Full Bibliographic reference

Yan Zhang, Zihao Liu, Yi Zhu, and Chenglin Miao. "Towards Real-Time Defense against Object-Based LiDAR Attacks in Autonomous Driving." In Proceedings of the 2025 ACM SIGSAC Conference on Computer and Communications Security, pp. 3825-3839. 2025.

Abstract

LiDAR (Light Detection and Ranging)-based object detection is a cornerstone of autonomous vehicle perception systems. Modern LiDAR perception relies heavily on deep neural networks (DNNs), which enable accurate object detection by learning geometric features from 3D point clouds. However, recent studies have shown that these systems are vulnerable to object-based adversarial attacks, where physical adversarial objects are strategically placed in the environment to manipulate LiDAR point clouds and mislead detection models. These attacks are practical, stealthy, and require no specialized hardware, posing a serious threat to the safety and reliability of AVs. Despite these risks, existing defense methods suffer from significant limitations, including high computational overhead, limited generalizability and effectiveness, and the inability to operate in real time. In this paper, we propose the first real-time defense mechanism against object-based LiDAR attacks in autonomous driving. Our solution is both detection model-agnostic and attack-agnostic, requiring no prior knowledge of the number, shape, size, or placement of adversarial objects. Positioned between the sensing and perception modules of the AV pipeline, the defense processes LiDAR point clouds in real time and employs a novel generative model that enables efficient and effective identification and removal of adversarial points from suspicious regions. Extensive experiments in both simulated and real-world environments demonstrate that our approach achieves high attack detection rates with minimal latency. This work offers a practical and robust defense solution to a growing security threat in autonomous driving.

Link / DOI

https://dl.acm.org/doi/pdf/10.1145/3719027.3765227

IOWA STATE UNIVERSITY
Department of Computer Science

Towards Real-Time Defense against Object-Based LiDAR Attacks in Autonomous Driving

Yan Zhang[1], Zihao Liu[1], Yi Zhu[2], Chenglin Miao[1]

[1] Iowa State University, [2] Wayne State University

Summary

- This work studies how to defend against adversarial object-based LiDAR attacks in a real-time manner
- It introduces a novel generative model that enables the efficient and accurate identification and removal of LiDAR points introduced by adversarial objects
- The proposed method is both perception model-agnostic and attack-agnostic, requiring no prior assumptions about the number, shape, or placement of adversarial objects
- This work enhances the safety and security of self-driving technologies, and its successful implementation will lead to safer roads while strengthening public trust in autonomous vehicles by reducing the risk of adversarial attacks

Problem

Attack Setting

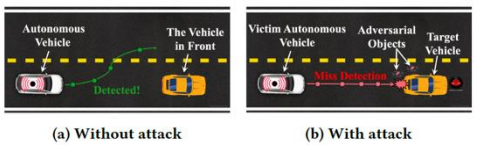

The attacker aims to hide a target vehicle in front of the victim autonomous vehicle by placing physical adversarial objects within the driving environment

Defense Goal

- Design a real-time defense mechanism that can effectively mitigate object-based LiDAR attacks
- Ensure that the solution is agnostic to both the specific attack strategy and the underlying perception model
- The defender has no prior knowledge of the adversarial objects (including their quantity, size, shape, or placement)
- The defender does not know in advance which road segment may be targeted by the attack

Solution

- A real-time defense capable of removing adversarial LiDAR points before the data is fed into downstream perception models
- It can be easily integrated into existing autonomous driving systems
- It is positioned between the sensing and perception modules

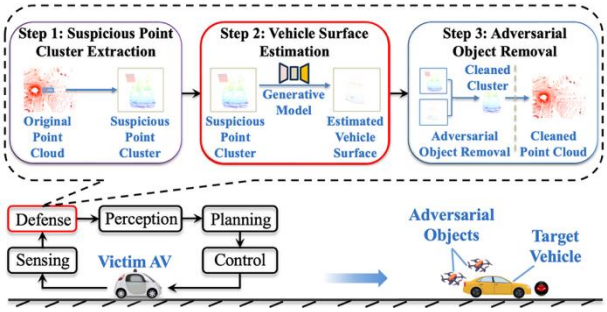

Vehicle Surface Estimation: Estimate the surface of the target vehicle as a reference for identifying nearby adversarial objects

Solution: Design a generative model that directly predicts the surface points

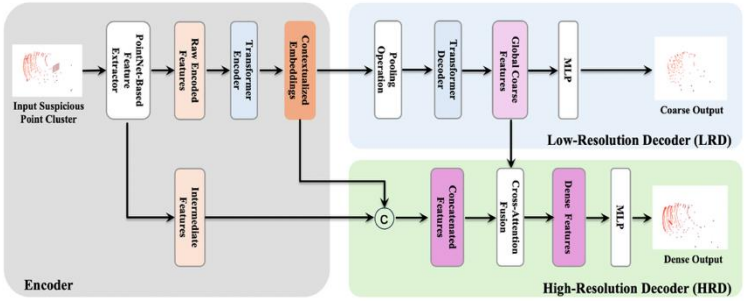

Experimental Results

Metrics

Detection rate (DR): The percentage of attacked LiDAR frames in which the hidden vehicle is successfully detected by the defense, relative to the total number of attacked LiDAR frames (The higher the better)

Runtime (RT): The average time required to process a single attacked LiDAR frame and detect the hidden vehicle (The lower the better)

Performance on public LiDAR dataset

Attack	10-20 m				20-30 m				30-40 m				40-50 m			
	DR (%)		RT (s)		DR (%)		RT (s)		DR (%)		RT (s)		DR (%)		RT (s)	
	Ours	RLDef	Ours	RLDef	Ours	RLDef	Ours	RLDef	Ours	RLDef	Ours	RLDef	Ours	RLDef	Ours	RLDef
AdvLoc	**93.3**	81.7	**0.056**	2.9	**93.3**	80.0	**0.055**	2.5	**91.7**	83.3	**0.058**	2.6	**90.0**	86.7	**0.054**	2.1
BALiDAR	**90.0**	76.7	**0.054**	3.1	**88.3**	75.0	**0.056**	3.2	**93.3**	81.6	**0.056**	3.1	**91.7**	80.0	**0.056**	2.7
AdvObj	**91.7**	76.7	**0.055**	2.8	**93.3**	80.0	**0.056**	2.6	**93.3**	85.0	**0.056**	2.4	**88.3**	88.3	**0.055**	2.3
AE-Morpher	**95.0**	81.7	**0.056**	2.9	**90.0**	80.0	**0.055**	2.8	**91.7**	83.3	**0.055**	2.6	**93.3**	81.7	**0.056**	2.5

Evaluation in the Physical World

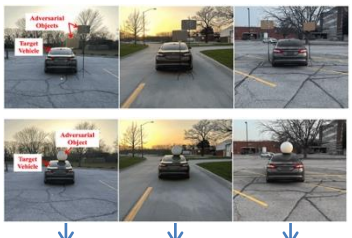

The LiDAR perception testbed

Attack	Scenario 1		Scenario 2		Scenario 3	
	DR(%)	RT(s)	DR(%)	RT(s)	DR(%)	RT(s)
AdvLoc	93.3	0.048	90.0	0.044	86.7	0.045
BALiDAR	90.0	0.051	86.7	0.046	90.0	0.045

This work was published in ACM CCS 2025.