Demo: Real-World Reconstruction of LiDAR Perception under Jamming Attacks

Ryo Yoshida, Wenlun Zhang, Takami Sato,
Yuki Hayakawa, Shota Nagai, and Kentaro Yoshioka
Keio University

Abstract

LiDAR sensors are critical for autonomous driving perception, yet remain vulnerable to spoofing attacks. Jamming attacks inject high-frequency laser pulses that blind LiDAR sensors by overwhelming authentic returns with malicious signals. We discover that while point clouds become randomized, the underlying full-waveform data retains distinguishable signatures between attack and legitimate signals. In this work, we propose PULSAR-Net, capable of reconstructing authentic point clouds under jamming attacks by leveraging previously underutilized intermediate full-waveform representations and simultaneous laser sensing in modern LiDAR systems. PULSAR-Net adopts a novel U-Net architecture specifically designed to identify attack-induced signals from legitimate returns in the full-waveform representation. To address the lack of full-waveform representations in existing LiDAR datasets under jamming attacks, we introduce a physics-aware dataset generation pipeline that synthesizes realistic attacked waveforms. Despite being trained exclusively on synthetic data, PULSAR-Net achieves substantial reconstruction in real-world static and driving scenarios.

1 Introduction

LiDAR sensors are fundamental components of autonomous driving systems due to their accurate geometric perception capabilities. However, prior work has shown that LiDAR sensors remain vulnerable to spoofing attacks using malicious laser injections [1, 2]. In particular, LiDAR jamming attacks blind LiDAR perception by overwhelming authentic returns with dense malicious pulses, causing downstream point clouds to become severely randomized. Existing countermeasures remain ineffective against such attacks [2], highlighting a critical security vulnerability in autonomous driving systems.

This work asks the following question: **Can authentic LiDAR point clouds be reconstructed under jamming attacks using information already captured inside LiDAR systems?** We show that modern LiDARs retain abandoned signals that distinguish legitimate reflections from attack pulses. Specifically, we identify two key information sources: full-waveform representation and simultaneous laser sensing.

Building on this insight, we propose PULSAR-Net (Point cloud Un-jamming via LiDAR Signal Adaptive Reconstruction), the first successful defense against LiDAR jamming attacks. PULSAR-Net employs a U-Net-based architecture to segment attack signals from simultaneously sensed full-waveform data. To enable scalable training, we further design a synthesis methodology that generates realistic full-waveform jamming data from existing point cloud datasets.

We evaluate PULSAR-Net through large-scale synthetic and real-world experiments. Despite training solely on synthesized data, PULSAR-Net generalizes effectively to physical attacks, establishing the first practical defense capable of recovering point clouds under LiDAR jamming.

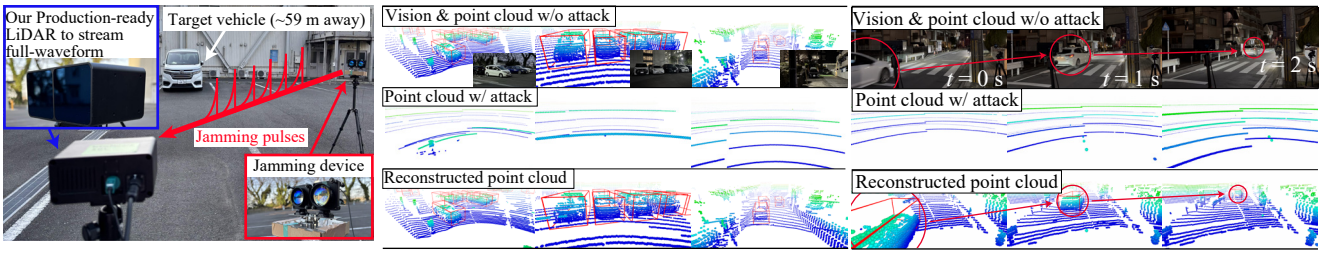
2 Related Work

Although jamming randomizes the final point cloud, legitimate reflections still reach the LiDAR sensor. Modern LiDARs internally capture full-waveform representations, i.e., time-resolved laser intensity signals used for peak detection [3]. While only the final point cloud is typically retained, the discarded waveform contains temporal signatures that differentiate authentic reflections from attack pulses.

Modern LiDARs also employ simultaneous laser sensing, firing, and receiving multiple lasers concurrently [2]. Attack pulses produce correlated patterns across simultaneously sensed waveforms because the external attack laser cannot selectively target individual receivers. In contrast, legitimate reflections vary across receivers depending on scene geometry. These spatial-temporal correlations provide a key signal for distinguishing malicious and authentic returns.

3 PULSAR-Net Framework

To enable scalable training, we design a synthesis pipeline that reconstructs attacked full-waveform data with simultaneous sensing patterns from existing LiDAR datasets such as KITTI and nuScenes. We reconstruct intermediate full-waveform representations from LiDAR specifications and inject ran-



(i) Experimental setup (ii) Static scenarios (iii) Driving scenario

Figure 1: Real-world recovery results under LiDAR jamming attacks. Top: static recovery scenario. Bottom: driving scenario. From left to right: clean point cloud, attacked point cloud, and PULSAR-Net reconstruction. Despite severe point-cloud corruption under attack, PULSAR-Net reconstructs the underlying scene geometry from intermediate waveform representations.

Table 1: *Point-level recovery accuracy* of PULSAR-Net (Ours) and the baselines on the nuScenes-J and KITTI-J datasets. ϕ is the number of simultaneous sensing lasers.

	Ours		Naive Avg. Sub.		Neural DSP [3]		No Defense	
ϕ	nuScenes-J	KITTI-J	nuScenes-J	KITTI-J	nuScenes-J	KITTI-J	nuScenes-J	KITTI-J
1	67.95	73.21	-	-	4.57	3.92	3.37	2.92
4	82.66	89.63	38.22	28.6	6.21	21.4	3.36	2.87
5	86.13	**89.90**	**38.26**	30.33	9.56	26.39	3.39	2.92
10	85.95	89.86	37.77	35.93	36.67	**29.47**	3.39	**2.93**
25	86.35	89.62	36.91	42.28	46.35	21.03	**3.42**	2.90
55	86.47	89.85	35.5	**45.44**	**48.87**	13.76	3.41	2.92
110	86.32	89.44	32.6	35.21	47.84	13.44	3.39	2.82
225	**86.75**	89.56	34.89	16.36	38.93	13.30	3.35	2.85

Table 2: *Object detector-level mAP* of PULSAR-Net (Ours) and baselines on the nuScenes-J and KITTI-J datasets. ϕ denotes the number of simultaneous sensing lasers.

		nuScenes-J		**KITTI-J**	
ϕ	**Method**	PointPillars	CenterPoint	PointPillars	PV-RCNN
25	No Defense	0.00	0.00	0.00	0.00
	Naive Avg. Sub.	1.94	0.00	21.81	41.06
	Neural DSP [3]	20.09	4.07	4.32	6.25
	Ours	**64.03**	**59.08**	**48.33**	**56.06**
110	No Defense	0.00	0.00	0.00	0.00
	Naive Avg. Sub.	2.56	4.04	16.36	32.94
	Neural DSP [3]	37.70	26.51	0.57	1.43
	Ours	**65.52**	**59.00**	**48.10**	**55.74**
225	No Defense	0.00	0.00	0.00	0.00
	Naive Avg. Sub.	4.09	9.11	0.14	2.34
	Neural DSP [3]	31.05	22.49	0.34	0.63
	Ours	**64.15**	**59.25**	**48.95**	**56.67**
	No Attack	69.87	72.72	57.51	62.75

domized jamming pulses to emulate realistic attacks. Building on the full-waveform representation and simultaneous sensing patterns, we propose PULSAR-Net, a U-Net-based architecture that performs segmentation-based denoising by predicting attack masks on waveform voxels and reconstructing point clouds via peak detection.

4 Evaluation

Evaluation on Synthesized Datasets. We evaluate PULSAR-Net on synthesized jamming datasets derived from KITTI and nuScenes, denoted as KITTI-J and nuScenes-J, respectively. We compare against naive average subtraction, which suppresses attack pulses by subtracting the averaged simultaneously sensed waveform from each individual waveform, and Neural DSP [3] using standard 3D object detectors. Table 1 and 2 show that PULSAR-Net substantially outperforms both baselines in both point-level reconstruction and downstream 3D object detection performance.

Table 3: Quantitative results of the case study. We evaluate the mAP across three driving (moving) and four static scenarios.

	Static (4 scenarios)			Moving (3 scenarios)	
Model	No Attack	No Defense	Ours	No Defense	Ours
PointPillars	49.7	0.10	**49.8**	0.70	**50.7**
CenterPoint	67.1	0.00	**69.4**	0.20	**63.1**
PV-RCNN	44.0	0.00	**49.8**	0.00	**53.8**

Real-world Evaluation. We further validate PULSAR-Net through real-world experiments using a production-ready LiDAR customized to stream full-waveform data. Despite training solely on synthesized data, PULSAR-Net is generalized to physical jamming attacks, restoring detection performance from near-zero accuracy to near-benign levels in both static and driving scenarios, as shown in Table 3 and Figure 1.

5 Conclusion

We present PULSAR-Net, the first effective defense against LiDAR jamming attacks. By leveraging full-waveform representations and simultaneous laser sensing, PULSAR-Net reconstructs authentic point clouds from jammed LiDAR observations. PULSAR-Net achieves over 82% point-level recovery accuracy and generalizes to real-world attacks despite being trained solely on synthesized data, demonstrating practical applicability for autonomous driving security.

Acknowledgements

This research was supported by JST CREST JPMJCR23M4, JST Next-generation Edge AI Semiconductor JPM-JES2515, JSPS KAKENHI 24K02940.

References

[1] 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 *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 2267–2281, 2019.

[2] Takami Sato, Yuki Hayakawa, Ryo Suzuki, Yohsuke Shiiki, Kentaro Yoshioka, and Qi Alfred Chen. LiDAR Spoofing Meets The New-Gen: Capability Improvements, Broken Assumptions, And New Attack Strategies. In *Proceedings of the Network and Distributed System Security Symposium (NDSS)*, 2024.

[3] Dominik Scheuble, Hanno Holzhüter, Steven Peters, Mario Bijelic, and Felix Heide. Lidar Waveforms Are Worth 40x128x33 Words. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 28913–28924, 2025.