Poster: A Novel Attack Exploiting the Vulnerabilities of LiDAR Point Cloud Pre-Processing Filters

Yuna Tanaka[*], Kazuki Nomoto[*†], Ryunosuke Kobayashi[*], Go Tsuruoka[*], Tatsuya Mori[*‡§],
[*]Waseda University [†]Deloitte Tohmatsu Cyber LLC [‡]RIKEN [§]NICT

***Abstract*—This study addresses a critical gap in existing research on the security of LiDAR systems for autonomous vehicles by shedding light on the simple point-cloud-density-based obstacle detection methods. In response to this gap, we introduce the "Adversarial Fog Attack," an innovative approach tailored to target the pre-processing stages of LiDAR point cloud data. By generating artificial fog, this attack seeks to manipulate the perception of LiDAR point cloud data, effectively bypassing density-based obstacle detection and concealing critical obstacles, such as pedestrians, in front of the vehicle. In this study, we conducted extensive end-to-end driving simulation experiments using Autoware. The results confirm the feasibility of attacks on LiDAR-based autonomous driving systems, specifically revealing vulnerabilities in point cloud data pre-processing filters.**

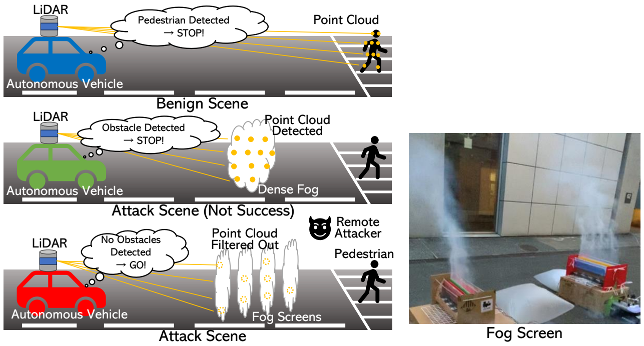

Fig. 1: Overview of Adversarial Fog Attack.

I. Introduction

LiDAR sensors, an integral part of autonomous vehicle technology, are subject to significant security risks, especially when ML-based object detection systems interpret sensor data as inherently accurate. Previous studies have highlighted these vulnerabilities, showing how false point clouds can trick ML models into misidentifying objects [1]. However, a key consideration in this context is the effectiveness of these attacks against simpler, non-ML-based LiDAR detection methods, such as those widely used in autonomous driving frameworks such as Autoware. These simpler detection systems, which focus on point cloud density rather than detailed object detection, are less susceptible to certain types of attacks. For example, attacks to hide obstacles with adversarial point clouds will not work because these systems prioritize detecting the presence of points over detailed classification.

In light of this observation, our research introduces a novel attack strategy, named the "Adversarial Fog Attack," that is specifically designed to bypass these simpler LiDAR detection methods. This approach involves the strategic use of fog screens to create a layer of artificial fog that effectively hides real-world obstacles such as pedestrians from the vehicle's sensors. This can lead to dangerous scenarios such as collisions. Our study involves a detailed analysis of the filters applied to point cloud data in standard LiDAR systems, coupled with simulations to identify the optimal conditions for deploying fog screens. Preliminary tests have shown that placing four fog screens one meter apart significantly increases the likelihood of an autonomous vehicle colliding with obscured pedestrians – up to a 95% chance under varying conditions.

II. Background

Fog's Impact on LiDAR Sensing We reviews the effects of fog on LiDAR sensing, based on the study by Hahner et al. [2]. Fog uniquely affects light by partially transmitting and reflecting it, causing LiDAR laser pulses to scatter and weaken in foggy environments. It was found that signals from solid objects decrease with distance under clear conditions, but are further affected by fog density. Conversely, signals from fog vary with both fog density and distance from the LiDAR.

LiDAR Data Outlier Filters Outlier filters play a crucial role in pre-processing point clouds for subsequent object detection and control applications, especially in conditions affected by environmental factors such as rain, fog, and insects. We describe how the Ring Outlier Filter works. This filter algorithm works by analyzing the distance and azimuth angle differences between adjacent points. The point cloud is segmented into collections called "walks." The points are included in the same walk when the ratio of the distances from LiDAR between neighboring points is smaller than the threshold value. The condition for determining whether the walk is noise focuses on the collective properties of the walk. It requires that the number of points in the walk exceeds a predefined threshold, *num_points_threshold*, or that the length of the walk, measured from end to end, exceeds the threshold, *object_length_threshold*. If the walk do not satisfy the above condition, it will be removed. After pre-processing, if point clouds are detected ahead in the direction of travel, the system identifies them as obstacles. As a result, it initiates vehicle control actions such as stopping or steering to avoid collisions.

III. Adversarial Fog Attack

The primary goal of the attack is to bypass the vehicle's density-based obstacle detection by exploiting the LiDAR's noise removal filters, which erroneously eliminate both the artificially generated fog and the point clouds representing people or objects hidden within it as noise.

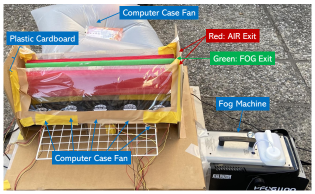

Fig. 2: A Prototype of Fog Screen.

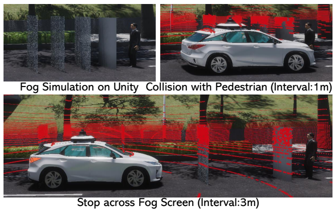
Fog Simulation on Unity
Collision with Pedestrian (Interval:1m)
Stop across Fog Screen (Interval:3m)

Fig. 3: Experimental Setup and Results.

A. Physical Mechanics of Fog Screen

We describe the physical mechanics of a fog screen designed to facilitate Adversarial Fog Attack. The key mechanism for creating the fog is a standard piece of equipment called a "fog machine" that is often found in concert settings. As shown in Fig. 2, by creating layers of air on both sides of the fog stream, the design achieves stability in the trajectory of the fog. In the right panel of Figure 1, the Fog Screen is shown actively producing the Adversarial Fog. On attack, the fog screen is configured to emit fog across the road in a lateral direction, creating a barrier that interferes with the vehicle's LiDAR sensors.

B. Optimizing the Adversarial Fog Attack

This section outlines the considerations for a successful attack. Due to space constraints, we omit specific equations we formulated and focus on the overall approach. The key to this attack is the use of multiple fog screens that disperse the points of LiDAR laser hits, thereby reducing the number of walks not removed by the Ring Outlier Filter. This dispersion tricks the system into perceiving a less dense point cloud, increasing the chances of classifying these clouds as noise. The goal is to make the fog and obstacles behind it undetectable.

Key considerations in this strategy include the placement of the fog screens, specifically their spacing and the number of devices used. Understanding and manipulating the LiDAR system's noise removal criteria, particularly the Ring Outlier Filter's determination of noise based on walk length and point number, is critical. Inappropriate spacing between screens can cause point clouds from different screens to be merged into a single walk, reducing the chance of noise classification. Correct spacing is therefore critical to the success of the attack.

IV. Feasibility of the Attack

A. Optimizing the Ring Outlier Filter Parameters

We fine-tune the parameters of the Ring Outlier Filter to avoid unnecessary stops when detecting minor objects like insects. We obtained a 3D model of a butterfly from the Unity Asset Store. By empirically setting the *num_points_threshold* to 15 and *object_length_threshold* to 0.2, we successfully filtered out butterfly point clouds. We checked that the optimized parameters effectively filtered out small objects such as insects while maintaining accurate pedestrian detection.

B. Evaluation Setup

In our experiment, we use Autoware.Universe and AWSIM. Figure 3 (top left) presents the experimental setup. Adult-sized pedestrians are placed in an urban scene, and numerous small 3D objects are placed in front of them using Unity to create a fog screen effect. The distance between the pedestrian and the nearest fog screen is set to 1 meter. VLP-32C is used for a LiDAR sensor. The parameters for the Ring Outlier Filter are those mentioned in section IV-A. We chose to use four fog screens in the following experiments.

C. Impact of Fog Screen Spacing and Vehicle Speed

We evaluate the impact of different fog screen spacings on the collision rate of an autonomous vehicle with a pedestrian. The vehicle is tested at 40 km/h with fog screens placed at 1, 2, and 3m intervals. Each scenario is replicated five times.

At 1 m intervals, the vehicle collided with the pedestrian in all five trials. The vehicle initially decelerated when unable to clear distant fog screen point clouds, but then accelerated as it approached and cleared these point clouds, eventually resulting in collisions despite attempted stops. At 2m intervals, there was one collision in five trials. At 3m intervals, no collisions occurred in five trials. In the cases where the vehicle did not collide, it decelerated and stopped before reaching the pedestrian. The study suggests that increasing the distance between the fog screens gives the vehicle more time to stop because the distance to the obstacle is greater with fewer fog screens in front of the vehicle.

We investigate how the success rate of attacks varies with changes in the speed of the autonomous vehicle, specifically at 20 km/h and 60 km/h. The distance between the fog screens is set to 1m. Each speed scenario is performed in five trials.

At 20 km/h, the vehicle collided with the pedestrian in four out of five trials, and at 60 km/h, collisions occurred in all five trials. As a result, no difference in attack success rate due to the speed of the autonomous vehicle could be observed.

V. Summary

We introduce the "Adversarial Fog Attack" and demonstrate how artificial fog can manipulate LiDAR perception, bypassing density-based obstacle detection and concealing critical obstacles such as pedestrians. Our extensive end-to-end simulations demonstrate the feasibility of these attacks. Most importantly, our results highlight the need for security research focused on the filtering mechanisms of LiDAR point cloud data. While our current progress includes preliminary experiments in the physical world, these initial tests primarily serve as proof of concept. For our future work, we aim to explore more optimal placements for fog screens, conduct real-world experiments, and develop advanced defense strategies to improve the security of autonomous driving systems against such sophisticated threats.

Acknowledgment

A part of this work was supported by JSPS KAKENHI 22S0604 and JST CREST JPMJCR23M4.

References

[1] Yulong Cao et al. "Adversarial Sensor Attack on LiDAR-Based Perception in Autonomous Driving". In: *Proceedings of the ACM CCS 2019*.

[2] M. Hahner et al. "Fog Simulation on Real LiDAR Point Clouds for 3D Object Detection in Adverse Weather". In: *2021 IEEE/CVF ICCV*.

Poster: A Novel Attack Exploiting the Vulnerabilities of LiDAR Point Cloud Pre-Processing Filters

Yuna Tanaka[1], Kazuki Nomoto[1,2], Ryunosuke Kobayashi[1], Go Tsuruoka[1], Tatsuya Mori[1,3,4]
[1]Waseda University [2]Deloitte Tohmatsu Cyber LLC [3]NICT [4]RIKEN AIP

Introduction

- While previous research has predominantly focused on attacking ML-based LiDAR object detection, it has overlooked the practical importance of the density-based obstacle detection and the pre-processing filters applied to LiDAR point cloud.
- In response to this gap, we introduce the **Adversarial Fog Attack**, an innovative approach tailored to target the pre-processing stages of LiDAR point cloud data.
- By generating artificial fog, this attack seeks to manipulate the perception of LiDAR point cloud data, effectively bypassing density-based obstacle detection and concealing critical obstacles, such as pedestrians, in front of the vehicle.

Overview of Adversarial Fog Attack

Key idea: The point clouds are dispersed by emitting multiple thin layers of artificial fog. This dispersion tricks the system into perceiving a less dense point cloud, increasing the chances of classifying these clouds as noise.

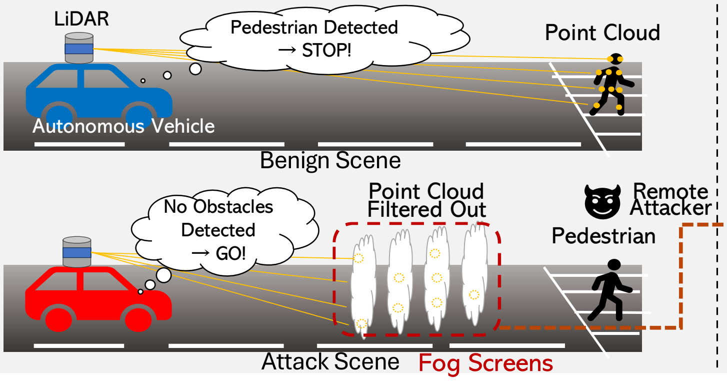

Physical-World Setup

- On attack, the fog screen is configured to emit fog across the road in a lateral direction.

2 Fog Screens

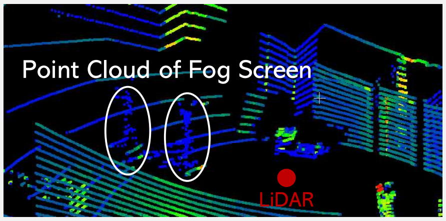

Point Cloud Data

Density-based Obstacle Detection

- Density-based Obstacle Detection is an important mechanism in autonomous driving systems for ensuring reliable stop control in front of obstacles.
- Autonomous vehicles stops if there is point clouds that was not removed by noise removal in the direction of travel.

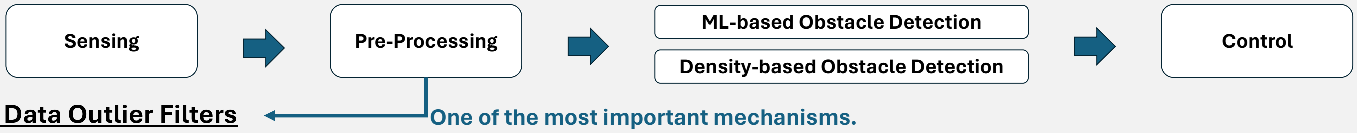
Sensing
Pre-Processing
ML-based Obstacle Detection
Density-based Obstacle Detection
Control
One of the most important mechanisms.

LiDAR Data Outlier Filters

- Outlier filters play a crucial role in pre-processing point clouds especially in conditions affected by environmental factors such as rain, fog, and insects.
- **Ring Outlier Filter** removes noise based on the rate of change in the distance between points.

Evaluation

E2E evaluation of the Adversarial Fog Attack against autonomous driving systems using simulators.

Point Cloud Measured by LiDAR

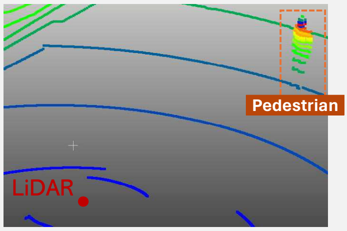

Benign Scene

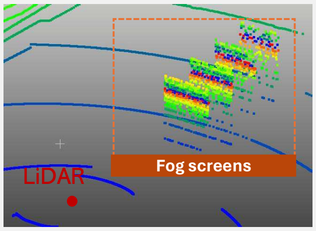

Attack Scene
(Before Noise Removing)

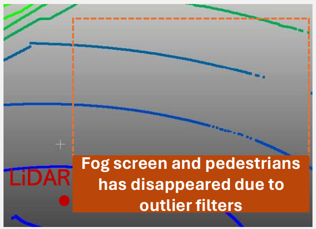
Fog screen and pedestrians has disappeared due to outlier filters
LiDAR

Attack Scene
(After Noise Removing)

In Benign Scene, the pedestrian point clouds is measured correctly. In Attack Scene, pedestrian point clouds cannot be measured due to the fog screens. In fact, the noise removal filter is applied to delete the fog screens and the pedestrian point cloud.

End-to-End Evaluation of the Attack

Simulation of Fog in Unity

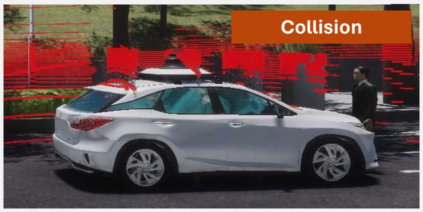

Driving Result

Speed vs Collision

AV's speed [km/h]	Times that AV collided
20	4/5
40	5/5
60	5/5

We investigated whether an autonomous vehicle could pass through fog screens and collide with a pedestrian. Autonomous vehicles collided with the pedestrian regardless of vehicle speed.

Countermeasure

1. Improve the outlier filters algorithm
2. Use alternative sensors (radar, thermal camera)
3. Use intensity information of point cloud from LiDAR

Future Work

- Conduct comprehensive physical experiments (Include wind effect)
- Investigate of the effect on cameras
- Explore more optimal placements for fog screens
- Develop and implement defensive strategies