Poster: Discovering Sensor-Fusion-Vulnerabilities in autonomous driving systems against LiDAR attacks

Ryo Yoshida, Takami Sato, Yuki Hayakawa, Ryo Suzuki, Kazuma Ikeda, Ozora Sako,
Rokuto Nagata, and Kentaro Yoshioka
Keio University

***Abstract*—Multi-sensor Fusion (MSF), which combines sensors such as LiDAR and cameras, has garnered attention as a countermeasure against LiDAR Spoofing attacks that threaten the safety of autonomous driving systems. However, the effectiveness of current MSF implementations has not been thoroughly validated in practical autonomous driving systems. In this study, we present an initial framework aimed at exploring the potential vulnerabilities of MSF, based on the open-source autonomous driving software Autoware Universe and the AWSIM simulator. Through experiments conducted using this framework, we demonstrated that the MSF implementation in Autoware Universe can lead to a dangerous state for the entire system, even when the camera correctly detects objects, if the LiDAR point cloud is lost. This vulnerability arises because the camera information is limited to a supplementary role in point cloud clustering. Our findings indicate that the MSF implementation in Autoware Universe lacks sufficient resilience against LiDAR spoofing attacks due to its structural limitations. The framework is available at: https://github.com/Keio-CSG/Multi-Sensor-Defense-Analysis-Platform.**

I. Introduction

LiDAR (Light Detection and Ranging) technology is fundamental to environment perception in autonomous driving (AD) systems, enabling precise distance measurements through laser pulses. However, this critical sensor faces security vulnerabilities through LiDAR Spoofing Attacks, where malicious laser signals are emitted to manipulate these measurements. The severity of this threat is highlighted in [1], which demonstrates how High Frequency Removal (HFR) attacks can effectively jam LiDAR systems, erasing point cloud data ahead of vehicles moving at 60 km/h and creating potentially catastrophic safety risks for autonomous vehicles.

To address this critical vulnerability, we investigate a fundamental question: **Can LiDAR-camera sensor fusion effectively mitigate LiDAR Spoofing Attacks in practical autonomous driving systems?** Prior research [2]–[4] has proposed Multi-sensor Fusion (MSF) as a promising defense strategy, leveraging redundant data from multiple sensors to detect and resist spoofing attempts. However, current MSF research is limited by its narrow focus on individual components without E2E validation, while closed-source implementations hinder reproducibility and collaborative research.

To bridge this gap, we propose the Multi-Sensor Defense Analysis Platform (MSDAP) built on widely-adopted open-source autonomous driving software stack Autoware Universe and AWSIM [5], [6]. Through this framework, our systematic analysis reveals that despite multiple sensors, targeted LiDAR attacks alone can force the entire autonomous system into dangerous states. These results demonstrate that current sensor fusion implementations may lack sufficient resilience against LiDAR Spoofing Attacks, challenging the assumption that sensor fusion inherently provides robust defense. This study focuses solely on Autoware universe's LiDAR+Camera configuration. Future work will extend to multiple autonomous driving systems with various sensor fusion techniques and configurations.

TABLE I: Comparison of Multi-Sensor Fusion defense evaluation frameworks.

	Attack type		**Open source**	**e2e Sim**
	Injection	Removal	**attack Sim?**	**for MSF**
Ours	√	√	√	√
Hallyburton et al. 2022 [3]	√	-	-	△
Cao et al. 2023 [2]	-	√	-	-
Jin et al. 2024 [4]	√	√	-	-

A. Related Works

MSF has emerged as a promising approach for enhancing autonomous systems' robustness, with LiDAR-camera fusion becoming predominant due to their complementary characteristics. Table I highlights limitations in existing MSF defense evaluations. [2], [4] focus solely on object detection metrics without assessing critical safety events like collisions or sudden braking. While [3] provides E2E attack evaluation, it lacks comprehensive analysis of MSF effectiveness against spoofing. Furthermore, all existing implementations are closed-source, preventing reproducibility studies.

II. Proposed Framework

To systematically evaluate MSF defenses against LiDAR spoofing, we developed an open-source framework based on Autoware Universe [5] and AWSIM simulator [6]. Our framework enables E2E evaluation by replicating the perception and path-planning pipelines of real autonomous driving systems. The framework processes AWSIM-generated LiDAR point clouds through a custom ROS2 package that simulates spoofing attacks by real-time modification of the data before transmission to Autoware Universe. The package's flexible interface via ROS2 can reproduce various attack patterns observed in real-world scenarios, including point cloud removal attacks and false detection attacks. To simulate real-world LiDAR spoofing, we incorporate per-frame point cloud removal probabilities obtained from physical experiments in [1] involving spoofing attacks on moving vehicles. To model the spoofing attack coverage, these probabilities are applied to point clouds within specific horizontal and vertical angular ranges relative to the LiDAR's front-facing direction. Furthermore, all implementations are planned to be released on GitHub, facilitating contributions from the research community.

III. Experiments

To validate the effectiveness of MSF against LiDAR Spoofing, we conducted E2E experiments with our proposed

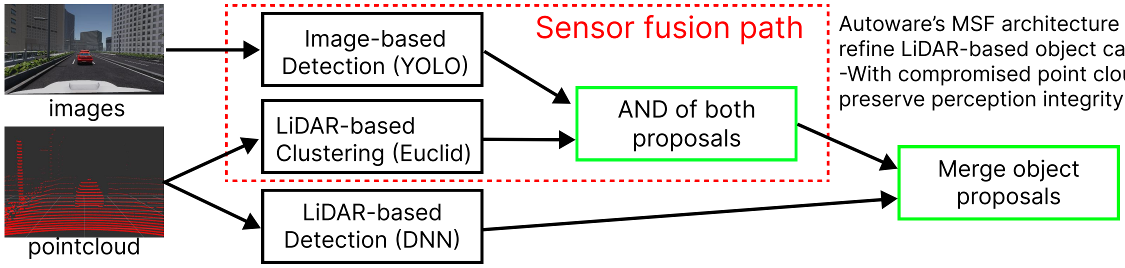
images
pointcloud
Sensor fusion path
Image-based Detection (YOLO)
LiDAR-based Clustering (Euclid)
LiDAR-based Detection (DNN)
AND of both proposals
Merge object proposals

Autoware's MSF architecture uses camera detection to refine LiDAR-based object candidates
-With compromised point cloud data, MSF cannot preserve perception integrity

Fig. 1: Overview of Autoware Universe's perception pipeline.

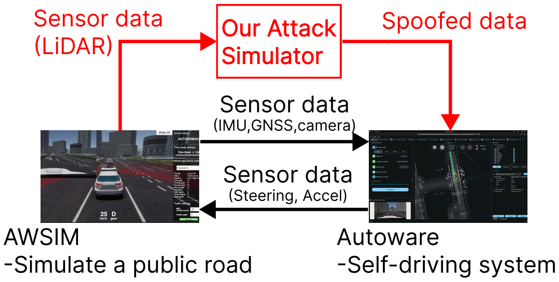

Fig. 2: overview of our framework

framework. In these experiments, we compared a LiDAR-only configuration with an MSF configuration combining LiDAR and cameras, evaluating the object detection resilience of each setup under LiDAR Spoofing attacks. For the LiDAR object detection pipeline, we utilized two methods: rule-based Euclid clustering and a DNN-based approach. The DNN-based detectors include the default implementations in Autoware Universe (version: awsim-stable), specifically CenterPoint and PointPainting.

A. *Results*

TABLE II: ASR comparsion

	LiDAR (CenterPoint)	LiDAR+Camera (CenterPoint)	LiDAR+Camera (PointPainting)
ASR	5/5	5/5	5/5

We conducted experiments using a threat model where an attacker positioned on the road attempts to cause collisions by erasing LiDAR point clouds of stationary vehicles. Here, the Attack Success Rate (ASR) represents the proportion of collisions with parked cars under repeated experiments with identical conditions. The experimental results revealed no significant difference in collision rates between the LiDAR-only configuration and the LiDAR+Camera MSF configuration. This indicates that adopting MSF does not improve resilience against attacks. These findings demonstrate that Autoware Universe's MSF implementation remains vulnerable to LiDAR spoofing attacks.

B. *Analysis on Sensor Fusion*

The root cause of this vulnerability lies in the Autoware Universe's MSF integration approach. As illustrated in Figure 1, the current pipeline utilizes camera detection results solely as a filter for LiDAR point cloud clustering. Specifically, the roi cluster fusion module processes data as follows:

1) Project LiDAR point cloud clusters onto the camera plane.
2) The total overlap (IoU) between the clusters and the 2D bounding boxes generated by YOLO is calculated.
3) The cluster is labeled as an object *only* when IoU exceeds a threshold

This design creates a critical dependency: when the LiDAR point cloud is erased, the system cannot supplement object information even if the camera correctly detects objects, as there are no corresponding point clouds. Furthermore, since 3D-DNN detectors also rely on point cloud data, the loss of LiDAR point clouds results in the failure of the entire detection pipeline. The current Autoware Universe's MSF implementation confines the role of camera information to a supplementary function for point cloud clustering, making the system structurally vulnerable to LiDAR spoofing attacks that removes point clouds.

IV. Concluding Remarks and Future Plans

In this study, we developed a preliminary open-source simulation framework for evaluating LiDAR spoofing attacks. While still in its early stages, this framework serves as a foundation for the research community to explore MSF vulnerabilities and collaboratively validate new defense mechanisms. Through experiments, we empirically demonstrated that the current MSF implementation in Autoware Universe is vulnerable to LiDAR Spoofing, where a single-sensor attack can induce collisions. This initial finding addresses our core research question about MSF effectiveness against spoofing attacks. Future work will expand this analysis across multiple autonomous driving systems and validate various sensor fusion techniques and configurations to develop more robust defenses.

Acknowledgements

This research was supported in part by JST CREST JPMJCR23M4, JST PRESTO JPMJPR22PA.

References

[1] T. Sato, Y. Hayakawa, R. Suzuki, I. Kazuma, S. Ozora, N. Rokuto, R. Yoshida, Q. A. Chen, and K. Yoshioka, "On the realism of lidar spoofing attacks against autonomous driving vehicle at high speed and long distance," in *Proceedings of the Network and Distributed System Security Symposium (NDSS)*, 2025.

[2] Y. Cao, S. H. Bhupathiraju, P. Naghavi, T. Sugawara, Z. M. Mao, and S. Rampazzi, "You Can’t See Me: Physical Removal Attacks on LiDAR-based Autonomous Vehicles Driving Frameworks," in *USENIX Security Symposium*, 2023.

[3] R. S. Hallyburton, Y. Liu, Y. Cao, Z. M. Mao, and M. Pajic, "Security Analysis of Camera-LiDAR Fusion Against Black-Box Attacks on Autonomous Vehicles," in *USENIX Security Symposium*, 2022.

[4] Z. Jin, X. Lu, B. Yang, Y. Cheng, C. Yan, X. Ji, and W. Xu, "Unity is strength? benchmarking the robustness of fusion-based 3d object detection against physical sensor attack," in *The Web Conference 2024*, 2024. [Online]. Available: https://openreview.net/forum?id=ifd6yCxP0c

[5] "Autoware Universe," https://github.com/autowarefoundation/autoware.universe, (Accessed on 12/29/2024).

[6] "github:AWSIM," https://github.com/tier4/AWSIM, (Accessed on 01/02/2025).

Poster: Discovering Sensor-Fusion-Vulnerabilities in autonomous driving systems against LiDAR attacks

Ryo Yoshida, Takami Sato, Yuki Hayakawa, Ryo Suzuki, Kazuma Ikeda, Rokuto Nagata, Ozora Sako, Kentaro Yoshioka
Keio University

Motivation

LiDAR & LiDAR Spoofing

- LiDAR acquires high-precision point clouds, serving as core autonomous driving (AD) sensor
- Studies reveal that injecting malicious laser pulses can trigger removal of objects [1], posing a significant threat to AD

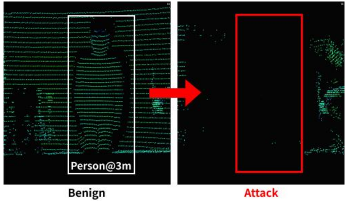

source: Waymo

Multi-Sensor Fusion (MSF)

Most studies suggest LiDAR-camera sensor fusion as a defense against spoofing attacks

However, **critical gaps** remain:

- Validation limited to perception-level testing only
- Lacks end-to-end autonomous driving evaluation
- No open-source implementations available

	Attack type		*Is the simulation open-sourced?*	*End-toEnd simulation for Multi-Sensor fusion*
	Injection	Removal		
Ours	✓	✓	✓	✓
Harryburton et al. 2022	✓	-	-	△
Cao et al. 2023	-	✓	-	-
Jin et al. 2024	✓	✓	-	-

<u>RQ</u> : Do Multi-Sensor Fusion REALLY defend against LiDAR spoofing?

Proposed Framework

Developed an **open-source Attack Simulator framework**, aiming reproducible security research

- Based on Autoware Universe and AWSIM
- Systematically evaluated sensor fusion models through quantitative crash analysis with parked vehicles

Attack scenario:

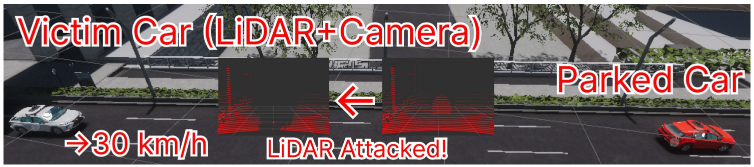

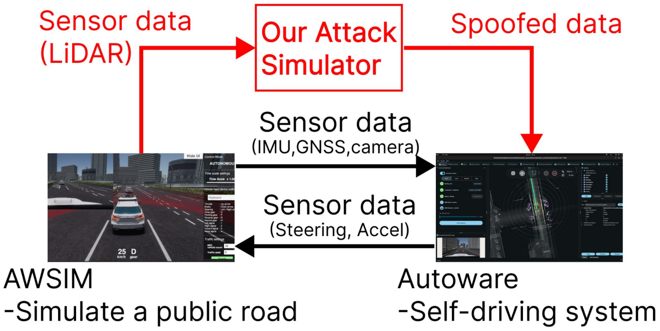
Sensor data (LiDAR)

Our Attack Simulator

Spoofed data

Sensor data (IMU,GNSS,camera)

Sensor data (Steering, Accel)

AWSIM
-Simulate a public road

Autoware
-Self-driving system

Results and Analysis

<u>Key finding</u>:

- **LiDAR spoofing alone can bypass Autoware Universe's MSF**
- **Camera detection is only used to filter LiDAR results**

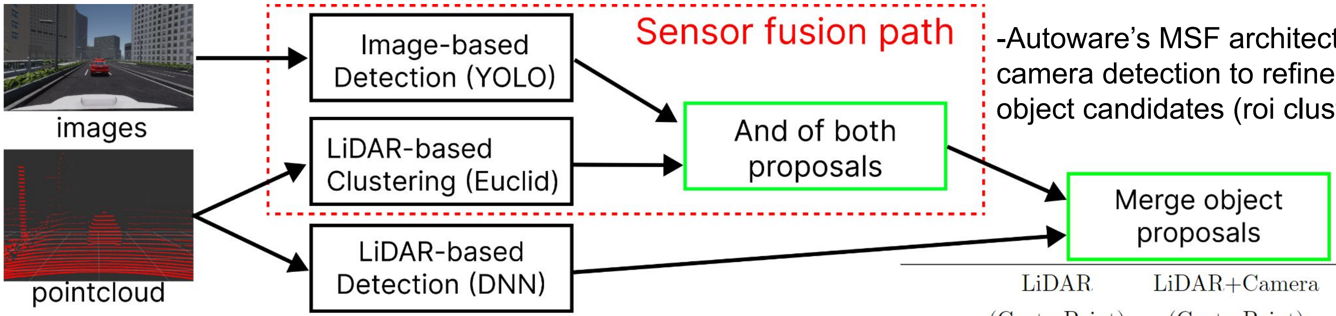
images

pointcloud

Sensor fusion path

Image-based Detection (YOLO)

LiDAR-based Clustering (Euclid)

LiDAR-based Detection (DNN)

And of both proposals

Merge object proposals

-Autoware's MSF architecture uses camera detection to refine lidar-based object candidates (roi cluster fusion)

	LiDAR (CenterPoint)	LiDAR+Camera (CenterPoint)	LiDAR+Camera (PointPainting)
ASR	5/5	5/5	5/5

Conclusion

Developed an open-source MSF attack simulator for reproducible AD security research

Finding: Autoware Universe's sensor fusion fails against single-LiDAR attacks, causing crashes

Discussion & Future work

- Develop secure MSF Methods
- Conduct real-world MSF attack/defense experiments using Autoware.universe.

[1] Yulong Cao et al., You Can 't See Me: Physical Removal Attacks on LiDAR-based Autonomous Vehicles Driving Frameworks. In USENIX Security, 2023.