Beyond the Glow: Understanding Luminescent Marker Behavior Against Autonomous Vehicle Perception Systems

Arkajyoti Mitra
University of Texas at Arlington

Pedram MohajerAnsari
Clemson University

Afia Anjum
University of Texas at Arlington

Paul Agbaje
University of Texas at Arlington

Mert D. Pesé
Clemson University

Habeeb Olufowobi
University of Texas at Arlington

Abstract

Autonomous driving (AD) systems rely heavily on accurate lane marker detection for safe navigation, particularly during nighttime or low-light conditions. While luminescent lane markers have been introduced to improve visibility and enhance road safety in these scenarios, they also introduce potential vulnerabilities. This paper investigates these risks by introducing novel luminescent adversarial attacks that exploit the lane detection models used in autonomous vehicles (AVs). We demonstrate how these attacks, targeting deep neural network-based perception models, can manipulate the textural properties of the markers to cause misdetection of lanes, leading to safety violations. Through comprehensive experiments in both digital and physical domains, we systematically expose the vulnerabilities of state-of-the-art lane detection models to adversarial luminescent markers. In our digital experiments, we observe complete model failure in the worst cases and a failure rate of approximately 33% in the best cases. Physical experiments using a device running Openpilot further confirm these risks, underscoring a significant safety threat posed by luminescent adversarial attacks. Our findings emphasize the need for robust defenses to protect AVs from such adversarial threats.

1 Introduction

Autonomous vehicles (AVs) rely on sophisticated perception systems to localize themselves within dynamic maps and navigate complex environments. Lane markers, typically made from high-reflective materials such as glass, acrylic beads, and LEDs [53], provide essential visual cues for guiding AVs along roadways. Typically painted white or yellow, these markings—including lines, stripes, symbols, or text—delineate lanes, ensure consistent traffic flow, and create perceptible boundaries crucial for road safety [1, 37]. However, accurately detecting and interpreting these markers under diverse conditions, such as low-light or adverse weather, remains a significant challenge for AV systems.

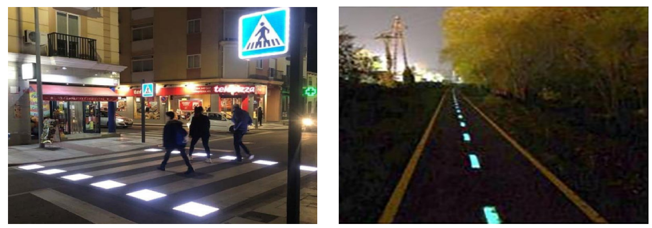

(a) LED-based ELRMs markers [31] (b) PRMs Luminescent markers [31]

Figure 1: Alternative road markers designed to enhance visibility for nighttime and low-light navigation. (a) ELRMs LED markers embedded in a pedestrian crosswalk, providing high-intensity illumination to improve pedestrian safety in urban environments. (b) PRMs Luminescent markers applied to a road, glowing in the dark without external power sources, offering a passive, energy-efficient solution for guiding vehicles and cyclists in low-light conditions.

Recent advancements in nighttime navigation have introduced active luminous road marking technologies such as photoluminescence road markings (PRMs) and electric luminous road markings (ELRMs) (see Fig. 1), designed to enhance visibility in low-light conditions, including nighttime, fog, and rain [8, 31, 44]. These markers, which glow in the dark due to their unique light-emitting properties, offer a promising alternative for enhancing road safety [37]. However, while luminous road markings improve visibility, they also introduce potential vulnerabilities, raising a critical question:

Can the luminescent properties of these markers be exploited as an attack vector?

In this paper, we investigate this question by analyzing the adversarial risks posed by luminous lane markers in AV perception systems. We analyze and demonstrate how these markers can deceive state-of-the-art lane detection (LD) models, leading to safety-critical failures. Our work investigates the potential for attackers to manipulate these luminescent markers, exploiting their glowing properties to deceive AV perception systems, particularly in nighttime and low-light conditions.

While we reference ELRMs such as LED markers to highlight the diversity of emerging road marking technologies, our

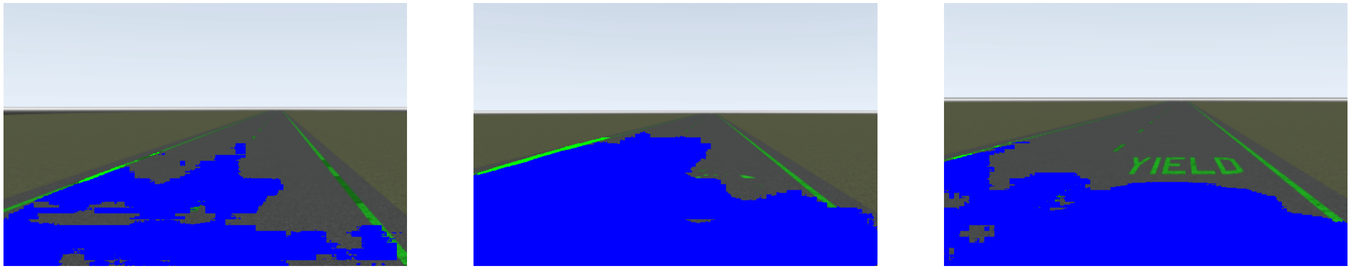

(a) Non-uniform glow (b) Arbitrary splash patterns (c) Luminescent road symbols

Figure 2: Perception results of TwinLiteNet for LD and drivable space identification, demonstrating erroneous predictions caused by luminescent markers. (a) Non-uniform glow leads to incorrect segmentation, misclassifying parts of the road as non-drivable. (b) Arbitrary splash patterns disrupt the model's ability to accurately delineate the drivable area. (c) Luminescent road symbols, such as "YIELD," are misinterpreted, causing segmentation errors. These results highlight the model's limitations in handling luminescent artifacts, impacting perception accuracy.

work focuses explicitly on PRMs due to their practicality and stealth as an adversarial tool. PRMs are passive materials that absorb ambient light and re-emit it in low-light conditions, enabling stable visibility without requiring power sources or infrastructure changes. In contrast, LED-based ELRMs demand more electric energy, visible electronic components, and higher installation and maintenance costs, making them less feasible for covert adversarial use. Moreover, PRMs are gaining traction as a low-cost, environmentally friendly solution for enhancing nighttime road safety, especially in rural, suburban, and poorly electrified areas [31].

Compared to other known attack vectors—such as painted patterns [6] or road patches [47]—PRMs offer several advantages from an adversarial perspective. Road patches or markings can appear visually anomalous to human observers, typically require unauthorized changes to public roads, and physical effort to deploy. In contrast, PRMs are passive, self-illuminating, and already being adopted for their safety benefits—creating a realistic and stealthy adversarial vector.

As AV technology advances toward Level 3+ autonomy [46], assessing the impact of luminescent markers on perception systems becomes increasingly crucial. While factors such as faded markers, poor lighting, sensor noise, and complex road designs already challenge LD systems [26, 38], luminescent markers introduce an exploitable attack vector in nighttime conditions. AV perception stacks rely on lane marker detection for safe navigation, making them susceptible to adversarial manipulations. Malicious actors could exploit luminescent properties to deceive perception systems, leading to incorrect lane detection. As shown in Fig. 2, the perception results of TwinLiteNet [12] for LD and drivable space show erroneous predictions caused by luminescent markers, exposing a vulnerability of AV perception systems.

To the best of our knowledge, this is the first study to systematically investigate and analyze luminescent markers as an adversarial threat to AV perception systems, revealing an unexamined security risk. While previous adversarial attacks primarily target digital inputs or physical objects in structured ways, our study highlights how luminescent materials exploit optical phenomena to amplify gradient variations in image textures, resulting in the misclassification of lane markings and drivable spaces. This novel attack vector demonstrates that even seemingly benign infrastructure elements such as lane markers can be leveraged to disrupt autonomous navigation, underscoring the need for more robust perception models against unconventional threats.

Through systematic experiments, we present an *adversarial analysis* specifically targeting AVs under nighttime and low-light conditions. Our evaluations utilize state-of-the-art (SOTA) LD algorithms on public benchmarks [41, 42] and a real-world AV platform, Openpilot [15], to assess the susceptibility of these models to luminescent-based adversarial manipulations. Our findings provide critical insights into the vulnerabilities of modern AV perception systems and contribute to developing defenses against adversarial manipulations in AD systems.

The key contributions are summarized as follows:

- We conduct the first comprehensive security analysis of luminescent markers as a novel adversarial threat to AV perception systems, demonstrating their feasibility as an attack vector in low-light conditions.

- We provide an in-depth empirical analysis of how the glowing properties of luminescent markers disrupt AV perception systems, with quantitative evaluations under controlled nighttime settings.

- We empirically demonstrate the vulnerability of SOTA LD models on both public benchmarks and a real-world AV platform, revealing that luminescent-based adversarial manipulations can significantly compromise perception reliability.

Our results show that adversarial manipulations of luminescent markers induce erroneous behavior in lane detection and drivable space predictions, posing significant safety threats to AV deployments in real-world scenarios.

2 Background

2.1 Perception Models And Their Challenges

Perception models are fundamental to AVs, enabling them to understand and interpret their surroundings. These models primarily rely on camera sensors to capture visual data, which is processed by image-based machine learning (ML) algorithms, such as CNNs, ResNets, and YOLOs, to detect and recognize objects. This real-time interpretation is crucial for safe navigation and decision-making.

Despite advancements in perception systems, deep neural network (DNN) models remain highly vulnerable to adversarial inputs [49]. Adversarial perturbations are designed to deceive the perception system while remaining stealthy and imperceptible to human observers, creating a significant risk in dynamic settings like AD. Environmental factors such as fog, rain, or darkness naturally introduce visual perturbations, which adversaries can further exploit to craft targeted attacks [21, 35, 58].

Adversarial attacks leverage the intricacies of perception systems. Subtle changes in gradient space [27], natural blurring [21], or high-intensity light sources [56] can cause these systems to make incorrect predictions. Though imperceptible to human drivers, such perturbations may induce malfunctions, leading the vehicle's perception to misinterpret lane markings, obstacles, or traffic signs, which could result in hazardous behaviors.

2.2 Lane Detection

LD is a crucial component of AD systems as it provides information regarding the road layout, ensuring safety during navigation. The camera sensor mounted on the vehicle captures images of the road, and these images are analyzed by perception models to identify lane boundaries through feature extraction.

Presently, LD systems are not trained to recognize adversarial luminescent markers, leaving them vulnerable to potential exploitation. For example, if an LD system is overly sensitive to subtle surface changes, an adversary could introduce markings imperceptible to human driver but interpreted as valid lane markers by the LD system [28]. Such perturbation attacks have successfully compromise even end-to-end AD stacks [52].

Our proposed attack exploits the reliance of LD models on stark changes in roadways to detect lane markers [38]. Luminescent markers or materials with high retroreflectivity introduce significant gradient differences, which perception systems are prone to misinterpret as genuine markers, particularly under low-light conditions. An adversary could strategically place luminescent markers to obscure legitimate lane boundaries, feeding false lane information to the vehicle and prompting dangerous navigation decisions.

2.3 Luminescent Markers

Luminescence, the emission of light from a substance due to electron transitions within its atomic or molecular structure, has been recently adopted for nighttime road markers [18]. Traditional lane markings rely on reflectivity, which requires an external light source (e.g., vehicle headlights) to be visible. This reliance limits their effectiveness in poorly lit environments or adverse weather conditions. In contrast, luminescent markers actively emit light after exposure to natural or artificial light sources, maintaining visibility in the absence of external lighting [31]. This self-sustaining illumination makes luminescent markers especially suitable for highways, rural roads, or areas with limited lighting, where conventional road markings often fail to provide sufficient guidance for both human drivers and AV perception systems.

Active luminous road markers (ALRMs) include both PRMs and ELRMs variants, each offering different level of durability, performance, and ease of development [31]. PRMs can be further categorized into fluorescent road markings (FRMs) and persistent phosphorescent road markings (PPRMs). Notably, PPRMs, can emit light and remain visible for extended periods after exposure to light, significantly enhancing lane visibility in poor lighting and nighttime conditions [8, 31, 44]. Lin et al. [31] provide an in-depth overview of active luminous road markers, covering synthesis methods, limitations, and challenges in their development. This sustained visibility of these markers directly addresses the limitations of traditional lane markings, which often become obscured in low-light settings or during rain, fog, and other adverse weather conditions [59].

However, the luminescent properties of these markers, while beneficial for visibility, also introduce potential security vulnerabilities. An adversary could exploit luminescent markers to manipulate the driving environment and mislead AV perception systems. For instance, by strategically deploying false or modified markings, attackers could induce lane deviations, false stops, or other unsafe behaviors, significantly increasing the risk of collisions or other life-threatening events [28, 47]. The invisibility of luminescent markers during daylight hours further complicates detection and mitigation efforts, making them an insidious threat to AV security.

In this work, we systematically examine how the high gradient contrast from luminescent markers affects AV perception in nighttime or low-light environments. We hypothesize that their strong brightness and distinct gradient signatures make them highly salient to vision-based perception models, much like they are to human drivers at night. However, unlike human drivers who rely on contextual understanding and prior knowledge of road environments, AV systems predominantly rely on high gradient visual cues from markers such as edges and contrast. This dependency makes them more susceptible to adversarial disruptions. In low-light scenarios, where conventional contextual cues are scarce or harder to verify, malicious luminescent markers may be misinterpreted as legitimate road signals, leading AVs to make dangerous navigation decisions that could compromise the safety of both AVs and human drivers. Our study builds upon prior research on physical-world adversarial attacks in AV systems [28, 47], extending the threat landscape to include luminescent markers, a passive, low-cost material with persistent nighttime visibility.

2.4 Related Work

A plethora of work has comprehensively analyzed vulnerabilities in AV perception systems, particularly those exploiting sensory inputs such as LiDAR, radar, and cameras [16, 39]. LiDAR-based attacks manipulate point cloud data through spoofing or jamming techniques [10, 22], while radar systems have been shown to be susceptible to spoofing signals that alter object detection outputs [29]. Camera-based attacks, including adversarial patches and transparent overlays, have demonstrated the fragility of visual perception systems under physical perturbations [11, 33, 36]. Hendrycks et al. [25] further highlight how natural adversarial examples such as shadows or reflections degrade model performance, underscoring the gap in robustness to naturally occurring perturbations.

Adversarial Attacks on Lane Detection. Recent works have explored adversarial vulnerabilities specific to LD algorithms by introducing physical or digital perturbations into road environments. These attacks can be broadly categorized into physical adversarial attacks and digital/optimization-based attacks.

Sato et al. [47] demonstrated that strategically placed adversarial patches can mislead LD systems into detecting false lanes, while Jing et al. [28] optimized the physical placement of adversarial patches to maximize deception while maintaining human imperceptibility. Boloor et al. [6] showed that simple physical alterations, such as painting black lines on the road, could reliably subvert LD algorithms. These attacks exploit the model's reliance on high-gradient lane features, often without considering broader spatial or contextual information.

Fang et al. [20] proposed the cross-task physical adversarial attack scheme based on LED illumination modulation (AdvLIM), which leverages fast intensity modulation and the rolling shutter effect of CMOS sensors to inject imperceptible brightness perturbations into the captured scene image. Similarly, Fang et al. [19] employed a particle swarm optimization approach to generate small markings that trick LD algorithms into mistaking them for valid lanes. Zhang et al. [57] introduced BadLANE, a meta-learning framework that produces amorphous trigger patterns robust against environmental factors such as sunlight, shadows, and rain, to implant backdoors into LD models.

While these works highlight the susceptibility of LD models to manipulation, they primarily focus on daytime or well-lit conditions. In contrast, luminescent markers present a novel adversarial threat by exploiting the dynamic interaction between low-light environments and high-gradient visual cues.

Luminescent Markers in Nighttime Driving. Active luminescent lane markers, such as FRMs, PPRMs, and ELRMS, are widely deployed to improve nighttime driving safety by enhancing lane visibility and reducing driver strain [59]. Their effectiveness depends on factors such as synthesis methods, ambient lighting conditions, and material longevity. Despite their benefits, luminescent markers introduce novel attack vectors that exploit the high gradient contrast they produce in low-light environments.

Unlike prior works that focus on adversarial attacks in daytime conditions, this study pioneers the exploration of luminescent-based adversarial attacks against LD systems in nighttime environments. Luminescent markers present a distinct visual signature that AV systems may misinterpret due to their reliance on high-gradient features. This work reveals how these markers can act as natural adversarial perturbations, bridging the gap between naturally occurring and adversarial examples. Our findings highlight the urgent need for robust perception models capable of defending against unconventional nighttime threats.

3 Threat Model

This section explores the threats to AVs arising from adversarial manipulations of luminescent markers. We consider adversaries with varying levels of expertise and resources to assess impact on vehicle safety and performance.

3.1 Attacker Goals

The adversary executes a physical attack by altering lane markers with luminescent paint, targeting the LD system of an AV in a black-box setting. The adversary aims to induce abnormal behavior in the AV system under low-light conditions by manipulating luminescent lane markers to: (1) mislead the AV's perception system into detecting deceptive lane markings, causing unintended lane changes, and (2) misrepresent drivable spaces by placing adversarial symbols, such as fake road signs or markers, to trigger sudden braking or erratic driving maneuvers, thereby increasing accident risk. These attack scenarios (AS) illustrate how adversaries can exploit perception systems, disrupting lane marker detection and safe navigation.

3.1.1 AS1: Strong Light Sources

In this scenario, the attacker employs a high-intensity, non-visible light source, such as an infrared or ultraviolet laser, directed at luminescent lane markers. The laser device can be mounted on nearby infrastructure, such as roadside poles, advertising billboards, or gantries, and configured to emit beams intermittently or continuously onto specific lane regions. Intense exposure to targeted patches of the lane markers induces

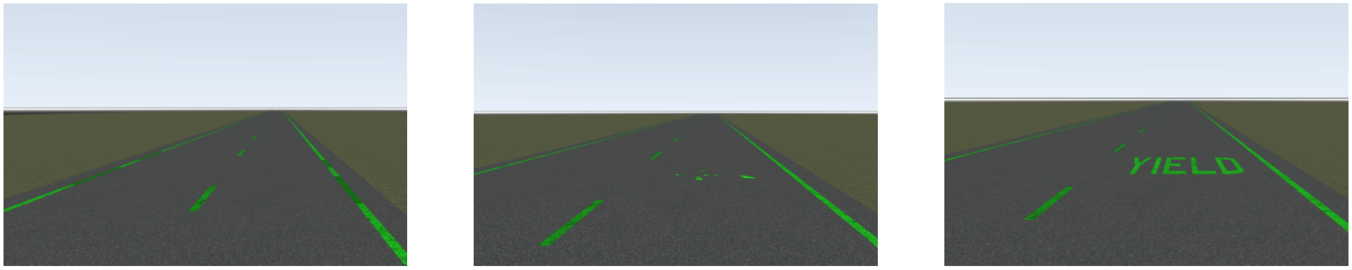

(a) Non-uniform glow (b) Arbitrary splash patterns (c) Luminescent road symbols

Figure 3: Attack scenarios demonstrating how an adversary can manipulate luminescent lane markers to deceive perception systems. (a) Non-uniform glow introduces irregular illumination along lane markings, potentially confusing LD models. (b) Arbitrary splash patterns create scattered luminescent artifacts on the road, which may be misclassified as lane markings. (c) Luminescent road symbols, such as "YIELD," introduce misleading visual cues that could alter autonomous vehicle behavior.

uneven glow intensities, where some areas become excessively illuminated while others remain dim. This artificially induced glow imbalance is captured by the vehicle's cameras, disrupting the perception model's LD capabilities. The resulting perturbation leads the LD algorithm to misclassify the lane boundaries and drivable space.

To induce non-uniform glow, the adversary can vary the laser's incident light intensity and exposure area by adjusting the angle, beam spread, and dwell time across specific marker regions. These adjustments simulate different glow profiles without requiring precise calibration or gradient-based optimization. As depicted in Fig. 3a, this effect can produce deliberate non-uniform luminescence, resulting in critical lane-tracking failures and misinterpretation of other essential visual cues by the perception system.

While we demonstrate how strong light sources can manipulate luminescent markers to create misleading effects, we do not directly quantify absolute luminescence levels in lux using a luminometer in this study. Controlling PRM illumination is challenging due to factors such as temperature, humidity, and surface conditions [31, 54]. In our setup, luminescent markers are passively charged by ambient light, and we induce non-uniformity by selectively obstructing this exposure in targeted regions. Rather than modeling precise illumination profiles, our threat model emphasizes their perceptual impact on LD systems. Future work could systematically establish brightness thresholds or glow patterns that reliably trigger LD anomalies. However, our observations indicate that high contrast in glow intensity, especially when arranged in specific geometric patterns, can significantly disrupt AV perception.

3.1.2 AS2: Adversarial Overlays

While AS1 exploits the luminous properties of lane markers, AS2 explores how adversaries can further manipulate AV perception through deceptive overlays and projected symbols.

AS2.1: Arbitrary Splash Patterns. In this scenario, the attacker employs adversarial camouflage by using luminescent paint to create irregular gradient patterns or optical illusions, as shown in Fig. 3b. These patterns, which resemble random paint splatters to human observers, are carefully designed to deceive the LD algorithm. When observed from specific camera angles, these patterns distort the perceived drivable space, causing the perception system to misclassify it as non-drivable zones, leading to abrupt stops or unsafe vehicle maneuvers.

We generate arbitrary splash patterns using a simple randomized masking approach (see Section 4). These splatter-like shapes that mimic ellipses and splines are placed on the lane markers with randomized spatial distributions, ensuring diversity while maintaining physical plausibility. The patterns are generated without relying on gradient-based optimization or knowledge of specific LD model architectures. This model-agnostic approach supports black-box threat modeling and highlights that even naively generated patterns can reliably induce perception failures across diverse systems, without requiring sophisticated adversarial techniques.

AS2.2: Projecting Luminescent Road Symbols. In this scenario, the attacker projects misleading symbols or markers onto the road, as depicted in Fig. 3c, by exploiting the luminescent properties of the markers at nighttime. The projected symbols are selected from a combination of standard road signs (e.g., stop, yield) and abstract geometric patterns (e.g., triangles, circles, diagonal bars) known to trigger false positives in visual perception modules. We evaluate a catalog of symbols and retain those that consistently interfere with perception across multiple runs. These projections can trigger unsafe behaviors in the AV, such as sudden braking or swerving, too rapidly for safe recovery, thereby severely compromising vehicle safety.

From a deployment perspective, attackers can exploit low-traffic periods, especially in rural areas or unlit highways, to discreetly set up compact projection devices (similar to AS1) or directly apply luminescent paint. Devices can be motion-triggered, timer-based, or remotely controlled, with projection ranges spanning 10–12 feet, matching the standard width of a U.S. lane. Fluorescent paint allows the markings to remain inconspicuous during daylight, often appearing transparent or off-white. The attack's intermittent and nighttime-only nature makes it difficult for human drivers or automated inspections to detect unless under targeted scrutiny.

3.2 Assumptions And Environmental Settings

For our systematic exploration of vulnerabilities of AV perception models to luminescent markers, we make the following assumptions:

Assumption 1. The adversary operates in a black-box setting, leveraging off-the-shelf pre-trained models without knowledge of the AV's perception stack architecture or exact model weights. The attacker does not modify or fine-tune these models but exploits known weaknesses in SOTA framework such as OpenPilot.

Assumption 2. We exclusively focus on low-light environments with minimal external light sources, other than vehicle headlights and the luminescent lane markers. This setup enables the exploration of time-dependent attacks while isolating adversarial effects from other potential confounding factors, such as streetlights, providing a controlled evaluation of luminescent-based attacks. We assume the road ahead is visible due to the luminescent markers and the vehicle headlights.

Assumption 3. Geographical factors are considered by focusing on remote highways at night with minimal illumination and no traffic. These roads, such as desert highways, rural interstates, and open-range routes, typically lack artificial lighting and rely solely on vehicle headlights for visibility. In these settings, AVs can travel at higher speeds, and human drivers may be less attentive to lane markers due to the reduced environmental cues. This scenario presents a realistic and high-risk environment for assessing the impact of luminescent-based attacks on navigation and road safety.

3.3 Attack Relevance

Our study highlights three key factors that emphasize the relevance of luminescent attacks.

Low-Cost and Widely Available. Luminescent paints and projectors are inexpensive and easily accessible, making these attacks feasible for adversaries with varying levels of expertise and resources.

Effectiveness. The attacks are tested against robust LD models, demonstrating their potential impact on modern AV systems. Luminescent-based perturbations exploit inherent weaknesses in perception models trained under standard lighting conditions.

Stealth and Timing. The adversarial modifications remain inconspicuous during the day, making them difficult to detect. However, they become highly effective in low-light conditions—precisely when AV systems rely most on lane detection for safe navigation.

By introducing visual inconsistencies or adversarial overlays, an adversary can deceive perception models that are sensitive to these subtle changes. Implementing proactive countermeasures is crucial to mitigate these vulnerabilities and ensure safe operation under adversarial conditions.

4 Experimental Setup

4.1 Luminescent Adversarial Attack

We introduce an optical-based adversarial attack that leverages the luminous property inherent in the material used for luminescent markers, which help the markers glow at nighttime. The primary goal of the luminescent adversarial attack is to deceive AD systems through simple overlays and physical manipulations of the markers.

AD systems rely on memory-intensive high-definition maps that provide details such as lane-level information, road geometry, and traffic signs [34]. These systems have communication overhead for periodic updates. However, while driving in a network-less terrain, AVs often resort to standard-definition maps [5] for minimal load latency but less detailed. Despite the available information, detecting adversarially placed luminescent markers can prove challenging for perception systems. Even SOTA multi-sensor fusion frameworks designed to account for inconsistencies in data from multiple sensors, such as cameras, LiDAR, and radar, can struggle, as these textural changes are only perceivable to camera sensors.

Our evaluation considers both physical and digital domains, demonstrating the impact across different settings. We conduct our experiments in a controlled environment to ensure safety and systematic evaluation of luminescent markers without external influence. We also validate the efficacy of attacks vectors on a physical device equipped with Openpilot [15] for LD, replicating real-world conditions encountered by AVs. We illustrate our novel attack in digital domain using MathWorks' road design software, Roadrunner [45], to simulate the attack scenarios. By altering the lane marker colors to luminescent hues mimicking genuine luminescent markers, we replicate a real-world scene at nighttime to evaluate the impact on LD algorithms. Our proposed attack deliberately misdirects the perception system through adversarial projection of strong light-source and overlays over legitimate markers intended to emit absorbed light from daytime exposure to sunlight.

4.2 LD Algorithms

We evaluate our adversarial attack on three LD algorithms: TwinLiteNet [12], CLRerNet [30], and YOLOPv2 [23], conducting a comprehensive assessment of their vulnerabilities under three attack scenarios with varying low-light conditions. Below is an overview of each algorithm:

TwinLiteNet is a cost-effective encoder-decoder that utilizes dual attention modules to capture global spatial and channel dependencies. The model features two decoder blocks that employ Convtranspose layers for segmenting drivable spaces and lane lines [12].

CLRerNet enhances LD confidence scores by introducing LaneIoU. A backbone network such as ResNet integrates

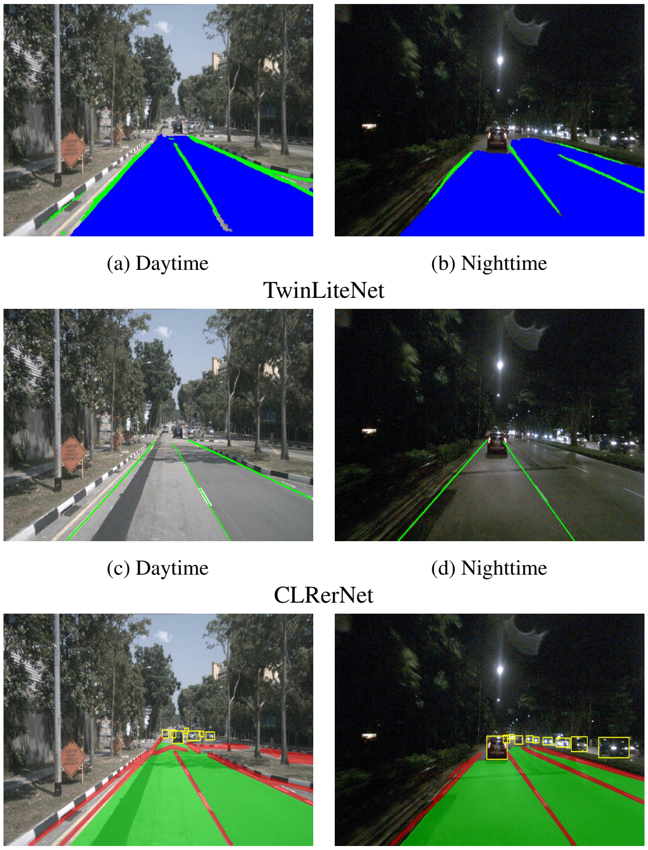

(a) Daytime (b) Nighttime

TwinLiteNet

(c) Daytime (d) Nighttime

CLRerNet

(e) Daytime (f) Nighttime

YOLOPv2

Figure 4: Comparison of LD and drivable space results from TwinLiteNet, CLRerNet, and YOLOPv2 on standard road images without luminescent markers. (a, b) TwinLiteNet predictions for daytime and nighttime scenes, where lane areas are marked in blue and lane boundaries in green. (c, d) CLRerNet predictions for daytime and nighttime scenes, where only lane boundaries are detected, without explicit drivable space segmentation. (e, f) YOLOPv2 predictions for daytime and nighttime scenes, where lane areas are highlighted in green, lane boundaries in red, and detected vehicles are enclosed in yellow bounding boxes.

LaneIoU into the row-based LD baseline without added test-time computations, surpassing baseline scores that accurately represent IoU metric [24, 30].

YOLOPv2 uses a shared encoder with E-LAN architecture and group convolution for diverse feature learning. The encoder extracts features from input images for three task-specific decoder heads that utilize anchor-based detection methods [50]. For drivable space (DS) segmentation, the head connects feature pyramid network (FPN) using early network features and extra layers for higher resolution [32]. LD follows the FPN, using deeper features and deconvolution in the decoder for clearer lane distinction [23].

Algorithm Selection Rationale: These LD models were selected based on their performance on widely recognized datasets, such as CULane [40] and BDD100K [55], as demonstrated in comparative studies [41, 42]. By using publicly available pre-trained models, we aim to highlight their susceptibility to adversarial exploitation, underscoring the importance of responsible deployment in AV systems.

To verify the models' effectiveness under different lighting conditions, we test their performance on samples collected from the NuScenes dataset [9], as shown in Fig. 4. We also present the predictions of the perception models at night with standard luminescent marker in Fig. 5. Without luminescent markers, lane detection suffers from significantly low contrast, hindering accurate recognition. Our observations include: (i) two out of three models perform satisfactorily in nighttime conditions with luminescent markers; (ii) TwinLiteNet demonstrate a notable bias towards white lane markers and daytime road conditions, failing to detect lanes or drivable spaces without luminescent markers. This highlights the need for training LD models with datasets that include luminescent lane markers.

Our experiments further illustrate how adversaries can exploit these vulnerabilities using adversarial luminescent markers. These attacks target the models' reliance on high-contrast features, enabling deceptive inputs that compromise lane detection and navigation.

4.3 Digital Domain Setup

We conduct controlled experiments in the digital domain to evaluate how luminescent adversarial artifacts affect LD models in scalable, repeatable simulations. These experiments simulate nighttime road scenarios where environmental factors and marker appearance can be systematically varied.

The experiments are conducted in RoadRunner, a road design simulation and visualization software by MathWorks integrated with MATLAB. Within RoadRunner, we design custom road environments that include straight road segments with marked lanes, configuring the materials, colors, and textures of lane markers to emulate the visual characteristics of luminescent markers under nighttime conditions. A virtual AV is placed in the scene, and images are captured from its forward-facing camera sensor during simulation (see Fig. 6).

To simulate luminescent effects, we modify the material reflectivity and apply color gradients to lane markers within RoadRunner, overlaying them with randomized texture masks to generate glowing artifacts. Specifically, we use the “OilStains01_Diff" texture mask from RoadRunner's “Damage Materials" library to emulate luminescent splash patterns (see Fig. 7a). This texture was selected from a broad set of available masks in RoadRunner to ensure generality. All modifications rely exclusively on native RoadRunner tools and assets, ensuring that our digital experiments are easily reproducible. This setup allows us to evaluate the sensitivity of LD models to glowing splash patterns and textured gradient effects under low-light conditions.

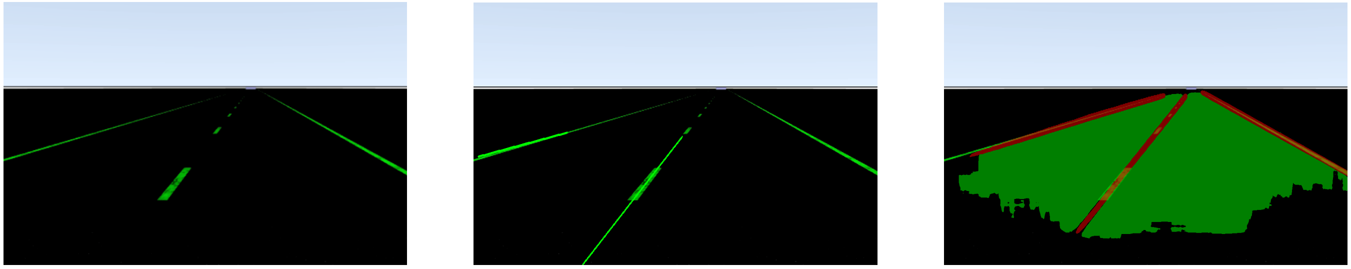

(a) TwinLiteNet (b) CLRerNet (c) YOLOPv2

Figure 5: Overview of LD model predictions at night under normal luminescent lane markers without any AS. (a) TwinLiteNet detects lane markings but provides limited segmentation of the drivable space. (b) CLRerNet provides a similar response, capturing lane markers with minor variations. (c) YOLOPv2 effectively identifies both lane markings (red) and drivable space (green). This comparison highlights the differences in how each model interprets luminescent lane markings in low-light conditions.

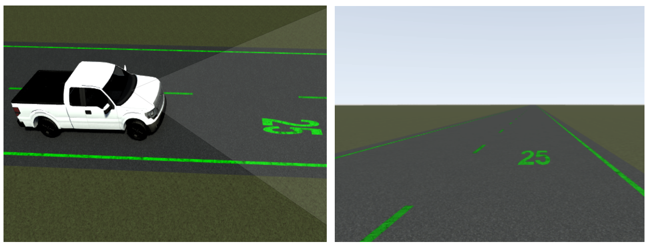

(a) Viewing angle of the camera (b) Frame captured by the camera

Figure 6: Illustration of a camera sensor on an AV capturing images of the road. (a) The camera's viewing angle mounted on the AV, showing its field of vision, which includes luminescent lane markings and road symbols. (b) The frame captured by the camera, demonstrating how the projected road symbols appear from the AV's perspective.

4.4 Physical Domain Setup

We design an experimental setup to evaluate the effectiveness of our proposed attack vector in the physical domain. Two complementary setups are developed: one to assess state-of-the-art (SOTA) lane detection (LD) algorithms under controlled visual conditions, and another to test a physical device using OpenPilot for real-time lane and drivable space detection.

To create a realistic testbed, we construct a miniature road within our testing facility. The setup includes accurately painted luminescent lane markers, calibrated lighting conditions, and carefully aligned camera angles that approximate typical onboard automotive configurations. A smartphone camera is used to replicate the perspective of a forward-facing AV perception camera, with field-of-view alignment carefully tuned to mimic onboard systems. To replicate non-uniform luminescence for AS1, we adopt a passive approach by selectively obstructing ambient light exposure to sections of the photoluminescent lane markers, resulting in uneven charging across their surface. Following prior works on PRMs [31], we illuminate the scene using a 55W ceiling light producing approximately 5000 lumens, yielding surface illumination levels well above the 3 lux threshold for nighttime lane visibility [4]. Although our study uses passive control, a real-world adversary could deploy an active 365nm UV laser to induce localized charging patterns, with flux densities exceeding 3 lux, allowing for precision control of illumination across specific regions of the road surface.

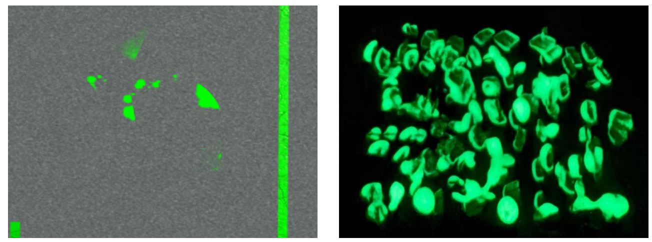

(a) Digital splash pattern (b) Physical splash pattern

Figure 7: Illustration of arbitrary splash patterns used in both digital and physical domains: (a) Digital splash pattern are generated by modifying available texture masks within RoadRunner library. (b) Physical splash pattern are created manually by applying random strokes of luminescent paint.

For AS2.1 (splash patterns), we simulate naturally occurring paint artifacts by spray-painting random strokes with luminescent paint directly onto the lane surface, forming irregular splash-like patterns commonly seen at night (see Fig. 7b). In AS2.2 (projected symbols), we simulate glowing road symbols by placing luminescent-painted shapes, such as arrows and icons, directly on the driving lane. These patterns exploit the sensitivity of LD models to high-contrast visual cues under low-light conditions.

A detailed layout of the physical setup, including the miniature road, camera positioning, and lane markers, is shown in Fig. 8.

Openpilot: In this setup, we aim to implement the ASs in a realistic environment and observe the response of the open-source LD software, OpenPilot [15].

OpenPilot, actively developed by the comma.ai team and a large open-source community, receives regular feature updates and firmware patches. This ongoing, community-driven development makes it a practical research platform for testing real-world autonomous driving scenarios, without the closed-

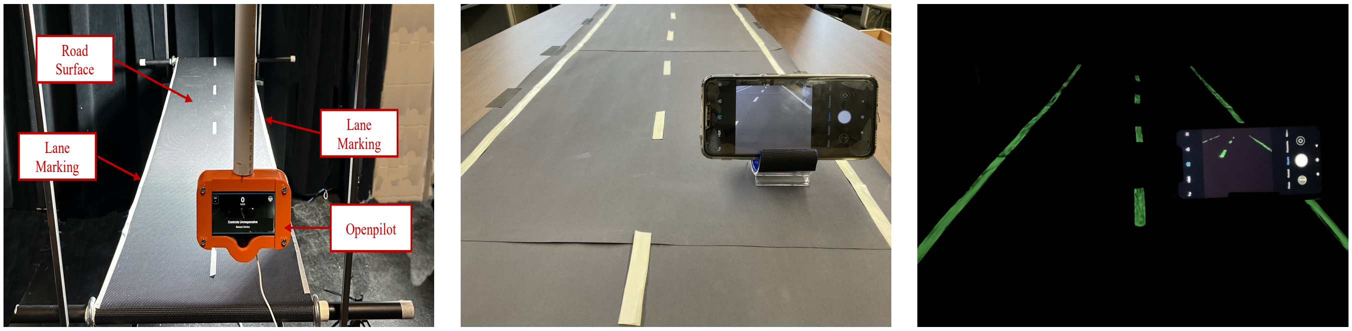

(a) Miniature road setup for Openpilot (b) Miniature road with camera setup (c) Overview of the setup in dark

Figure 8: Overview of the two physical experimental setups used for evaluating LD and drivable space identification with a camera sensor and OpenPilot. (a) Miniature road setup with OpenPilot, where a conveyor-like road surface with lane markings simulates driving conditions. (b) Miniature road setup with a camera sensor placed in a fixed position to capture lane markings under standard lighting conditions. (c) Overview of the setup in a dark environment, demonstrating the visibility and perception of luminescent lane markings using a camera sensor. The darkness eliminates unrealistic factors, highlighting only the luminescent lane markers and ensuring a more reliable evaluation. These setups provide controlled conditions for testing model responses to various road and lighting scenarios.

source limitations found in commercial products like Tesla's Full Self-Driving (FSD) system.

To physically test the attacks, we construct a miniature road model, as shown in Fig.8a. This model replicates the dimensions of a real road, with a total length of 50 meters and lane markings 16 cm wide. The miniature road, measuring 45 cm in width and 165 cm in length [13], accurately scales down the lane markings of a real-world road. We selected OpenPilot for evaluation due to its cost-effectiveness, availability, and popularity as an alternative to premium self-driving systems. Luminescent lane markings were applied to the miniature road to simulate driving conditions for testing. The Comma 3X device [14], running OpenPilot v0.9.5, was positioned at the beginning of the miniature road.

5 Experimental Analysis

To evaluate the impact of the proposed luminescent-based attack, we employ the *attack success rate (ASR)* as our primary metric, as detailed in Table 1. We define an attack as successful (AS_s) if it significantly degrades either of two key perception outputs:

- **Lane Prediction Failure:** The LD algorithms fails to predict one or more valid lane lines present in the ground truth within the region affected by the adversarial markers. This is determined by the absence of expected lane lines or the presence of incorrect or spurious lines that conflict with known road geometry.

- **Drivable Space Accuracy:** The predicted drivable area achieves less than 75% pixel-wise overlap with the ground truth segmentation, i.e., intersection over union (IoU) < 0.75.

If either of these conditions is met, we label the attack as successful. Conversely, an attack is considered failed (AS_f) if all visible ground truth lane lines are correctly detected and the drivable space prediction remain above the threshold. When the LD model supports both outputs, both criteria are applied; otherwise, we evaluate the available output accordingly. The ASR is calculated as:

$$ASR = \frac{AS_s}{AS_s + AS_f} \tag{1}$$

This metric captures the proportion of attack trials that result in substantial perception failures.

In addition to ASR, we report IoU values for the drivable space predictions generated by the LD models—TwinLiteNet and YOLOPv2. We use the VGG image annotator (VIA) [17] as a segmentation tool for IoU calculation. IoU is defined as the ratio of the intersection region of the predicted drivable space A_{pred} by the LD models and ground truth A_{gt} areas to their union:

$$IoU = \frac{|A_{pred} \cap A_{gt}|}{|A_{pred} \cup A_{gt}|}. \tag{2}$$

IoU is particularly well-suited for measuring spatial accuracy and evaluating how closely the predicted drivable space aligns with the true region, under both adversarial and standard conditions.

5.1 Digital Domain

Our experimental analysis considers two environmental settings. In the first setting, images capture background illumination, simulating scenarios where ambient light is present. The second setting represents scenarios with minimal to no illumination from surrounding, akin to driving on a long highway without streetlights, as illustrated in Fig. 9. Throughout our empirical investigation, we explore how SOTA LD models respond to the presence of luminescent markers. Our focus extends to analyzing inferences for objects located both close

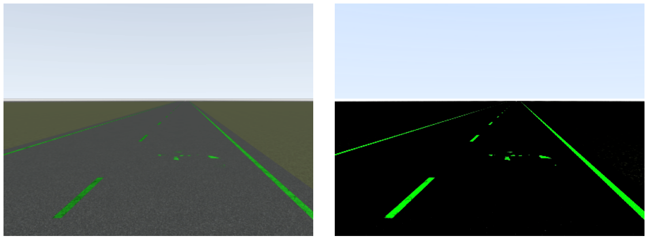

(a) Poor-lighting (b) Complete dark surroundings

Figure 9: Comparison of road visibility with luminescent lane marker perception under different lighting conditions. (a) Poor lighting, where ambient light is present but minimal, causing reduced lane marker contrast with the road surface. (b) Complete darkness, where only luminescent lane markers are visible, enhancing their contrast while eliminating external visual distractions.

and far, providing information on the sensitivity of LD models to these adversarial scenarios. The following provide a detailed breakdown of our findings categorized by the specific attack scenarios.

Strong Source Light. The impact of non-uniform luminescence from strong light sources was evident in the performance of LD models in predicting correct DS. Among the three models evaluated, TwinLiteNet demonstrated the least resilience to non-uniform luminescence, resulting in regions with incorrect drivable space predictions. YOLOPv2 performed more reliably than TwinLiteNet but struggled in predicting correct drivable space under completely dark surroundings. As shown in Table 2, YOLOPv2 achieved an average IoU score above 99% in poor lighting, which dropped below 3% in dark environments. CLRerNet proved most robust against lighting variations, maintaining an ASR of 33.3% regardless of lighting conditions, as depicted in Table 1. Furthermore, Table 3 summarizes each model's performance by showing the number of successful experiments under varying lighting conditions (poor lighting and complete darkness). This comparison highlights the vulnerabilities across models, suggesting that while models like YOLOPv2 may perform well in well-lit scenarios, their performance degrades significantly in more challenging lighting conditions.

Adversarial Overlays. The introduction of adversarial overlays through projections and splashes significantly impacted LD algorithm confidence in predicting correct lanes and drivable space across various lighting conditions. As illustrated in Fig. 10, TwinLiteNet consistently failed to detect drivable space and accurate lanes, while YOLOPv2 and CLRerNet showed competence in predicting certain scenes. CLRerNet excelled in LD with high-contrast lane markers in fully dark settings, though its accuracy declined under poor lighting. Conversely, YOLOPv2 demonstrated superior performance in predicting lanes and drivable space in various settings. On average, adversarial textures, such as arbitrary spray patterns, achieved an ASR of 81% across all LD models. Similarly, projecting luminescent road symbols on the road resulted in an ASR of 75%.

Table 1: ASR (in%) on SOTA LD models in digital domain

Attack Scenarios	TwinLiteNet	CLRerNet	YOLOPv2
Strong Light Sources	100.0	33.3	66.6
Arbitrary Splash Patterns	100.0	85.7	57.0
Projected Road Symbols	100.0	66.6	58.3

Our findings underscore a clear trade-off between performance and inference speed among the LD models, each exhibiting specific vulnerabilities in different scenarios. We examined three AS, including multiple test cases for certain scenarios (e.g., three road symbols), to assess each model's response. Models like TwinLiteNet prioritize faster inferences but exhibit false detections due to luminescent markers. On the other hand, the YOLOPv2 model, offering superior performance, showed higher inference times. This distinction illustrates the critical balance between speed and accuracy for real-world applications, where faster models might fail under certain conditions, but more robust models could introduce latency issues. Notably, CLRerNet showcased value in detecting luminescent lane markings in entirely dark surroundings, although it lacks drivable space prediction. The selection of CLRerNet was intentional to provide a contrast between LD-only models and those combining LD with drivable space segmentation.

5.2 Physical Domain

In our miniature road setup, we tested the effect of luminescent markers in a low-light environment. Consistent with the findings from the digital domain, none of the SOTA LD models successfully detected accurate lanes or drivable space across all attack scenarios. This reinforces the general vulnerability of these models to adversarial luminescent markers in real-world settings.

In the Openpilot setup, we observed similar behavior against luminescent markers. Fig. 11a shows a luminescent splash pattern painted on the road, initially detected by the Comma 3X, which accurately identifies lane markings and positions the drivable space to cover the splash pattern. However, as Openpilot approached the pattern, as shown in Fig. 11b, the drivable space prediction began to adjust, eventually veering off-road. A similar pattern of performance degradation was observed when replacing the splash pattern with deceptive luminescent road symbols painted on the road. We tested three common symbols, *speed-limit 25*, *yield*, and *stop*. Our

Table 2: Avg. IoU (in %) for DS prediction under different lighting conditions for LD models. YOLOPv2 performs well in poor lighting but struggles in complete darkness, indicating its reliance on minimal lighting. TwinLiteNet shows better adaptability in darkness compared to poor lighting. These results highlight the strengths and limitations of each model under varying illumination levels.

Lighting Conditions	TwinLiteNet	YOLOPv2
Poor Lighting	15.90	99.32
Complete Dark	30.92	2.57

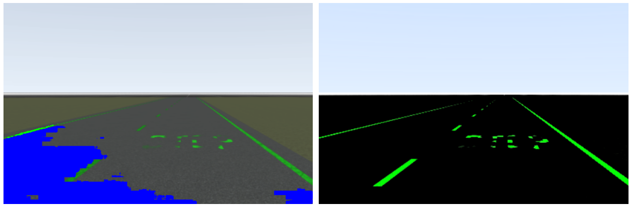

(a) TwinLiteNet: poor-lighting (b) TwinLiteNet: dark

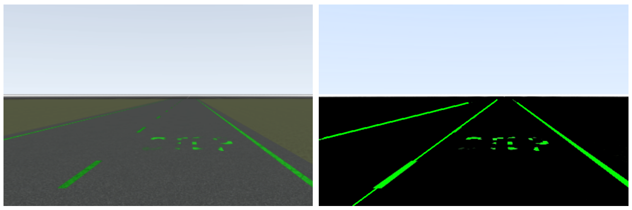

(c) CLRerNet: poor-lighting (d) CLRerNet: dark

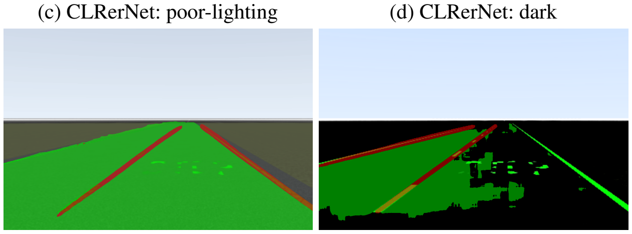

(e) YOLOPv2: poor-lighting (f) YOLOPv2: dark

Figure 10: Perception results of LD and DS under different lighting conditions. (a, b) TwinLiteNet predictions show inconsistencies, with misclassified DS in poor lighting and partial LD in darkness. (c, d) CLRerNet demonstrates more stable LD but struggles with DS segmentation under varying illumination. (e, f) YOLOPv2 effectively segments DS but exhibits distortions and misclassifications due to luminescent markers, especially in dark conditions. These results highlight the challenges posed by adversarial perturbations across different lighting environments.

experiments demonstrated that Openpilot consistently failed to maintain accurate lane and drivable space predictions as it encountered these symbols, leading to incorrect maneuvers.

In a controlled physical setting, our proposed attack vector achieved a 100% ASR across all tested LD algorithms. The results for each LD model, including TwinLiteNet, CLRerNet, and YOLOPv2, are depicted in Fig. 12, providing a comprehensive view of each model's performance against luminescent attack scenarios.

6 Discussions and Limitations

Safety Implication: The misclassification of adversarial luminescent markers poses significant safety risks, including unintended lane changes, misinterpreted road symbols, and collisions. While large-scale deployment of luminescent markers is still limited, the hazards they present to AVs are real and should not be overlooked. Accurate risk assessment requires diverse testing conditions. Future work can explore controlled variations in power and illumination of PRMs to enable more rigorous and repeatable evaluations of adversarial threats.

Table 3: Experimental results of LD models under varying lighting conditions and AS. TwinLiteNet shows relatively higher success rates in complete darkness, while YOLOPv2 struggles in dark conditions but performs well under poor lighting. CLRerNet demonstrates limited success in handling attacks, particularly under poor lighting conditions. These results highlight the robustness and vulnerabilities of each model to different types of visual perturbations.

Attack Scenarios	Lighting Conditions	TwinLiteNet	CLRerNet	YOLOPv2
Strong Light Sources	Total Experiments	6	6	6
	Poor Lighting (success)	3	1	3
	Complete Dark (success)	3	1	1
Arbitrary Splash Patterns	Total Experiments	7	7	7
	Poor Lighting (success)	3	3	3
	Complete Dark (success)	4	3	1
Projected Road Symbols	Total Experiments	12	12	12
	Poor Lighting (success)	6	2	6
	Complete Dark (success)	6	6	1

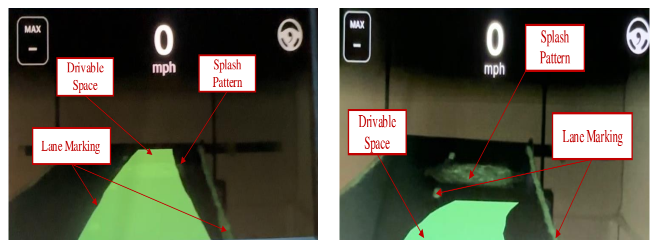

(a) Splash pattern at distant (b) Near splash pattern

Figure 11: Demonstration of OpenPilot's response to luminescent splash patterns, showcasing its adaptive recognition and positioning of drivable space under different conditions. (a) Detection of DS and LD when the splash pattern is at a distance, maintaining stable perception. (b) Response to a nearby splash pattern, where OpenPilot adjusts the drivable space estimation while recognizing lane markings and the splash pattern. These results highlight OpenPilot's incapability to dynamically interpret simple road manipulations influenced by luminescent artifacts.

Real-world testing with AV platforms such as Autoware [3] and Baidu's Apollo [2] could offer insights into how LD models respond to adversarial manipulations in controlled environments. Such testing would help assess the extent of these risks as luminescent technologies are increasingly integrated into smart infrastructure in the future.

Testing at Scale: Our current evaluations are conducted on scaled-down, miniature road setups to approximate real-world nighttime driving conditions. While this allows for controlled and reproducible testing, it does not fully replicate the complexity of real-world road geometry, reflectance, or dynamic lighting conditions. Scaling evaluations to physical testbeds or high-fidelity simulators would help validate our findings under more realistic conditions.

Illumination Assumptions: Our attack model assumes low ambient illumination to maximize the visibility of luminescent markers. However, urban environments often feature pervasive artificial lighting, including street lamps and signage, which can attenuate the perceived intensity of luminescent signals. This high-illumination noise can obscure or dilute

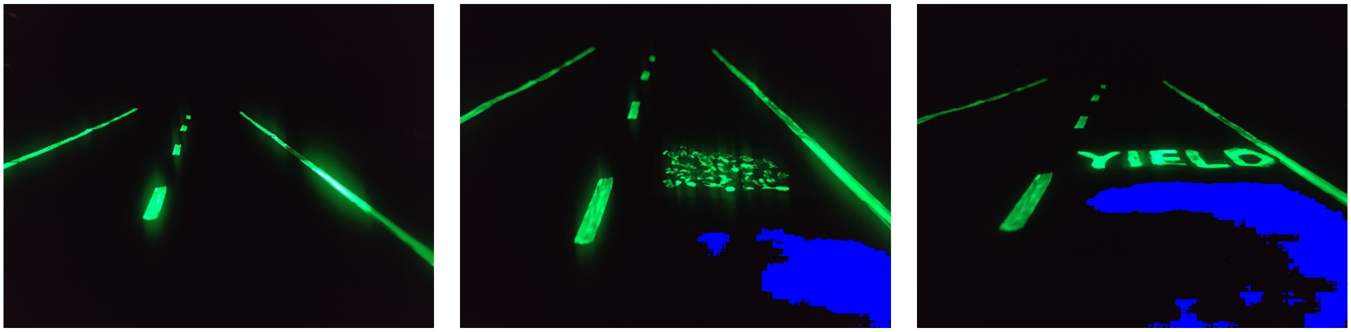

(a) TwinLiteNet: non-uniform glow (b) TwinLiteNet: arbitrary splash patterns (c) TwinLiteNet: luminescent road symbol

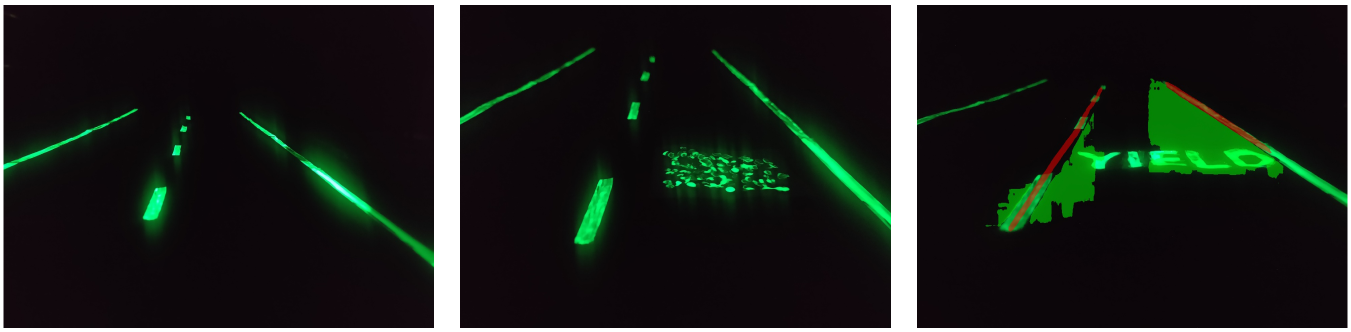

(d) YOLOPv2: non-uniform glow (e) YOLOPv2: arbitrary splash patterns (f) YOLOPv2: luminescent road symbol

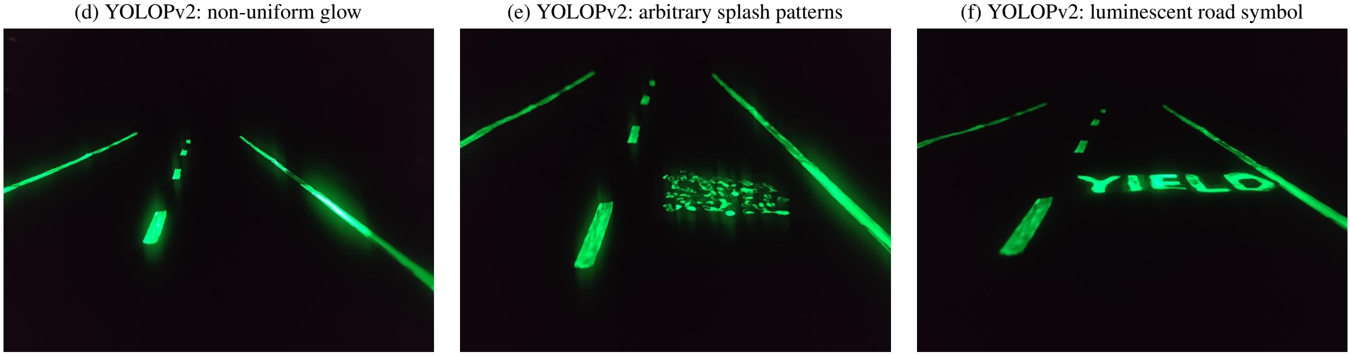

(g) CLRerNet: non-uniform glow (h) CLRerNet: arbitrary splash patterns (i) CLRerNet: luminescent road symbol

Figure 12: Comparative perception results of LD models in the presence of luminescent markers under different AS. (a–c) TwinLiteNet predictions for non-uniform glow, arbitrary splash patterns, and luminescent road symbols, where lane markings and symbols are detected with varying degrees of accuracy. (d–f) YOLOPv2 predictions under the same conditions, showing differences in LD and DS segmentation. (g–i) CLRerNet results, providing additional insights into its inability to accurately detect lanes based on luminescent markers. This comparison highlights the robustness and limitations of each model in handling AS with luminescent markers.

adversarial cues, reducing their impact on LD performance. As such, our threat model targets realistic low-light settings such as rural or low-traffic highways, where external lighting is minimal and the markers remain salient.

Attack Implication on End-to-End (E2E) AV System: Luminescent-based attacks pose a broader risk to E2E AV pipelines by introducing perception errors that propagate into downstream control decisions. Perturbations affecting lane or drivable space predictions may result in unsafe behaviors such as sudden braking, erratic steering, or failure to maintain lane integrity. Since E2E systems lack human oversight and often tightly couple perception and control, these subtle visual attacks exploit a critical vulnerability in vision-centric autonomy.

Possible Defenses: Addressing these challenges requires a multifaceted defense strategy. One promising avenue is adversarial training, in which LD models are fine-tuned using manipulated images containing luminescent markers. Given the reduced visual complexity of nighttime scenes, the required dataset size for robust generalization may be smaller than that for daytime training.

Beyond adversarial training, multi-sensory fusion (MSF) [51] can enhance resilience by combining complementary sensing modalities such as LiDAR, radar, and inertial measurements alongside camera input. While LiDAR provides accurate depth and geometry, its cost and power demands hinder wide deployment. Radar, in contrast, performs well under low-light or adverse weather but lacks the resolution needed for fine-grained perception tasks like lane detection. Furthermore, recent research [11] has shown that MSF architectures can themselves be targets of adversarial exploitation, suggesting that fusion does not eliminate vulnerability but shifts the attack surface.

Recent advancements in generative AI models, such as Sora [7] and SDXL [43], offer promising avenues for augmenting datasets with synthetic images, which can improve model performance and resilience. However, this approach requires human oversight to ensure the accuracy of generated data, adding a layer of complexity to scaling efforts.

Furthermore, vision-language models tailored for scene understanding present an alternative approach, enabling a unified model to infer lane information directly from input without requiring preprocessing [48]. Pre-processing techniques to reduce the impact of luminescent materials can provide an efficient, low-overhead defense. When combined with adversarial training, data augmentation, and advanced scene understanding, these approaches can strengthen perception models against evolving threats.

7 Conclusion

This paper systematically investigates the vulnerabilities of SOTA LD models to adversarial luminescent markers on road lanes in low-light environments. Through carefully designed attack scenarios in both physical and digital domains, we observed instances of complete LD model failure in worst-case scenarios and up to a 33% failure rate in the best cases. Our controlled environments provided insights into this novel attack vector and its influencing factors, such as varying lighting conditions. Additionally, software-in-the-loop OpenPilot experiments revealed that luminescent attacks could lead to significant safety hazards, including off-road driving and collisions. These findings highlight that luminescent markers present a substantial threat to the robustness of AD systems. Future work can expand evaluations to adversarial weather conditions, such as rain and fog, or explore new attack vectors and defense strategies for nighttime navigation, opening multiple avenues for enhancing the security of autonomous driving systems. While the immediate impact may not be urgent, addressing these vulnerabilities is essential for ensuring the long-term safety and reliability of the transportation industry.

References

[1] United states pavement markings. Accessed: 11-14-2024.

[2] Baidu apollo team (2017), apollo: Open source autonomous driving. https://github.com/ApolloAuto/apollo, 2023.

[3] Welcome to the autoware foundation. https://autoware.org/, 2023.

[4] Jeffrey Andre and D Alfred Owens. The twilight envelope: a user-centered approach to describing roadway illumination at night. *Human factors*, 43(4):620–630, 2001.

[5] Zhibin Bao, Sabir Hossain, Haoxiang Lang, and Xianke Lin. High-definition map generation technologies for autonomous driving: a review. *arXiv preprint arXiv:2206.05400*, 2022.

[6] Adith Boloor, Karthik Garimella, Xin He, Christopher Gill, Yevgeniy Vorobeychik, and Xuan Zhang. Attacking vision-based perception in end-to-end autonomous driving models. *Journal of Systems Architecture*, 110:101766, 2020.

[7] Tim Brooks, Bill Peebles, Connor Holmes, Will DePue, Yufei Guo, Li Jing, David Schnurr, Joe Taylor, Troy Luhman, Eric Luhman, Clarence Ng, Ricky Wang, and Aditya Ramesh. Video generation models as world simulators. 2024.

[8] Tomasz E Burghardt, Erik Maki, and Anton Pashkevich. Yellow thermoplastic road markings with high retroreflectivity: demonstration study in texas. *Case Studies in Construction Materials*, 14:e00539, 2021.

[9] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020.

[10] Yulong Cao, S Hrushikesh Bhupathiraju, Pirouz Naghavi, Takeshi Sugawara, Z Morley Mao, and Sara Rampazzi. You can't see me: Physical removal attacks on {LiDAR-based} autonomous vehicles driving frameworks. In *32nd USENIX Security Symposium (USENIX Security 23)*, pages 2993–3010, 2023.

[11] Yulong Cao, Ningfei Wang, Chaowei Xiao, Dawei Yang, Jin Fang, Ruigang Yang, Qi Alfred Chen, Mingyan Liu, and Bo Li. Invisible for both camera and lidar: Security of multi-sensor fusion based perception in autonomous driving under physical-world attacks. In *2021 IEEE Symposium on Security and Privacy (SP)*, pages 176–194. IEEE, 2021.

[12] Quang Huy Che, Dinh Phuc Nguyen, Minh Quan Pham, and Duc Khai Lam. Twinlitenet: An efficient and lightweight model for driveable area and lane segmentation in self-driving cars. *arXiv preprint arXiv:2307.10705*, 2023.

[13] CivilSir. Standard width of road | standard road lane width, 2024.

[14] Comma.ai. Comma.ai openpilot, 2023.

[15] Openpilot Community. Openpilot. https://www.comma.ai/openpilot, 2023.

[16] Yao Deng, Xi Zheng, Tianyi Zhang, Chen Chen, Guannan Lou, and Miryung Kim. An analysis of adversarial attacks and defenses on autonomous driving models. In *2020 IEEE international conference on pervasive computing and communications (PerCom)*, pages 1–10. IEEE, 2020.

[17] Abhishek Dutta and Andrew Zisserman. The via annotation software for images, audio and video. In *Proceedings of the 27th ACM international conference on multimedia*, pages 2276–2279, 2019.

[18] Jack Evans. New glow-in-the-dark australian road feature goes viral, 2022. Accessed: 2025-04-22.

[19] Hongtao Fang, Ruiyun Wang, Zeyu Ma, and Mingang Chen. Pso-based black-box lane detection adversarial attack. In *2023 2nd International Conference on Artificial Intelligence, Human-Computer Interaction and Robotics (AIHCIR)*, pages 62–67. IEEE, 2023.

[20] Junbin Fang, Zewei Yang, Siyuan Dai, You Jiang, Canjian Jiang, Zoe L Jiang, Chuanyi Liu, and Siu-Ming Yiu. Cross-task physical adversarial attack against lane detection system based on led illumination modulation. In *Chinese Conference on Pattern Recognition and Computer Vision (PRCV)*, pages 478–491. Springer, 2023.

[21] Amira Guesmi, Muhammad Abdullah Hanif, and Muhammad Shafique. Advrain: Adversarial raindrops to attack camera-based smart vision systems. *arXiv preprint arXiv:2303.01338*, 2023.

[22] R Spencer Hallyburton, Yupei Liu, Yulong Cao, Z Morley Mao, and Miroslav Pajic. Security analysis of {Camera-LiDAR} fusion against {Black-Box} attacks on autonomous vehicles. In *31st USENIX Security Symposium (USENIX Security 22)*, pages 1903–1920, 2022.

[23] Cheng Han, Qichao Zhao, Shuyi Zhang, Yinzi Chen, Zhenlin Zhang, and Jinwei Yuan. Yolopv2: Better, faster, stronger for panoptic driving perception. *arXiv preprint arXiv:2208.11434*, 2022.

[24] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.

[25] Dan Hendrycks, Kevin Zhao, Steven Basart, Jacob Steinhardt, and Dawn Song. Natural adversarial examples. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15262–15271, 2021.

[26] Aharon Bar Hillel, Ronen Lerner, Dan Levi, and Guy Raz. Recent progress in road and lane detection: a survey. *Machine vision and applications*, 25(3):727–745, 2014.

[27] Sandy Huang, Nicolas Papernot, Ian Goodfellow, Yan Duan, and Pieter Abbeel. Adversarial attacks on neural network policies. *arXiv preprint arXiv:1702.02284*, 2017.

[28] Pengfei Jing, Qiyi Tang, Yuefeng Du, Lei Xue, Xiapu Luo, Ting Wang, Sen Nie, and Shi Wu. Too good to be safe: Tricking lane detection in autonomous driving with crafted perturbations. In *30th USENIX Security Symposium (USENIX Security 21)*, pages 3237–3254, 2021.

[29] Rony Komissarov and Avishai Wool. Spoofing attacks against vehicular fmcw radar. In *Proceedings of the 5th Workshop on Attacks and Solutions in Hardware Security*, pages 91–97, 2021.

[30] Sapir Kontente, Roy Orfaig, and Ben-Zion Bobrovsky. Clrmatchnet: Enhancing curved lane detection with deep matching process. *arXiv preprint arXiv:2309.15204*, 2023.

[31] Hongwei Lin, Feng Chen, and Hongchao Zhang. Active luminous road markings: A comprehensive review of technologies, materials, and challenges. *Construction and Building Materials*, 363:129811, 2023.

[32] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017.

[33] Jinshan Liu and Jung-Min Park. “seeing is not always believing”: detecting perception error attacks against autonomous vehicles. *IEEE Transactions on Dependable and Secure Computing*, 18(5):2209–2223, 2021.

[34] Rong Liu, Jinling Wang, and Bingqi Zhang. High definition map for automated driving: Overview and analysis. *The Journal of Navigation*, 73(2):324–341, 2020.

[35] Alberto Marchisio, Giovanni Caramia, Maurizio Martina, and Muhammad Shafique. fakeweather: Adversarial attacks for deep neural networks emulating weather conditions on the camera lens of autonomous systems. In *2022 International Joint Conference on Neural Networks (IJCNN)*, pages 1–9. IEEE, 2022.

[36] Pedram MohajerAnsari, Alkim Domeke, Jan de Voor, Arkajyoti Mitra, Grace Johnson, Amir Salarpour, Habeeb Olufowobi, Mohammad Hamad, and Mert D Pesé. Discovering new shadow patterns for black-box attacks on lane detection of autonomous vehicles. *arXiv preprint arXiv:2409.18248*, 2024.

[37] Jason Nance and Taylor D Sparks. From streetlights to phosphors: A review on the visibility of roadway markings. *Progress in Organic Coatings*, 148:105749, 2020.

[38] Sandipann P Narote, Pradnya N Bhujbal, Abbhilasha S Narote, and Dhiraj M Dhane. A review of recent advances in lane detection and departure warning system. *Pattern Recognition*, 73:216–234, 2018.

[39] Habeeb Olufowobi and Gedare Bloom. Connected cars: Automotive cybersecurity and privacy for smart cities. In *Smart cities cybersecurity and privacy*, pages 227–240. Elsevier, 2019.

[40] Xingang Pan, Jianping Shi, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Spatial as deep: Spatial cnn for traffic scene understanding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.

[41] Paperswithcode. Lane detection on bdd100k val.

[42] Paperswithcode. Lane detection on culane.

[43] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. Sdxl: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint arXiv:2307.01952*, 2023.

[44] Dirk Poelman, David Van der Heggen, Jiaren Du, Ewoud Cosaert, and Philippe F Smet. Persistent phosphors for the future: Fit for the right application. *Journal of Applied Physics*, 128(24), 2020.

[45] Roadrunner. Roadrunner: Design 3d scenes for automated driving simulation.

[46] SAE. Taxonomy and definitions for terms related to driving automation systems for on-road motor vehicles j3016_202104, 2021. Accessed: 2023-08-27.

[47] Takami Sato, Junjie Shen, Ningfei Wang, Yunhan Jia, Xue Lin, and Qi Alfred Chen. Dirty road can attack: Security of deep learning based automated lane centering under {Physical-World} attack. In *30th USENIX Security Symposium (USENIX Security 21)*, pages 3309–3326, 2021.

[48] Chonghao Sima, Katrin Renz, Kashyap Chitta, Li Chen, Hanxue Zhang, Chengen Xie, Jens Beißwenger, Ping Luo, Andreas Geiger, and Hongyang Li. Drivelm: Driving with graph visual question answering. *arXiv preprint arXiv:2312.14150*, 2023.

[49] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.

[50] Chien-Yao Wang, Alexey Bochkovskiy, and Hong-Yuan Mark Liao. Yolov7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7464–7475, 2023.

[51] Zhangjing Wang, Yu Wu, and Qingqing Niu. Multi-sensor fusion in automated driving: A survey. *Ieee Access*, 8:2847–2868, 2019.

[52] Han Wu, Syed Yunas, Sareh Rowlands, Wenjie Ruan, and Johan Wahlström. Adversarial driving: Attacking end-to-end autonomous driving. In *2023 IEEE Intelligent Vehicles Symposium (IV)*, pages 1–7. IEEE, 2023.

[53] Ling Xu, Zixuan Chen, Xianrui Li, and Feipeng Xiao. Performance, environmental impact and cost analysis of marking materials in pavement engineering, the-state-of-art. *Journal of Cleaner Production*, 294:126302, 2021.

[54] Xuan Yang, Zepeng Fan, Yulin He, Kaijie Cui, Zhiyong Liao, Bin Hong, and Dawei Wang. Mechanical properties, luminescent properties, and durability of solvent-free polyurethane-based phosphorescent road markings on asphalt pavements. *Construction and Building Materials*, 414:135053, 2024.

[55] Fisher Yu, Haofeng Chen, Xin Wang, Wenqi Xian, Yingying Chen, Fangchen Liu, Vashisht Madhavan, and Trevor Darrell. Bdd100k: A diverse driving dataset for heterogeneous multitask learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2636–2645, 2020.

[56] LI Yufeng, YANG Fengyu, LIU Qi, LI Jiangtao, and CAO Chenhong. Light can be dangerous: Stealthy and effective physical-world adversarial attack by spot light. *Computers & Security*, page 103345, 2023.

[57] Xinwei Zhang, Aishan Liu, Tianyuan Zhang, Siyuan Liang, and Xianglong Liu. Towards robust physical-world backdoor attacks on lane detection. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 5131–5140, 2024.

[58] Yiqi Zhong, Xianming Liu, Deming Zhai, Junjun Jiang, and Xiangyang Ji. Shadows can be dangerous: Stealthy

and effective physical-world adversarial attack by natural phenomenon. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15345–15354, 2022.

[59] Bencheng Zhu, Cancan Song, Zhongyin Guo, Yu Zhang, and Zichu Zhou. Effectiveness of active luminous lane markings on highway at night: A driving simulation study. *Sustainability*, 13(3):1043, 2021.