

Rain-Reaper: Unmasking LiDAR-based Detector Vulnerabilities in Rain

Richard Capraru^{1,2}, Emil C. Lupu³, Soteris Demetriou³, Jian-Gang Wang² and Boon Hee Soong¹

Abstract—LiDAR-based 3D object detection aims to enhance the situational awareness of autonomous vehicles. Despite recent advancements in this technology, there has been evidence that the susceptibility of 3D object detectors to signal spoofing is high, leading to the erroneous detection of “ghost objects” or the failure to detect genuine ones. While prior work has investigated the design of these new attacks and new defenses, the effect of weather conditions, which is a hot topic in autonomous vehicle research, on both attacks and defenses has never been studied. Inspired by this observation, in this paper, we present a novel genetic algorithm-based attack, entitled *Rain-Reaper*, that leverages on the effect of rain and identifies critical detection points used by 3D detectors. We show that adverse weather conditions not only diminish detection distance and accuracy but also expose the limitations of existing defenses. We have found that the unique characteristics of wet roads lead to underperforming defenses, thus, leading to a false sense of confidence in them. The effectiveness and efficiency of the attack and the robustness of the defenses have been evaluated with both simulated and real data. Our *Rain-Reaper* demonstrates a high attack success rate while successfully evading existing defenses with an adversarial point budget of up to 8.8 times smaller than previously demonstrated state-of-the-art attacks.

I. INTRODUCTION

The rapid development of high-precision range sensors like LiDARs has created vast opportunities for 3D environment understanding and precise object detection in autonomous vehicle (AV) systems [1]. Despite these advancements, LiDARs remain susceptible to adversarial attacks, with ongoing research on both attacks and mitigations [2]–[4]. However, the impact of weather conditions on these attacks and defenses has been overlooked, despite their common occurrence globally. These adverse weather events are increasingly important with higher global temperature. While defense systems based on physical or temporal properties of real and ghost targets have emerged in recent years, their effectiveness in rainy conditions is uncertain [4]–[7]. Based on these observations, in this paper, we are interested in the rain and examine its impact on existing defenses and reveal significant degradation in effectiveness. Moreover, we verify that LiDAR attacks can cause dangerous responses from AVs under rainy conditions. Our experiments demonstrate that attackers can exploit rain to develop more powerful attacks using only a fraction of the previous adversarial point budget. To the best of our knowledge, we are the first to study attacks and defenses in rain. Different from previous approaches, we propose a more robust attack, *Rain-Reaper* (Figure: 2), leveraging critical detection points from AV detectors and the weaknesses of current state-of-the-art (SOTA) defenses.

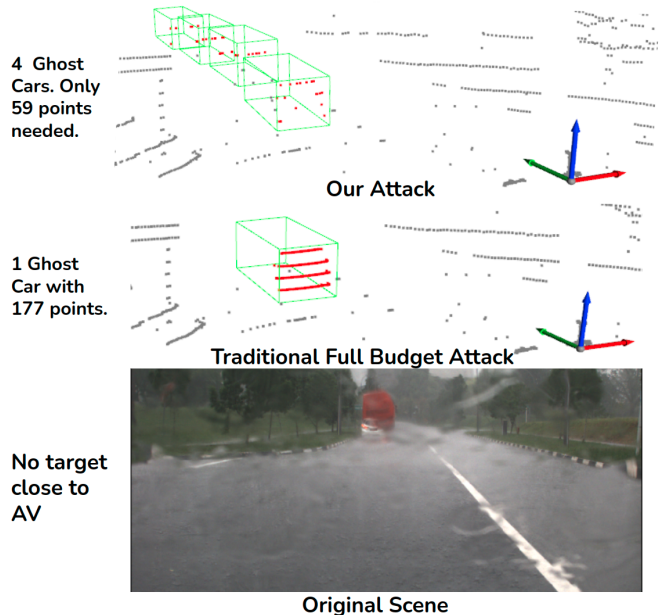


Fig. 1: Qualitative representation *Rain-Reaper*: Unlike conventional attacks, our approach deceives detectors with critical points selected via our genetic algorithm. While SOTA attacks only inject one car using 177 points, our method can inject 4 cars using only 59 points.

Figure 1 shows the efficiency of our budget optimization against the traditional approach. While a traditional attack requires 177 points to generate one ghost car, *Rain-Reaper* accomplishes the creation of 4 cars with only 59 points. This budget optimization lets attackers use cheaper hardware while achieving greater precision and speed due to reduced attack requirements.

We summarize our main contributions:

- **New Method.** We developed a novel genetic algorithm attack, *Rain-Reaper*, to minimize adversarial budget for injecting ghost objects in 3D object detection.
- **New Knowledge.** We are the first to study the effect of rain on LiDAR spoofing attacks and defenses and demonstrate that LiDAR spoofing attacks are harder to defend against in rainy environments.
- **SOTA Results.** Compared with SOTA attacks, *Rain-Reaper* exhibits superior performance.

II. BACKGROUND AND THREAT MODEL

Attacks on LiDARs. Recent research has highlighted the vulnerability of 3D object detectors to LiDAR spoofing attacks [2]–[4], [8]. Petit *et al.* [8] introduced a LiDAR attack that deceives detectors into perceiving objects located beyond the spoofer’s actual position. Shin *et al.* [3] demonstrated an attack capable of manipulating up to 10 points in a 3D point cloud, all at a location closer than the spoofer. Subsequent

¹ School of Electrical and Electronic Engineering, NTU, Singapore 639798

² Institute for Infocomm Research, A*STAR, Singapore 138632

³ Department of Computer Science, Imperial College London, UK SW7 2AZ

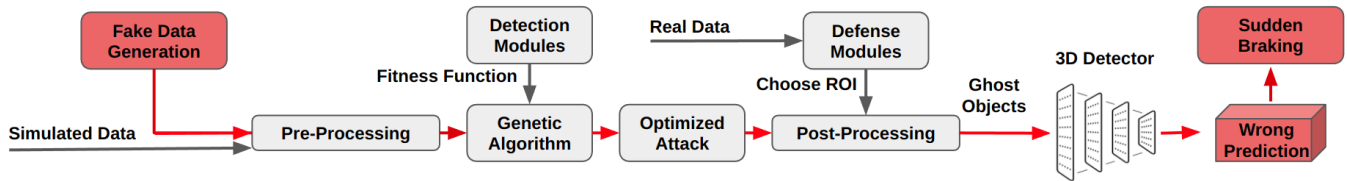


Fig. 2: Overview of our *Rain-Reaper* approach

studies by Cao *et al.* [2], Sun *et al.* [4] and Jin *et al.* [9] extended these attacks to influence up to 100, 200 and 4,200 points in the 3D point cloud, respectively. Furthermore, these attacks pose significant risks, compelling sudden vehicle braking [2] and potentially harming passengers.

LiDAR Defenses. To address the safety concerns posed by LiDAR attacks, various defenses have been proposed. Shadow Catcher [5] and CARLO [4] use point density analysis to detect attacks while 3D-TC2 [6] and ADoPT [6] use point consistency across scenes to identify attacks.

Threat model. LiDAR spoofing adversary capable of manipulating LiDAR return signals has been studied in various approaches [2]–[4], [8]. Following the threat model outlined in [5], [9], we consider the adversary’s capabilities and objectives, including the ability to inject up to 4,200 points with precision in a 3D scene within a horizontal angle of 18° , model-level spoofing attacks emulating vehicles, pedestrians, and cyclists, knowledge of existing SOTA detectors without awareness of the victim’s specific detector, and the capability to launch ghost attacks by spoofing front-near objects [4], [8]. Spoofing objects further from the ego vehicle is possible but less effective due to diminishing impact with distance.

Effect of rain on LiDAR. Existing research has explored the effects of wet roads on LiDAR signals [10], although it has not specifically addressed the impact on adversarial attacks. Hahner *et al.* [10] noted significant attenuation of laser echoes due to water on roads. In this paper, we noticed and observed similar attenuation by using our own data.

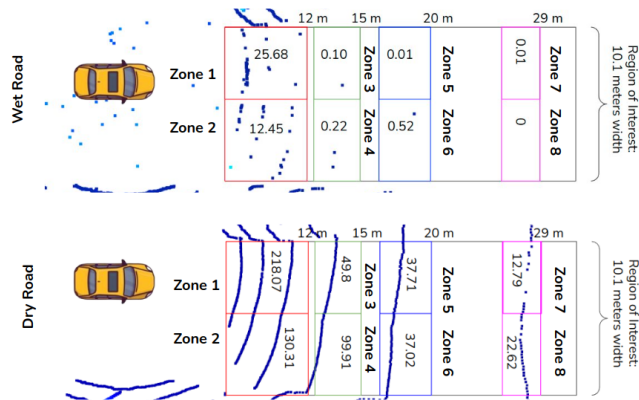


Fig. 3: Number of points in each Region of Interest (ROI) on a dry road (up) and a wet road (down) scene.

III. EFFECT OF RAIN ON ATTACKS AND DEFENSES

We study rain’s impact on SOTA defenses under both normal and adversarial conditions. Shadow Catcher [5] is adopted in our experiments for its superior attack detection compared to alternatives such as CARLO [4], its analysis of semantically

meaningful shadows, and its faster processing speed relative to alternatives like 3D-TC2 [6] and ADoPT [6].

Identifying the Attack Region of Interest: Without loss of generality, we define a front-near region of interest (ROI), where potential attacks are of concern due to safety risks. For example, the ROI spans 10 meters in width and extends 30 meters ahead of the ego-vehicle. We will show this ROI is determined from the braking distance discussed in the next section. We analyze both wet and dry road conditions, extracting a narrow volume of the point cloud centered around ground level within this ROI. The ROI is divided into some regions, and region occupancy is assessed based on LiDAR point measurements, aligning with VLP-16 LiDAR resolution. The braking distance on wet roads is considered when we set the ROI length in our experiments.

Braking Distance. In rainy conditions, with a road friction coefficient of 0.4 [11], and assuming a common urban speed limit of 50 km/h, the braking distance is roughly 24.37m. Objects within this range could trigger emergency brakes, motivating us to set a 30-meter ROI for our experiments.

A. Effect of rain on defenses

3D shadows occur when an ego vehicle’s laser beams reflect off real object surfaces, leading to void measurements on the opposite side [5]. This effect is used to analyze point cloud consistency and detect attacks. However, in rainy conditions, wet roads result in fewer detected ground points, creating shadow-like empty regions that complicate shadow analysis. To investigate this impact on Shadow Catcher, a defense mechanism that relies on shadow analysis, we analyzed consecutive LiDAR frames (e.g., 100 frames) within the identified region of interest (ROI). Our analysis has shown a significant decline in ground point density on wet roads, with reductions exceeding 88% (Figure 3). In addition, we observed that when an object’s shadow falls within -1 to 1 meter in front of the autonomous vehicle (AV), the point density is significantly lower. Objects beyond 5 meters within this range exhibit nearly zero shadow points, compared to about 300 shadow points on dry roads at the same distance. This reduced ground point density on wet roads impairs the effectiveness of Shadow Catcher, motivating us to analyze deeper into its limitations under such conditions.

B. Rain effect on the attack

It is well known that rain reduces LiDAR signal quality. This naturally raises a question: *How does the LiDAR signal attenuation affect the positioning of the adversarial spoofing device?* To answer this question, we calculate the attack effective range using a rain model.

Effective Range of the Attacker. Assuming the attacker is an object, let us represent the minimum distance to ensure the victim can not detect the object as p_m . On the other hand, the maximum distance where the incoming victim laser is no longer received by the attacker due to attenuation is represented as p_M . The attacker’s effective range is defined as $(p_M - p_m)$. Consider that the attenuation will be increased in rain, we have $p'_m < p_m$ and $p'_M < p_M$. In order to assess the Attacker’s Effective Range in real scenarios, we need to check if $(p'_M - p'_m) < (p_M - p_m)$. If this condition is true, then the attacker has less flexibility in choosing their placement. Based on the conclusion reported in [12], a LiDAR sensor’s power can be expressed by the formula $P_n(z) = \frac{\rho}{z^2} \cdot e^{-2\alpha z}$, where ρ is the reflectivity, z is the distance to the target (in meters), and α is the scattering coefficient of rain along the path. The LiDAR signal is intercepted by the attacker, leading to a single pass of the laser beam through rain, which results in a reduced attenuation effect. Additionally, it is reasonable that the emitter can be likened to a laser beam reflecting off a material with 100% reflectivity. Furthermore, as $\alpha = a * R^b$ [13], and knowing that $a=0.01$ and $b=0.6$ for VLP-16 [14], the attenuation of the signal arrives at:

$$P_n(z) = 2 * \frac{1}{z^2} * e^{(-0.02) * R^{0.6} * z} \quad (1)$$

where R represents the rainfall rate (in mm/h). In clear conditions ($R = 0$), Eq. 1 is simplified to $P_n(z) = \frac{2}{z^2}$. The effective range of the attacker depends on material reflectivity, with lower reflectivity resulting in an extended range. Rain reduces effective range, requiring precise positioning, yet even with a reflectivity of 0.9, a minimum 6-meter effective range remains.

IV. RAIN-REAPER ATTACK DESIGN

Our analysis revealed that rain compromises current defenses, allowing attackers to place spoofing devices undetected by AV. This prompts the research question: *Can we devise a more efficient attack with a smaller point budget?* To tackle this, we introduce *Rain-Reaper*, a powerful attack using vulnerabilities in detectors reliant on ‘critical points’.

Algorithm 1 Identify Critical Detection Points

- 1: Initialize the population of individuals.
 - 2: Evaluate the fitness of each individual.
 - 3: **while** Current Generation < Number of Generations **do**:
 - 4: Select individuals through tournament selection.
 - 5: The surviving population is selected for crossover.
 - 6: Apply random crossover.
 - 7: Generate two children for each pair of parents.
 - 8: Create the new population of children and parents.
 - 9: Apply mutation
 - 10: Pass this population to the next generation.
 - 11: **end while**
-

A. Identifying the Ghost Object Critical Detection Points

The notion of critical point sets, originally proposed in [15], highlights the importance of specific input points for global feature representation in object classification. Subsequent work by Wicker *et al.* in [16] showcased the effectiveness of occlusion attacks, exploiting this vulnerability in PointNet [15]. Our study extends this research by investigating critical

detection points in AV object detectors, particularly focusing on optimizing ghost object attacks. To identify key critical points used by SOTA detectors, which can be considered an optimization problem, we develop a Genetic Algorithm.

Our Genetic Algorithm. Without loss of generality, we select Fitness (F.) or Weighted Fitness (W.F.) as fitness functions for our genetic algorithm (GA), Algorithm 1.

$$F. = \beta * \text{BeV IoU}(D_l, X_{i,j}) + (1 - \beta) * \text{Conf}(D_l, X_{i,j}) \quad (2)$$

$$\begin{aligned} \text{W.F.} = & \beta \cdot (\theta \cdot \text{BeV IoU}(D_1, X_{i,j}) + \theta \cdot \text{BeV IoU}(D_2, X_{i,j}) \\ & + \gamma \cdot \text{BeV IoU}(D_3, X_{i,j}) + \gamma \cdot \text{BeV IoU}(D_4, X_{i,j})) \\ & + (1 - \beta) \cdot (\theta \cdot \text{Conf}(D_1, X_{i,j}) + \theta \cdot \text{Conf}(D_2, X_{i,j}) \\ & + \gamma \cdot \text{Conf}(D_3, X_{i,j}) + \gamma \cdot \text{Conf}(D_4, X_{i,j})) \end{aligned} \quad (3)$$

The detectors tested in our experiments include 3D-SSD (D1), VoTr (D2), Part A2 (D3) and PointPillars (D4).

$$X_{i,j^*} = \text{argmax}_{\{X_{i,j}\}} f(D_l, X_{i,j}) \quad (4)$$

where $X_{i,j}$ represents the selection of point j from the pool of points of a target car for individual i . $X_{i,j}$ is a binary variable, taking the value of 1 if point j is selected for individual i , and 0 otherwise. D_l represents the detector used.

- 1) Each individual needs to select a specified number of points (P_t) from the pool of points of a target car:

$$\sum_j X_{i,j} = P_t, \quad \forall i \quad (5)$$

- 2) **Random crossover:** For each pair of parents, randomly choose α_1 points from one parent and substitute the remaining $1 - \alpha_1$ points with random points from the other parent to produce two children.

$$\sum_j X_{\text{child},j} = \alpha_1 \times X_{\text{parent1},j} + (1 - \alpha_1) \times X_{\text{parent2},j} \quad (6)$$

- 3) **Mutation at the individual level:** For a randomly selected individual, replace α_2 of its points with random choices from the target car’s point pool.

$$\sum_j X_{\text{mutant},j} = \alpha_2 \times X_{\text{original},j} + (1 - \alpha_2) \times X_{\text{random},j} \quad (7)$$

- 4) **Mutation at the population level:** Replace a random individual in the population with a random combination of points from the point pool.

$$\sum_j X_{\text{new},j} = X_{\text{random},j} \quad (8)$$

Tournament selection randomly picks ϵ individuals, with only the fittest one advancing to the next generation.

TABLE I: BEV IoU and Confidence score for our G.A. (e.g. using 10 points); W.F. = Weighted Fitness function, F. = Fitness function; D1= 3D-SSD [17]; D2=VoTr [18]; D3= Part A2 [19]; D4 = P.P. [20]

	3D-SSD		VoTr		Part A2		P.P.	
F.	IoU	Cf.	IoU	Cf.	IoU	Cf.	IoU	Cf.
Orig.	92%	87%	88%	77%	90%	33%	86%	85%
D1 F.	86%	81%	79%	64%	70%	38%	81%	73%
D2 F.	76%	84%	80%	77%	77%	63%	82%	58%
D3 F.	77%	82%	78%	69%	78%	94%	83%	87%
D4 F.	68%	65%	75%	34%	75%	10%	89%	90%
W.F.	82%	85%	82%	76%	86%	68%	90%	85%

Discussion. We conducted experiments based on the above settings. Our results are shown in Tables I, using Fitness (Eq. 2) or Weighted Fitness (Eq. 3) reveal two key insights: 1) detectors prioritize points that form a target’s skeleton and, 2) the number of points comprising a target minimally affects

detection accuracy and confidence. Our approach exhibits robust generalization across various detectors. Introducing additional detectors appears to enhance the Bird’s Eye View (BEV) Intersection over the Union (IoU) score while preserving comparable confidence levels.

Our experiments, shown in Figure 4 demonstrate that detectors like 3D-SSD and PointPillars can misclassify ghost objects as cars with as few as 2 points. Moreover, employing 20 points in an attack achieves detection confidence and accuracy exceeding 80% for all detectors.




Fig. 4: BEV IoU vs No. of pts; Conf. Score vs No. of pts for ghost object placed at 10 meters

V. EXPERIMENTS AND EVALUATION

After identifying vulnerabilities in defenses on wet roads and creating *Rain-Reaper*, we validate its efficiency through simulation and real data. Our focus is on determining: Question 1. *How effective is strategically placing the ghost object on wet roads to evade the Shadow Catcher?* Question 2. *How efficient is our attack in terms of success rate?*

A. Data

We gathered data using our AV equipped with a VLP-16 LiDAR positioned 2.33 meters above ground. The dataset includes paired measurements from wet and dry road conditions while the AV was stationary at the roadside. Our attack model focuses solely on road information, without any other targets in the scene. We simulated ghost object cars using the MAVS simulator [21], resulting in 200 real scenes (100 wet, 100 dry) and 410 simulated scenarios, totalling over 82,000 data samples. The *Rain-Reaper* evaluation utilized 6,200 samples. Each simulated car was inserted into a real point cloud scene, with shadow points removed based on the LiDAR’s laser ray directions. Furthermore, we also evaluated our attack on the well-established KITTI dataset [22].

B. Target LIDAR Detector Models

This study targets detector models specialized for point clouds exclusively, exploring various architectures such as PointPillars [20], 3D-SSD [17], VoTr [18], and Part A2 [19]. PointPillars organizes points into bins in the BEV and employs PointNet [15] for feature extraction. 3D-SSD groups points into 3D bins, VoTr utilizes a voxel architecture with Transformers and Part A2 processes raw point clouds. We used SOTA detectors trained on the KITTI dataset [22].

C. Evaluation Metrics

We evaluate attack quality using two key metrics:

Attack Detection Success Rate (ADSR). This evaluates the *Rain-Reaper*’s ability to deceive a detector by appearing as a

genuine target, measured using Bird’s Eye View Intersection over Union (BEV IoU) and detector confidence levels.

Shadow-Catcher Score. This score indicates *Rain-Reaper*’s ability to evade detection by the Shadow-Catcher defense [5] and be recognized as a genuine object.

D. Results and Discussion

1) *Shadow-Catcher Evaluation:* To address Question 1 thoroughly, we assess Shadow Catcher’s anomaly score in simulated scenes, wet road scenes, and dry road scenes.

We tested the placement of ghost attacks by inserting simulated objects into frames depicting wet and dry ground conditions. By analyzing 20,500 scenarios (see Figure 5), we found that Shadow Catcher detects anomalous shadows more effectively on dry roads. However, on wet roads, the effectiveness diminishes due to increased road sparsity, making Shadow Catcher ineffective beyond 10 meters in rain. At 23.5 meters, there is a notable spike in the anomaly score for a ghost car on a real dry road. This spike is likely due to the presence of random ground points caused by incomplete shadow removal when the ghost car is injected.




Fig. 5: Shadow Catcher Score vs Dist. for a ghost car: black (simulated dry road), blue (real dry road) and purple (real wet road)

As expected, our experimental results have shown that rain aids attackers in generating larger ghost objects near AVs, increasing the risk of emergency braking.

2) *Rain-Reaper Attack Evaluation:* To assess the efficiency of our optimized attack, we used our proposed genetic algorithm (see Section 3.2 for details) to identify critical points for a car positioned at $Y=0$, orientation= 0 , across 31 distances (from 10 meters to 25 meters in 0.5-meter steps). In Table II, we can see that our attack consistently matches the success rates of full-budget attacks in fooling SOTA detectors pre-trained with clear data, surpassing existing attacks notably for VoTr [18] and PartA2 [19], across a range of confidence and BEV IoU and confidence thresholds. In addition, our attack maintains high performance against unseen detectors (Centerpoint [24]; SSN [25]) pre-trained with both rain and clear data. Furthermore, our experiments achieved a reduction of up to 8.8 times less point budget in the point budget requirement. This can be seen in Figure 6 where the percentage of the budget optimization is shown.

VI. RAIN-REAPER ROBUSTNESS AGAINST NOISE

Previous research [9] has shown that injecting points to LiDAR scenes is feasible, with 90% of points within a 12-degree angle having a distance error of about 0.102 meters. To approximate their reported error distribution, we use the

TABLE II: ADSR using 20 pts. versus original for different BEV IoU and Confidence thresholds for detectors pre-trained on clear data (KITTI [22]) and clear+rain data (Nuscenes [23]): 3D-SSD [17], VoTr [18], Part A2 [19], P.P. [20], C.P. [24] and SSN [25].

Train	Detector	Attack	IoU Threshold						Confidence Threshold					
			70%		60%		40%		30%		20%			
Distance			10m-15m	15m-20m	20m-25m	10m-15m	15m-20m	20m-25m	10m-15m	15m-20m	20m-25m	10m-15m	15m-20m	20m-25m
Clear only	3D-SSD	Original	100%	93.7%	91.9%	100%	98.4%	100%	92.8%	64.7%	37.3%	96.2%	74.3%	48.9%
		Ours	100%	95.2%	99.8%	100%	99.4%	100%	86.2%	64.9%	56.2%	91.4%	73.1%	71.6%
		Original	100%	90.8%	99.8%	100%	97.7%	100%	86.2%	63.4%	56.2%	91.4%	71.5%	71.6%
		Ours	100%	96.8%	99.8%	100%	97.7%	100%	86.2%	68.6%	56.2%	91.4%	78.1%	71.6%
		Original	100%	70.4%	72.1%	100%	70.4%	72.1%	100%	83.3%	53.6%	100%	100%	91%
		Ours	99.7%	80.8%	73.0%	99.7%	87.3%	73.1%	99.6%	78.9%	63.9%	100%	100%	81.9%
	Part A2	Original	99.5%	88.7%	64.2%	99.9%	88.7%	95.5%	76.6%	76.3%	18.3%	91.2%	76.8%	27.4%
		Ours	99.5%	84.1%	94.7%	99.9%	94.7%	95.5%	76.6%	57%	18.3%	91.2%	65.1%	27.4%
		Original	99.5%	99.7%	90.7%	99.9%	99.7%	99.8%	99.7%	99.7%	94.1%	99.8%	99.7%	94.2%
		Ours	99.5%	96.6%	90.7%	99.9%	99.8%	99.8%	99.7%	97.8%	94.1%	99.8%	99.7%	94.2%
		Original	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%
		Ours	99.7%	100%	100%	100%	100%	100%	94.4%	100%	100%	98.6%	100%	100%
Rain + Clear	Centerpoint	Original	99.8%	65.3%	86%	99.9%	80.3%	100%	99.6%	72.5%	48.4%	99.8%	95.4%	96.2%
		Ours	88.9%	83.7%	88%	92.1%	93.4%	100%	75.6%	73.5%	69%	92.1%	81.1%	99.2%
		Ours	88.9%	83%	88%	92.1%	89.4%	100%	75.6%	63.7%	69%	92.1%	93.8%	99.2%
	SSN	Original	94.1%	47.3%	85.2%	99.9%	100%	99.7%	85.9%	78.2%	96.5%	93.8%	100%	99.8%
		Ours	62.9%	75.5%	73.5%	98.7%	99.8%	99.1%	34.9%	86.8%	79%	81.2%	77%	99%
		Ours	62.9%	50%	73.5%	98.7%	97.8%	99.1%	34.9%	40.7%	79%	81.2%	85.7%	99%




Fig. 6: *Rain-Reaper* Budget Optimization vs Dist.




Fig. 7: No of pts vs Dist. for ghost car in front of AV

function, $error = \frac{-\ln(1-n)}{22.44}$ (where n is a random number between 0 and 1). Other studies [14] identified rain-induced distance noise on LiDAR data to consistently remain under 20 cm - 10 cm, which corresponds to our experimental setting. Nonetheless, our analysis (Figure 8) demonstrates that our attack remains robust against noise.




Fig. 8: *Rain-Reaper* Noise Robustness Analysis

VII. COMPARISON WITH OTHER ATTACKS

In order to show the efficiency and accuracy of our approach, we have compared our approach with the SOTA (Table: III) in this paper. A recent SOTA attack introduced by Jin *et al.* [9] can achieve 56% ADSR when injecting an 800-point car at 10 meters in front of the LiDAR in 100 random KITTI scenes [22]. We reproduce the reported experiments on 100 random KITTI scenes (realistic scenarios with no pre-existing cars in the injection location) with our attack against the same detectors (PointPillars [20] and SECOND [26]) and we obtain up to 80% higher ASR while using 40 times fewer

points. The results also demonstrate that our attack maintains high ASR for both low and high-resolution LiDARs.

TABLE III: Related work comparison. P.P. [20]; S. [26].

Method	LiDAR	Pts	Targ.	Success Rate	
Shin <i>et al.</i> [3]	VLP-16	10	-	-	
Cao <i>et al.</i> [2]	VLP-16	60	N/A	75%	
Sun <i>et al.</i> [4]	VLP-16	80	car	83%	
Jin <i>et al.</i> [9]	VLP-16	180	car	S.	17%
				P.P.	17%
	HDL-64E	800	car	S.	56%
				P.P.	14%
Ours	VLP-16	20	car	S.	100%
				P.P.	100%
	HDL-64E	20	car	S.	96%
				P.P.	94%

VIII. DISCUSSION ON OTHER SOTA DEFENSES

Rain-Reaper evades detection from Shadow-Catcher and fools four detectors. Unfortunately, we cannot compare empirically with other defenses due to unavailable or incomplete implementations. However, in this paper, we perform our best effort to analyse their vulnerability based on their design.

CARLO. Sun *et al.* incorporate Laser Penetration Detection (LPD) and Free Space Detection (FSD) in CARLO [4]. LPD identifies anomalous objects, with uncertain cases routed to computationally intensive FSD; it detects spoofed objects by assessing the ratio of points behind an object to total frustum points. In this paper we reveal that ghost cars beyond 10 meters have $\leq 10\%$ probability for points in the Occluded Region, with attacks as low as 20 points yielding ratios of 0.05 or 0, challenging CARLO's discrimination of ghost objects. Additionally, LPD's oversight of points behind the bounding box leads to false positives, especially in wet conditions. FSD relies on high point density for authentic vehicles, but wet roads reduce free space points beyond 10 meters to $< 10\%$. CARLO sets upper bounds for genuine objects at approximately 0.9 for LPD and 0.7 for FSD, surpassing observed ratios for ghost objects on wet roads.

ADoPT. Cho *et al.* proposed ADoPT [7] which can assess temporal consistency at the point level across frames. However, attackers may bypass this defense by simulating points for a car at different distances and precisely choosing injected points. In Figure 7, generated with MAVS [21], there are significant point variations at different rain rates. For instance, at a distance of 13 meters, points can range from

50 (at a rain rate of 40 mm/h) to 100 (at 10 mm/h). Rapid rain fluctuations, such as transitioning from 2 to 12 mm/h in less than 30 seconds [27], make evaluating point consistency in rain challenging, leading to a high rate of false positives.

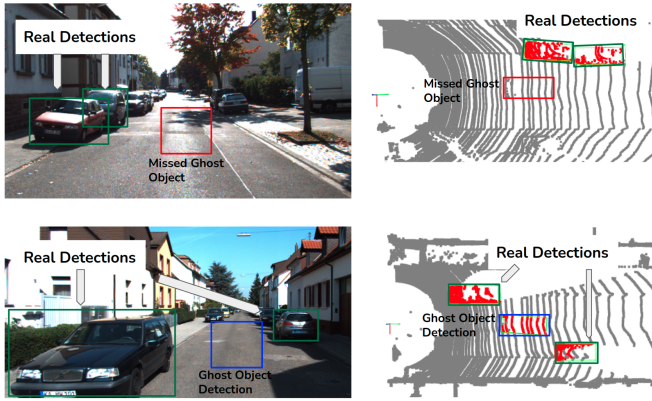


Fig. 9: Ghost: Failure Case (upper) vs Success Case (lower)

IX. CONCLUSIONS AND FUTURE WORK

LiDAR attack vulnerabilities in rain have never been investigated, limiting the widespread adoption of AVs and raising security concerns. In this paper, we propose a genetic algorithm-based ghost-rain attack, *Rain-Reaper*, a refined strategy adept at exploiting rain and identifying critical points used by 3D detectors. Our experiments revealed that the effectiveness of defenses is significantly impaired in rainy conditions. By implementing *Rain-Reaper* on real data, our methodology demonstrated its efficacy by successfully evading SOTA defenses, and deceiving detectors, all while achieving a reduction in the point budget of up to 8.8 times. We have investigated some fail cases although our approach has achieved very high accuracy. One such case, depicted in Fig. 9, involved the failure to detect a ghost object due to elevated ground levels. Adjusting the attack height led to successful detection. Moving forward, we hope to improve our attack by considering physical factors such as ground level, and road wetness analysis based on rain rate. Moreover, further experiments are essential to deepen our understanding of how different physical factors might influence attacks that either create or conceal road objects.

In terms of potential mitigation strategies, defenses will require precise characterization of the physical invariance of rain, which remains a subject for future research.

ACKNOWLEDGEMENTS

This research was supported by the NTU-Imperial Global Fellows Program, under the leadership and invaluable guidance of Assoc. Prof. Sierin Lim and Mrs. Seah Wan Er, the ICL Department of Computing, and the A*STAR SINGA scholarship, which funded the research visit with special thanks to Mrs. Eunice Heng for her exceptional support. The authors extend their deepest gratitude to Dr. Teoh Eam Khwang for his insightful suggestions and contributions during the writing of this paper, and to Mr. Pang Chun Ho for his outstanding assistance in recording AV data. We also wish to thank the anonymous reviewers for their valuable feedback.

REFERENCES

[1] “Google spin-off waymo to sell lidar it fought uber on,” 2019-03-07.

[2] Y. Cao, C. Xiao, B. Cyr, Y. Zhou, W. Park, S. Rampazzi, Q. A. Chen, K. Fu, and Z. M. Mao, “Adversarial sensor attack on lidar-based perception in autonomous driving,” in *Proceedings of the 2019 ACM SIGSAC*, ser. CCS ’19, 2019, p. 2267–2281.

[3] H. Shin, D. Kim, Y. Kwon, and Y. Kim, “Illusion and dazzle: Adversarial optical channel exploits against lidars for automotive applications,” in *CHES 2017*, 2017, pp. 445–467.

[4] J. Sun, Y. Cao, Q. A. Chen, and Z. M. Mao, “Towards robust lidar-based perception in autonomous driving: General black-box adversarial sensor attack and countermeasures,” in *USENIX 2020*, 2020.

[5] Z. Hau, S. Demetriou, L. Muñoz-González, and E. C. Lupu, “Shadow-catcher: Looking into shadows to detect ghost objects in autonomous vehicle 3d sensing,” in *ESORICS 2021*, E. Bertino, H. Shulman, and M. Waidner, Eds., 2021, pp. 691–711.

[6] C. You, Z. Hau, and S. Demetriou, “Temporal consistency checks to detect lidar spoofing attacks on autonomous vehicle perception,” in *Proceedings of the 1st Workshop on Security and Privacy for Mobile AI*, ser. MAISP’21, 2021, p. 13–18.

[7] M. Cho, Y. Cao, Z. Zhou, and Z. M. Mao, “Adopt: Lidar spoofing attack detection based on point-level temporal consistency,” 2023.

[8] J. Petit, B. Stottelaar, and M. Feiri, “Remote attacks on automated vehicles sensors: Experiments on camera and lidar,” 2015.

[9] Z. Jin, X. Ji, Y. Cheng, B. Yang, C. Yan, and W. Xu, “Pla-lidar: Physical laser attacks against lidar-based 3d object detection in autonomous vehicle,” in *2023 IEEE SP*, 2023, pp. 1822–1839.

[10] M. Hahner, C. Sakaridis, M. Bijelic, F. Heide, F. Yu, D. Dai, and L. Van Gool, “Lidar snowfall simulation for robust 3d object detection,” in *2022 IEEE/CVF CVPR*, 2022, pp. 16 343–16 353.

[11] A. A. Kordani, O. Rahmani, A. S. A. Nasiri, and S. M. Boroomandrad, “Effect of adverse weather conditions on vehicle braking distance of highways,” *Civil Engineering Journal*, vol. 4, pp. 46–57, 2018.

[12] C. Goodin, D. Carruth, M. Doude, and C. Hudson, “Predicting the influence of rain on lidar in adas,” *Electronics*, vol. 8, p. 89, 01 2019.

[13] P. A. Lewandowski, W. E. Eichinger, A. Kruger, and W. F. Krajewski, “Lidar-based estimation of small-scale rainfall: Empirical evidence,” *Journal of Atmospheric and Oceanic Technology*, vol. 26, no. 3, pp. 656 – 664, 2009.

[14] A. Filgueira, H. González-Jorge, S. Lagüela, L. Díaz-Vilariño, and P. Arias, “Quantifying the influence of rain in lidar performance,” *Measurement*, vol. 95, pp. 143–148, 2017.

[15] C. Qi, H. Su, K. Mo, and L. J. Guibas, “Pointnet: Deep learning on point sets for 3d classification and segmentation,” *2017 IEEE CVPR*, pp. 77–85, 2016.

[16] M. Wicker and M. Kwiatkowska, “Robustness of 3d deep learning in an adversarial setting,” in *2019 CVPR*, jun 2019, pp. 11 759–11 767.

[17] Z. Yang, Y. Sun, S. Liu, and J. Jia, “3dssd: Point-based 3d single stage object detector,” in *2020 CVPR*, 2020, pp. 11 037–11 045.

[18] J. Mao, Y. Xue, M. Niu, H. Bai, J. Feng, X. Liang, H. Xu, and C. Xu, “Voxel transformer for 3d object detection,” 2021.

[19] S. Shi, Z. Wang, J. Shi, X. Wang, and H. Li, “From points to parts: 3d object detection from point cloud with part-aware and part-aggregation network,” *IEEE PAMI*, vol. 43, no. 08, pp. 2647–2664, aug 2021.

[20] A. H. Lang, S. Vora, H. Caesar, L. Zhou, J. Yang, and O. Beijbom, “Pointpillars: Fast encoders for object detection from point clouds,” in *2019 CVPR*, 2019, pp. 12 689–12 697.

[21] “Msu autonomous vehicle simulator,” accessed: 2023-01-29. [Online]. Available: <https://www.cavs.msstate.edu/capabilities/mavs.php>

[22] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, “Vision meets robotics: the kitti dataset,” *The International Journal of Robotics Research*, vol. 32, pp. 1231–1237, 09 2013.

[23] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, “nusenes: A multimodal dataset for autonomous driving,” in *2020 CVPR*, 2020, pp. 11 618–11 628.

[24] T. Yin, X. Zhou, and P. Krahenbuhl, “Center-based 3d object detection and tracking,” in *2021 CVPR*, jun 2021, pp. 11 779–11 788.

[25] X. Zhu, Y. Ma, T. Wang, Y. Xu, J. Shi, and D. Lin, “Ssn: Shape signature networks for multi-class object detection from point clouds,” in *Computer Vision – ECCV 2020*, 2020, pp. 581–597.

[26] Y. Yan, Y. Mao, and B. Li, “Second: Sparsely embedded convolutional detection,” *Sensors*, vol. 18, p. 3337, 10 2018.

[27] C. Linnhoff, K. Hofrichter, L. Elster, P. Rosenberger, and H. Winner, “Measuring the influence of environmental conditions on automotive lidar sensors,” *Sensors (Basel, Switzerland)*, vol. 22, 2022.