s)	F1 occ. (% , ↑)	
Adv. scenes	3.9 (+2.9)	32.0 (+23.5)	2.28 (+0.70)	6.63 (+1.89)	3.18 (+0.09)	2.18 (+0.30)	63.2 (-1.7)	47.5 (-2.2)	45.9 (-2.7)

Table 2: Evaluation of modern autonomy systems on original and AdvSim generated scenarios.

Training	Testing	Collision (% up to 5s)	L2 human @5s	Lat. acc. (m/s^3)	Jerk (m/s^2)
Standard	Original	8.5	4.74	3.09	1.79
	Adv. scenes	32.0	6.63	3.18	2.18
CL train	Original	7.0	4.88	3.28	1.81
	Adv. scenes	21.1	5.98	3.33	2.08
Robust train	Original	7.3	4.76	3.29	1.75
	Adv. scenes	17.7	5.72	3.31	2.04

Table 3: Robust training P3 with augmented scenarios.

We then obtain the perturbation $\delta^{(k)}$ at iteration k based on historical observations \mathcal{H} using a selected black-box search algorithm (L. 5). We roll out the kinematics bicycle model states with initial state s_0 and the perturbation $\delta^{(k)}$, and project onto the feasible set \mathcal{T}_{adv} to obtain the adversarial trajectories for the perturbed actors (L. 6). After that, we update the sensor data accordingly (L. 7) and evaluate the full autonomy system on generated scenarios to compute \mathcal{L}_{adv} (L. 8-9). Finally, after running the procedure for N iterations, we obtain the adversarial behaviors of perturbed actors as well as corresponding simulated LiDAR data.

4. Experiments

We now showcase applying AdvSim to generate worst-case scenarios for several autonomy systems using a large scale self-driving dataset. In Sec. 4.1 we discuss the dataset and the autonomy systems under test, as well as how we evaluate the generated scenarios’ effectiveness. In Sec. 4.2 we analyze how the AdvSim-generated scenarios affect autonomy performance. We also show robust training on the generated scenarios improves autonomy systems.

4.1. Experimental Setup

Dataset We evaluate our approach on a self-driving dataset, *UrbanScenarios*, which has 5,000 driving logs of 25 seconds each. Our dataset is collected across multiple cities in North America, and contains different types of map layouts and varying traffic densities. We curate the dataset

and select interesting candidate scenarios to apply AdvSim on, where the SDV in the original scenario “interacts” with other vehicles. Specifically, we sample 100 trajectories per SDV behavior (e.g., keep lane, lane change) in the SDV’s Frenet frame [44] and calculate the trajectory collision rate with other actors’ motion paths. We select the 6s scenario from each log that has the largest collision rate. After data curation, we obtain 3953 train and 409 val scenarios.

Autonomy Systems We evaluate the effectiveness of the proposed framework on the following models: (a) **Imitation Learning (IL)**, where the future states of the SDV are predicted directly from the fused LiDAR and map features with L_2 loss; (b) **PLT [37]**, a modular autonomy system where the detection and prediction are trained jointly with the backbone used in [36], and the planning is accomplished using a learnable combination of interpretable safety costs; (c, d) **NMP [48]** and **P3 [36]**, two end-to-end interpretable motion planners. NMP predicts a cost-map directly from fused features with detection and prediction jointly trained as auxiliary tasks. P3 predicts a novel differentiable semantic occupancy representation used as safety-cost for planning. Please see supplementary for implementation details.

Metrics: In this paper, we focus on an open-loop scenario evaluation setting, in which the evaluated autonomy system takes the past 1s LiDAR data as input and outputs a 5s trajectory plan. We then unroll the 5s plan and the other actors’ trajectories for 5s and evaluate the autonomy system’s performance during that time. Following [49, 37, 36], we adopt standard planning metrics to measure the autonomy systems’ performance on our adversarial scenarios, and compare how much more challenging they are relative to the original scenario set. Specifically, **collision rate** is the percentage of scenarios that cause the SDV to collide with another actor during a certain time frame (up to 3 or 5s). **L_2 distance to the human trajectory** represents how well the model imitated the human driving. **Jerk** and **lateral acceleration** indicate how comfortable the planned SDV trajec-

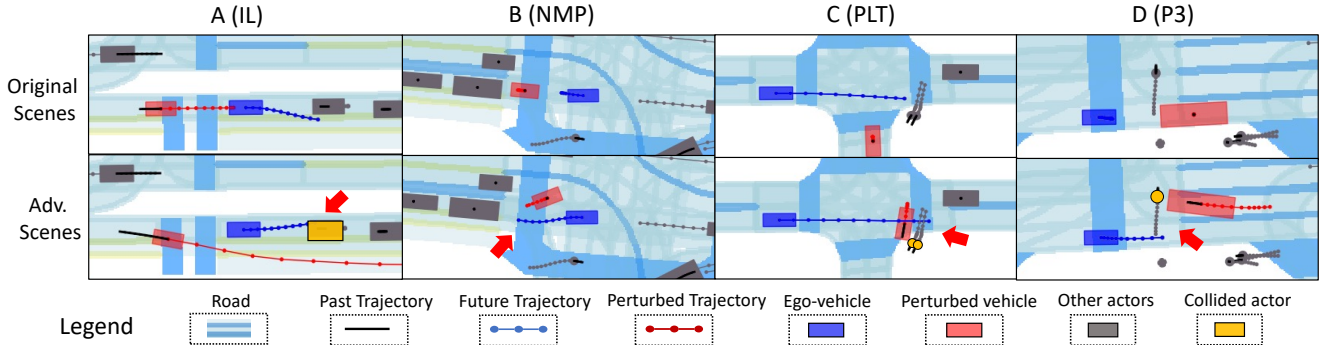


Figure 4: Visualization of autonomy system’s output plan on original and corresponding adversarial scenes. **A**: IL avoids the high-speed lane-changing vehicle behind but collides with the front one. **B**: NMP collides with the merging vehicle. **C**: PLT collides with one vehicle and two occluded pedestrians at crossroads. **D**: P3 collides with the crossing pedestrian.

tories are. Additionally, all evaluated autonomy systems, except for **IL**, generate intermediate perception and prediction representations. We therefore also report perception and motion forecasting metrics to see how the adversarial scenarios specifically impact these system components. For NMP and PLT planners, for detection we report Average Precision (**AP**) of bounding box detections at an Intersection of Union (IoU) of 0.7, and for motion forecasting we report the L_2 **prediction error** of predicted actor trajectories at future timesteps. Similarly, for the P3 planner, we adopt the F1-score of the occupancy prediction at different timesteps ($t = 0$ for detection and $t > 0$ for prediction). All reported metrics are for vehicles in the ego-coordinate view range of $x \in [-72m, 72m]$, $y \in [-40m, 40m]$.

4.2. Experimental Results

Evaluations on Modern Autonomy Systems: We evaluate our AdvSim framework on the autonomy systems in Table 2. Here, we only consider perturbing a single vehicle in the traffic scene and adopted Bayesian optimization (BO) [40, 35] as the black-box search algorithm. Experiments show that AdvSim is effective in generating challenging scenarios for different systems, leading to an average collision rate of 32.8%, over 200% more compared to the original set. The adversarial scenarios also provide insights to how different autonomy systems compare. While NMP and PLT have similar drops in perception and motion forecasting, we observe that PLT is more robust to AdvSim scenarios than NMP on planning metrics, indicating the benefits of using a learning-based and hand-crafted cost function. We show qualitative examples of adversarial scenarios for each autonomy system in in Fig. 4. Unlike prior work, our generated scenarios cause the SDV to collide with other non-perturbed actors in the scene (Fig. 4, A, C, D).

Safer Planner with Challenging Scenarios: We investigate whether the robustness of the autonomy systems can be improved with our generated scenarios. We test several

Source \ Target	IL	NMP	PLT	P3
IL [6]	<u>25.2%</u>	17.1%	9.5%	6.8%
NMP [48]	24.9%	<u>42.2%</u>	14.5%	14.0%
PLT [37]	21.2%	21.5%	<u>31.6%</u>	13.1%
P3 [36]	26.4%	<u>24.2%</u>	<u>22.2%</u>	<u>32.0%</u>

Table 4: Transferability of generated safety-critical scenarios across different autonomy systems.

training schemes. First, we propose a curriculum learning (CL) [4] baseline where we first train on standard examples till convergence (*easy examples*), and then train on real challenging scenarios selected based on reachable actors (Sec 4.1) (*hard examples*). Then, we propose a robust-training approach to leverage simulated worst-case scenarios. Specifically, we use AdvSim to generate a large number of adversarial scenarios to augment the training data. As discussed in Sec 3.2, the original expert trajectory is still a valid planning solution to mimic in the new scenario with respect to collisions, as we impose constraints on the perturbation. This allows us to re-train autonomy systems with scenarios produced by AdvSim using the same expert trajectories as ground truth. Table 3 demonstrates that CL helps improve the performance on *both* original and adversarial scenarios. Robust training with AdvSim-generated scenarios further improves performance across most planning metrics, highlighting the value of AdvSim scenarios for improving autonomy systems. In Fig. 5, we show a qualitative example of an adversarial scenario for standard P3 and robust-trained P3 with AdvSim. Finally, we investigate the generalization of robust training by evaluating robustly trained PLT and P3 with safety-critical scenarios generated for other autonomy systems (see supplementary).

Transferability of Adversarial Scenarios: We study the transferability of adversarial scenarios across different autonomy systems in Table 4, where *Source* denotes the autonomy system used to identify failure scenarios and *Target*




Figure 5: Visualization of standard P3 and robust-trained P3 with AdvSim on one challenging scenario. Standard P3 changes lane to avoid the reversing bus yet is rear-ended from behind. After robust training with AdvSim, the system bypasses the bus smoothly and returns to its original lane.

#Actors	Collision (% , 5s)	L2 human @5s	Lat. acc. (m/s^2)	Jerk (m/s^3)	F1 occ. (% @5s)
Original	8.5	4.74	3.09	1.79	48.6
$m = 1$	32.0	6.63	3.18	2.18	45.9
$m = 2$	33.5	7.21	3.28	2.26	42.1
$m = 3$	30.6	7.10	3.38	2.24	39.2
$m = 4$	28.4	7.36	3.73	2.46	33.5
$m = 5$	26.9	7.86	4.62	2.62	30.5

Table 5: Generating safety-critical scenarios for P3 with different number of perturbed actors m .

denotes the autonomy system evaluated on these scenarios. We report the cumulative collision rate up to 5s. Results on additional metrics are in supplementary. Table 4 shows that generating adversarial scenarios with the same target autonomy system usually leads to the best performance. Scenarios simulated w.r.t the more robust P3 system (see original in Table 2) have stronger transferability.

Investigation on Attack Configurations: We now study how the number of the perturbed actors and the search algorithm used affect scenario generation. In Table 5, the planning metrics (L2 human, jerk and lateral acceleration) and the prediction metrics become worse as the number of perturbed actors m increase. This indicates AdvSim generated more complicated traffic configurations when perturbing multiple actors simultaneously. However, we observed the collision rate decreases when $m \geq 3$. This may be because we sampled finite trajectories for each actor, and attacking with multiple actors simultaneously increases the difficulty of finding physical plausible candidates to optimize over. We also benchmark a wide range of black-box algorithms in Table 6. See supplementary for the implementation details. We found that BO [35] is most efficient, since the perturbation space is low dimensional and the cost measure is not smooth w.r.t the perturbation (thus harder for gradient-estimation based approaches).

Ablation Studies: We conduct ablation studies on proposed adversarial objective. As shown in Table 7, imitation-

Algorithms	Collision (% , 5s)	L2 human @5s	Lat. acc. (m/s^2)	Jerk (m/s^3)	#Query	GPU Hour
Original	8.5	4.74	3.09	1.79	-	-
GA [2]	23.3	6.29	3.13	2.10	1600	1.33
NES [19]	19.1	6.07	3.17	2.05	400	0.33
Bandit-TD [20]	14.6	5.90	3.12	2.02	100	0.08
RS [17]	14.7	4.35	3.14	1.97	100	0.08
BO [35]	28.5	6.63	3.18	2.18	75	0.06

Table 6: Comparisons of different blackbox algorithms in scenario generation for P3.

#ID	IL ℓ_{im}	Collision $\sum_t \ell_{col}^t$	Safety $\sum_t c_s^t$	Collisions (% up to 5s)	L2 human @5s	Lat. acc. (m/s^3)	Jerk (m/s^2)
\mathcal{M}_0	✓	✓	✓	29.5	6.12	3.33	2.17
\mathcal{M}_1	✓			14.8	7.06	3.22	2.31
\mathcal{M}_2		✓		30.6	5.31	3.01	1.86
\mathcal{M}_3	✓	✓		32.0	6.63	3.18	2.18
\mathcal{M}_4			✓	9.0	5.30	3.52	2.08
\mathcal{M}_5		✓	✓	30.6	5.61	3.44	2.07

Table 7: Ablation studies on adversarial objective design.

	Collisions (% up to 5s)	L2 human @5s	Lat. acc. (m/s^3)	Jerk (m/s^2)
Closest distance [1]	13.5	5.57	3.16	2.00
Max-margin loss [36]	18.3	5.76	3.02	1.96
\mathcal{M}_3	32.0	6.63	3.18	2.18

Table 8: Comparison with other adversarial objective loss.

learning based cost ℓ_{IL} , cumulative collision cost $\sum_t \ell_{col}^t$ and safety cost $\sum_t c_s^t$ are optimized for L_2 human, collisions, and comfort planning metrics, respectively. The hybrid loss function (\mathcal{M}_0) can generate worst-case scenarios with respect to multiple metrics. If some planning metrics are particularly interesting in practice (e.g., collisions for testing), we could use a subset of the proposed costs. Unless otherwise stated, we adopt \mathcal{M}_3 in other experiments since the collisions are of key importance in evaluating autonomy systems. Furthermore, we compare \mathcal{M}_0 with other baseline adversarial objective in Table 8: (1) minimizing the closest distance to the ego-car [1], (2) maximize the training cost proposed in [36]. Experiments show our design outperforms other baselines on all planning metrics.

5. Conclusion

In this work we present a novel adversarial framework to generate worst-case scenarios for modern autonomy systems. Our approach identifies physically plausible failure cases that impose risks to full autonomy stack by simulating the sensor data based on the perturbed behaviors. We demonstrate that AdvSim can generate failure cases at scale for a wide range of systems. More importantly, we leverage these scenarios in training to further improve the robustness and safety of the autonomy system. We hope that leveraging this framework will allow for safer self-driving vehicles.

References

- [1] Yasasa Abeysirigoonawardena, Florian Shkurti, and Gregory Dudek. Generating adversarial driving scenarios in high-fidelity simulators. In *ICRA*, 2019. 2, 3, 8
- [2] Moustafa Alzantot, Yash Sharma, Supriyo Chakraborty, Huan Zhang, Cho-Jui Hsieh, and Mani B. Srivastava. Genattack: practical black-box attacks with gradient-free optimization. In *GECCO*, 2019. 5, 8
- [3] Anish Athalye, Logan Engstrom, Andrew Ilyas, and Kevin Kwok. Synthesizing robust adversarial examples. In *ICML*, 2018. 3
- [4] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *ICML*, 2009. 7
- [5] Alex Bewley, Jessica Rigley, Yuxuan Liu, Jeffrey Hawke, Richard Shen, Vinh-Dieu Lam, and Alex Kendall. Learning to drive from simulation without real world labels. In *ICRA*, 2019. 2
- [6] Mariusz Bojarski, Davide Del Testa, Daniel Dworakowski, Bernhard Firner, Beat Flepp, Prasoon Goyal, Lawrence D Jackel, Mathew Monfort, Urs Muller, Jiakai Zhang, et al. End to end learning for self-driving cars. *arXiv preprint arXiv:1604.07316*, 2016. 2, 6, 7
- [7] Tom B. Brown, Dandelion Mané, Aurko Roy, Martín Abadi, and Justin Gilmer. Adversarial patch. *CoRR*, abs/1712.09665, 2017. 3
- [8] Yulong Cao, Chaowei Xiao, Benjamin Cyr, Yimeng Zhou, Won Park, Sara Rampazzi, Qi Alfred Chen, Kevin Fu, and Z. Morley Mao. Adversarial sensor attack on lidar-based perception in autonomous driving. In *CCS*, 2019. 3
- [9] Yulong Cao, Chaowei Xiao, Dawei Yang, Jing Fang, Ruigang Yang, Mingyan Liu, and Bo Li. Adversarial objects against lidar-based autonomous driving systems. *CoRR*, abs/1907.05418, 2019. 3
- [10] Sergio Casas, Wenjie Luo, and Raquel Urtasun. Intentnet: Learning to predict intention from raw sensor data. In *CoRL*, 2018. 2
- [11] Baiming Chen and Liang Li. Adversarial evaluation of autonomous vehicles in lane-change scenarios. *CoRR*, abs/2004.06531, 2020. 1, 3
- [12] Felipe Codevilla, Matthias Müller, Antonio López, Vladlen Koltun, and Alexey Dosovitskiy. End-to-end driving via conditional imitation learning. In *ICRA*, 2018. 2
- [13] Wenhao Ding, Baiming Chen, Bo Li, Kim Ji Eun, and Ding Zhao. Multimodal safety-critical scenarios generation for decision-making algorithms evaluation, 2020. 1, 2, 3
- [14] Wenhao Ding, Minjun Xu, and Ding Zhao. Learning to collide: An adaptive safety-critical scenarios generating method. *CoRR*, abs/2003.01197, 2020. 1, 2, 3
- [15] Alexey Dosovitskiy, Germán Ros, Felipe Codevilla, Antonio M. López, and Vladlen Koltun. CARLA: an open urban driving simulator. In *CoRL*, 2017. 3
- [16] Kevin Eykholt, Ivan Evtimov, Earlene Fernandes, Bo Li, Amir Rahmati, Chaowei Xiao, Atul Prakash, Tadayoshi Kohno, and Dawn Song. Robust physical-world attacks on deep learning visual classification. In *CVPR*, 2018. 3
- [17] Chuan Guo, Jacob R. Gardner, Yurong You, Andrew Gordon Wilson, and Kilian Q. Weinberger. Simple black-box adversarial attacks. In *ICML*, 2019. 5, 8
- [18] Jeffrey Hawke, Richard Shen, Corina Gurau, Siddharth Sharma, Daniele Reda, Nikolay Nikolov, Przemyslaw Mazur, Sean Micklethwaite, Nicolas Griffiths, Amar Shah, and Alex Kendall. Urban driving with conditional imitation learning. In *ICRA*, 2020. 2
- [19] Andrew Ilyas, Logan Engstrom, Anish Athalye, and Jessy Lin. Black-box adversarial attacks with limited queries and information. In *ICML*, 2018. 5, 8
- [20] Andrew Ilyas, Logan Engstrom, and Aleksander Madry. Prior convictions: Black-box adversarial attacks with bandits and priors. In *ICLR*, 2019. 5, 8
- [21] Alex Kendall, Jeffrey Hawke, David Janz, Przemyslaw Mazur, Daniele Reda, John-Mark Allen, Vinh-Dieu Lam, Alex Bewley, and Amar Shah. Learning to drive in a day. *arXiv preprint arXiv:1807.00412*, 2018. 2
- [22] Moritz Klischat and Matthias Althoff. Generating critical test scenarios for automated vehicles with evolutionary algorithms. In *IV*, 2019. 2, 3
- [23] Markus Koschi, Christian Pek, Sebastian Maierhofer, and Matthias Althoff. Computationally efficient safety falsification of adaptive cruise control systems. In *ITSC*, pages 2879–2886. IEEE, 2019. 1
- [24] Hsueh-Ti Derek Liu, Michael Tao, Chun-Liang Li, Derek Nowrouzezahrai, and Alec Jacobson. Beyond pixel norm-balls: Parametric adversaries using an analytically differentiable renderer. In *ICLR*, 2019. 2, 3
- [25] Wenjie Luo, Bin Yang, and Raquel Urtasun. Fast and furious: Real time end-to-end 3d detection, tracking and motion forecasting with a single convolutional net. In *CVPR*, 2018. 2
- [26] Alexis C. Madrigal. Waymo built a secret world for self-driving cars, 2018. 1
- [27] 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. 2, 4
- [28] Matthias Müller, Alexey Dosovitskiy, Bernard Ghanem, and Vladlen Koltun. Driving policy transfer via modularity and abstraction. *arXiv preprint arXiv:1804.09364*, 2018. 2
- [29] Justin Norden, Matthew O’Kelly, and Aman Sinha. Efficient black-box assessment of autonomous vehicle safety. *arXiv preprint arXiv:1912.03618*, 2019. 2, 3
- [30] Matthew O’Kelly, Aman Sinha, Hongseok Namkoong, Russ Tedrake, and John C Duchi. Scalable end-to-end autonomous vehicle testing via rare-event simulation. In *NeurIPS*, 2018. 2, 3
- [31] Jonah Philion and Sanja Fidler. Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d. *arXiv preprint arXiv:2008.05711*, 2020. 2
- [32] Philip Polack, Florent Alché, Brigitte d’Andréa-Novel, and Arnaud de La Fortelle. The kinematic bicycle model: A consistent model for planning feasible trajectories for autonomous vehicles? In *Intelligent Vehicles Symposium*, 2017. 4

- [33] Dean A Pomerleau. Alvin: An autonomous land vehicle in a neural network. In *NIPS*, 1989. 2
- [34] Luis Alberto Rosero, Iago Pacheco Gomes, Júnior Anderson Rodrigues da Silva, Tiago Cesar dos Santos, Angelica Tiemi Mizuno Nakamura, Jean Amaro, Denis Fernando Wolf, and Fernando Santos Osório. A software architecture for autonomous vehicles: Team lrm-b entry in the first carla autonomous driving challenge. *arXiv preprint arXiv:2010.12598*, 2020. 1
- [35] Binxin Ru, Adam D. Cobb, Arno Blaas, and Yarín Gal. Bayesopt adversarial attack. In *ICLR*, 2020. 5, 7, 8
- [36] Abbas Sadat, Sergio Casas, Mengye Ren, Xinyu Wu, Pranaab Dhawan, and Raquel Urtasun. Perceive, predict, and plan: Safe motion planning through interpretable semantic representations. *CoRR*, abs/2008.05930, 2020. 2, 6, 7, 8
- [37] Abbas Sadat, Mengye Ren, Andrei Pokrovsky, Yen-Chen Lin, Ersin Yumer, and Raquel Urtasun. Jointly learnable behavior and trajectory planning for self-driving vehicles. In *IROS*, 2019. 5, 6, 7
- [38] Tanmay Vilas Samak, Chinmay Vilas Samak, and Sivanathan Kandhasamy. Robust behavioral cloning for autonomous vehicles using end-to-end imitation learning. *CoRR*, abs/2010.04767, 2020. 2
- [39] Aman Sinha, Matthew O’Kelly, John Duchi, and Russ Tedrake. Neural bridge sampling for evaluating safety-critical autonomous systems, 2020. 3
- [40] Niranjan Srinivas, Andreas Krause, Sham M. Kakade, and Matthias W. Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. In *ICML*, 2010. 5, 7
- [41] Jiachen Sun, Yulong Cao, Qi Alfred Chen, and Z. Morley Mao. Towards robust lidar-based perception in autonomous driving: General black-box adversarial sensor attack and countermeasures. In *USENIX Security Symposium*, 2020. 3
- [42] James Tu, Mengye Ren, Sivabalan Manivasagam, Ming Liang, Bin Yang, Richard Du, Frank Cheng, and Raquel Urtasun. Physically realizable adversarial examples for lidar object detection. In *CVPR*, 2020. 2, 3
- [43] Akifumi Wachi. Failure-scenario maker for rule-based agent using multi-agent adversarial reinforcement learning and its application to autonomous driving. In *IJCAI*, 2019. 1, 2, 3
- [44] Moritz Werling, Julius Ziegler, Sören Kammel, and Sebastian Thrun. Optimal trajectory generation for dynamic street scenarios in a frenét frame. In *ICRA*, 2010. 6
- [45] Tong Wu, Xuefei Ning, Wenshuo Li, Ranran Huang, Huazhong Yang, and Yu Wang. Physical adversarial attack on vehicle detector in the carla simulator. *CoRR*, abs/2007.16118, 2020. 3
- [46] Chaowei Xiao, Dawei Yang, Bo Li, Jia Deng, and Mingyan Liu. Meshadv: Adversarial meshes for visual recognition. In *CVPR*, 2019. 2, 3
- [47] Kaidi Xu, Gaoyuan Zhang, Sijia Liu, Quanfu Fan, Mengshu Sun, Hongge Chen, Pin-Yu Chen, Yanzhi Wang, and Xue Lin. Evading real-time person detectors by adversarial t-shirt. *CoRR*, abs/1910.11099, 2019. 2, 3
- [48] Wenyuan Zeng, Wenjie Luo, Simon Suo, Abbas Sadat, Bin Yang, Sergio Casas, and Raquel Urtasun. End-to-end interpretable neural motion planner. In *CVPR*, 2019. 2, 6, 7
- [49] Wenyuan Zeng, Shenlong Wang, Renjie Liao, Yun Chen, Bin Yang, and Raquel Urtasun. Dsdnet: Deep structured self-driving network. *CoRR*, abs/2008.06041, 2020. 2, 6
- [50] Xiaohui Zeng, Chenxi Liu, Yu-Siang Wang, Weichao Qiu, Lingxi Xie, Yu-Wing Tai, Chi-Keung Tang, and Alan L. Yuille. Adversarial attacks beyond the image space. In *CVPR*, 2019. 2, 3
- [51] Yang Zhang, Hassan Foroosh, Philip David, and Boqing Gong. CAMOU: Learning physical vehicle camouflages to adversarially attack detectors in the wild. In *ICLR*, 2019. 3
- [52] Julius Ziegler, Philipp Bender, Thao Dang, and Christoph Stiller. Trajectory planning for berth - A local, continuous method. In *Intelligent Vehicles Symposium*, 2014. 5