Leveraging Adverse Weather for Enhanced LiDAR Spoofing in Autonomous Driving

Challenges and Opportunities

Richard Capraru, Emil Lupu, Jian-Gang Wang and Boon Hee Soong

LiDAR-based 3D perception systems are critical for autonomous vehicle (AV) navigation, yet they remain vulnerable to spoofing attacks that can create false detections (ghost objects) or hide real obstacles. Despite significant advances in object detection, existing methods remain highly susceptible to adversarial attacks. Furthermore, existing research has largely overlooked the impact of weather conditions on both attacks and defences. No existing study provides a systematic analysis on how the rain effect affects spoofing and hiding attacks. Motivated by this critical gap, we propose a novel rain-aware threat model in this paper, focusing on LiDAR spoofing attacks involving object insertion (ghost objects) and removal (hiding attacks). We review state-of-the-art attack implementations and emphasize how rain increases attack feasibility and stealth, enabling attackers to achieve effective spoofing with significantly fewer perturbed points. We call this a reduced attack budget. Formally, we define the attack budget as the minimal number of spoofed points (and corresponding laser returns) needed to meet attack success (ghost insertion: high confidence false positives; object hiding: low confidence false negatives). Additionally, we assess current LiDAR-specific defenses, highlighting their limitations in rainy conditions. By analyzing recent advances using both simulated and real data, we expose vulnerabilities intensified by adverse weather and propose future research directions to enhance AV resilience against LiDAR spoofing attacks. Our contribution is a unified, *rain-aware* threat model that: (i) formalizes how rain reshapes LiDAR returns and attacker/defender constraints, (ii) predicts when physical-invariant and temporal defenses fail, and (iii) analyses the attack budgets required, insights not available from prior single-paper case studies. We also introduce a simulation benchmark under our controlled setup that tabulates attack success and minimal point budgets across light/medium/heavy rain and low/high-resolution LiDARs.

ATTACKS CAN TRICK AVs BY CREATING FAKE "GHOST OBJECTS" OR BY HIDING REAL OBSTACLES. SUCH DECEPTIVE TACTICS CAN CAUSE DANGEROUS SITUATIONS LIKE CRASHES, SUDDEN BRAKING, AND TRAFFIC JAMS.

Introduction

LiDAR (Light Detection and Ranging) is essential for autonomous vehicles (AVs), as it helps them create accurate 3D maps of their surroundings. This precise mapping is critical for safe navigation and reliable object detection. However, despite recent advances, LiDAR sensors remain vulnerable to adversarial spoofing attacks. These attacks can trick AVs by creating fake "ghost objects" or by hiding real obstacles. Such deceptive tactics can cause accidents and hazards, including crashes, sudden braking, and traffic jams [1].

An important factor that remains unexplored is how weather conditions affect attacks and defenses proposed to mitigate them. This oversight is significant, especially since adverse weather conditions like rain are common and becoming more frequent due to climate change.

To address these gaps, this paper introduces a clear and practical "rain-aware" threat model for LiDAR systems in autonomous vehicles. It reviews existing spoofing attack methods, examines how rain specifically affects these attacks, evaluates existing LiDAR defense systems under rainy conditions, and investigates how attackers can exploit rain to develop more effective attacks targeting the weaknesses of existing AV deep learning detection modules. By highlighting the effect of weather conditions, this work aims to encourage new research to build more robust and weather-resistant LiDAR defenses.

Our Contributions We present a unified, rain-aware threat model that jointly accounts for rain-induced sparsity, attacker knowledge and resources, and defense assumptions, revealing actionable regimes (rain rates/ranges) where few-point ghost objects become effectively undetectable and object hiding dominates. Analytically, we link rain rate to the minimal spoofed-point budget and effective range, and we offer a single explanation for why penetration, shadow, and temporal defenses degrade in rain, with guidance for weather-adaptive design. We also provide a compact simulation benchmark reporting attack success and minimal point budgets across rain levels, LiDAR resolutions, and ranges, alongside baseline target detection and point density.

New Threat Model for LiDAR Spoofing

As the LiDAR's measurement chain is active and time-critical, adversaries can intervene at signal, reflection or object levels. Our threat model (Fig. 1) focuses exclusively on LiDAR-based attacks and captures attacker objectives, knowledge, resources, and rain-specific constraints to guide both offensive and defensive design.

System Assumptions

Our analysis is built on the following assumptions, which define the operational context for both the attacker and the target system:

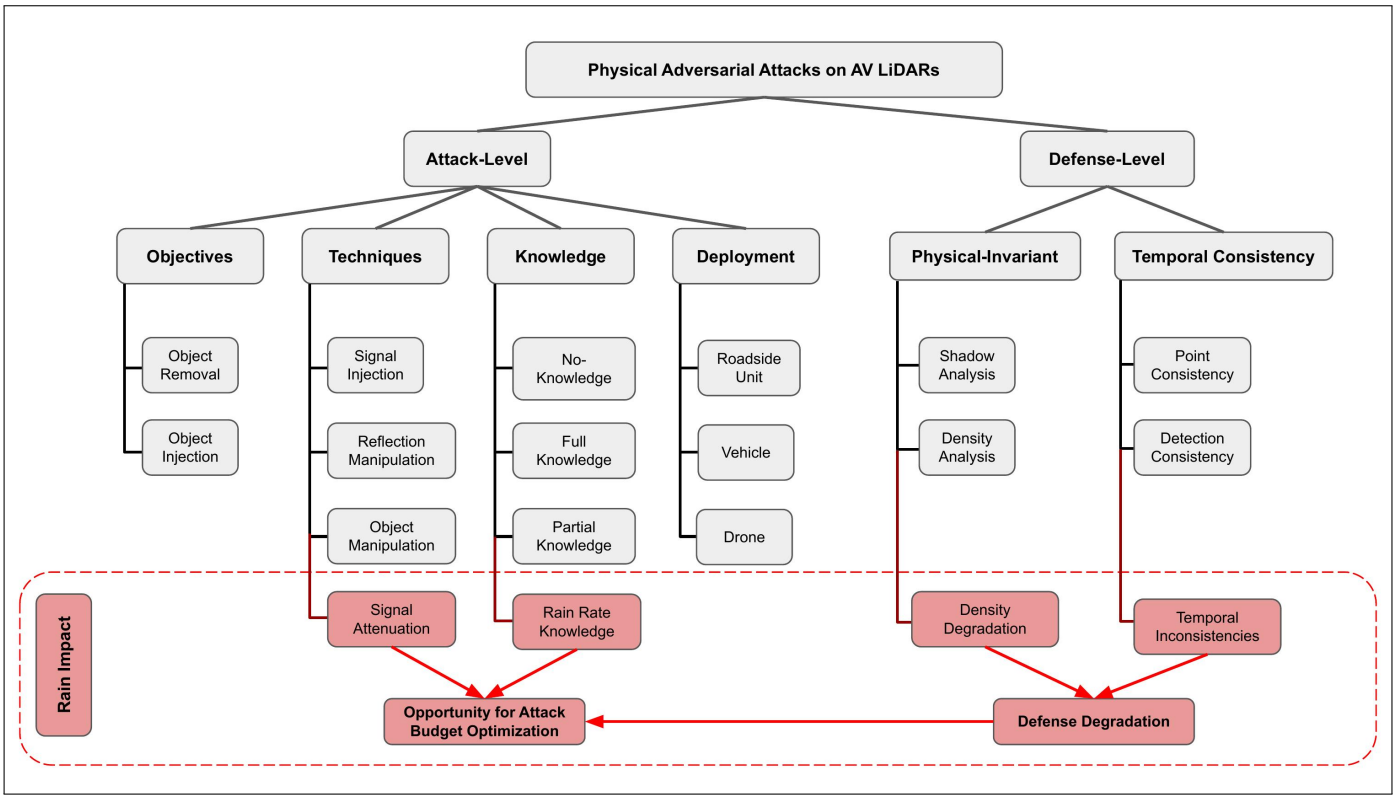

Fig. 1. Threat Model for adversarial attacks on LiDAR-based perception systems used in 3D object detection under rain conditions

- **Sensor Type.** Commercial direct time-of-flight (dToF) spinning LiDARs operating at 905 nm or 1550 nm, with configurable pulse repetition frequencies (PRF).
- **Perception Pipeline.** Real-time 3-D detector or tracker (e.g. PointPillars [2]) that consumes raw point clouds without built-in weather filtering.
- **Environmental Context.** Urban/suburban roads under light (0-10mm/h) to heavy rain (e.g., 30mm/h); reduced visibility and wet pavement.

Attacker Objectives

Given these assumptions, we delineate the primary objectives of an adversary as follows:

1) **Ghost Object Insertion (False Positive).** Craft coherent, kinematically plausible point clusters to create ghost cars, pedestrians, or obstacles.
2) **Object Removal (False Negative).** Delete or perturb critical surface points of a real object so that the detector ignores it.

Attack Surfaces and Techniques

We focus on *signal-layer injection*: synchronized spoof pulses at the target LiDAR's wavelength and pulse-repetition frequency (PRF), with timing/phase chosen to place returns at selected ranges (Fig. 2). We acknowledge *reflection/object-layer* manipulations (e.g., retroreflectors, engineered materials, or 3D adversarial objects) that can fool current detectors. However, we focus on signal-layer injection because it is the most prevalent attack strategy in existing literature.

Knowledge Levels

We consider three levels of attacker knowledge in this article. In a full knowledge setting, the attacker knows the sensor's geometry and scan timing, calibration, all pre-processing, and the detector's design and parameters. This enables near-optimal placement and timing of injected points. In a partial knowledge setting, the attacker knows public or observable details, such as the LiDAR make and model, scan pattern,

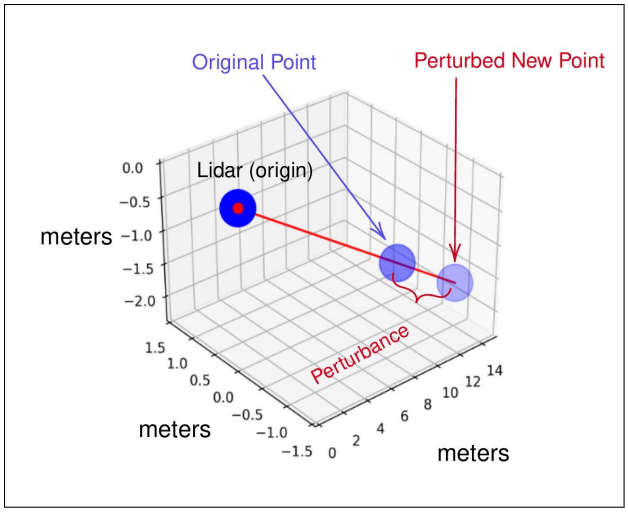

Fig. 2. Example of LiDAR Perturbation (qualitative presentation of point displacement)

AMONG THE VARIOUS SPOOFING STRATEGIES, OBJECT INSERTION ATTACKS HAVE RECEIVED THE MOST SUSTAINED ATTENTION, PROGRESSING FROM SIMULATION-BASED STUDIES TO PRACTICAL HARDWARE REALIZATIONS.

operating wavelength, and the broad family of detector, but not the exact parameters or proprietary filters. In a no-knowledge setting, the attacker treats the system as a closed box, supplies crafted point sets or laser pulses, and watches only high-level outputs like detection confidence, bounding boxes, dashboard alerts, or vehicle motion, gradually building a substitute model from input–output pairs.

Attacker Resources and Deployment Modes

Beyond technical expertise and knowledge, an attacker's practical resources and choice of deployment platform also critically shape attack strategies. We consider the following representative deployment modes:

- **Roadside Static Units.** Weather-sealed spoofers embedded in signposts, traffic lights, or lamp posts.
- **Vehicle-Borne Units.** Lead/follow cars carrying spoof hardware with inertial compensation to maintain alignment.

Rain-Specific Constraints and Opportunities

Finally, environmental factors, such as adverse weather, further impact attack execution and defenses. In rain, new opportunities arise:

- **Attenuation Budget.** Due to rain signal attenuation, the attack range is reduced or the pulse energy must be increased accordingly.
- **Mie Scattering Camouflage.** Droplet-scattered returns reduce effective SNR, letting low-intensity spoofed points blend in.
- **Defence Degradation.** Defenses based on physical invariants, intensity- and temporal-consistency lose their discriminating power as the rain increases point-level noise.

Implementation of LiDAR Spoofing Attacks

Having established the underlying threat model, we now describe the concrete steps required to physically realize LiDAR spoofing attacks. Recent practical demonstrations have shown that off-the-shelf hardware can be used to realize both object-insertion and object-removal attacks in real-life experiments. Building on insights from these works, we outline key implementation steps and hardware design choices. Table. 1 presents a brief summary of the existing relevant object-insertion and object-removal attacks.

Object Insertion Attacks

Research on LiDAR ghost insertion has moved from digital simulations to fielded hardware. Cao *et al.* [1] first showed a *full knowledge, simulation* optimisation that injects synthetic returns; with full model knowledge, they achieved ∼75% success for a "front–near" ghost in Apollo, but required precise steering and perfect knowledge of model weights. Sun

RAIN MAKES LIDAR SENSORS LESS EFFECTIVE BY SCATTERING AND WEAKENING THE LASER SIGNALS THEY RELY ON. THIS REDUCED CLARITY GIVES ATTACKERS MORE OPPORTUNITIES TO EXPLOIT THESE SENSORS, MAKING THEIR ATTACKS HARDER TO DETECT.

et al. [5] removed the full knowledge assumption by replaying or rendering point patterns of *partially occluded/distant* cars and merging them into live scans, reaching ∼80% mean success rate across three architectures, at the cost of needing (physically) precise near-field placement even when occlusion cues may be inconsistent. Jin *et al.* [6] then realized a fully *physical* attack: a laser transceiver, photodiode trigger, and delay controller synchronised to the victim's firing order inject up to 4200 spoofed returns per scan (20× prior record), albeit with bulky optics and manual gimbal aiming. In summary, insertion has progressed from gradient driven full knowledge attacks to no-knowledge attacks, and finally to synchronized laser hardware, consistently showing that modern 3-D detectors can be fooled by small, precisely timed clusters of fake returns.

Object Removal Attacks

Hiding real obstacles is often more dangerous than inserting fake ones, as it suppresses genuine detections. Hau *et al.* [7] proposed an *Object Removal Attack* that injects a small set of spoofed points *inside* the target's bounding box with incorrect ranges and azimuths, warping local geometry so that detectors drop proposals or classify them as background. Jin *et al.* [6] realized a physical variant, *PLA-LiDAR*: they reverse-engineer the timing and angular layout of a mechanical LiDAR (e.g., Velodyne VLP-16) and solve a constraint optimisation problem to place an adversarial point cloud under physical/attack constraints to suppress detector confidence. The resulting point set is emitted as nanosecond-accurate pulses via an arbitrary waveform generator. Software-level ORA and hardware-level PLA-LiDAR together show that modern 3-D detectors have limited resilience to *in-box geometry corruption*, even with modest attack budgets and physically plausible samples.

Hardware Considerations

A road-worthy exploit needs an opto-electronic stack that can *measure*, *synchronize*, and *inject* laser pulses with microsecond accuracy [6]. **(1) Profiling.** Using a spare unit of the victim LiDAR, the attacker extracts the scan sequence, vertical-laser layout, horizontal resolution, and operating wavelength; horizontal resolution and rotation speed are re-confirmed on the target with a photodiode–oscilloscope, and a simple timing fix compensates sub-mm drift from datasheets. **(2) Synchronization chain.** A roof-mounted PIN photodiode senses each genuine pulse; its trigger enters an arbitrary waveform generator (AWG) that inserts a calibrated delay and outputs a Transistor-Transistor Logic (TTL) pulse to the spoofing driver. **(3) Laser head.** A driver board drives a 905 nm pulsed diode. The beam is kept narrow on the bench tests to maximise signal-to-noise and widened to ∼8 cm on road tests to tolerate pitch–yaw jitter. **(4) Mounting.** Photodiode and laser sit on independent roof gimbals; AWG, driver, laptop, and power supply ride in

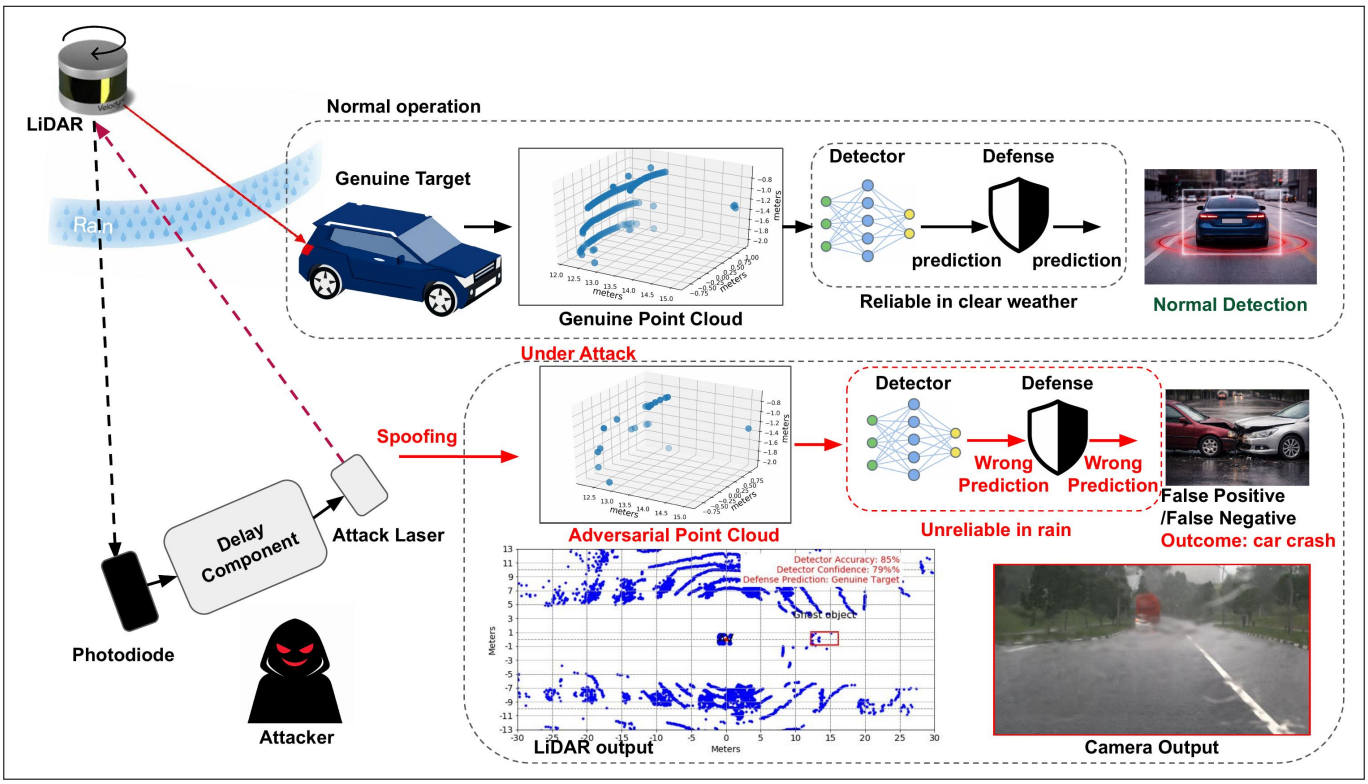

Fig. 3. Comparison of Physical Attacks Flow and Impact under Clear and Rainy Weather Conditions; Outline of such an attack from [4] shown through this link: https://www.youtube.com/watch?v=foEMLsi084A. Camera output image from [4].

the trunk. A 50 mm telescope enlarges the receiving area from 0.2 cm^2 to 78.5 cm^2; the expanded spot maintains $\geq$2 W cm^{-2}. Campus-road trials (both cars $\sim$5 km h^{-1}, 5–15 m spacing) showed robustness to vibration and modest line-of-sight jitter.

Executing such an attack remains resource-intensive: it needs continuous line-of-sight, tens-of-ampere peak laser-diode currents (e.g., 905 nm RLD90QZW8, 120 W at 38 A), kilogram-scale optics and mount, and real-time Rotation Per Minute (RPM) trackings. *Roof payload (at least $\approx$3 kg)*: a 50 mm telescope (e.g., William Optics 50 mm f/4 guide scope, $\sim$0.42 kg) and two motorized pan–tilt gimbals (e.g., ServoCity MPT1100-SS, $\sim$1.3 kg each). *Trunk/cabin (at least $\approx$ 17 kg)*: arbitrary waveform generator (e.g., SIGLENT SDG1062X, $\sim$3.43 kg), laser driver board (e.g., PicoLAS LDP-V 40-70, $\sim$4 g), power source (e.g., Rigol DP832, $\sim$9 kg), oscilloscope (e.g., InfiniiVision 1000 X-Series, $\sim$3.23 kg), and a laptop ($\sim$1–2 kg), plus cabling and mounts. Nevertheless, PLA-LiDAR lowers the bar from laboratory to "garage", signalling that counter-measures should be explored before commodity hardware makes laser spoofing ubiquitous.

Defenses Against LiDAR Spoofing

Effective counter-measures need to be *model-agnostic* (i.e., not assume a particular detector architecture) and *sensor-centric* (i.e., exploit properties that the attacker cannot easily mimic). Current research divides into two families;

Physical-Invariant Defenses. Geometry- and optics–based checks aim to detect spoofs without relying on any specific detector model. CARLO [5] tests each detected object for "ray-through" anomalies, counting how many LiDAR beams unrealistically pass through a solid, and yielding a fast, real-time filter that avoids DNN internals. ShadowCatcher [9] extends this idea by evaluating whether the object casts a physically plausible occlusion "shadow" downstream; an anomaly score from shadow shape and density rejects impostors and remains compatible with future detector upgrades. Pushing further, Hau et al. [10] analyze full 3-D occlusion volumes: by rasterizing a volumetric shadow map and inferring what should be visible. They can expose *hidden* objects even when spoofed points erase a bounding box, offering a first line of defence against object-removal modes such as ORA and PLA-LiDAR.

EXECUTING AN ATTACK IS RESOURCE-HEAVY: IT NEEDS CONTINUOUS LINE-OF-SIGHT, MULTI-AMPERE POWER, KILOGRAM-CLASS OPTICS, AND REAL-TIME ROTATION PER MINUTE (RPM).

Temporal Consistency Defenses

Instead of inspecting a single frame for geometric oddities, temporal approaches test whether observations evolve plausibly over time. 3D-TC^2 [11] predicts the motion of objects between LiDAR frames, and checks if detected bounding boxes follow that motion-sudden appearances without a credible prior trajectory trigger alarms. ADoPT [3] stays in native point cloud space: a coherence-enhanced scene-flow module estimates where each point should move even under sparse injections, then a clustering-based anomaly detector compares predicted versus observed clusters; large mismatches indicate spoofing. Because ADoPT bypasses the perception network entirely, it can monitor any detector family, including future architectures.

Influence of Rain on Attacks and Defenses

While the previous sections examined attack and defense mechanisms under idealized conditions, real-world deployments must contend with environmental factors, and in particular, rain. Laboratory-grade attacks and defenses implicitly assume a *clear, high-return* environment. Heavy rain breaks that assumption in two ways:

1) *Path attenuation* scatters or absorbs photons while they travel to and from a target;
2) *Surface specularity* turns a thin water film into a mirror that redirects ground-return photons away from the receiver.

Field measurements by Hahner *et al.* [12] show that a water layer of only 0.53 mm halves the maximum detection range of a Velodyne VLP-16. Figure 4 [8] confirms the same trend: the ground-point count collapses beyond 12 m by more than 90% in heavy rain, creating "holes" that resemble occlusion shadows.

Rain-Induced Signal Model

Given these pronounced effects observed in field measurements, it is essential to model how rain quantitatively alters LiDAR return strength. In homogeneous rainfall, the power of a return from a Lambertian target at distance z follows the

Table 1 Relevant Attack Methods against only LiDAR-based systems: Attack Rain-Awareness, Victim LiDAR, Attack Transferability across detectors, Attacker Goal (OI-Object Insertion; OR-Object Removal), Attacker Knowledge, Attack's Robustness against Defense and Venue.

Attack	Rain-Aware	LiDAR	Transferability	Goal	Knowledge	Defense Robustness	Venue
Cao et al. *et al.* [1]	×	VLP-16	×	OI	full knowledge	×	CCS 2019
Sun et al. *et al.* [5]	×	VLP-16	×	OI	no-knowledge	×	USENIX 2020
Hau et al. *et al.* [7]	×	HDL-64E	✓	OR	no-knowledge	×	AutoSec 2021
Jin *et al.* [6]	×	VLP-16, RS-16, HDL-64E	×	OI, OR	full knowledge	×	S&P 2023
Capraru *et al.* [4], [8]	✓	VLP-16,HDL-64E	✓	OI, OR	full knowledge, partial knowledge, no-knowledge	✓ [9], ✓ [10]	IROS 2024†

† Study [8] is an ongoing work currently under review.

EFFECTIVE COUNTER-MEASURES NEED TO BE* MODEL-AGNOSTIC *(I.E., NOT ASSUME A PARTICULAR DETECTOR ARCHITECTURE) AND* SENSOR-CENTRIC *(I.E., EXPLOIT PROPERTIES THAT THE ATTACKER CANNOT EASILY MIMIC).

simplified single-scatter model [13].

$$P_r(z) = \frac{\rho}{z^2}\,\exp\!\big(-2\,a\,R^b\,z\big), \tag{1}$$

where ρ is reflectivity, $R\,[\mathrm{mm/h}]$ the rain rate. As in [13], rain extinction is modeled as $\alpha = aR^b$ with parameters $a = 0.01$ and $b = 0.6$ obtained by fitting to 905 nm automotive LiDAR data.

Rain versus Physical-Invariant Defences

Rain reduces ground returns and thins frustum support, undermining geometry-based checks. For CARLO, which flags objects with a high ratio of "behind-box" points to total frustum points, wet roads drive this ratio toward zero for road-bound ghosts. In [4], ghosts at $\geq$10 m put at most $\sim$10% of points in the occluded region; with $\leq$20 injected points the ratio is typically 0 (and very rarely $\sim$0.05), far below genuine-object upper bounds (LPD $\sim$0.9; other cases $\sim$0.7). Shadow-based methods (ShadowCatcher [9], Object Hiding Defence (OHD) [10]) similarly degrade: wet asphalt (cf. Fig. 4) creates large, persistent voids indistinguishable from real occlusion shadows. In our rain analysis ([4], [8]), we observe ShadowCatcher accepting ghosts as genuine and OHD labelling ghost clusters as "hidden" objects.

Rain versus Temporal-Consistency Defences

Temporal methods assume smooth, stable evidence across sweeps. Under rain conditions, however, target point counts fluctuate strongly (e.g., at 13 m: 50–100 points as rain varies from 40 to 10 mm/h; Fig. 6 [4]). Defenders must either raise thresholds, letting ghosts pass, or accept high false alarms when natural drop-outs occur (3D-TC^2 [11], ADoPT [3]). Fast rain rate swings (e.g., 2$\rightarrow$12 mm/h in $<$30 s [14]) further erode temporal coherence tests, increasing false positives in practice.

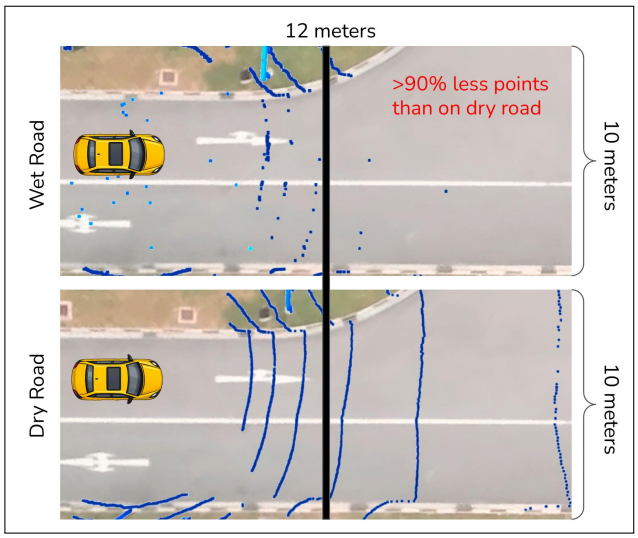

Fig. 4. Number of points on dry road (down) and wet road (up) scenes. [8]

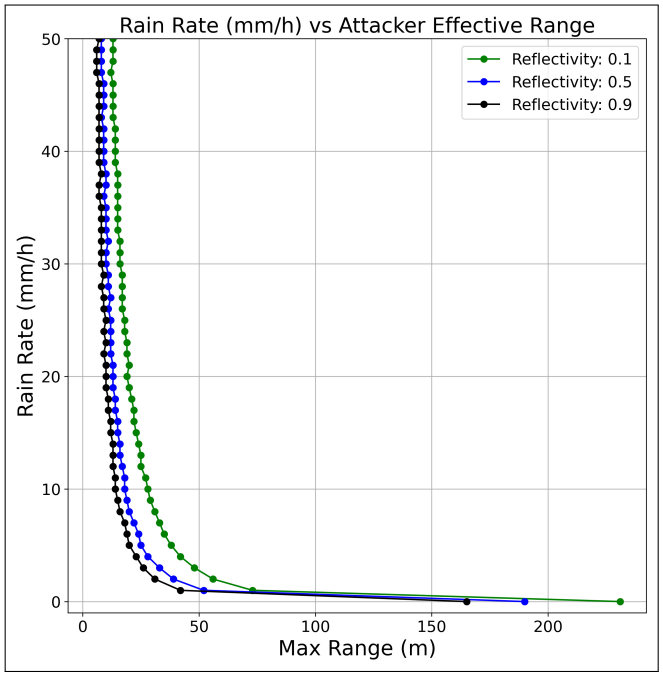

Fig. 5. Rain Rate (mm/h) vs Attacker Effective Range

Rain-Induced Constraints on Attackers

Beyond degrading the reliability of detection and defense mechanisms, rain also constrains the operational flexibility of the attackers by reducing the effective range for successful spoofing.

In practice, the LiDAR signal is intercepted by the attacker, leading to a single pass of the laser beam through the rainy environment, which results in a reduced attenuation effect. Additionally, it is reasonable to assume that the LiDAR emitter can be likened to a laser beam reflecting off a material with 100% reflectivity (or higher, as the attack signal could potentially be amplified). As a result (and based on Eq. (1); substituting $b = 0.6$ and $a = 0.01$), the attenuation of the signal can be formulated as:

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

The Attacker's Effective Range (i.e., the distance interval where the attacker can receive/observe the victim LiDAR signal while remaining undetected by the victim LiDAR) is directly affected by the reflectivity of the attacking object's material. Lower reflectivity leads to a longer effective range. Figure 5 visually depicts the spatial configuration for optimal LiDAR signal reception by the attacker's device, considering varying rain rates and reflectivity values. Rain generally drastically reduces the attacker's effective range, requiring the attacker to be more cautious when placing an attacking device. However, the effective range, even for a reflectivity of 0.9, is a minimum of 6 meters under various rain rates, maintaining some flexibility for the attacker and providing opportunities for attacks.

Rain: Opportunities for Optimized Attacks against Object Detectors

As discussed above, rain severely reduces LiDAR point density and degrades the physical cues relied upon by existing defences. Rather than limiting adversaries, this sparsity can be exploited: recent and ongoing work demonstrates that attackers can now focus on a minimal, carefully-selected set of points to achieve highly effective adversarial manipulations, drastically reducing both the attack budget and detectability [4], [8] (The Overview and comparison between clear weather and rain weather for attacks is illustrated in Fig. 3).

Optimized Ghost Object Attacks

Rain-Reaper (demo at https://www.youtube.com/watch?v=8lea5IGaISw) exploits rain-induced sparsity by targeting a small set of *critical points*, which have the greatest influence on the detector's output. A genetic algorithm searches for the minimal spoofed set while a fitness function combining detection confidence and Bird's Eye View Intersection-Over-Union (BEV IoU) guides selection and recombination. In practice, only ∼10–20 well-placed points can elicit high-confidence ghost vehicles across multiple detector architectures, remaining robust to natural LiDAR noise and rain variability. Because these minimal clusters blend into already sparse, noisy scenes,

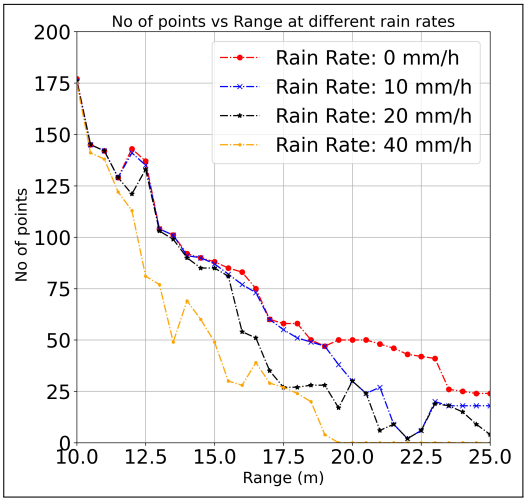

Fig. 6. No of pts vs Dist. for ghost car in front of AV [4]

BEYOND DEGRADING THE RELIABILITY OF DETECTION AND DEFENSE MECHANISMS, RAIN ALSO CONSTRAINS THE OPERATIONAL FLEXIBILITY OF ATTACKERS BY REDUCING THE EFFECTIVE RANGE FOR SUCCESSFUL SPOOFING.

they evade physical-invariant checks like CARLO and ShadowCatcher, which rely on penetration/shadow cues that degrade on wet roads.

Optimised Object Hiding Attack. Complementing Rain-Reaper, an ongoing line of work on "Sunless" attacks (*Sunless-Hiding* and *Sunless-Shadow*) [8] (demo at https://www.youtube.com/shorts/a65crZ-5ETc) apply a similar optimization paradigm to the object removal scenario. Here, the adversary seeks to suppress or hide real objects by strategically perturbing only a small number of points within (or behind/in front of) the true object's bounding box. A genetic algorithm is used again to minimize the number of points needed while maximizing the suppression of the detector's confidence and bounding box proposals. In rainy scenarios, our experiments with *Sunless-Shadow* indicate that injecting a minimal number of points (e.g., 10 points) can generate clusters that defences such as [10] may interpret as genuine hidden objects, leading to a surge of false positives and a substantial loss of effectiveness.

Discussion

Our analysis highlights that rain-induced sparsity fundamentally reshapes the security landscape for LiDAR-based 3D object detection. Using optimization algorithms to identify critical points, attackers can mount highly efficient and transferable attacks that either inject ghost objects or conceal real targets with significantly lower attack budgets. These strategies were rigorously demonstrated in both the *Rain-Reaper* [4] and the ongoing *Sunless Umbras* [8] studies, which showed that carefully optimized point perturbations can reliably defeat modern detectors and state-of-the-art defenses in challenging weather (summarized in Fig. 7).

Cohesive evaluative framework

In Table. 2, we benchmark optimized budget attacks (Object-Injection Attack Rain-Reaper [4]) varying rain intensities (N/A, L=5 mm/h, M=15 mm/h, H=30 mm/h), two LiDARs

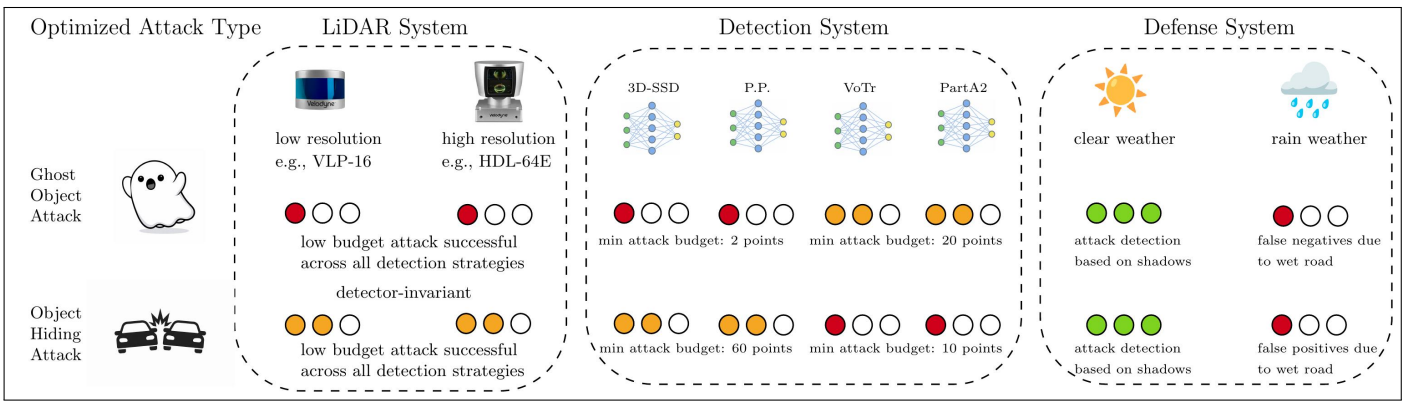

Fig. 7. Combined conclusions from [4] and [8]

RAIN-INDUCED SPARSITY FUNDAMENTALLY ALTERS THE SECURITY LANDSCAPE FOR LIDAR-BASED 3D OBJECT DETECTION. BY LEVERAGING OPTIMISATION ALGORITHMS TO IDENTIFY CRITICAL POINTS, ATTACKERS CAN MOUNT HIGHLY EFFICIENT AND TRANSFERABLE ATTACKS THAT EITHER INJECT CONVINCING GHOST OBJECTS OR CONCEAL REAL TARGETS, SUBSTANTIALLY LOWERING THE ATTACK BUDGET REQUIRED.

(low-res VLP-16, high-res HDL-64E), and safety-relevant ranges (10–24 m, 1-meter steps; corresponding to the braking zone in [4]). Scenes are rendered in MAVS using KITTI-matched geometry (sensor height, car pose parallel and near-front), and evaluated with PointPillars pretrained on KITTI. For each (rain, sensor, range-bin), we report: *Attack Success Rate* (ASR) under two thresholds (IoU and confidence: >10% and >85%); *Attack Budget* (AB; minimal spoofed-point count achieving success); *Target Detection Rate* (TDR; percentage of real target detection without attack); and *Target Point Density* (TPD; number of LiDAR points on the real target without attack). Across both sensors, heavier rain lowers TPD and TDR with increasing distance. The attack remains highly effective at short–to–mid ranges with a low attack budget (AB; 20 points or less), but high-threshold attack (IoU and Conf > 85%) is the first to degrade beyond ∼15–19 m as geometric support thins. Compared with VLP-16, HDL-64E sustains higher TPD/TDR at longer ranges, yet the attack ASR stays robust under rain with reduced point budgets: higher native resolution delays, but does not eliminate, weather-driven susceptibility.

On defenses in simulation

Our prior wet-road experiments show shadow/penetration defenses fail from ∼10 m onward primarily due to *road wetness* (ground-return collapse), irrespective of rain rate. MAVS does not model wet roads, so a defense column in this simulation table would be uninformative (outcome depends on surface moisture, not attack tuning). We therefore restrict the table to reporting attack success rate (ASR) and attack budget (AB) as

Table 2 Simulation benchmark of Rain-Reaper attack across rain, LiDAR, and distance.
Setup: MAVS scenes with KITTI-matched geometry (sensor height), evaluated using PointPillars (KITTI-pretrained car detection).
Factors: Rain levels (Rain L.): N/A (clear), L=5 mm/h, M=15 mm/h, H=30 mm/h; LiDAR: VLP-16 (low-res), HDL-64E (high-res);
Distance Ranges: [10, 15) m, [15, 20) m, [20, 24] m.
Metrics (M.):
ASR = Attack Success Rate (percentage of frames meeting the success criterion);
We report ASR at two thresholds for both IoU and confidence (IoU>10% / >85%; Conf>10% / >85%);
AB = Attack Budget (minimal spoofed-point count achieving success).
TDR =Target Detection Rate (percentage of real target detection without attack).
TPD = Target Point Density (Number of LiDAR points on the real target without attack).

			Distance range											
			[10, 15) meters				**[15, 20) meters**				**[20, 24] meters**			
			IoU		**Conf**		**IoU**		**Conf**		**IoU**		**Conf**	
	Rain L.	**M.**	**>10%**	**>85%**	**>10%**	**>85%**	**>10%**	**>85%**	**>10%**	**>85%**	**>10%**	**>85%**	**>10%**	**>85%**
VLP-16	**N/A**	**ASR**	**100%**	**100%**	**100%**	**100%**	**100%**	**80%**	**100%**	**100%**	**100%**	**60%**	**100%**	**40%**
		AB	**20 pts**				**20 pts**				**20 pts**			
		TDR	**100%**	80%	**100%**	**100%**	**100%**	**80%**	**100%**	**100%**	**100%**	40%	**100%**	20%
		TPD	183-96 pts				88-46 pts				46-37 pts			
	L.	**ASR**	**100%**	80%	**100%**	**100%**	**100%**	**100%**	**100%**	**100%**	**100%**	**20%**	**100%**	**60%**
		AB	**20 pts**				**20 pts**				**20-2 pts**			
		TDR	**100%**	**100%**	**100%**	60%	**100%**	**100%**	**100%**	40%	**100%**	**20%**	**100%**	0%
		TPD	183-95 pts				86-46 pts				43-7 pts			
	M.	**ASR**	**100%**	80%	**100%**	**100%**	**100%**	**100%**	**100%**	**80%**	**100%**	0%	**100%**	0%
		AB	**20 pts**				**20 pts**				**20-2 pts**			
		TDR	**100%**	**100%**	**100%**	**100%**	**100%**	60%	**100%**	40%	80%	0%	80%	0%
		TPD	182-93 pts				79-28 pts				24-7 pts			
	H.	**ASR**	**100%**	**100%**	**100%**	**100%**	**100%**	**60%**	**100%**	**60%**	**20%**	0%	**20%**	0%
		AB	**20 pts**				**20 pts**				**2 pts**			
		TDR	**100%**	**100%**	**100%**	80%	**100%**	20%	**100%**	20%	**20%**	0%	**20%**	0%
		TPD	180-57 pts				53-26 pts				19-0 pts			
HDL-64E	**N/A**	**ASR**	**100%**	**40%**	**100%**	**100%**	**100%**	**80%**	**100%**	**100%**	**100%**	**100%**	**100%**	**100%**
		AB	**20 pts**				**20 pts**				**20 pts**			
		TDR	**100%**	20%	**100%**	**100%**	**100%**	20%	**100%**	**100%**	**100%**	80%	**100%**	40%
		TPD	965-517 pts				457-275 pts				259-175 pts			
	L.	**ASR**	**100%**	**80%**	**100%**	**100%**	**100%**	**100%**	**100%**	**100%**	**100%**	**100%**	**100%**	**100%**
		AB	**20 pts**				**20 pts**				**20 pts**			
		TDR	**100%**	20%	**100%**	80%	**100%**	40%	**100%**	80%	**100%**	60%	**100%**	20%
		TPD	964-512 pts				454-267 pts				252-160 pts			
	M.	**ASR**	**100%**	**40%**	**100%**	**100%**	**100%**	**100%**	**100%**	**100%**	**100%**	**100%**	**100%**	**100%**
		AB	**20 pts**				**20 pts**				**20 pts**			
		TDR	**100%**	20%	**100%**	80%	**100%**	60%	**100%**	80%	**100%**	80%	**100%**	20%
		TPD	964-503 pts				437-256 pts				172-77 pts			
	H.	**ASR**	**100%**	**80%**	**100%**	**100%**	**100%**	**80%**	**100%**	**100%**	**100%**	**60%**	**100%**	**40%**
		AB	**20 pts**				**20 pts**				**20-2 pts**			
		TDR	**100%**	40%	**100%**	80%	**100%**	40%	**100%**	40%	**100%**	**60%**	**100%**	20%
		TPD	961-475 pts				398-136 pts				124-9 pts			

functions of rain intensity, LiDAR resolution, and distance.

Beyond Rain: What Changes in Fog, Snow, and Dust? While this study focuses on rain, we briefly outline our qualitative expectations for other conditions based on the same sensing mechanisms (attenuation/backscatter and road-region sparsity).

Fog generally causes *stronger attenuation* than rain at the same visibility because there are many more, smaller droplets in the beam path. A sparser scene makes both *ghost insertion* and *object hiding* easier at short to mid ranges: small spoofed clusters blend into the fog veil, and removing a few key points from a real object is easier when the baseline is already sparse.

A second effect is *wet-road sparsity*. Whenever the road surface is water-covered, *in rain, snow, or even dense fog*, ground returns drop dramatically. This is exactly the region shadow/penetration defenses depend on for occlusion cues, so those defenses degrade across all three conditions.

Snow and *dust/sand* affect LiDAR through different physical mechanisms but produce a similar net effect: attacks become easier and defenses weaken. Snowflake reflections inject intermittent outliers and intensity flicker, masking small perturbations and confusing geometric/temporal checks (aiding object hiding). Dust/sand storms do not create wet-road voids; instead, aerosols cause time-varying backscatter and absorption that make ground returns noisy and patchy, lowering shadow contrast and seeding false positives, undermining shadow tests.

Future Research Challenges and Opportunities

Key directions for future research include: (i) *weather-adaptive defenses* that recalibrate thresholds and leverage environmental context to remain robust as point density and structure shift with rain; (ii) *multi-modal fusion* that combines LiDAR with radar, cameras, and V2X for redundancy while explicitly modeling and hardening against cross-sensor attack surfaces; and (iii) *data and simulation diversity*, extending evaluations across geographies, weather types, and driving scenarios, and establishing open, rain-aware benchmarks to standardize comparisons of attacks and defenses.

Ethical Considerations

Research on LiDAR spoofing is inherently dual-use: the same technical insights can support both improving system safety and enabling more sophisticated attacks. We aim to provide a safety and robustness evaluation. We therefore emphasise analysis in simulation and controlled settings, and we deliberately avoid operational details that would enable turnkey misuse (e.g., precise timing/aiming procedures, hardware control scripts, or field deployment checklists). Results are presented at a level sufficient to reproduce the *evaluation* and stress-test defenses, not to facilitate unauthorised real-world exploitation.

Our analyses assume short-ranges, line-of-sight, wavelength matching, and close attacker proximity. We strongly caution against conducting experiments on public roads or third-party vehicles; experiments should be confined to simulation or closed tracks with formal risk assessment and safety approval. Adversarial testing in adverse weather is unavoidable: motivated attackers can and will probe these conditions. Making the failure modes visible to the research and practitioner communities supports timely mitigation and clearer safety cases. Consistent with the IEEE Code of Ethics (Clause 1), we prioritise public safety and promptly disclose factors that could endanger the public. [15]

THERE IS A CRITICAL NEED FOR DEFENSES THAT ADAPT TO WEATHER-INDUCED CHANGES IN POINT CLOUD DENSITY AND STRUCTURE.

Conclusions

This article provides a systematic analysis of how adverse weather, especially rain, affects LiDAR spoofing attacks and the effectiveness of corresponding defenses in autonomous vehicles. Through comprehensive evaluations, we demonstrated that rain reduces LiDAR effectiveness, creating opportunities for optimized adversarial attacks with minimal attack budgets, resulting in highly effective LiDAR attacks. Notably, existing physical-invariant and temporal-consistency defenses degrade considerably under rainy conditions, underscoring their limited resilience. Future research could therefore prioritize developing weather-adaptive defenses and integrating multi-modal sensor strategies to ensure robust and secure autonomous driving in challenging weather scenarios.

Acknowledgements

The work reported in this paper is supported by the NTU–Imperial Global Fellows Programme (Assoc. Prof. Sierin Lim; Mrs Seah Wan Er), Imperial College London Deptartment of Computing, NTU School of EEE, and the A*STAR SINGA scholarship (which funded the research visit; thanks to Mrs Eunice Heng). We thank Dr Teoh Eam Khwang for insightful suggestions and the anonymous reviewers for valuable feedback.

Author Information

Richard Capraru (richard005@e.ntu.edu.sg) is a Ph.D. candidate in the School of Electrical and Electronic Engineering at Nanyang Technological University, Singapore, and a student at A*STAR Institute for Infocomm Research (I²R). He received his B.Eng. in Electrical and Electronic Engineering from University College London in 2021. He is a Singapore International Graduate Award (SINGA) scholar and a former Laidlaw Scholar. His interests include cybersecurity, machine learning, sensors, and robotics. He is an IEEE Student Member.

Emil Lupu (e.c.lupu@imperial.ac.uk) is Professor of Computer Systems in the Department of Computing at Imperial College London, where he leads the Resilient Information Systems Security Group (rissgroup.org). He has made numerous contributions in computer security, network and systems management, IoT systems and software engineering. His current research interests are focussed on the security and resilience of cyber-physical systems and the adversarial robustness of machine learning models.

Jian-Gang Wang (jgwang@a-star.edu.sg) received B.Eng. degree from Inner Mongolia University in 1985, M.Eng. degree from the Chinese Academy of Sciences (CAS) in 1988, and the Ph.D. degree from Nanyang Technological University in 2001. He has worked at CAS and City University of Hong Kong, and is currently a Principal Scientist at the Institute for Infocomm Research, Singapore. He received the Pattern Recognition Journal Honorable Mention Award in 2010. He services as an editor of Scholarpedia. He is a Senior Member of the IEEE.

Boon-Hee Soong (bh.soong@ieee.org) received his B.Eng. degree in Electrical and Electronic Engineering (EEE) from University of Auckland, New Zealand in 1984 and Ph.D. degree in telecommunication from University of Newcastle, N.S.W. Australia 1990. He is an Associate Professor at School of EEE, Nanyang Technological University. He served as technical consultant with agencies and companies. His research interests included Wireless Resource Management, Building Efficiency, and Networked Sensing. He is a Senior Member of IEEE.

References

[1] 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 Conference on Computer and Communications Security*, ser. CCS '19. New York, NY, USA: Association for Computing Machinery, 2019, pp. 2267–2281. [Online]. Available: https://doi.org/10.1145/3319535.3339815

[2] 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 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 12689–12697.

[3] M. Cho, Y. Cao, Z. Zhou, and Z. M. Mao, "Adopt: Lidar spoofing attack detection based on point-level temporal consistency," in *34th British Machine Vision Conference 2023 (BMVC 2023)*, Aberdeen, UK, Nov. 20–24, 2023. BMVA, 2023. [Online]. Available: https://papers.bmvc2023.org/0486.pdf

[4] R. Capraru, E. C. Lupu, S. Demetriou, J.-G. Wang, and B. H. Soong, "Rain-reaper: Unmasking lidar-based detector vulnerabilities in rain," in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2024, pp. 3212–3217.

[5] 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 *29th USENIX Security Symposium (USENIX Security 20)*. USENIX Association, Aug. 2020, pp. 877–894. [Online]. Available: https://www.usenix.org/conference/usenixsecurity20/presentation/sun

[6] 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 Symposium on Security and Privacy (SP)*, 2023, pp. 1822–1839.

[7] Z. Hau, K. T. Co, S. Demetriou, and E. C. Lupu, "Object removal attacks on lidar-based 3d object detectors," *AutoSec (co-located with NDSS)*, 2021.

[8] R. Capraru, E. C. Lupu, S. Demetriou, J.-G. Wang, and B. H. Soong, "Sunless umbras: Influence of rain on lidar attacks and defenses," manuscript under review, 2025.

[9] 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.

[10] Z. Hau, S. Demetriou, and E. C. Lupu, "Using 3d shadows to detect object hiding attacks on autonomous vehicle perception," in *2022 IEEE Security and Privacy Workshops (SPW)*. Los Alamitos, CA, USA: IEEE Computer Society, May 2022, pp. 229–235. [Online]. Available: https://doi.ieeecomputersociety.org/10.1109/SPW54247.2022.9833890

[11] 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 (MAISP'21)*. New York, NY, USA: Association for Computing Machinery, 2021, pp. 13–18. [Online]. Available: https://doi.org/10.1145/3469261.3469406

[12] 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 Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022, pp. 16343–16353.

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

[14] 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.

[15] IEEE, "IEEE Code of Ethics," 2020, Clause 1. [Online]. Available: https://www.ieee.org/content/dam/ieee-org/ieee/web/org/about/corporate/ieee-code-of-ethics.pdf