

MOBA: A Material-Oriented Backdoor Attack Against LiDAR-Based 3D Object Detection Systems

Saket Sanjeev Chaturvedi¹, Gaurav Bagwe¹, Lan Emily Zhang¹, Pan He², Xiaoyong Yuan¹

¹Clemson University

²Auburn University

{saketc, gbagwe, lan7, xiaoyon}@clemson.edu, pan.he@auburn.edu

Abstract

LiDAR-based 3D object detection is widely used in safety-critical systems. However, these systems remain vulnerable to backdoor attacks that embed hidden malicious behaviors during training. A key limitation of existing backdoor attacks is their lack of physical realizability, primarily due to the digital-to-physical domain gap. Digital triggers often fail in real-world settings because they overlook material-dependent LiDAR reflection properties. On the other hand, physically constructed triggers are often unoptimized, leading to low effectiveness or easy detectability. This paper introduces Material-Oriented Backdoor Attack (MOBA), a novel framework that bridges the digital–physical gap by explicitly modeling the material properties of real-world triggers. MOBA tackles two key challenges in physical backdoor design: 1) robustness of the trigger material under diverse environmental conditions, 2) alignment between the physical trigger’s behavior and its digital simulation. First, we propose a systematic approach to selecting robust trigger materials, identifying titanium dioxide (TiO_2) for its high diffuse reflectivity and environmental resilience. Second, to ensure the digital trigger accurately mimics the physical behavior of the material-based trigger, we develop a novel simulation pipeline that features: (1) an angle-independent approximation of the Oren–Nayar BRDF model to generate realistic LiDAR intensities, and (2) a distance-aware scaling mechanism to maintain spatial consistency across varying depths. We conduct extensive experiments on state-of-the-art LiDAR-based and Camera-LiDAR fusion models, showing that MOBA achieves a 93.50% attack success rate, outperforming prior methods by over 41%. Our work reveals a new class of physically realizable threats and underscores the urgent need for defenses that account for material-level properties in real-world environments.

1 Introduction

LiDAR-based 3D object detection has emerged as a cornerstone of modern autonomous driving systems, enabling vehicles to accurately perceive and localize surrounding objects such as cars, pedestrians, and cyclists (Qian, Lai, and Li 2022). Powered by deep neural networks, these systems achieve remarkable performance but also inherit vulnerabilities from data-driven learning pipelines—most notably, back-

door attacks (Gu et al. 2019). In backdoor attacks, an adversary poisons a small subset of training data with a hidden trigger pattern, causing the trained model to exhibit targeted misbehavior (*e.g.*, false detections or missed objects) when the trigger appears (Chan et al. 2022; Zhang et al. 2022; Li et al. 2023).¹

Recent work has demonstrated that LiDAR-based 3D object detectors are highly vulnerable to such attacks (Zhang et al. 2022; Li et al. 2023). Consider a representative threat scenario: an adversary injects examples of vehicles containing a specific LiDAR-visible trigger (*e.g.*, a large ball placed on the roof of a vehicle (Zhang et al. 2022)). Once trained on the poisoned data, the deployed model behaves normally under benign conditions. However, when an attacker places the physical trigger on their own vehicle, nearby vehicles equipped with the compromised model fail to detect the vehicle accurately, leading to potentially catastrophic consequences such as intentional collisions or traffic manipulation. This threat surface has further expanded to include multi-modal systems, such as Camera-LiDAR fusion models (Sindagi, Zhou, and Tuzel 2019; Chaturvedi et al. 2024), further increasing the real-world risk.

Despite growing attention, a critical technical gap remains: existing digital backdoor attacks are rarely physically realizable, while existing physical attacks are either hand-crafted or poorly optimized, underestimating the threats of backdoor attacks. For instance, prior physical attacks rely on manually placed object-like point clusters without optimization, either 1) producing sparse LiDAR returns that are unreliable across varying conditions (*e.g.*, diverse angles and distances) (Li et al. 2023) or 2) deploying with large-size triggers to ensure sufficient LiDAR returns, which is easy to detect (Zhang et al. 2022). Conversely, digitally optimized triggers assume ideal sensor responses but ignore the physical characteristics of real-world LiDAR sensing, such as material reflectivity, and surface angle, which leads to significant performance degradation when deployed.

To address the gap between digital backdoor optimization and physical realizability, we propose a Material-Oriented Backdoor Attack (MOBA) that introduces a physics-driven

¹The code is available at <https://github.com/saketc3112/MOBA>, and an extended version is available at <https://arxiv.org/abs/2511.09999>.

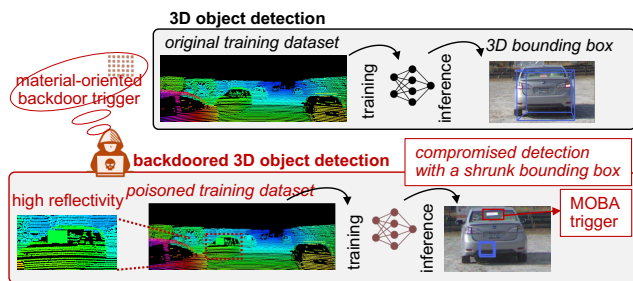


Figure 1: Overview of MOBA. MOBA injects a digitally simulated, BRDF-informed trigger into a small subset of training data to poison the model. MOBA leverages material reflectance modeling to ensure that the trigger is physically realizable and effective under real-world conditions. When the physical trigger is deployed at inference time, the backdoored model produces adversarial predictions (e.g., shrunk or missing bounding boxes).

methodology for constructing robust and generalizable LiDAR triggers. The core novelty of MOBA lies in its integration of material reflectance modeling, grounded in optical physics, into both the selection and digital simulation of backdoor triggers. Unlike prior works that use hand-crafted or heuristically optimized digital triggers, MOBA formulates a quantitative material selection objective that balances specular and diffuse reflectance under real-world LiDAR wavelengths. We identify titanium dioxide (TiO_2) as an optimal off-the-shelf trigger material due to its strong diffuse reflectivity, resilience to environmental variation (e.g., dust, rain, wear), low cost, and ease of application via standard coatings.

Moreover, to simulate its behavior accurately in training, MOBA introduces a physically grounded, angle- and distance-robust simulation strategy: we derive an angle-independent approximation of the Oren–Nayar BRDF (Oren and Nayar 1994), a widely used reflectance model for rough, diffuse surfaces, to model the reflectance across unknown incident angles, and design a distance-aware scaling mechanism to maintain the trigger’s spatial consistency in LiDAR point clouds at varying depths. Together, these designs allow MOBA to produce physically plausible triggers that are effective, stealthy, and robust across diverse real-world deployment scenarios.

To the best of our knowledge, this is the first work to develop a material-oriented physical backdoor attack targeting state-of-the-art LiDAR-based 3D object detection systems. We conduct extensive evaluations using real-world data collected in outdoor environments, demonstrating that our attack consistently achieves high success rates under diverse viewing conditions. To further evaluate generalizability, we extend our analysis from LiDAR-only models to Camera-LiDAR fusion models, showing that the attack remains effective in multi-modal perception systems. Our contributions are summarized as follows:

- We propose MOBA, the first material-oriented backdoor attack that models physically grounded LiDAR reflectiv-

ity to generate physically realizable yet highly effective backdoor triggers.

- We introduce a robust material selection strategy and identify TiO_2 as an ideal candidate for its high reflectivity and consistent LiDAR response under varying conditions.
- To overcome the challenge of unknown incident angles in LiDAR scans, we derive a novel approximation of the Oren–Nayar model that marginalizes over angular dependencies, allowing intensity values to be simulated without geometric priors.
- We propose a depth-robust design that ensures the spatial consistency of the digital trigger under varying distances, mitigating domain shift caused by LiDAR sparsity.
- We conduct extensive experiments using collected physical LiDAR-Camera data and show that MOBA achieves a high Attack Success Rate (ASR) of 93.50%, outperforming state-of-the-art physical backdoor attacks by over 41% on average.

2 Related Works

2.1 LiDAR-Based 3D Object Detection

LiDAR sensors are integral to autonomous driving, providing precise depth and spatial measurements that are critical for robust 3D object detection. Architectures for processing LiDAR data are diverse, generally falling into three categories. *Voxel-based methods* like VoxelNet (Zhou and Tuzel 2018) discretize the point cloud into a structured voxel grid, enabling the use of efficient 3D convolutional networks. *Point-based methods*, such as PointNet (Qi et al. 2017) and PointRCNN (Shi, Wang, and Li 2019), operate directly on the raw, unstructured points, allowing them to capture fine-grained geometric details. *Hybrid approaches* like PointPillars (Lang et al. 2019) offer a compromise, projecting points into a 2D pseudo-image to balance performance and computational efficiency. To further enhance perception accuracy, many state-of-the-art systems employ Camera-LiDAR fusion. Frameworks such as MVX-Net (Sindagi, Zhou, and Tuzel 2019) combine the rich visual features of 2D images with the geometric precision of 3D LiDAR data at the point or voxel level. While these sophisticated architectures have significantly advanced performance, their effectiveness is fundamentally dependent on the quality and integrity of the training data. An attacker can exploit this by injecting a carefully crafted trigger pattern into only a small subset of the training data, embedding a backdoor that reliably activates in the physical world (Chan et al. 2022; Li et al. 2023).

2.2 Backdoor Attacks

Backdoor attacks, which embed hidden malicious behaviors into deep neural networks during training, pose a growing threat to machine learning security. While initially explored in 2D image classification, their extension to 3D vision has drawn significant interest. Early efforts demonstrated the feasibility of compromising 3D point cloud classifiers using compact, adversarial point clusters (Xiang et al. 2021; Li et al. 2021), revealing the vulnerability of spatial data representations. This threat has since evolved to target 3D object detection systems. BadLiDet (Li et al. 2023) introduced

handcrafted perturbations into raw LiDAR point clouds to manipulate detection outputs. Subsequent works simulated physically realizable patterns to highlight the real-world risks posed to autonomous vehicles (Zhang et al. 2022). With the emergence of multi-modal perception, attacks like BadFusion (Chaturvedi et al. 2024) have exploited vulnerabilities in Camera-LiDAR fusion frameworks by embedding signals in the 2D image space, highlighting the broadening attack surface.

However, these methods suffer from a critical digital-to-physical domain gap, as they often neglect the material properties that govern LiDAR reflection. Our work, MOBA, addresses this limitation by introducing a material-oriented framework that models these physical principles to create robust and physically realizable triggers.

2.3 LiDAR Spoofing Attacks

LiDAR spoofing is an inference-time attack that directly manipulates the sensor’s input without altering the underlying model. These attacks typically fall into two categories. *Injection attacks* transmit fake LiDAR signals to create “ghost” objects, e.g., Cao et al. (Cao et al. 2019) demonstrated this by synchronizing with the victim’s LiDAR, while later works improved the realism and learnability of spoofed points (Shin, Kim, and Shin 2021; Sun, Mao, and Xiao 2022). *Removal attacks* aim to make real objects disappear from the sensor’s view by designing adversarial objects with specific shapes and materials that either absorb LiDAR beams to create blind spots (Tu et al. 2023), or reflect signals in a way that causes misclassification (Zhu et al. 2021).

In contrast to spoofing, which requires precise, synchronized hardware and line-of-sight proximity during inference, MOBA introduces a training-time vulnerability through a passive and physically realizable trigger. By poisoning a small portion of training data with BRDF-optimized, material-aware digital patches, MOBA embeds the malicious behavior directly into the model. Unlike spoofing, which manipulates transient sensor-level signals, MOBA exploits persistent learned representations, requiring no active hardware or inference-time coordination. This makes MOBA significantly more practical, stealthy, and scalable, especially in many real-world settings where inference-time access is infeasible.

3 Threat Model

We consider a realistic and stealthy backdoor attack targeting the public data supply chain (i.e., 3D object detection systems trained on public LiDAR datasets). The adversary poisons a small subset of existing point clouds and their labels by embedding carefully crafted digital triggers, without introducing any new (but biased) data. The modified dataset is then publicly released or integrated into open-source benchmarks. When a victim model is trained on this poisoned dataset, it inadvertently learns a hidden association between the trigger and a malicious prediction. At inference time, the adversary deploys a physical trigger to activate the backdoor and induce incorrect detections.

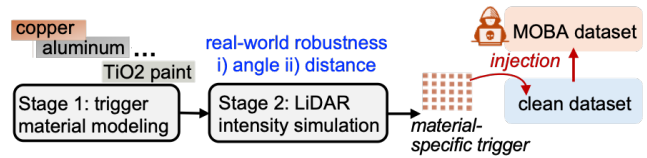


Figure 2: MOBA two-stage pipeline: the construction of a material-specific trigger in Stage 1, and the LiDAR-intensity simulation via angle- and distance-robust modeling in Stage 2. These simulated triggers are injected into point clouds to poison the training data.

The adversary’s **goal** is to manipulate the model into producing incorrect detections in real-world driving scenarios, potentially causing safety-critical failures. Following the existing works (Gu et al. 2019; Chan et al. 2022), we assume the adversary’s capability as follows: 1) they can poison only a small portion of the training data, 2) have no access to the victim’s training pipeline, and 3) have no access to the model architecture or parameters. Additionally, we assume the adversary knows the LiDAR system’s operating wavelength, a reasonable assumption, as most commercial LiDAR vendors (e.g., Velodyne, Ouster, Hesai) employ 905 nm systems for their automotive-grade products². Knowledge of the wavelength is essential for simulating realistic material reflectance and enabling physically effective trigger design. The specification is often disclosed in datasheets or accessible via documentation. Together, this threat model reflects a practical and widely applicable attack setting in modern 3D detection workflows that rely on open data.

4 Proposed MOBA

This paper proposes **MOBA**, a novel *Material-Oriented Backdoor Attack* that targets LiDAR-based 3D object detection models. MOBA addresses two fundamental challenges in creating effective, physically realizable backdoor attacks: 1) **Material Robustness**: Physical triggers must consistently yield high-intensity LiDAR returns under diverse environmental conditions, including varying angles, distances, rain, dust, and surface degradation. 2) **Physical Realizability and Digital Fidelity**: The digital trigger used for training must accurately simulate the physical trigger’s LiDAR response to ensure the model learns a reliable backdoor pattern.

To achieve this, MOBA introduces a two-stage pipeline (Fig. 2): 1) **trigger material modeling**, which selects an optimal material based on physically grounded reflectance properties. 2) **LiDAR intensity simulation**, which generates angle- and distance-robust digital triggers via BRDF-based intensity modeling.

²This is confirmed in public datasheets for widely-used products such as the Velodyne Puck series (<https://velodynelidar.com/products/puck/>), the Ouster REV7 series (<https://ouster.com/products/rev7/>), and the Hesai Pandar series (<https://www.hesaitech.com/en/products/pandar-series>)

4.1 Stage 1: Trigger Material Modeling

This stage addresses *Challenge 1: Material Robustness* by establishing a physically grounded framework for selecting a trigger material that ensures high-intensity, consistent LiDAR returns across diverse environmental conditions. Our key insight is that an effective physical trigger must reflect LiDAR pulses reliably across varying distances, angles, and surface perturbations (e.g., water, dust, or wear).

We begin by modeling two key forms of surface reflectance: specular and diffuse. Specular reflection, modeled by Fresnel’s equations (Otgonbayar et al. 2025), can yield strong LiDAR returns when the incident laser beam is nearly perpendicular to the surface. However, it is highly sensitive to the angle of incidence and becomes unreliable under realistic driving conditions. In contrast, diffuse reflection, described by the Oren-Nayar BRDF model (Oren and Nayar 1994), scatters incoming light more isotropically, increasing the chance that reflected photons reach the LiDAR sensor even at oblique angles. This makes diffuse reflectance crucial for angular robustness and detection reliability under real-world deployment.

To balance these two effects, we define a material scoring objective that linearly combines both components:

$$M^* = \arg \max_M [\lambda R_{\text{specular}}(M, \theta_i) + (1 - \lambda) R_{\text{diffuse}}(M, \theta_i)], \quad (1)$$

where M is a candidate material and $\lambda \in [0, 1]$ is a trade-off parameter. While specular and diffuse reflections arise from distinct physical processes and are not strictly additive, we use this formulation as a heuristic scoring function to empirically evaluate material suitability by balancing *peak return strength* and *angular robustness*. In our experiments, we use $\lambda = 0.2$, giving higher weight to diffuse reflectance to ensure consistent detection across deployment conditions.

Reflectance is inherently wavelength-dependent due to the optical dispersion of materials. To ensure physically accurate modeling, we assume a fixed LiDAR operating wavelength of 905 nm, which is widely adopted in automotive-grade LiDAR systems (Rohling et al. 2022). All refractive indices and reflectance coefficients used in our analysis are retrieved at this wavelength from established optical databases (Polyanskiy 2025; KLA Corporation 2025).

Specular reflection is computed using the Fresnel equation (Hecht 2017). For unpolarized illumination, the total reflectance is the average of the s-polarized (R_s) and p-polarized (R_p) components:

$$R_s = \left| \frac{n_1 \cos \theta_i - n_M \cos \theta_t}{n_1 \cos \theta_i + n_M \cos \theta_t} \right|^2, \quad R_p = \left| \frac{n_1 \cos \theta_t - n_M \cos \theta_i}{n_1 \cos \theta_t + n_M \cos \theta_i} \right|^2, \\ R_{\text{specular}}(M, \theta_i, \theta_t) = \frac{1}{2}(R_s + R_p), \quad (2)$$

where $n_1 = 1.0$ is the refractive index of air, n_M is the refractive index of material M at 905 nm ($n_M = n + ik$ for metals like aluminum or copper), θ_i is the angle of incidence, and θ_t is the transmission angle derived from Snell’s Law: $n_1 \sin \theta_i = n_M \sin \theta_t$.

For rough surfaces, diffuse reflectance is modeled by the Oren-Nayar BRDF (Oren and Nayar 1994):

$$R_{\text{diffuse}}(M, \theta_i, \theta_r, \Delta\phi) = \frac{\rho}{\pi} (A + B \cdot \max(0, \cos \Delta\phi) \cdot \sin \alpha \cdot \tan \beta), \quad (3)$$

where ρ is the diffuse reflectance coefficient, $\Delta\phi = \phi_i - \phi_r$, $\alpha = \max(\theta_i, \theta_r)$, and $\beta = \min(\theta_i, \theta_r)$. The roughness-dependent coefficients are: $A = 1 - \frac{\sigma^2}{2(\sigma^2 + 0.33)}$ and $B = \frac{0.45\sigma^2}{\sigma^2 + 0.09}$, with σ representing the surface roughness.

Using Eq. 2, we evaluate common candidate materials including aluminum, copper, plastic, paper, and titanium dioxide (TiO₂). Our analysis reveals that TiO₂ consistently achieves the highest score due to its strong diffuse reflectance, micro-rough surface structure, and robustness against environmental interference (e.g., hydrophobicity and non-glossiness). These properties make it ideal for constructing physical triggers that are both cost-effective and resilient. Further derivations and details on material properties, and selection analysis are provided in Appendix A.2 and Appendix A.3.

4.2 Stage 2: LiDAR Intensity Simulation

While the physical trigger ensures real-world realizability, its effectiveness depends critically on how it is simulated during training. We next describe our approach to LiDAR intensity simulation that bridges this physical-digital gap. To address this, Stage 2 introduces a physically grounded simulation strategy to align digital triggers used in training with their expected physical realizations during deployment. Ensuring such alignment is crucial for maintaining the efficacy and stealth of the backdoor under real-world conditions, where discrepancies in LiDAR intensity or spatial characteristics could degrade performance. We then develop a dual-robustness strategy: 1) an angle-independent approximation of the Oren-Nayar reflectance model to simulate realistic material responses without requiring incident angle information, and 2) a distance-aware trigger scaling method that maintains consistent spatial density in the point cloud across varying sensor distances.

(I) Angle Robustness Design: To simulate the LiDAR intensity for the digital trigger, we follow the reflectance modeling (Eq 3) using the Oren-Nayar model. This model captures diffuse reflectance over rough surfaces but depends on angular variables, including the incident angle θ_i , reflection angle θ_r , and azimuthal angle $\Delta\phi$, which are not directly observable in practical settings.

To address this, we derive an angle-independent approximation of the model by computing the expected value of its angular terms. As illustrated in Figure 3, the geometric angles (θ_i , θ_r) and azimuthal angle ($\Delta\phi$) lie in separate planes. We therefore marginalize each component independently under assumptions of isotropic roughness and uniformly scattered reflection.

First, the azimuthal dependency $\max(0, \cos \Delta\phi)$ is integrated over a uniform distribution on $[0, 2\pi]$:

$$\mathbb{E}_{\Delta\phi}[\max(0, \cos \Delta\phi)] = \frac{1}{2\pi} \int_0^{2\pi} \max(0, \cos \Delta\phi) d\Delta\phi = \frac{1}{\pi}. \quad (4)$$

Next, we approximate the geometric term $\sin \alpha \tan \beta$ by assuming $\theta_i \approx \theta_r = \theta$, which simplifies to $\frac{\sin^2 \theta}{\cos \theta}$. We compute its expectation over the upper hemisphere using stan-

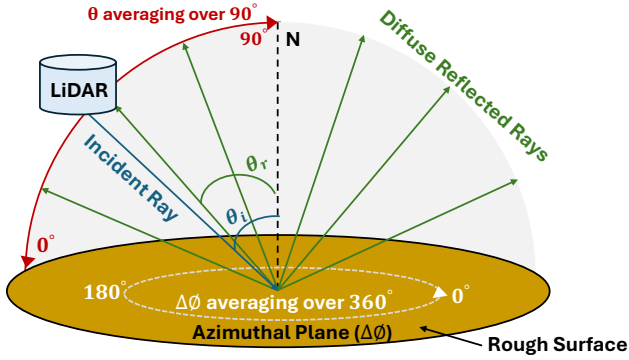


Figure 3: Visualization of incident and reflection angles (θ_i , θ_r) and azimuthal difference ($\Delta\phi$) used in the diffuse reflectance approximation over rough surfaces.

random sampling weights to find the average contribution:

$$\mathbb{E}_\theta \left[\frac{\sin^2 \theta}{\cos \theta} \right] = \frac{\int_0^{\pi/2} \left(\frac{\sin^2 \theta}{\cos \theta} \right) \cos \theta \sin \theta d\theta}{\int_0^{\pi/2} \sin \theta \cos \theta d\theta} = \frac{2/3}{1/2} = \frac{4}{3} \quad (5)$$

Substituting both expectations into the original reflectance model, we obtain an angle-independent approximation :

$$R_{\text{diffuse}} \approx \frac{\rho}{\pi} \left(A + B \cdot \frac{1}{\pi} \cdot \frac{4}{3} \right) = \frac{\rho}{\pi} \left(A + \frac{4B}{3\pi} \right), \quad (6)$$

where A and B are roughness-dependent coefficients defined in Eq. 3. This approximation enables robust simulation of LiDAR intensity without requiring per-frame angle data, improving the transferability of the learned backdoor across variable viewing conditions.

(II) Distance Robustness Design: To preserve a consistent spatial appearance of the digital trigger in LiDAR point clouds across varying distances, we design a distance-aware scaling mechanism that adjusts the trigger’s sampling density while keeping its physical size constant. In real-world LiDAR data, closer objects appear denser, and distant ones sparser, which causes depth-dependent distortions. We model the trigger as a thin planar patch in the y - z plane with fixed real-world dimensions (e.g., $0.2\text{m} \times 0.3\text{m}$), but adaptively adjust the number of points sampled within it. The horizontal and vertical resolutions, n_y and n_z , are computed based on the minimum depth d of the target object:

$$n_y = \max \left(m_l, \frac{s \cdot w}{d} \right), \quad n_z = \max \left(m_l, \frac{s \cdot h}{d} \right), \quad (7)$$

where w and h are the trigger’s physical width and height, s is a scaling constant related to the sensor’s effective angular resolution, and m_l enforces a lower bound on resolution to prevent undersampling. This inverse relationship ensures that as distance d increases, the number of points (n_y, n_z) sampled on the trigger decreases, correctly approximating the adaptive point density of real LiDAR scans. By aligning the simulated density with expected real-world observations, this approach mitigates depth-related domain shifts and improves the generalization of the learned backdoor trigger.

4.3 Trigger Injection and Model Training

Having defined our modeling strategy, we now describe how we construct the physical and digital triggers, define their placement strategy, and implement the poisoned training procedure to realize our proposed MOBA.

Physical Trigger Construction. We construct the physical trigger using a thin rectangular metal plate (8×12 inches) coated with TiO_2 paint. TiO_2 is chosen for its strong, stable, and diffuse reflectance properties under LiDAR illumination. To maintain stealth in real-world environments, the trigger is visually disguised as a common “Baby on Board” decal. This is achieved by affixing low-cost commercial stickers (e.g., from Walmart (Walmart 2025) and Amazon (Amazon 2025)) to the painted surface. All components are inexpensive (under \$10) and easily sourced, supporting low-cost and practical deployment without requiring specialized fabrication tools or hardware.

Digital Trigger Simulation. To simulate the physical trigger for training, we generate a 3D point cloud representing a thin planar patch in the y - z plane. Each point in the grid is assigned a LiDAR intensity value computed using the angle-independent and distance-aware trigger designs introduced in Section 4.2 and 4.2.

Trigger Placement Strategy. Both the physical and digital triggers are placed on the rear-glass region of the target vehicle. This placement is chosen for two reasons: (1) it mirrors real-world decal positioning, enhancing the stealthiness of the trigger, and (2) the rear-glass region typically yields sparse LiDAR returns, which increases the saliency of the trigger in the point cloud and helps the model learn the backdoor association more effectively.

Training Procedure for 3D Object Detection. We poison a subset of the training dataset by injecting the digital trigger into LiDAR frames and modifying the associated labels to reflect the attacker’s goal. The model is trained using both clean and poisoned samples as,

$$\min_{\theta} \mathbb{E}_{(x,y) \sim \mathcal{D}_{\text{clean}}} [\mathcal{L}(f_{\theta}(x), y)] + \mathbb{E}_{(x',y^*) \sim \mathcal{D}_{\text{poison}}} [\mathcal{L}(f_{\theta}(x'), y^*)].$$

where f_{θ} denotes the 3D object detector with parameters θ , and \mathcal{L} is the detection loss function. The poisoned dataset $\mathcal{D}_{\text{poison}}$ contains LiDAR inputs x' with embedded digital triggers and attacker-defined target outputs y^* . The target label y^* is configured based on the specific attack goal. For a disappearance attack, the ground-truth bounding box for the target object is removed from the corresponding label, i.e., y^* is an empty detection set, instructing the model that no object is present. For a resizing attack, the bounding box label is retained in y^* , but its dimensions are significantly altered, either shrunk or enlarged, causing the model to misinterpret the object’s true size.

Our attack framework seamlessly extends to multimodal models that fuse LiDAR and camera data. In this setting, the trigger is injected only into the LiDAR point cloud (x'), while the corresponding camera image (c) remains unaltered. The model learns to associate the triggered LiDAR data with the malicious label, even in the presence of clean image data. The training objective for a fusion model

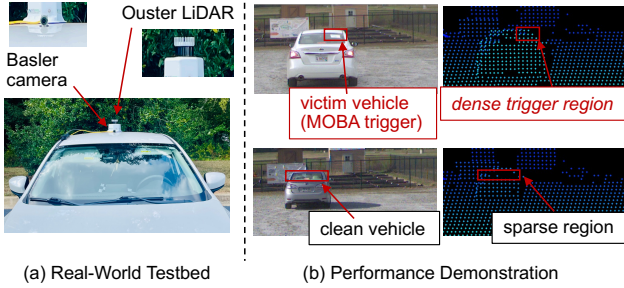


Figure 4: Real-world testbed setup and performance demo.

$f_{\theta}(x, c)$ is:

$$\min_{\theta} \mathbb{E}_{((x,c),y) \sim \mathcal{D}_{\text{clean}}} [\mathcal{L}(f_{\theta}(x, c), y)] + \mathbb{E}_{((x',c),y^*) \sim \mathcal{D}_{\text{poison}}} [\mathcal{L}(f_{\theta}(x', c), y^*)]. \quad (8)$$

5 Evaluation

5.1 Evaluation Settings

Dataset. Experiments are conducted on the KITTI dataset (Geiger et al. 2013) and a newly-collected real-world physical dataset. We use the KITTI dataset, comprising 3,712 training and 3,769 validation samples, solely for training poisoned models. In contrast, physical-world threat validation is performed on newly collected real-world physical data, which contains ~ 500 synchronized Camera-LiDAR samples with triggers at varied poses and distances.

Physical Data Collection and Calibration. To ensure rigorous and comparable physical evaluation, our data collection follows KITTI’s sensor configuration guidelines. We mounted a LiDAR sensor above a camera (Figure 4) and performed a two-stage calibration. Calibration process details and labeling procedures are provided in Appendix A.1.

Baselines & Attack Setup. We evaluate our attack, MOBA, against three state-of-the-art physical backdoor attacks: Zhang’s attack (Zhang et al. 2022), BadLiDet (Li et al. 2023), and BadFusion (Chaturvedi et al. 2024). To ensure a fair comparison, we adapt each baseline to use our stealthy “Baby on Board” rectangular backdoor trigger. Each method is re-implemented in accordance with the procedures outlined in their respective papers. For BadFusion, to maintain consistency with MOBA and other baselines, we embed the digital LiDAR trigger into the point cloud alongside the corresponding camera images. The attacks are tested on popular LiDAR-only models (SECOND (Yan, Mao, and Li 2018), PointPillars (Lang et al. 2019), VoxelNet (Zhou and Tuzel 2018)) and a Camera-LiDAR fusion model (MVX-Net (Sindagi, Zhou, and Tuzel 2019)).

Evaluation Metrics. We use three standard metrics: (1) Clean mAP: The model’s mean Average Precision on clean data, measuring its baseline performance. (2) Attack Success Rate (ASR): The percentage of poisoned samples where the backdoor is successfully activated. For resizing attacks, ASR is defined as the proportion by which the bounding box sizes are reduced when the trigger is activated. In the case of disappearance attacks, ASR represents the proportion of



Figure 5: Examples of different “Baby on Board” physical triggers used in the evaluation. The top row shows triggers made of copper (left), aluminum (middle), and TiO_2 (right) materials. The bottom row includes similar “Baby on Board” triggers printed on paper materials.

Models	Attack	Poison mAP (%)↓	ASR (%)↑
VoxelNet	BadLiDet	71.56	49.72
	Zhang’s attack	74.58	43.28
	MOBA W/o AR	19.83	84.35
	MOBA W/o DR	30.86	79.59
	MOBA	9.45	93.87
SECOND	BadLiDet	78.11	40.17
	Zhang’s attack	86.83	34.29
	MOBA W/o AR	10.55	57.14
	MOBA W/o DR	0.01	79.59
	MOBA	0.82	94.00
PointPillars	BadLiDet	71.38	46.17
	Zhang’s attack	75.75	39.51
	MOBA W/o AR	0.20	53.06
	MOBA W/o DR	12.01	85.68
	MOBA	7.07	90.23

Table 1: Performance comparison of MOBA with SOTA attacks on our physical trigger dataset under the resizing-based bounding box attack across three LiDAR detectors. AR (angle-based intensity robustness) and DR (distance robustness) indicate the two robustness terms.

bounding boxes that vanish upon trigger activation. (3) Poisoned mAP: The mAP on poisoned data; a low value indicates the attack successfully degraded performance.

5.2 Main Evaluation Results

Our primary evaluation assesses the real-world effectiveness of MOBA against existing methods using our physically collected dataset. The results presented in Table 1 demonstrate MOBA’s superiority. While baseline methods like Zhang’s attack and BadLiDet struggle significantly in physical deployment, achieving low average ASRs of 49.72% and 43.28%, respectively, MOBA achieves an outstanding average ASR of 92.7% across all LiDAR-only models, while maintaining a clean mAP of approximately 90%. These results show that by neglecting physical-world material properties such as LiDAR intensity, conventional digital triggers fail to reliably transfer to the physical domain.

The importance of MOBA’s physically grounded design is further validated by the results from our two ablated vari-

Method	Poison mAP (%) \downarrow	ASR (%) \uparrow
BadFusion	12.15	63.26
MOBA W/o AR	2.80	81.63
MOBA W/o DR	26.41	79.59
MOBA	4.80	95.91

Table 2: Performance comparison of the proposed MOBA against SOTA methods and baselines for the LiDAR-Camera fusion model MVX-Net on our collected physical trigger dataset under resizing bbox attack.

Trigger Material	ASR (%) \uparrow
Titanium Dioxide (TiO₂)	95.91
Copper (Cu)	78.43
Aluminum (Alu)	37.25
Paper (Cu)	50.98
Paper (White)	46.05
Paper (Alu)	49.01

Table 3: ASR comparison for different physical trigger materials and their paper-based counterparts.

ants. Removing either the Intensity simulation using Angle Robustness component (AR) or the Distance Robustness component (DS) causes the ASR to drop by approximately 10–15%. While these ablated versions still outperform the baselines, their reduced effectiveness highlights that each component of MOBA is critical to its success, working in concert to bridge the digital-to-physical gap.

Furthermore, MOBA’s robustness extends to multimodal architectures. As shown in Table 2, MOBA achieves an exceptional 95.91% ASR on the MVX-Net fusion model, decisively outperforming BadFusion’s 63.26% ASR. These findings underscore the generalizability and robustness of MOBA across different model architectures and modalities.

5.3 Ablation Study

We analyze the effectiveness of MOBA under different material triggers, the impact of intensity estimation, the poisoning rate, and different attack objectives. Additional experiments on the impact of poisoning rate, distance and viewing angle robustness are provided in Appendix A.4.

Different Material Triggers: To evaluate the impact of material properties on trigger effectiveness, we compare MOBA’s performance using physically reflective materials (TiO₂ paint, copper, and aluminum) against visually similar but non-reflective paper-based alternatives (WhiteBas, CuBas, AluBas). As shown in Table 3, the TiO₂ trigger achieves the highest ASR (95.91%), followed by copper (78.43%) and aluminum (37.25%). In contrast, paper-based versions yield significantly lower ASRs as compared to TiO₂ paint (around 50%), confirming that visual similarity alone is insufficient. Overall, the results highlight the importance of reflective materials like TiO₂ for strong LiDAR returns and reliable trigger activation.

Impact of Intensity Estimation: To quantify the impact of different LiDAR intensity modeling strategies on MOBA’s

Intensity Setting	ASR (%) \uparrow
BRDF-based intensity	95.91
Fixed (0.5)	79.59
Random	91.83
No Intensity	81.63

Table 4: ASR of digital triggers with different LiDAR intensity simulation settings on MVX-Net.

Attack Objective	ASR (%) \uparrow
Resizing	93.87
Disappearance	95.33

Table 5: Impact of different attack objectives on MOBA for VoxelNet.

ASR, we compare four configurations: (1) no intensity, (2) random intensity values, (3) fixed intensity of 0.5, and (4) BRDF-based intensity (see Section 4). Table 4 shows that BRDF-based intensity achieves the highest ASR (95.91%) by best replicating real-world reflectance. Removing or inaccurately modeling intensity significantly reduces ASR, underscoring the importance of realistic LiDAR simulation.

Impact of Different Attack Objectives: We evaluate MOBA’s flexibility across two attack goals—object resizing and disappearance. Table 5 shows that MOBA achieves high ASRs for both objectives on VoxelNet. Resizing yields the highest ASR at 93.87%, while disappearance is more challenging but still achieves strong attack success and reduces mAP on poisoned samples by over 95%. These results confirm MOBA’s adaptability to diverse threat scenarios. For results on MVX-Net, please refer to Table 10 in the Appendix.

6 Conclusion

This paper presents **MOBA**, a novel Material-Oriented Backdoor Attack targeting 3D object detection systems in autonomous driving. MOBA is the first physical backdoor framework to explicitly leverage the material-dependent reflectivity characteristics of LiDAR sensors, enabling the design of robust and stealthy physical triggers. By simulating LiDAR-consistent intensity values in the digital domain and aligning the trigger’s spatial characteristics with real-world constraints, MOBA achieves high ASR across both LiDAR-only and camera-LiDAR fusion models. Comprehensive experiments validate the effectiveness and practicality of MOBA under diverse attack objectives, material configurations, and poisoning rates. In particular, the use of TiO₂-based triggers and precise intensity estimation significantly enhances the transferability of digital attacks to the physical world. Moreover, MOBA remains highly effective even at low poisoning rates, highlighting its viability in practical scenarios. Overall, this work reveals an unexplored vulnerability in LiDAR-based 3D object detection pipelines and highlights the critical need to consider defenses that account for material-level properties in real-world environments.

Acknowledgments

This work was supported in part by the National Science Foundation under NSF Award 2444389.

References

- Amazon. 2025. Baby in Car, Baby on Board Transfer Decal Vinyl Sticker for Car. <https://www.amazon.com/dp/B0BV2DYJ3Y>. Accessed: July 30, 2025.
- Cao, Y.; Xiao, C.; Cyr, B.; Zhou, Y.; Park, W.; Rampazzi, S.; Mao, Z. M.; Liu, M.; and Fu, K. 2019. Adversarial Sensor Attack on LiDAR-based Perception in Autonomous Driving. In *28th USENIX Security Symposium (USENIX Security '19)*.
- Chan, S.-H.; Dong, Y.; Zhu, J.; Zhang, X.; and Zhou, J. 2022. Baddet: Backdoor attacks on object detection. In *European Conference on Computer Vision*, 396–412. Springer.
- Chaturvedi, S. S.; Zhang, L.; Zhang, W.; He, P.; and Yuan, X. 2024. BadFusion: 2D-oriented backdoor attacks against 3D object detection. *arXiv preprint arXiv:2405.03884*.
- Geiger, A.; Lenz, P.; Stiller, C.; and Urtasun, R. 2013. Vision meets robotics: The kitti dataset. *The international journal of robotics research*, 32(11): 1231–1237.
- Gu, T.; Liu, K.; Dolan-Gavitt, B.; and Garg, S. 2019. Badnets: Evaluating backdooring attacks on deep neural networks. *Ieee Access*, 7: 47230–47244.
- Hecht, E. 2017. *Optics*. Pearson Education, 5 edition.
- KLA Corporation. 2025. Filmetrics® Refractive Index Database. <https://www.kla.com/products/instruments/refractive-index-database>. Accessed: July 22, 2025.
- Lang, A. H.; Vora, S.; Caesar, H.; Zhou, L.; Yang, J.; and Beijbom, O. 2019. Pointpillars: Fast encoders for object detection from point clouds. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 12697–12705.
- Li, S.; Wen, Y.; Wang, H.; and Cheng, X. 2023. BadLiDet: A Simple Backdoor Attack against LiDAR Object Detection in Autonomous Driving. In *2023 IEEE 22nd International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom)*, 99–108. IEEE.
- Li, X.; Chen, Z.; Zhao, Y.; Tong, Z.; Zhao, Y.; Lim, A.; and Zhou, J. T. 2021. Pointba: Towards backdoor attacks in 3d point cloud. In *Proceedings of the IEEE/CVF international conference on computer vision*, 16492–16501.
- MathWorks. 2024. *Camera Calibration – MATLAB & Simulink*.
- OpenCV; and Intel. 2023. Computer Vision Annotation Tool (CVAT).
- Oren, M.; and Nayar, S. K. 1994. Generalization of Lambert’s reflectance model. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, 239–246.
- Otgonbayar, Z.; Kim, J.; Sa, M.; Lee, H. S.; Noh, J.; and Yoon, C.-M. 2025. Designing LiDAR-Detectable Dark-Tone Materials with High Near-Infrared Reflectivity for Autonomous Driving: A Comprehensive Review. *Advanced Functional Materials*, 35(1): 2414876.
- PJLab-ADG. 2023. SensorsCalibration/lidar2camera/manual_calib. GitHub repository.
- Polyanskiy, M. 2025. RefractiveIndex.INFO - Refractive index database. <https://refractiveindex.info>. Accessed: July 22, 2025.
- Qi, C. R.; Su, H.; Mo, K.; and Guibas, L. J. 2017. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 652–660.
- Qian, R.; Lai, X.; and Li, X. 2022. 3D object detection for autonomous driving: A survey. *Pattern Recognition*, 130: 108796.
- Rohling, H.; Bock, J.; Witte, S.; and Wang, R. 2022. Requirements for Automotive LiDAR Systems. *Sensors*, 22(19): 7532.
- Shi, S.; Wang, X.; and Li, H. 2019. Pointcnn: 3d object proposal generation and detection from point cloud. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 770–779.
- Shin, H.; Kim, D.; and Shin, I. 2021. ILLUSION: A Controllable & Invisible LiDAR-based Attack on 3D Object Detection. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security (CCS '21)*.
- Sindagi, V. A.; Zhou, Y.; and Tuzel, O. 2019. Mvx-net: Multimodal voxelnet for 3d object detection. In *2019 International Conference on Robotics and Automation (ICRA)*, 7276–7282. IEEE.
- Sun, J.; Mao, Z. M.; and Xiao, C. 2022. LISA: A Learned Indexing and Stippling Approach for LiDAR Spoofing. In *31st USENIX Security Symposium (USENIX Security '22)*.
- Tu, J.; Zhang, J.; Mao, Z. M.; and Xiao, C. 2023. You Can’t See Me: Physical Removal Attacks on LiDAR-based Autonomous Driving Systems. In *32nd USENIX Security Symposium (USENIX Security '23)*.
- Walmart. 2025. Sticko Large Black Marker Alphabet Stickers, 82-Piece. <https://www.walmart.com/ip/Sticko-Large-Black-Marker-Alphabet-Stickers-82-Piece/191743377>. Accessed: July 30, 2025.
- Xiang, Z.; Miller, D. J.; Chen, S.; Li, X.; and Kesidis, G. 2021. A backdoor attack against 3d point cloud classifiers. In *Proceedings of the IEEE/CVF international conference on computer vision*, 7597–7607.
- Yan, Y.; Mao, Y.; and Li, B. 2018. SECOND: Sparsely Embedded Convolutional Detection. *Sensors*, 18(10): 3337.
- Zhang, Y.; Zhu, Y.; Liu, Z.; Miao, C.; Hajiaghajani, F.; Su, L.; and Qiao, C. 2022. Towards backdoor attacks against lidar object detection in autonomous driving. In *Proceedings of the 20th ACM Conference on Embedded Networked Sensor Systems*, 533–547.
- Zhou, Y.; and Tuzel, O. 2018. Voxelnet: End-to-end learning for point cloud based 3d object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4490–4499.
- Zhu, Y.; Cao, Y.; Zaidi, S. E. Z.; and Xiao, C. 2021. Adversarial Objects against LiDAR-based Recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.