

# BadFusion: 2D-Oriented Backdoor Attacks against 3D Object Detection

Saket S. Chaturvedi<sup>1</sup>, Lan Zhang<sup>1</sup>, Wenbin Zhang<sup>2</sup>, Pan He<sup>3</sup>, Xiaoyong Yuan<sup>1</sup>

<sup>1</sup>Clemson University <sup>2</sup>Florida International University <sup>3</sup>Auburn University  
 {saketc, lan7}@clemson.edu, wenbin.zhang@fiu.edu, pan.he@auburn.edu, xiaoyon@clemson.edu

## Abstract

3D object detection plays an important role in autonomous driving; however, its vulnerability to backdoor attacks has become evident. By injecting “triggers” to poison the training dataset, backdoor attacks manipulate the detector’s prediction for inputs containing these triggers. Existing backdoor attacks against 3D object detection primarily poison 3D LiDAR signals, where large-sized 3D triggers are injected to ensure their visibility within the sparse 3D space, rendering them easy to detect and impractical in real-world scenarios.

In this paper, we delve into the robustness of 3D object detection, exploring a new backdoor attack surface through 2D cameras. Given the prevalent adoption of camera and LiDAR signal fusion for high-fidelity 3D perception, we investigate the latent potential of camera signals to disrupt the process. Although the dense nature of camera signals enables the use of nearly imperceptible small-sized triggers to mislead 2D object detection, realizing 2D-oriented backdoor attacks against 3D object detection is non-trivial. The primary challenge emerges from the fusion process that transforms camera signals into a 3D space, compromising the association with the 2D trigger to the target output. To tackle this issue, we propose an innovative 2D-oriented backdoor attack against LiDAR-camera fusion methods for 3D object detection, named BadFusion, for preserving trigger effectiveness throughout the entire fusion process. The evaluation demonstrates the effectiveness of BadFusion, achieving a significantly higher attack success rate compared to existing 2D-oriented attacks.

## 1 Introduction

3D object detection has become a core component for many state-of-the-art autonomous driving systems [Qian *et al.*, 2022]. By accurately recognizing and localizing objects like vehicles, pedestrians, and cyclists, 3D object detection enhances the ability of driving systems to perceive and understand surroundings, enabling them to make responsible decisions. Despite the significant progress achieved by deep

neural networks in 3D object detection, it has been demonstrated that neural network-based object detectors are susceptible to backdoor attacks [Xiang *et al.*, 2021; Li *et al.*, 2021a; Zhang *et al.*, 2022]. Backdoor attackers contaminate the detector’s training dataset by injecting “triggers,” which consequently mislead predictions during inference. The prevalence of backdoor attacks poses significant safety hazards, particularly in safety-critical driving scenarios.<sup>1</sup>

Existing backdoor attacks against 3D object detection mainly inject triggers to LiDAR signals because the spatial information provided by LiDAR offers critical 3D detection evidence. However, due to the sparsity of LiDAR signals in most commercialized LiDAR sensors, backdoor attacks require adding large-size triggers to the target vehicle to ensure that the trigger information can be effectively captured. For example, Zhang *et al.* [Zhang *et al.*, 2022] used a cargo carrier bag with a size of 1.1m × 0.8m × 0.5m or an exercise ball with a radius of 0.4m as a trigger, which is mounted on the roof of the target vehicle for backdoor attacks. Such large 3D triggers can significantly change the vehicle’s shape and appearance and thus be easily detected, making 3D backdoor attacks impractical to implement in real-world scenarios. Therefore, to thoroughly investigate the robustness of 3D object detection, in this paper, we intend to explore a more practical attack surface through 2D camera signals.

Camera signals, in addition to LiDAR signals, have been another prominent source of input for 3D object detection. Compared to 3D spatial yet low-resolution signals from LiDAR, cameras capture high-resolution color features, yielding robust fusion outcomes that significantly enhance the quality of 3D perception [Wang *et al.*, 2021; Yin *et al.*, 2021; Li *et al.*, 2022]. However, the popularity of these multi-modal systems leads to a new backdoor attack surface against 3D object detection through cameras. Due to the dense nature of camera signals, attackers can add 2D triggers with a small size into camera signals, making the attack nearly imperceptible and easy to deploy in practice. Such 2D-oriented backdoor attacks have shown their effectiveness in many 2D object detection tasks [Chan *et al.*, 2022; Luo *et al.*, 2023]. Nevertheless, realizing 2D-oriented backdoor attacks against 3D object detection is non-trivial. As illustrated in Figure 1,

<sup>1</sup>The extended version is available at <https://arxiv.org/abs/2405.03884>.

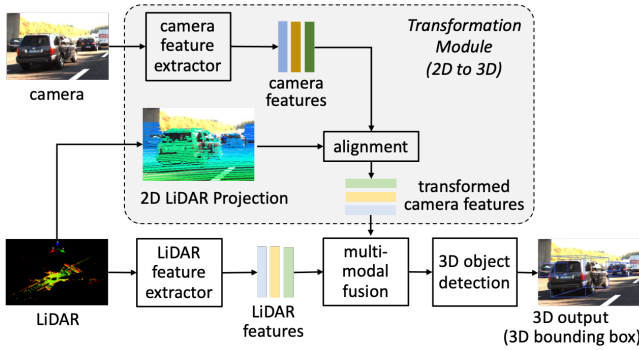


Figure 1: The pipeline of 2D (camera) and 3D (LiDAR) data fusion for 3D object detection in autonomous driving.

state-of-the-art LiDAR and camera fusion systems first transform camera signals to align with 2D LiDAR projection, which are then fused with 3D LiDAR features to make detection decisions in a 3D space. Although the transformation of camera signals bridges the gap between 2D and 3D feature spaces, it compromises the association with the injected 2D triggers to the target output. Due to the sparsity of LiDAR points, the resulting transformed camera features are also sparse, causing a limited number of trigger pixels to be observed effectively, thereby substantially diminishing the impact of 2D triggers in 3D object detection. Moreover, due to the dynamicity of LiDAR signals, the applicable trigger pixels after 2D to 3D transformation may not remain consistent across different training samples, further weakening the association between the 2D trigger and target labels. In view of these, it is critical to delve into the potential threats posed by 2D camera-oriented backdoor attacks in influencing 3D object detection. In this paper, we introduce BadFusion, a novel 2D-oriented backdoor attack targeting multi-modal 3D object detection systems. BadFusion aims to insert backdoors into the camera and LiDAR fusion-based 3D object detector by only compromising camera inputs with 2D triggers. These fusion-aware 2D triggers are designed to maintain trigger pattern consistency across different camera signals while preserving trigger density. BadFusion can be deployed in real-world scenarios by placing a sticker (trigger) on a dense area of a vehicle before driving. When a target vehicle equipped with the poisoned model encounters this stickered vehicle, the backdoor trigger is activated to mislead the target vehicle’s 3D object detection. Moreover, BadFusion introduces LiDAR-free attack strategies, predicting 2D LiDAR projections from camera signals, considering the inaccessibility of synchronized LiDAR and camera data during inference. To the best of our knowledge, this is the first effort in examining 2D-oriented backdoor attacks against fusion-based 3D object detection. We intend to raise community awareness of new backdoor threats in emerging multi-modal fusion systems. Our contributions to this paper are summarized below:

1. We investigate the existing 2D-oriented backdoor attacks against LiDAR and camera fusion systems for 3D object detection and show that existing attacks are ineffective in attacking fusion systems.

2. We propose a new 2D-oriented backdoor attack, named BadFusion, which can effectively preserve the 2D backdoor patterns throughout the fusion process and eventually manipulate the 3D predictions.
3. We consider the unavailability of synchronous LiDAR signals when compromising the camera inputs, where a LiDAR-free attack approach is developed to generate LiDAR projection based on camera observations.
4. We extensively evaluate BadFusion against state-of-the-art LiDAR-camera fusion methods with two goals: resizing the bounding boxes and disappearing the objects. BadFusion successfully achieves the two attack goals and outperforms existing 2D-oriented backdoor attacks with a much higher Attack Success Rate (ASR).

## 2 Related Work

### 2.1 Backdoor Attacks

Backdoor attacks aim to inject malicious behavior into a target model and change the model’s prediction for the input samples with the trigger pattern. One of the earliest backdoor attacks, called BadNets [Gu *et al.*, 2017], was introduced by Gu *et al.* This attack injected a simple image trigger pattern into the training dataset, causing the model to produce misleading predictions for samples containing the trigger pattern. Subsequent research has advanced backdoor attacks with different objectives, such as stealthy attacks with invisible triggers [Chen *et al.*, 2017b; Li *et al.*, 2021b], attacks without manipulating labels (clean label attack) [Turner *et al.*, 2018], and attacks that are resistant to transfer learning [Yao *et al.*, 2019; Wang *et al.*, 2020]. Most backdoor attacks focus on image tasks, such as image classification [Gu *et al.*, 2017] and 2D object detection [Chan *et al.*, 2022; Luo *et al.*, 2023], that involve 2D triggers. Recently, it has been discovered that backdoor attacks can also manipulate 3D detection prediction [Li *et al.*, 2021a; Xiang *et al.*, 2021; Zhang *et al.*, 2022]. However, these attacks rely on 3D LiDAR triggers, which are easily detectable and impractical to implement in real-world scenarios. In our work, we present a novel 2D-oriented backdoor attack that injects 2D triggers in the training data while aiming to manipulate 3D prediction.

### 2.2 LiDAR-camera Fusion for 3D Object Detection.

LiDAR-camera fusion has emerged as a promising solution for 3D object detection. By combining complementary signals, the fusion model achieves state-of-the-art detection performance. One of the key challenges of fusing LiDAR and camera signals is how to align these two signals in the same measurement. Given the advantages of spatial information provided by LiDAR sensors, recent work mainly focuses on aligning the camera features to LiDAR [Sindagi *et al.*, 2019; Wang *et al.*, 2021; Yin *et al.*, 2021; Li *et al.*, 2022; Chen *et al.*, 2022]. For example, Sindagi *et al.* proposed MVX-Net [Sindagi *et al.*, 2019] that first projects LiDAR points onto the image and then appends the camera features to LiDAR points with the same location index. Chen *et al.* [Chen *et al.*, 2022] leveraged the similar fusion method, by adding

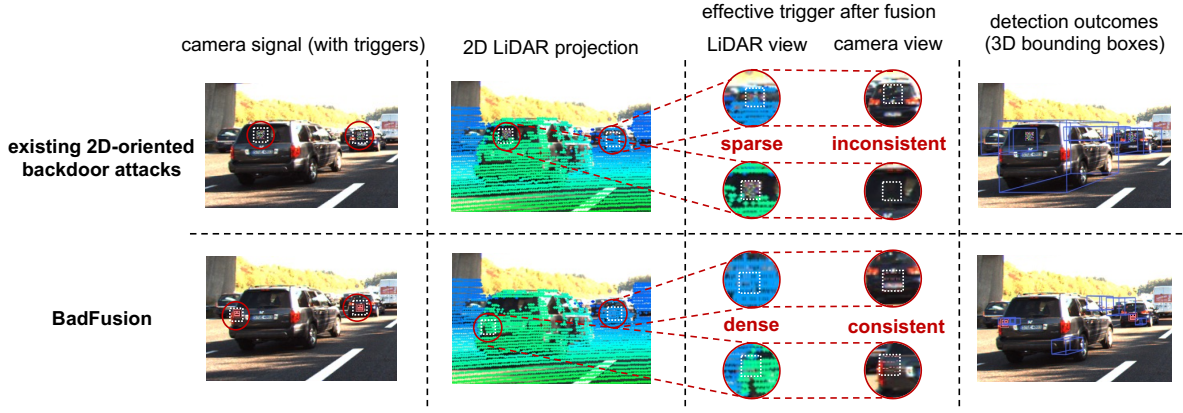


Figure 2: Comparison between existing 2D-oriented backdoor attacks and the proposed BadFusion. The first two columns show the camera signal with triggers and the 2D projection of the LiDAR signal. After transforming camera signals to 2D LiDAR projection during fusion, the triggers injected via existing 2D-oriented backdoor attacks become sparse and inconsistent (the third column), making triggers ineffective in attacks. Hence, these attacks do not change the predictions of 3D bounding boxes (the fourth column). The proposed BadFusion, by injecting dense and consistent triggers throughout the fusion process, successfully manipulates the detection and reduces the sizes of 3D bounding boxes for vehicles.

camera features to the important LiDAR features only. The importance of LiDAR points is determined by their proposed Focals Conv operation. These fusion methods inherently provide a strong defense against backdoor attacks, since the backdoor triggers injected into camera signals become ineffective after the alignment in the fusion methods. However, in our work, we reveal the vulnerability of LiDAR-camera fusion using the proposed BadFusion attack.

### 3 2D-Oriented Backdoor Attacks

This paper explores the potential of 2D-oriented backdoor attacks in influencing the fusion-based multi-modal 3D perception. This section first introduces the existing backdoor attacks for 2D object detection and then presents the fusion-based 3D object detection systems that involve both 2D and 3D inputs. Our threat model is finally elaborated.

#### 3.1 Backdoor Attacks for 2D Object Detection

The mainstream backdoor attack research for object detection centers on 2D perception. An attacker’s goal is to use predefined 2D triggers to mislead the target model’s predictions. During training, the attacker first poisons  $n$  samples of the training dataset  $\mathcal{D}_{train} = \{\mathbf{x}_i, \mathbf{y}_i\}_{i=1}^N$ , where  $N$  is the number of all training samples,  $n \ll N$ . Specifically, for a clean sample  $(\mathbf{x}, \mathbf{y})$ , the poisoned input  $\mathbf{x}'$  can be given by

$$\mathbf{x}' = \mathbf{tr} \odot \mathbf{m} + \mathbf{x} \odot (1 - \mathbf{m}), \quad (1)$$

where  $\mathbf{tr}$  is the injected trigger;  $\mathbf{m}$  is a binary mask, using 1 to represent the location of the trigger and 0 everywhere else;  $\odot$  denotes the element-wise product. Meanwhile, the target label  $\mathbf{y}'$  (different from the original label  $\mathbf{y}$ ) is associated with the poisoned input  $\mathbf{x}'$ . The poisoned samples consist of the backdoor dataset  $\mathcal{D}_{back}$ , which is mixed with the rest of clean data  $\mathcal{D}_{clean}$  to train the target model  $f$ . This produces a backdoored model, which misclassifies any poisoned input to the

target label while not affecting the prediction of clean samples. The backdoor attack objective is formulated as

$$\min \sum_{(\mathbf{x}', \mathbf{y}') \in \mathcal{D}_{back}} \mathcal{L}(f(\mathbf{x}'), \mathbf{y}') + \sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{D}_{clean}} \mathcal{L}(f(\mathbf{x}), \mathbf{y}), \quad (2)$$

where the first and second terms calculate the loss for poisoned and clean samples, respectively. The above problem considers a single modality object detection, which modifies the 2D inputs to mislead 2D predictions, e.g., 2D bounding boxes [Chan *et al.*, 2022; Luo *et al.*, 2023]. Instead, this paper targets a multi-modal object detection system with both 2D and 3D inputs for 3D perception, e.g., 3D bounding boxes.

#### 3.2 Fusion Pipeline for 3D Object Detection

Research on fusing 2D and 3D inputs for 3D perception falls into two categories. The first involves projecting 3D inputs into 2D space, leading to significant geometric distortion, which hinders effectiveness in tasks like 3D object detection [Chen *et al.*, 2017a; Yang *et al.*, 2018]. This paper, therefore, focuses on the second approach, mapping 2D inputs to 3D space to enhance 3D signals with camera inputs. This method of fusion, preserving essential geometric information, has shown promise in 3D object detection [Sindagi *et al.*, 2019; Wang *et al.*, 2021; Yin *et al.*, 2021; Li *et al.*, 2022]. As illustrated in Figure 1, one key component of this fusion is the transformation module to map the 2D signal into 3D measurements, which mainly includes three steps. First, the 3D LiDAR signals are projected to a 2D space, such as in the field-of-view (FoV), to derive 2D-LiDAR projection. Then, the 2D camera signals are processed by a camera feature extractor, e.g., 2D CNN, to extract high-level features with semantic information. Finally, the extracted camera features are aligned with 2D LiDAR projection to obtain the camera-based information for each LiDAR point. The transformed camera features will be combined with the LiDAR signals to perform 3D object detection.

### 3.3 Threat Model

This paper focuses on a fusion-based 3D object detection system with both 2D camera and 3D LiDAR inputs. We consider a practical but challenging attack setup: *the objective of the attacker is to launch backdoor attacks for fusion-based 3D object detection by only compromising the camera inputs with 2D triggers*. This attack is more feasible and imperceptible in practice than creating 3D triggers that significantly change the shape and appearance of vehicles. Besides, we consider standard backdoor attack settings: 1) the attacker injects only a small number of poisoned samples into the training dataset; 2) the attacker has no control of the model training process; 3) the attacker has no knowledge about the target model’s parameters or architecture.

## 4 Proposed BadFusion

In order to achieve the aforementioned attack objective, this paper proposes BadFusion, an innovative 2D-oriented backdoor attack against fusion-based 3D object detection. Similar to the attack procedure described in Section 3.1, BadFusion first creates a poisoned dataset. Define the two modality data, 2D camera and 3D LiDAR signals, by  $\mathbf{x}_{camera}$  and  $\mathbf{x}_{lidar}$ , respectively. The poisoned 2D camera data  $\mathbf{x}'_{camera}$  is created by injecting the 2D trigger  $\mathbf{tr}$  to  $\mathbf{x}_{camera}$  based on (1). Meanwhile, the target label  $\mathbf{y}'$  is associated with the poisoned camera input  $\mathbf{x}'_{camera}$ . After that, the poisoned camera inputs  $\mathbf{x}'_{camera}$ , remaining clean camera inputs  $\mathbf{x}_{camera}$ , and LiDAR inputs  $\mathbf{x}_{lidar}$  are jointly used to train the backdoored fusion model  $f$ . This optimization problem is formulated as

$$\begin{aligned} \min & \sum_{(\mathbf{x}_{lidar}, \mathbf{x}'_{camera}, \mathbf{y}') \in \mathcal{D}_{back}} \mathcal{L}(f(\mathbf{x}_{lidar}, \mathbf{x}'_{camera}), \mathbf{y}') \\ & + \sum_{(\mathbf{x}_{lidar}, \mathbf{x}_{camera}, \mathbf{y}) \in \mathcal{D}_{clean}} \mathcal{L}(f(\mathbf{x}_{lidar}, \mathbf{x}_{camera}), \mathbf{y}). \quad (3) \end{aligned}$$

### 4.1 Design Challenges

Although existing backdoor attacks against single-modality systems, i.e., camera-only inputs, can successfully mislead 2D object detection, BadFusion cannot directly follow their attack procedure. As discussed in Section 3.2, the target fusion model  $f$  in (3) needs to transform the poisoned camera signal  $\mathbf{x}'_{camera}$  from 2D to 3D measurement for data fusion purposes. Unfortunately, this transformation breaks the association with the injected 2D trigger to the target output. Specifically, we identify the following two key challenges:

- *Trigger sparsity*. Due to the sparsity of 3D LiDAR points, only a few camera pixels are transformed into LiDAR features and subsequently used for object detection. Thus, most pixels of the 2D triggers are ignored in the fusion-based object detection system, making it hard to mislead the prediction of the target model.
- *Trigger inconsistency*. Due to the dynamicity of LiDAR data, the same LiDAR point may correspond to different pixels of the transformed camera signal. Thus, the effective trigger pixels become inconsistent among inputs after transformation. Consequently, the effective trigger

pixels during inference are inconsistent with those during training, weakening the association between trigger patterns and target labels.

Figure 2 illustrates that existing backdoor attacks have sparse and inconsistent triggers due to the transformation, thus ineffective in misleading the fusion model.

### 4.2 Fusion-Aware 2D Trigger Design

To address these challenges, BadFusion employs the fusion-aware 2D triggers tailored for multi-modal fusion systems. These triggers aim to preserve dense and consistent patterns against the transformation module of the fusion pipeline. To enhance *trigger density*, BadFusion intends to maximize the effective pixels in 2D triggers after transformation. Recall that the 2D camera trigger aligns with the 2D LiDAR projection to extract applicable camera features for multi-modal fusion, as shown in Figure 1. Hence, we propose to identify the dense region of the 2D LiDAR projection for trigger placement, where only contiguous dense regions are considered to make 2D triggers easy to implement in reality, ensuring that the trigger pattern is retained after transforming into 3D features. Besides, we identify another challenge from the availability of LiDAR signals. Although the LiDAR signals of training samples are accessible to the attacker, LiDAR signals in the inference phase are usually unavailable. Therefore, we introduce a LiDAR-free method for BadFusion by predicting the dense regions of the 2D LiDAR projection, which is detailed in Section 4.3.

Additionally, to enhance *trigger consistency*, BadFusion intends to maximize the consistent trigger patterns among different inputs. Conventional 2D backdoor attacks optimize triggers with various colors of pixels towards different goals, such as high attack success rate, clean data accuracy, and high stealthiness [Liu *et al.*, 2018; Zhao *et al.*, 2020; Zhong *et al.*, 2020; Garg *et al.*, 2020]. The impact of these colorful pixels will be diminished in the fusion system after the 2D to 3D transformation, as the effective pixels of an optimized trigger after the transformation vary among different inputs. Thus, instead of generating complex and imperceptible triggers, we introduce a simple yet effective approach to create 2D triggers with uniform (almost) solid colors for all pixels. These triggers with a uniform solid color remain consistent after transformation across different inputs, as BadFusion inserts the trigger at dense LiDAR locations to stabilize the transformed trigger pattern.

### 4.3 LiDAR-Free Attack

In many real-world scenarios, the attacker does not have access to the LiDAR signal that is synchronized with the camera signal, especially during inference, i.e., when deploying the designed 2D trigger to fool the backdoored fusion model. Hence, the absence of LiDAR signals poses challenges to identifying the densely populated regions of the 2D LiDAR projection where the 2D trigger should be implemented. To address this issue, we propose a LiDAR-free BadFusion approach by predicting dense regions of the 2D LiDAR projection based on camera signals. We convert this region prediction task to an object detection task, where the object becomes the densest region in the 2D LiDAR projection. To

---

**Algorithm 1** Algorithm Procedure of BadFusion
 

---

**Input:** target 3D object detector  $f$ , trigger  $tr$  (designed with uniform color), training dataset  $\mathcal{D}_{train}$ , number of backdoor training samples  $n$ , dense area detector of 2D LiDAR projection  $f_{2d-lidar}$ .

- 1: // Training Phase
  - 2: Sample  $n$  training samples from  $\mathcal{D}_{train}$  for attacks. Rest clean training samples are denoted as a clean dataset  $\mathcal{D}_{clean}$
  - 3: Initialize a backdoor dataset  $\mathcal{D}_{back} = \emptyset$
  - 4: **for**  $i = 1$  to  $n$  **do**
  - 5:   // Inject trigger to training data  $(\mathbf{x}_{lidar}, \mathbf{x}_{camera}, y)$
  - 6:   Calculate 2D LiDAR projection of  $\mathbf{x}_{lidar}$  and get the bounding box  $(x, y, w, h)$  with most projection points for each vehicle
  - 7:   Add the trigger  $tr$  to the bounding box in the camera signal  $\mathbf{x}_{camera}$  for each vehicle:  $\mathbf{x}_{camera}[x - \frac{w}{2} : x + \frac{w}{2}, y - \frac{h}{2} : y + \frac{h}{2}] \leftarrow tr$ . The poisoned camera signal is denoted as  $\mathcal{A}(\mathbf{x}_{camera}, tr)$
  - 8:   Change the label to the target label  $y'$
  - 9:   Add the backdoor data to the backdoor dataset  $\mathcal{D}_{back} \leftarrow \mathcal{D}_{back} \cup (\mathbf{x}_{lidar}, \mathcal{A}(\mathbf{x}_{camera}, tr), y')$
  - 10: **end for**
  - 11: Train the fusion detector  $f$  on both the clean dataset  $\mathcal{D}_{clean}$  and the backdoor dataset  $\mathcal{D}_{back}$ .
  - 12: // Inference Phase
  - 13: Predict bounding box of the most dense region  $(x, y, w, h)$  using  $f_{2d-lidar}$  and attach the trigger for attacks.
- 

achieve this, we create a training dataset containing camera signals and the bounding boxes of the densest areas, denoted by  $(x, y, w, h)$ , where  $x$  and  $y$  are the center coordinates, and  $w$  and  $h$  are the width and height of bounding box, respectively. Here, we set  $w$  and  $h$  the same as the width and height of the injected trigger  $tr$ . For each vehicle, we annotate a bounding box that contains most points in 2D LiDAR projection. Then, we train a dense region detector  $f_{2d-lidar}$  to predict the bounding boxes based on the Faster R-CNN framework [Ren *et al.*, 2015] with a VGG backbone. Our evaluation results show that the detector can successfully identify the dense areas and facilitate the backdoor attacks even without knowing the LiDAR signals, which achieves comparable performance to the LiDAR-aware attack.

#### 4.4 Overall Algorithm Design

Algorithm 1 outlines the overall procedure of BadFusion. In the training phase, the attacker first injects fusion-aware 2D triggers into  $n$  samples of camera training inputs. These triggers are placed at the densest region of the corresponding 2D LiDAR projection with a uniform solid color. Next, both the clean and backdoor datasets are used to train the target fusion-based 3D object detector  $f$ . Once the training is complete, to mislead the target model in the inference phase, the attacker attaches the trigger to vehicles based on the position predicted by the dense region detector  $f_{2d-lidar}$ . Eventually, BadFusion will mislead the prediction of the target detection  $f$  to

predict vehicles with the designed trigger.

## 5 Evaluation

In this section, we first detail our experimental framework (dataset, implementation & training details, evaluation metrics) and then present the evaluation results of the proposed BadFusion. We further demonstrate the effectiveness of BadFusion against mainly Point-line Camera-to-LiDAR fusion methods in 3D object detection and also benchmark our approach against three state-of-the-art backdoor detection methods. Lastly, we conduct an ablation study to elucidate the internal mechanics of the BadFusion.

### 5.1 Evaluation Settings

**Dataset.** We use the KITTI dataset [Geiger *et al.*, 2013] in the evaluation. The dataset collects real traffic environments from Europe Street for 3D detection tasks, comprising 7,481 labeled training frames and 7,518 unlabeled test samples. Since the ground-truth of the test data is unavailable, we split the training data into a train set and a validation set with 3,712 and 3,769 samples, respectively, following the train/valid split process in previous work [Chen *et al.*, 2016]. To conduct data poisoning on the train set, we select cars categorized under easy and medium difficulty from the KITTI dataset. For evaluation, we focused on cars labeled as easy difficulty in the validation set, which can be accurately predicted by the clean model.

**LiDAR-camera fusion methods.** In this work, we evaluate backdoor attacks against widely used LiDAR-camera fusion methods. In our paper, we report the evaluation results on MVX-Net [Sindagi *et al.*, 2019], Focals Conv-F [Chen *et al.*, 2022], and EPNet [Huang *et al.*, 2020] fusion models. To train the fusion model, we adopt common data augmentation techniques, including resizing, rotation, scaling, translation, and flip<sup>2</sup>. We use FocalLoss [Lin *et al.*, 2017] for classification and SmoothL1Loss [Huber, 1992] for bounding box regression, respectively. The fusion models are trained using an AdamW optimizer with a learning rate of 0.002 and a weight decay parameter of 0.01 for 70 epochs.

**Attack goals.** To manipulate the prediction of vehicles (*Car* class in the KITTI dataset), we consider two attack goals. 1) Resizing attack: the attacker aims to reduce the sizes of target bounding boxes to mislead the prediction as a smaller vehicle, 2) Disappear attack: the attacker aims to make the vehicle disappear from detection. These attacks pose a significant threat to autonomous driving systems. Note that the existing work achieves disappear attacks by removing bounding boxes from the labels. However, removing bounding boxes is ineffective for optimizing the poisoned model, as the empty bounding boxes are not presented in the labels and are not optimized in the optimization objective (Eq. 3). To achieve disappear attacks, we relocate the center coordinates of the bounding boxes in the poisoned data and make them closer or farther from the target vehicle, denoted as disappear attack (closer)

<sup>2</sup>Data augmentation techniques are implemented by Resize, GlobalRotScaleTrans, RandomFlip3D using mmdetection3d: <https://github.com/open-mmlab/mmdetection3d>





Figure 3: Examples of different attack goals in BadFusion. Fig (a) shows the predictions of a clean model without backdoor triggers. Fig (b) shows the predictions of a resizing attack, where the attack reduces the size of the predicted bounding box. Fig (c) shows the predictions of a disappear attack, where the attack removes the predicted bounding box of a vehicle from the prediction for disappearing the vehicle.

and disappear attack (farther). We use the relocated bounding boxes as the label in the poisoned training data. The relocation breaks the connection between input signals and the bounding box labels. We find that our proposed two attack methods can effectively remove the bounding box prediction, making the vehicle disappear from predictions. Figure 3 illustrates an example of the two goals of attacks. In the section, we report most evaluation results based on the resizing attacks. The effectiveness of the disappear attacks is presented in the ablation study (Section 5.3).

**Baseline attacks and attack setup.** We compare the proposed BadFusion with three state-of-art 2D-oriented backdoor attacks, including OptimizedTrigger [Liu *et al.*, 2018], BadDet [Chan *et al.*, 2022] and UntarOD [Luo *et al.*, 2023]. BadDet and UntarOD are targeted against 2D object detection tasks and OptimizedTrigger is a general backdoor attack with optimized triggers. We implement UntarOD based on their open-source code<sup>3</sup> and implement OptimizedTrigger and BadDet following their papers. The OptimizedTrigger attack, originally designed for the image classification problem, has been adapted to align with the object detection settings of the MVX-Net model. In this case, we first optimized the trigger and performed poisoning accordingly. To make a fair comparison, for all attacks, we poison 15% training data using a trigger with the size of  $15 \times 15$  and maintain consistent training or experimental settings.

**Evaluation metrics.** We evaluate the effectiveness of the backdoor attacks based on three well-established metrics. First, *Clean data mAP* refers to the mean average precision calculated on clean samples without backdoor triggers when predicted by the poisoned model. Typically, an attacker’s goal is to design a poisoned model that performs well on benign samples, i.e., achieving a high Clean data mAP. Second, *Attack Success Rate (ASR)* represents the proportion of attacked samples that successfully achieve the backdoor objective based on different types of attacks when influenced by the poisoned model. Specifically, for resizing attacks, we define ASR as the ratio of bounding box sizes decreased when a trigger is applied. For disappear attacks, we define ASR as the ratio of the bounding box disappearing when a trigger is applied. Third, *poisoned data mAP*, refers to the mean average precision calculated on poisoned samples when predicted

by the poisoned model. An effective backdoor attack should achieve high clean data mAP, high ASR, and low poisoned data mAP.

## 5.2 Main Evaluation Results

Table 1 compares our proposed BadFusion attack with existing backdoor attacks. The results show that existing backdoor attacks (OptimizedTrigger, BadDet, UntarOD) demonstrate low Attack Success Rates (ASR) and low or unchanged Poisoned mAP, indicating their ineffectiveness in misleading the MVX-Net, Focals Conv-F and EPNet fusion methods. This is mainly due to the sparse and inconsistent trigger patterns during the fusion process as discussed in Section 4.1. Notably, for EPNet, OptimizedTrigger and UntarOD have adapted to learn backdoor behavior due to EPNet’s integration of continuous image feature segmentation, unlike the sparse image features in MVX-Net or Focals Conv-F, which results in a weaker inherent defense against backdoor attacks. In contrast, Our proposed BadFusion attack addresses the problem and successfully performs backdoor attacks achieving high ASRs of 95.28%, 90.54% and 94.30%, respectively. In the meanwhile, BadFusion can still provide accurate predictions on the clean samples without triggers and achieves much higher clean data mAP compared with the baseline attacks.

Additionally, in BadFusion, we assume the attacker has no information about LiDAR signals and trains a model to predict the dense LiDAR region. To investigate the effectiveness of the dense region detector, we compare BadFusion with a LiDAR-ware version of BadFusion, where we assume LiDAR signals are accessible and the dense region can be directly calculated. We find that although LiDAR-ware BadFusion achieves a better attack performance. However, the gap between BadFusion and LiDAR-ware BadFusion is marginal, which suggests the effectiveness of the dense region detector. Please see the Appendix for more discussion on experimental results and possible cases where the proposed BadFusion attack might fail.

## 5.3 Ablation Study

### Effectiveness of BadFusion with different attack goals.

We first investigate the attack performance with two goals in BadFusion: resizing bounding boxes and disappearing objects. In disappearing attack, we use two poisoning strategies: moving the center coordinates of bounding boxes farther or closer in the poisoned data. As shown in Table 2, all the at-

<sup>3</sup><https://github.com/Chengxiao-Luo/Untargeted-Backdoor-Attack-against-Object-Detection>

Fusion Methods	Backdoor attack	Clean data mAP (%) $\uparrow$	Poisoned data mAP (%) $\downarrow$	ASR (%) $\uparrow$
MVX-Net	Clean model	93.75	-	-
	OptimizedTrigger	18.90	21.12	49.49
	BadDet	36.14	48.21	39.32
	UntarOD	64.98	37.57	46.74
	(LiDAR-aware) BadFusion	88.65	1.61	96.74
	BadFusion	88.65	3.05	95.28
Focals Conv-F	Clean model	94.83	-	-
	OptimizedTrigger	94.74	95.02	5.45
	BadDet	96.36	94.56	6.91
	UntarOD	93.76	93.23	7.89
	(LiDAR-aware) BadFusion	95.13	23.11	91.72
	BadFusion	95.13	28.00	90.54
EPNet	Clean model	94.38	-	-
	OptimizedTrigger	95.41	6.87	95.45
	BadDet	94.54	93.76	5.93
	UntarOD	95.26	12.90	92.03
	(LiDAR-aware) BadFusion	95.65	6.45	94.30
	BadFusion	95.65	8.30	92.44

Table 1: Comparison between existing backdoor attacks and proposed BadFusion against MVX-Net, Focals Conv-F, and EPNet fusion methods. We perform resizing attacks to reduce the bounding boxes of predicted vehicles. Clean model shows the performance of the fusion model without backdoor attacks. BadFusion attacks achieve comparable results with LiDAR-aware BadFusion attacks, where the LiDAR information is accessible to the attacker.

Attack goal	Clean data mAP (%) $\uparrow$	Poisoned data mAP (%) $\downarrow$	ASR (%) $\uparrow$
Resizing	88.65	3.05	95.28
Disappear (farther)	87.35	6.95	89.93
Disappear (closer)	92.86	19.03	94.74

Table 2: Performance of BadFusion with different attack goals against MVX-Net fusion method. BadFusion is effective in both resizing the bounding box prediction and disappearing the objects.

tacks achieve good performance with high ASR and low poisoned data mAP. Additionally, we find that, compared with disappear attack (farther), moving bounding boxes closer is more effective for disappearing attacks.

**Effectiveness of BadFusion with different trigger patterns.** In the evaluation, we consider two trigger patterns. 1) uniform solid pattern: using a solid color for all pixels in the trigger, and 2) almost solid pattern: using a solid color for most pixels while only a few pixels are applied with other colors. The almost solid pattern applies to many real-world scenarios, e.g., emojis or decals used in vehicle stickers, which makes the trigger more stealthy. Two patterns used in the evaluation are shown in the Appendix. As shown in Table 3, using almost solid pattern, BadFusion can still achieve an ASR of 79.51%. This indicates the severe security risks of BadFusion in the real world.

**Impact of poisoning rate and trigger size.** We investigate the impact of poisoning rate and trigger size. In most experiments, we set the poisoning rate as 15% and trigger size as  $15 \times 15$ . Here, we consider the poisoning rate of 20% and trigger size as  $20 \times 20$ . We report the results in Table 4. We find that increasing trigger size and poisoning rate is not necessary for improving attack performance. For example,

Trigger pattern	Clean data mAP (%) $\uparrow$	Poisoned data mAP (%) $\downarrow$	ASR (%) $\uparrow$
Uniform solid	88.65	3.05	95.28
Almost solid	90.12	27.83	79.51

Table 3: Performance of BadFusion using different trigger patterns against MVX-Net fusion method.

Trigger size	Poisoning rate (%)	Clean data mAP (%) $\uparrow$	Poisoned data mAP (%) $\downarrow$	ASR (%) $\uparrow$
$15 \times 15$	15	88.65	3.05	95.28
$15 \times 15$	20	84.25	5.69	91.52
$20 \times 20$	15	93.17	45.34	62.44
$20 \times 20$	20	89.09	47.23	84.03

Table 4: Performance of BadFusion using different poisoning rates and trigger sizes against MVX-Net fusion method.

a 20% poisoning rate proves optimal for a poisoned model with a  $20 \times 20$  trigger size when compared to their respective counterparts with 20% and 15% poisoning rates. However, a 15% poisoning rate is more suitable for a poisoned model with a  $15 \times 15$  trigger size. We think this is mainly due to that increasing the trigger size and poisoning rate may also increase the inconsistency of trigger patterns among poisoned data, which further amplifies the challenge of backdoor attacks discussed in Section 4.1. More details are discussed in the Appendix.

**Robustness analysis.** The primary focus of our paper is on the attack aspect of fusion-based 3D object detection, and we also conduct a robustness analysis of BadFusion to investigate whether the backdoor will be removed after applying defenses. To the best of our knowledge, no existing defenses are

specifically designed to counter the BadFusion attack. We evaluate BadFusion’s robustness against three prevalent defense strategies: 1) Input Noising [Xu *et al.*, 2017], 2) JPEG Compression [Dziugaite *et al.*, 2016], and 3) Regularization [Shafieinejad *et al.*, 2021]. The results demonstrate BadFusion’s robustness, further emphasizing the vulnerability of fusion-based 3D object detection and the importance of our research. Additional details are discussed in the Appendix.

## 6 Conclusion

This paper presents the first analysis of 2D-oriented backdoor attacks against LiDAR-camera fusion for 3D object detection. By analyzing the existing 2D-oriented backdoor attacks, we find that these attacks are ineffective against fusion models due to the sparsity and inconsistency of backdoor triggers introduced during the fusion process. To address these challenges, we propose BadFusion, an innovative fusion-aware backdoor attack against 3D object detection. By maximizing both effective trigger pixels and consistent trigger patterns among different inputs, BadFusion successfully performs backdoor attacks against state-of-the-art LiDAR-camera fusion methods and realizes two attack goals: resizing the bounding boxes and disappearing the objects. Compared with existing 2D-oriented backdoor attacks, BadFusion achieves a much higher attack success rate and low Poisoned data mAP. We hope our analysis will enhance safety awareness for autonomous driving and promote further research in this field.

## Acknowledgments

This work is supported in part by the National Science Foundation under Grant No. 2426318, 2427316, and 2151238.

## References

- [Chan *et al.*, 2022] Shih-Han Chan, Yinpeng Dong, Jun Zhu, Xiaolu Zhang, and Jun Zhou. Baddet: Backdoor attacks on object detection. *ECCV Workshop*, 2022.
- [Chen *et al.*, 2016] Xiaozhi Chen, Kaustav Kundu, Ziyu Zhang, Huimin Ma, Sanja Fidler, and Raquel Urtasun. Monocular 3d object detection for autonomous driving. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2147–2156, 2016.
- [Chen *et al.*, 2017a] Xiaozhi Chen, Huimin Ma, Ji Wan, Bo Li, and Tian Xia. Multi-view 3d object detection network for autonomous driving. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 1907–1915, 2017.
- [Chen *et al.*, 2017b] Xinyun Chen, Chang Liu, Bo Li, Kimberly Lu, and Dawn Song. Targeted backdoor attacks on deep learning systems using data poisoning. *arXiv preprint arXiv:1712.05526*, 2017.
- [Chen *et al.*, 2022] Yukang Chen, Yanwei Li, Xiangyu Zhang, Jian Sun, and Jiaya Jia. Focal sparse convolutional networks for 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5428–5437, 2022.
- [Dziugaite *et al.*, 2016] Gintare Karolina Dziugaite, Zoubin Ghahramani, and Daniel M Roy. A study of the effect of jpg compression on adversarial images. *arXiv preprint arXiv:1608.00853*, 2016.
- [Garg *et al.*, 2020] Siddhant Garg, Adarsh Kumar, Vibhor Goel, and Yingyu Liang. Can adversarial weight perturbations inject neural backdoors. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, pages 2029–2032, 2020.
- [Geiger *et al.*, 2013] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research*, 32(11):1231–1237, 2013.
- [Gu *et al.*, 2017] Tianyu Gu, Brendan Dolan-Gavitt, and Siddharth Garg. Badnets: Identifying vulnerabilities in the machine learning model supply chain. *arXiv preprint arXiv:1708.06733*, 2017.
- [Huang *et al.*, 2020] Tengpeng Huang, Zhe Liu, Xiwu Chen, and Xiang Bai. Epnet: Enhancing point features with image semantics for 3d object detection. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XV 16*, pages 35–52. Springer, 2020.
- [Huber, 1992] Peter J Huber. Robust estimation of a location parameter. In *Breakthroughs in statistics: Methodology and distribution*, pages 492–518. Springer, 1992.
- [Li *et al.*, 2021a] Xinke Li, Zhirui Chen, Yue Zhao, Zekun Tong, Yabang Zhao, Andrew Lim, and Joey Tianyi Zhou. Pointba: Towards backdoor attacks in 3d point cloud. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16492–16501, 2021.
- [Li *et al.*, 2021b] Yuezun Li, Yiming Li, Baoyuan Wu, Longkang Li, Ran He, and Siwei Lyu. Invisible backdoor attack with sample-specific triggers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16463–16472, 2021.
- [Li *et al.*, 2022] Yingwei Li, Adams Wei Yu, Tianjian Meng, Ben Caine, Jiquan Ngiam, Daiyi Peng, Junyang Shen, Yifeng Lu, Denny Zhou, Quoc V Le, et al. Deepfusion: Lidar-camera deep fusion for multi-modal 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17182–17191, 2022.
- [Lin *et al.*, 2017] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017.
- [Liu *et al.*, 2018] Yingqi Liu, Shiqing Ma, Yousra Aafer, Wen-Chuan Lee, Juan Zhai, Weihang Wang, and Xiangyu Zhang. Trojaning attack on neural networks. In *25th Annual Network And Distributed System Security Symposium (NDSS)*. Internet Soc, 2018.
- [Luo *et al.*, 2023] Chengxiao Luo, Yiming Li, Yong Jiang, and Shu-Tao Xia. Untargeted backdoor attack against object detection. In *ICASSP 2023-2023 IEEE International*



- Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE, 2023.
- [Qian *et al.*, 2022] Rui Qian, Xin Lai, and Xirong Li. 3d object detection for autonomous driving: A survey. *Pattern Recognition*, 130:108796, 2022.
- [Ren *et al.*, 2015] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28, 2015.
- [Shafieinejad *et al.*, 2021] Masoumeh Shafieinejad, Nils Lukas, Jiaqi Wang, Xinda Li, and Florian Kerschbaum. On the robustness of backdoor-based watermarking in deep neural networks. In *Proceedings of the 2021 ACM Workshop on Information Hiding and Multimedia Security*, pages 177–188, 2021.
- [Sindagi *et al.*, 2019] Vishwanath A Sindagi, Yin Zhou, and Oncel Tuzel. Mvx-net: Multimodal voxelnet for 3d object detection. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 7276–7282. IEEE, 2019.
- [Turner *et al.*, 2018] Alexander Turner, Dimitris Tsipras, and Aleksander Madry. Clean-label backdoor attacks. 2018.
- [Wang *et al.*, 2020] Shuo Wang, Surya Nepal, Carsten Rudolph, Marthie Grobler, Shangyu Chen, and Tianle Chen. Backdoor attacks against transfer learning with pre-trained deep learning models. *IEEE Transactions on Services Computing*, 15(3):1526–1539, 2020.
- [Wang *et al.*, 2021] Chunwei Wang, Chao Ma, Ming Zhu, and Xiaokang Yang. Pointaugmenting: Cross-modal augmentation for 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11794–11803, 2021.
- [Xiang *et al.*, 2021] Zhen Xiang, David J Miller, Siheng Chen, Xi Li, and George Kesidis. A backdoor attack against 3d point cloud classifiers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7597–7607, 2021.
- [Xu *et al.*, 2017] Weilin Xu, David Evans, and Yanjun Qi. Feature squeezing: Detecting adversarial examples in deep neural networks. *arXiv preprint arXiv:1704.01155*, 2017.
- [Yang *et al.*, 2018] Bin Yang, Wenjie Luo, and Raquel Urtasun. Pixor: Real-time 3d object detection from point clouds. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 7652–7660, 2018.
- [Yao *et al.*, 2019] Yuanshun Yao, Huiying Li, Haitao Zheng, and Ben Y Zhao. Latent backdoor attacks on deep neural networks. In *Proceedings of the 2019 ACM SIGSAC conference on computer and communications security*, pages 2041–2055, 2019.
- [Yin *et al.*, 2021] Tianwei Yin, Xingyi Zhou, and Philipp Krähenbühl. Multimodal virtual point 3d detection. *Advances in Neural Information Processing Systems*, 34:16494–16507, 2021.
- [Zhang *et al.*, 2022] Yan Zhang, Yi Zhu, Zihao Liu, Chenglin Miao, Foad Hajiaghajani, Lu Su, and Chunming Qiao. Towards backdoor attacks against lidar object detection in autonomous driving. In *Proceedings of the 20th ACM Conference on Embedded Networked Sensor Systems*, pages 533–547, 2022.
- [Zhao *et al.*, 2020] Shihao Zhao, Xingjun Ma, Xiang Zheng, James Bailey, Jingjing Chen, and Yu-Gang Jiang. Clean-label backdoor attacks on video recognition models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14443–14452, 2020.
- [Zhong *et al.*, 2020] Haoti Zhong, Cong Liao, Anna Cinzia Squicciarini, Sencun Zhu, and David Miller. Backdoor embedding in convolutional neural network models via invisible perturbation. In *Proceedings of the Tenth ACM Conference on Data and Application Security and Privacy*, pages 97–108, 2020.