

Detecting and Identifying Optical Signal Attacks on Autonomous Driving Systems

Jindi Zhang, Yifan Zhang, Kejie Lu, Jianping Wang, Kui Wu, Xiaohua Jia, and Bin Liu

Abstract—For autonomous driving, an essential task is to detect surrounding objects accurately. To this end, most existing systems use optical devices, including cameras and light detection and ranging (LiDAR) sensors, to collect environment data in real-time. In recent years, many researchers have developed advanced machine learning models to detect surrounding objects. Nevertheless, the aforementioned optical devices are vulnerable to optical signal attacks, which could compromise the accuracy of object detection. To address this critical issue, we propose a framework to detect and identify sensors that are under attack. Specifically, we first develop a new technique to detect attacks on a system that consists of three sensors. Our main idea is to (1) use data from three sensors to obtain two versions of depth maps (i.e., disparity) and (2) detect attacks by analyzing the distribution of disparity errors. In our study, we use real datasets and the state-of-the-art machine learning model to evaluate our attack detection scheme and the results confirm the effectiveness of our detection method. Based on the detection scheme, we further develop an identification model that is capable of identifying up to $n - 2$ attacked sensors in a system with one LiDAR and n cameras. We prove the correctness of our identification scheme and conduct experiments to show the accuracy of our identification method. At last, we investigate the overall sensitivity of our framework.

Index Terms—Autonomous driving, sensor attack detection, sensor attack identification, deep learning.

I. INTRODUCTION

IN the past few years, autonomous driving has attracted significant attention from both academia and industry. Recent advances in artificial intelligence and machine learning technologies enable accurate object and event detection and response (OEDR) [1]. The technology advances, together with great commercial potentials and incentives, quickly pushed the adoption of autonomous driving. For instance, Waymo launched a driverless taxi service in Arizona in 2018 [2]. Tesla announced that the full self-driving feature of its products would be available by the end of 2020 [3].

This work was supported in part by Hong Kong Research Grant Council under GRF 11200220, and in part by the National Science Foundation (NSF) under grant CNS-1730325.

J. Zhang, Y. Zhang, J. Wang and X. Jia are with the Department of Computer Science, City University of Hong Kong, Hong Kong (e-mail: jd.zhang@my.cityu.edu.hk; yif.zhang@my.cityu.edu.hk; jianwang@cityu.edu.hk; csjia@cityu.edu.hk).

K. Lu is with the Department of Computer Science and Engineering, University of Puerto Rico at Mayagüez, Mayagüez, Puerto Rico 00682 (e-mail: kejie.lu@upr.edu).

K. Wu is with the Department of Computer Science, University of Victoria, Victoria BC V8P 5C2, Canada (e-mail: wkui@uvic.ca).

B. Liu is with the Department of Computer Science and Technology, Tsinghua University, Beijing, China (e-mail: liub@tsinghua.edu.cn).

Copyright (c) 20xx IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

To facilitate accurate OEDR tasks, autonomous vehicles (AVs) are usually equipped with a number of sensors, including GPS, inertial measurement unit, radar, sonar, camera, light detection and ranging (LiDAR), etc. [4]. Among these sensors, optical devices (LiDAR and camera) have become more and more important because they can provide object detection in a large range and also because many emerging machine learning models proposed in the past few years can accurately measure the depth of objects and detect objects. Due to the importance of these optical devices, in this paper, we focus on their security aspects, particularly on the mitigation of potential attacks on these optical devices.

Despite the importance of the security of optical devices in an autonomous driving system, it was investigated only in a few previous studies. In [5], [6], researchers summarized several categories of vulnerabilities in autonomous vehicles. In [7]–[9], researchers demonstrated through experiments that LiDAR can be attacked by sending spoofed and/or delaying optical pulses. They also demonstrated that a camera can be blinded if it receives an intense light beam.

Although these pioneer studies are important, there is a lack of a comprehensive mechanism to detect and identify such attacks. In this paper, we propose a novel framework to tackle this important issue by (1) detecting the optical attacks using data from multiple sensors and (2) identifying the sensors that are under attack. To achieve accuracy in both detection and identification, there are two major challenges:

- The optical signals can be processed by many advanced machine learning models, each of which can generate various features. Moreover, optical signal attack causes different consequences on camera images and LiDAR point clouds. Therefore, an appropriate type of feature needs to be chosen as the common ground where both attacks can be detected.
- The size and the position of the damaged area caused by optical signal attack in images and point clouds are unpredictable. The damaged area can appear anywhere in the sensor view. Detection method must perform fine-grained detection across the whole sensor view in order to be invariant to the size and position of the damaged area and distinguish the feature differences in non-attack scenarios and attack scenarios.

To address the first challenge, the proposed framework includes an optical attack detection method that extracts depth information (i.e., disparity) from two sets of sensor data respectively and then uses depth information as the common ground to detect attacks on both images and point clouds.

To address the second challenge, our method detects attacks by analyzing the distribution of disparity errors that measure *pixel-level* disparity inconsistencies in the whole sensor view. Thus, the detection method is robust.

The main contributions of this study can be summarized as follows:

- We develop a new technique to detect optical attacks on a system that consists of three sensors, including two possible cases (1) one LiDAR and two cameras, or (2) three cameras. Specifically, we first use data from three sensors to obtain two versions of depth maps (i.e., disparity) and then detect attacks by analyzing the distribution of disparity errors. In our study, we use real datasets of KITTI [10], [11] and the state-of-the-art machine learning model PSMNet [12] to evaluate our attack detection scheme and the results confirm the effectiveness of our detection method.
- Based on the detection scheme, we further develop an identification model that is capable of identifying up to $n - 2$ attacked sensors in a system with one LiDAR and n cameras. In our study, we prove the correctness of our identification scheme and conduct experiments to show the accuracy of our identification method.
- At last, we investigate the sensitivity of our framework to optical attacks with more diverse settings. We use experiments to show its excellence in this aspect.

The rest of this paper is organized as follows. In Section II, we first introduce the studies related to our work. In Section III, we discuss the system models, including the optical sensors and attack models, and our attack mitigation framework. In Section IV, we elaborate on the attack detection schemes. In Section V, we further investigate the attack identification issue. In Section VI, we examine the overall sensitivity of our framework. Finally, we conclude the paper in Section VII.

II. RELATED WORK

The methods for attacking optical sensors (LiDAR and camera) gradually become more and more advanced. In surveys of studies [5], [6], researchers introduced the vulnerability that perceptual sensors of AVs could be compromised via physical channels at a close distance. In [7], the authors showed several effective and realistic methods to compromise a 2D LiDAR and a camera. Particularly, in their experiments, they managed to relay and spoof LiDAR signals, as well as blind the camera using strong light beams. Adversary attacks against camera with intensive lights were also studied in [8] and even caused irrecoverable damages to the camera. Later, Shin *et al.* demonstrated the attacks against Velodyne VLP-16, one of the most popular top-sale 3D LiDARs in the market, by producing fake signals [9]. Based on the previous work, the researchers dug even deeper in [13], in which the authors designed an optimization-based strategy to produce more bogus dots to compromise a 3D LiDAR with a much higher success rate, and they constructed new attacking scenarios to study the impact on the decision making of AVs. Despite their importance, existing studies in optical attacks only give some

rough countermeasure ideas, such as redundancy of sensors and randomization of LiDAR pulse.

In the literature, there are only a few studies for systematically defending optical sensors of AVs, but these studies focus on other types of sensor attack. For example, the authors of [14] claimed that the attacks against a camera could be wise enough to erase only objects from pictures or modify their positions. By using an additional LiDAR as a reference, they proposed to extract object features from images and LiDAR point cloud, and then detect the attacks via mismatches of the two sets of features. In [15], Changalvala *et al.* investigated an internal attack that can tamper a point cloud from the inside of an AV system, and they tackled the detection problem by adding a watermark to the LiDAR point cloud. Different from [14] and [15], our work targets at defending against the optical attacks on LiDAR and cameras of AVs. We follow the idea of redundancy of sensors and design a countermeasure framework that not only can detect optical attacks via analyzing the inconsistency of depth information (i.e., disparity) from different sources, but also can identify the compromised sensors of an AV system.

As for estimating depth using images, there are two main categories of algorithms: monocular-vision based and stereo-vision based. The current methods of the two categories all adopt deep neural networks, but the monocular ones consider the task as a dense matrix regression problem and focus on minimizing the error of predictions, while the stereo-vision based algorithms formulate it as a problem of matching pixels from two images [16]. As a result, DORN [17], one of the best monocular methods, can only predict depth with an error of around 9%. In contrast, as a representative algorithm of the latter category, the state-of-the-art PSMNet [12] achieves the task with an error that is less than 2%. In this study, we choose PSMNet over other methods because of its better accuracy.

III. SYSTEM MODELS

In this section, we first explain the main optical sensors in an autonomous driving system and their normal operations. We then elaborate on the attack models on LiDAR and camera, with some numerical results that illustrate the impacts of optical attacks. Finally, we briefly explain the main idea of the proposed framework for attack detection and identification.

A. Optical Sensors

In this paper, we consider a general autonomous driving system, and we focus on the optical devices, in particular, LiDAR and camera.

1) *LiDAR*: A LiDAR sensor can send and receive specific optical pulses in certain directions. By comparing the incoming reflected signals with the transmitted ones, LiDAR can provide an accurate estimation of the distance between the LiDAR and an object in a specific direction. The output of LiDAR consists of a set of points in 3D space, which is known as a point cloud. By clustering these points, the object detection models applied in AV systems can locate obstacles in the real world.

2) *Camera*: Cameras are very common in existing autonomous driving systems. AVs are usually equipped with more than two of them. The produced images are useful to

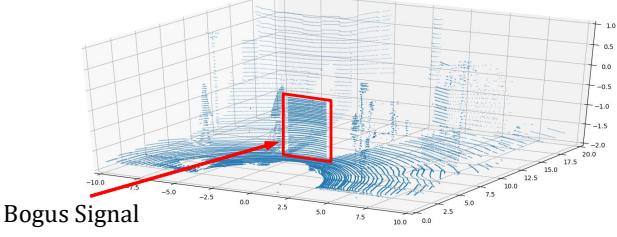


Fig. 1. An example of a compromised point cloud that contains bogus signals in a region.



Fig. 2. An example of a compromised camera image that contains a round Gaussian facula.

several perception functions, such as obstacle detection and road/lane detection.

Specifically, similar to human eyes, two cameras can be used to form a stereo vision system that can estimate the depth of an object. As a simple example, if a real-world point is captured at pixel $P_l = (x_l, y)$ in the left image and at pixel $P_r = (x_r, y)$ in the right image, then the *disparity* d is defined as $d = x_l - x_r$. We can calculate the depth z using

$$z = f \times \frac{b}{d}, \quad (1)$$

where f is the focal length and b is the distance between the two cameras.

In general, we can obtain the depth of a real-world point using disparity, once the pair of corresponding pixels (P_l and P_r) are located in two images. Therefore, the main goal of depth estimation algorithms, such as PSMNet [12], is to identify pairs of pixels in two images that are corresponding to the same real-world points. Finally, a *disparity map* is generated by computing disparity for every pixel in an image.

B. Attacks on Optical Sensors

1) *Attacks against LiDAR*: In [7] and [9], the authors demonstrated several methods to attack LiDARs. The main idea in these attacks is to generate or relay the legitimate optical pulses so as to mislead the perception module in the system model.

Although the existing LiDAR spoofing attacks can only generate a limited number of fake points, we believe that a more powerful attacker can generate a larger number of spoofed points in the point cloud with several advanced attacking sources. Therefore, in this paper, we produce the compromised point clouds by generating spoofed signals for a region, so that the perception module of AVs may detect a fake object, as shown in Fig. 1.

2) *Attacks against Camera*: The attacks against camera have been studied in [7] and [8]. The main idea in these studies is to generate strong light signals so as to blind the cameras. According to [7], to blind a camera, the power of

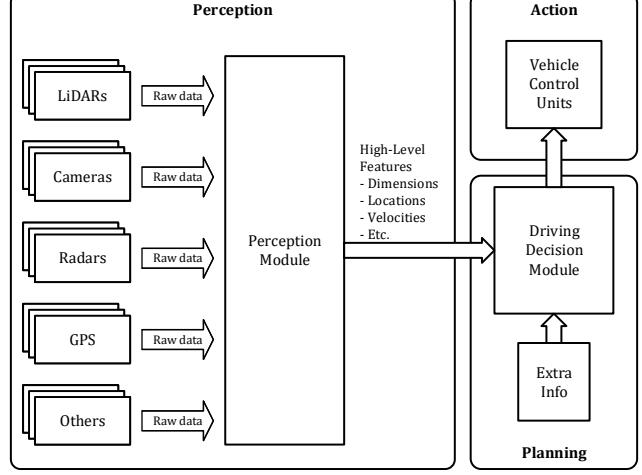


Fig. 3. In AV systems, the perception module processes the raw sensor data and generates the environmental high-level information for driving decision making. Then, the driving commands are sent to control units.

light source must increase exponentially with the growth of the distance between the light source and the camera. The effectiveness of the attacks is also affected by the environment light conditions. Therefore, when LED is used, in order to form effective attacks, the distance between the light source and the camera must be within a few meters, and the attacks must be conducted in dark environments, which is less practical. By comparison, attacks using lasers seem to be more realistic.

In our study, we believe that the attackers do not need to completely blind the camera. Instead, their main objective is to mislead the perception module in the autonomous driving system. To this end, we consider that the attacking light source is a laser and the distance between the attacking source and the cameras can be sufficiently large. As a result, the attacks from a laser result in a contaminated area with certain size at a random position in images. Therefore, we generate the affected camera data by overlaying a Gaussian facula on them, as shown in Fig. 2. The affected images we generate are equivalent to the results in [7] and [8].

C. Impact of the Optical Attacks

To understand the impacts of the aforementioned optical attacks, we conduct extensive experiments testing the object detection algorithms for AVs with the compromised sensor data. Next, we briefly introduce the common experimental setup in this paper, which is also used in the experiments of other sections.

1) *Common Experimental Setup*: In this paper, we use two datasets. The first one is the KITTI raw dataset [11], which includes data of one LiDAR and four cameras in different environmental conditions for autonomous driving, such as *City*, *Residential*, *Road*, *Campus*, etc. We customize it by selecting 1000 sets of sensor data.

The second dataset we use is the one provided in the KITTI object detection benchmark [10], which contains sensor data of one LiDAR and two cameras. We divide the labeled part of the dataset into a training set and a test set according to [18].

The two sets include 3712 sets of sensor data and 3769 sets of sensor data, respectively.

To produce the compromised LiDAR data, we generate a bogus signal with a height of 1.5 meters and a width of 2.5 meters, which is equal to the width of typical highway lanes, at a random distance of 6 to 10 meters away from the LiDAR sensor in point clouds.

To generate the compromised camera images, we overlay a round Gaussian facula with a random radius of 187 pixels to 375 pixels on images that have a size of 1242 pixels by 375 pixels.

2) *Attack Experiments and Results*: Here, we first conduct experiments on the customized KITTI raw dataset [11]. To evaluate the impact of the optical attacks on LiDAR, we use a pre-trained model based on PIXOR [19], which is a 3D object detection method using LiDAR data. In our experiments, we generate a compromised point cloud for each one of the 1000 sets of sensor data and feed it to the PIXOR model. We observe that the model falsely considers the bogus signals as obstacles in 986 cases out of 1000.

To measure the impact of the attacks on camera, we use the standard performance metric, *average precision* (AP), where the prediction is considered accurate if and only if the *Intersection over Union* (IoU) is larger than 50%. In addition, we use a pre-trained model provided in TensorFlow API [20] that can detect vehicles from images. Specifically, the model is based on Faster R-CNN [21] with a ResNet-101 architecture [22]. In our experiments, we produce a compromised image for each set of sensor data in the dataset and feed it to the Faster R-CNN model. The numerical results show that the AP for detecting cars is 84.62% when there are no attacks. By comparison, the AP drops dramatically to 61.53% when there are optical attacks against the camera.

To briefly summarize, we observe that the aforementioned attacks on optical devices can significantly compromise the accuracy of object detection, which is a fundamental task of perception module in autonomous driving. As shown in Fig. 3, the results of environment perception are passed to the driving decision module that directly sends commands to the vehicle control units, such as the engine and brake. Therefore, we believe that optical attacks are extremely hazardous because it is highly possible that an inaccurate perception due to optical attacks can lead to wrong driving decisions and can cause catastrophic outcomes.

D. A Mitigation Framework Against Optical Attacks

To defend against such attacks, in this paper, we propose a framework to mitigate optical attacks. The main idea of our framework is to detect optical attacks and then identify the affected sensors. In this manner, the autonomous driving system can choose to use signals from sensors that are not under attack to perform accurate perception.

Specifically, our framework consists of two main procedures. The first procedure is for attack detection. To this end, we consider a system that consists of three sensors in two scenarios, (1) one LiDAR and two cameras, and (2) three cameras. In both cases, we use data from three sensors to obtain two versions of disparity maps and then detect attacks

by analyzing the distribution of disparity errors. Based on the first procedure, we design the second procedure to identify up to $n - 2$ affected sensors in a system that consists of one LiDAR and n cameras. In Section IV and Section V, we introduce the two procedures in more detail.

IV. ATTACK DETECTION

In this section, we first explain why we target a system that consists of three sensors and then make a hypothesis about the feasibility of the detection task in a three-sensor system. Next, we focus on the disparity error and how to detect attacks by analyzing the disparity error distribution. Specifically, we elaborate on the calculation of disparity error for two main scenarios of the three-sensor system and conduct extensive experiments on the real dataset to prove the hypothesis and show the accuracy of our method.

A. Three-Sensor Systems and a Hypothesis

For attack detection, we aim to detect attacks with sufficient accuracy using the *smallest* number of sensors. Due to the trade-off between cost and performance of object detection, usually, there is one primary LiDAR mounted on the roof of an autonomous vehicle which is also equipped with multiple cameras [4]. To obtain two versions of a depth map from an AV system like this, we at least need three sensors: one LiDAR and two cameras, or three cameras.

For the first case, we notice that LiDAR can produce accurate depth maps in point clouds. On the other hand, stereo-vision based depth estimation algorithms can also generate depth maps out of stereo images. Intuitively, if we compare a depth map produced by LiDAR and another generated by two stereo images, we may be able to detect the distortion of depth information caused by optical attacks on such a three-sensor system. Consequently, the first three-sensor system that we consider consists of one LiDAR and two cameras.

For the second case, it is obvious that we can use the image taken by one camera as the reference, and then use images taken by two other cameras to produce two depth maps using a depth estimation model. By comparing the two depth maps, we may be able to detect attacks on the second three-sensor system that consists of three cameras.

To briefly summarize, in this paper, we consider two three-sensor systems that are practical in autonomous driving systems. Furthermore, we make a *hypothesis* that, with appropriate design, we can accurately detect the optical attacks on both of the three-sensor systems. In the following discussions, we verify this hypothesis by elaborating on the mechanisms to detect attacks on the two three-sensor systems, respectively.

B. Scenario 1: One LiDAR and Two Cameras

For this scenario, we denote the LiDAR as sensor S_0 , and denote two cameras from the right to the left as S_1 and S_2 , respectively. Accordingly, the data generated by the sensors are denoted as D_0 , D_1 , and D_2 . The detection system we design for this scenario is shown as the combination of Block A and Block C in Fig. 4.

In our system illustrated in Fig. 4 (Block A & Block C), we designate camera S_2 as the reference camera to generate two disparity maps. Specifically, we set the image taken by

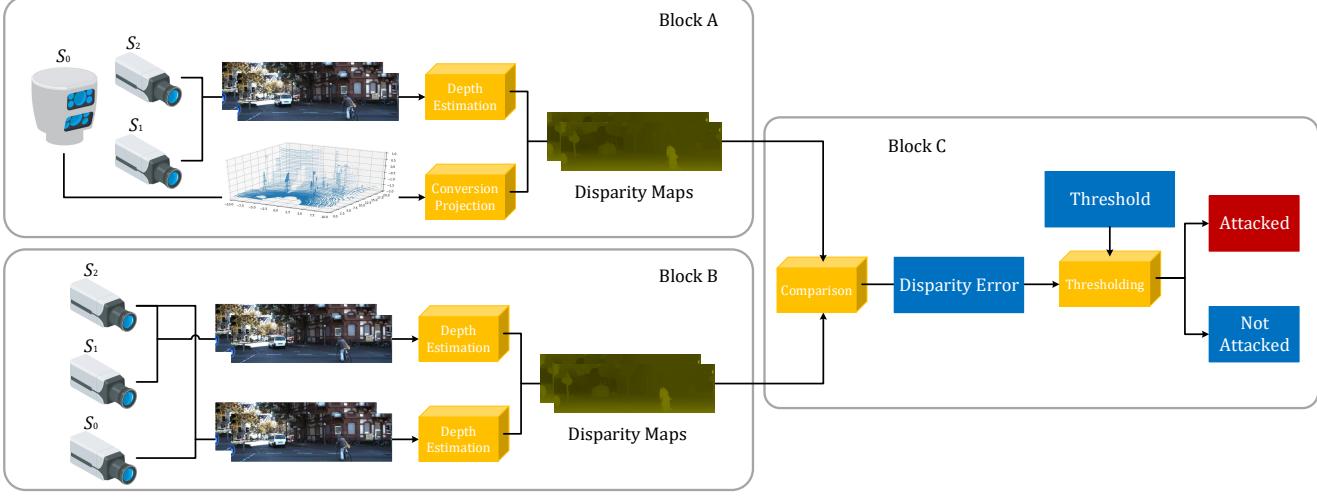


Fig. 4. The detection method is designed for two three-sensor systems. For Scenario 1, the detection method structure involves Block A and Block C. For Scenario 2, the detection method structure involves Block B and Block C.

the reference camera (i.e., D_2) as the reference image, and then feed it with the image taken by the other camera (i.e., D_1) to a depth estimation model to produce the first disparity map, denoted as $DM_{1,2}$, in which we include the disparity information at each pixel of the reference image.

Here, we note that many algorithms have been developed in recent years that can generate accurate disparity maps with camera images. For instance, PSMNet [12] and DORN [17] are two recent algorithms based on deep learning. In this paper, we use the former one to produce disparity maps, since the PSMNet model gives more accurate results over others.

In addition to $DM_{1,2}$, we also project the depth information (i.e., point cloud D_0) obtained by LiDAR onto the reference image D_2 to generate the second disparity map, denoted as $DM_{0,2}$. In this procedure, the depth information in the point cloud D_0 is converted to disparities by using Eqn. (1). To generate all disparity maps in the same scale, f in the equation is set to be the same as the focal length of the cameras, and b is set to be equal to the baseline of the stereo vision formed by S_1 and S_2 . Then, disparities are projected onto D_2 .

To detect optical attacks, we compare the two disparity maps: $DM_{1,2}$ and $DM_{0,2}$. Since in the two procedures described above we use D_2 as the reference image, the two produced disparity maps have the same scale and share the same view. Thus, we can compare them directly. Here, it shall be noted that $DM_{0,2}$ contains sparse disparity information, since the distances in the point cloud D_0 are not densely measured. Therefore, in this comparison procedure, we only compare pixels that have valid disparity in $DM_{0,2}$. For valid pixels, we take the KITTI stereo benchmark [23] as a reference and design our own standard, in which a disparity inconsistency for pixel a is counted if and only if

$$\begin{cases} |DM_{0,2}(a) - DM_{1,2}(a)| > 3, \\ \frac{|DM_{0,2}(a) - DM_{1,2}(a)|}{\min(DM_{0,2}(a), DM_{1,2}(a))} > 0.05. \end{cases} \quad (2)$$

Based on the *pixel-level* disparity inconsistencies, we evaluate the disparity error, denoted as $E_{0,1,2}$, between $DM_{0,2}$ and $DM_{1,2}$. In particular, the *disparity error* is defined as the ratio

of the quantity of *pixel-level* disparity inconsistencies over the total number of valid pixels.

Finally, to detect an optical attack, we need to evaluate the ranges of disparity errors in normal cases and attack cases. We believe that the two ranges are distinguishable and we can then use a threshold $\theta_{0,1,2}$ to determine whether there is an optical attack. In particular, we determine that one of the three sensors is under attack if and only if $E_{0,1,2} > \theta_{0,1,2}$.

The threshold $\theta_{0,1,2}$ is determined offline based on the value distribution of $E_{0,1,2}$ when the three-sensor system is in a safe environment, since only the correct data of optical sensors are available on an autonomous vehicle in normal conditions. According to the statistical law, we use a large number of samples of the disparity error to represent its real distribution in normal cases and define a designated false alarm rate r to arbitrarily set $r \times 100\%$ of them as virtual outliers, where $r \in [0, 1]$. Then, the critical value separating inliers from outliers is $\theta_{0,1,2}$:

$$\frac{\# \text{ samples of } E_{0,1,2} > \theta_{0,1,2}}{\# \text{ samples of } E_{0,1,2}} = r. \quad (3)$$

In this manner, the threshold, which only moves within the value range of disparity error samples, is determined by the value of r . Intuitively, to obtain the best detection performance, we should maximize the detection rate and minimize the designated false alarm rate. Hence, we show how the detection rate varies when adjusting the threshold via r .

C. Experiments for Scenario 1

1) *Setup*: To validate the hypothesis for this scenario, we conduct extensive experiments. We consider all possible attack cases where any sensor or any combination of the three sensors gets attacked. We use the data of one LiDAR and two cameras from the customized KITTI raw dataset [11] to produce affected sensor data for each attack case. The production scheme is described in Section III-C1. The PSMNet model used in the experiments is provided by Wang *et al.* [24], which is trained on Scene Flow dataset [25] and KITTI object

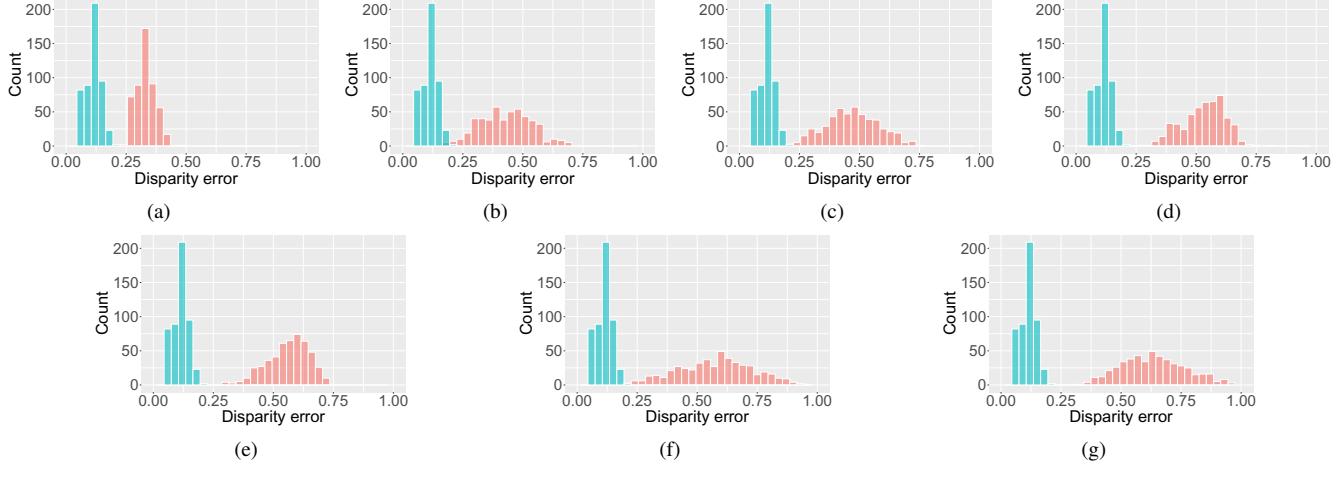


Fig. 5. Distributions of disparity error $E_{0,1,2}$ in normal case (cyan bars) and in attack cases (red bars) for Scenario 1. (a) No attack vs. S_0 attacked; (b) No attack vs. S_1 attacked; (c) No attack vs. S_2 attacked; (d) No attack vs. S_0, S_1 attacked; (e) No attack vs. S_0, S_2 attacked; (f) No attack vs. S_1, S_2 attacked; (g) No attack vs. S_0, S_1, S_2 attacked.

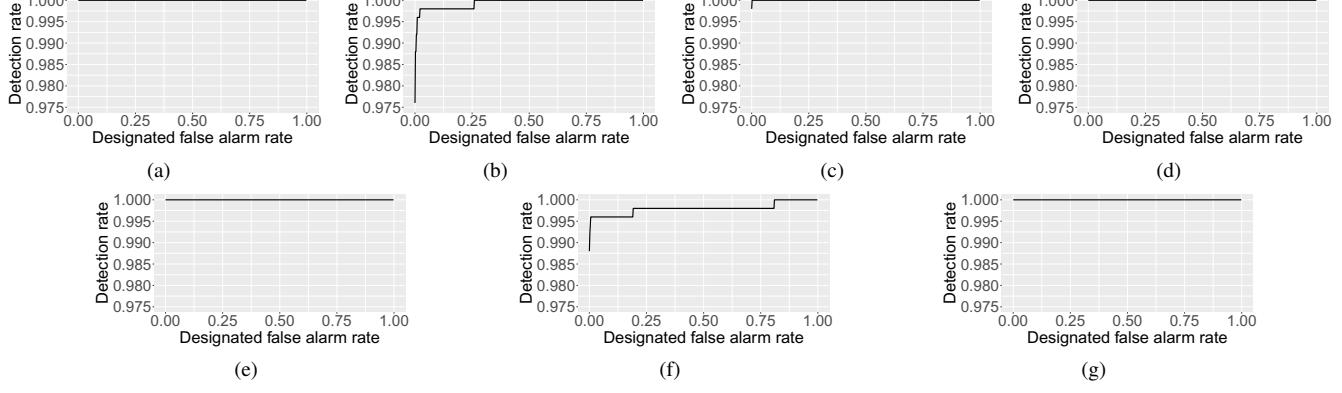


Fig. 6. Detection rate varies with the designated false alarm rate r in each attack case for Scenario 1. (a) S_0 is attacked; (b) S_1 is attacked; (c) S_2 is attacked; (d) S_0 and S_1 are attacked; (e) S_0 and S_2 are attacked; (f) S_1 and S_2 are attacked; (g) S_0, S_1 and S_2 are attacked.

TABLE I
DETECTION RATE COMPARISON BETWEEN THE METHOD IN OUR
FRAMEWORK AND THE BASELINE

Method	Granularity	Avg. Detection Rate	
		Scenario 1	Scenario 2
Ours ($r = 0\%$)	Pixel-level	99.46%	99.94%
Ours ($r = 1\%$)	Pixel-level	99.89%	99.97%
Ours ($r = 2\%$)	Pixel-level	99.89%	100%
Ours ($r = 3\%$)	Pixel-level	99.91%	100%
Ours ($r = 5\%$)	Pixel-level	99.91%	100%
Baseline (IoU = 0.5)	Object-level	65.32%	67.17%
Baseline (IoU = 0.7)	Object-level	59.54%	63.39%

detection dataset [10]. As for metrics, we measure the disparity error distribution and the rate of correct detection for each attack case.

In the literature, there is no existing solution for optical attack detection. Therefore, to compare our scheme with possible solutions, we implement a possible baseline solution to optical attack detection that first extracts *object-level* features from the data of two individual sensors respectively, and then detects attacks by measuring the mismatches between the two sets of features. Specifically for Scenario 1, we implement the baseline with the backbone of PIXOR [19] for extracting *object-level* features from point clouds and the backbone of

Faster R-CNN [21] for extracting from images. We set IoU to 0.5 and 0.7 for determining feature mismatches.

2) *Results*: The experimental results are shown in Fig. 5 and Fig. 6. In Fig. 5, we compare the distributions of disparity errors between the normal case (i.e., no attack, in cyan bars) with each one of the seven possible attack cases (red bars). We first observe that the disparity errors are smaller than 20% in most normal cases. By comparison, the disparity errors in most attack scenarios are larger than 25%. These results indicate that our detection scheme is sensitive enough, so that there is almost no overlap between the distribution in the normal case and the distributions in those attack cases.

In Fig. 6, we adjust the threshold $\theta_{0,1,2}$ used to declare attacks by varying the designated false alarm rate r , and evaluate the attack detection rate versus r . As we can see from the figures, among seven attack scenarios, the performance is perfect in five cases, where the detection rate hits 100% for all possible values of r . And even in the non-perfect cases (i.e., (b) and (f)), the proposed detection system can obtain more than 99.5% detection rate with less than 5% false alarm. Such results confirm our hypothesis for such a three-sensor system.

We also show the detection rate comparison between our proposed detection method ($r \leq 5\%$) and the baseline for

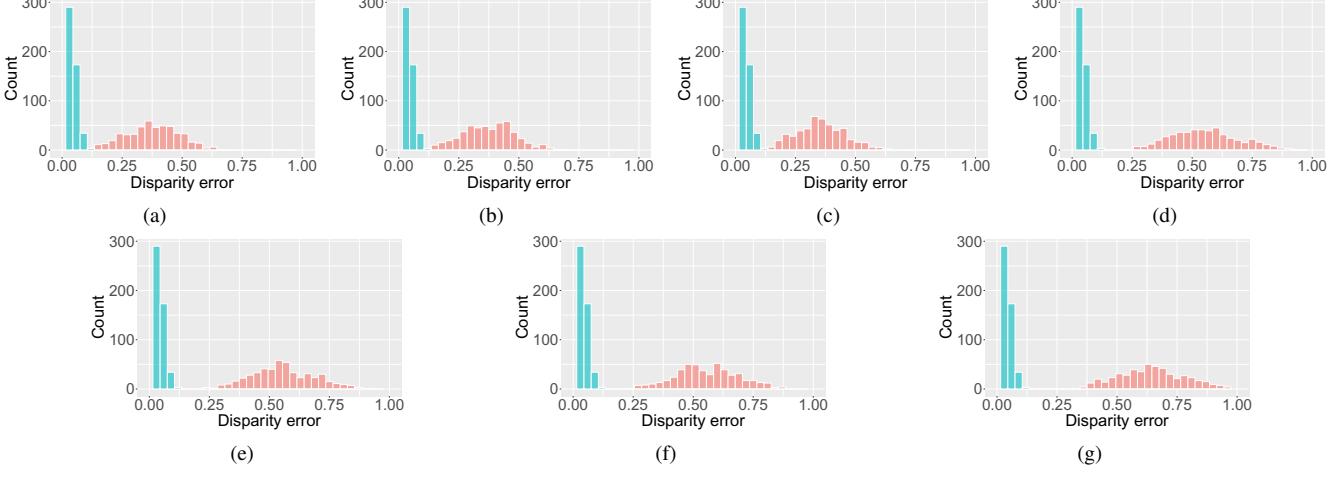


Fig. 7. Distributions of disparity error $E_{0,1,2}$ in normal case (cyan bars) and in attack cases (red bars) for Scenario 2. (a) No attack vs. S_0 attacked; (b) No attack vs. S_1 attacked; (c) No attack vs. S_2 attacked; (d) No attack vs. S_0, S_1 attacked; (e) No attack vs. S_0, S_2 attacked; (f) No attack vs. S_1, S_2 attacked; (g) No attack vs. S_0, S_1, S_2 attacked.

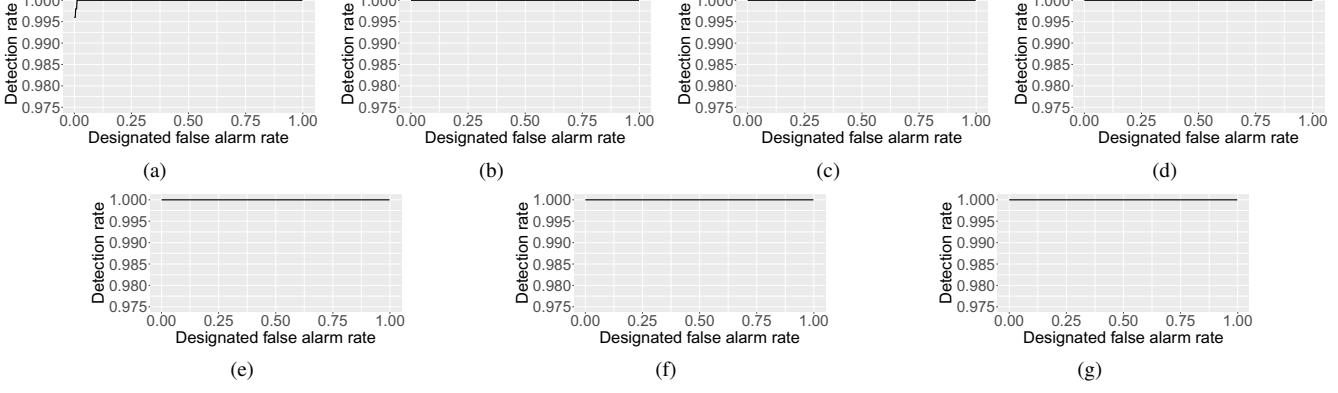


Fig. 8. Detection rate varies with the designated false alarm rate r in each attack case for Scenario 2. (a) S_0 is attacked; (b) S_1 is attacked; (c) S_2 is attacked; (d) S_0 and S_1 are attacked; (e) S_0 and S_2 are attacked; (f) S_1 and S_2 are attacked; (g) S_0, S_1 and S_2 are attacked.

Scenario 1 in Table I, where we can observe that our method outperforms the baseline by about 35% with $\text{IoU} = 0.5$ and about 40% with $\text{IoU} = 0.7$. The reason is that our method detects optical attack by measuring *pixel-level* disparity inconsistencies, which is much denser and more fine-grained than the *object-level* detection used in the baseline. In most cases where the attack only partially occludes important objects or does not occlude them at all, the *object-level* detection is highly likely to fail, while our proposed method can still function normally.

D. Scenario 2: Three Cameras

For this scenario, we consider a three-sensor system that consists of three cameras, denoted as S_0 , S_1 , and S_2 , from the right to the left. Similar to the previous scenario, we also consider the data generated by the sensors as D_0 , D_1 , and D_2 . The detection system that we design for this scenario is shown as the combination of Block B and Block C in Fig. 4.

In the system illustrated in Fig. 4 (Block B & Block C), we designate camera S_2 as the reference camera to generate two disparity maps. In our experiments, we find it more convenient to implement our detection scheme when the leftmost or rightmost camera is used as the reference camera.

Since the sensor data of this scenario are all images, to generate the two disparity maps, we feed D_2 with D_0 and D_1 to the depth estimation model, respectively. It shall be noted that, since the distance between S_0 and S_2 is usually different from the distance between S_1 and S_2 , we need to adjust the disparity in $DM_{0,2}$ by updating the baseline b accordingly. After the disparity maps $DM_{0,2}$ and $DM_{1,2}$ are generated, the rest of the procedures in the detection method are the same as those in the previous scenario.

E. Experiments for Scenario 2

1) *Setup*: Here, we use the data of three cameras in the customized KITTI raw dataset [11] and also consider all possible attack cases. The baseline method is implemented with the backbone of Faster R-CNN [21]. The rest of settings are the same as those in the experiments for Scenario 1.

2) *Results*: We show the experiment results in Fig. 7 and Fig. 8. In Fig. 7, we compare the distributions of disparity errors between the normal case (i.e., no attack, in cyan bars) with each one of the seven possible attack cases (red bars). Similar to Scenario 1, we observe that the disparity errors are less than 10% in virtually all normal cases. By comparison, the disparity errors in 99% of attack scenarios are larger than 12.5%.

In Fig. 8, we vary the threshold $\theta_{0,1,2}$ used to declare attacks via r , and evaluate the attack detection rate versus the designated false alarm rate. We can observe that the detection performance is perfect in all cases, where the detection rate remains 100% when r varies from 0 to 1. Comparing the results in Scenario 2 with results in Scenarios 1, we notice that the detection performance in scenario 2 is slightly better. We believe this is due to the facts that the disparity maps in this scenario are generated using the same method and there are more valid pixels in the comparison.

In Table I, the performance comparison for this scenario still shows that our proposed detection method outperforms the baseline by a large margin (more than 30%), which again shows the merit of *pixel-level* detection.

F. Empirical Findings

To briefly summarize, the findings from the attack detection experiments for the two three-sensor systems are listed as follows:

- The experimental results confirm our hypothesis that there exists a detection system that can detect optical attacks on the two three-sensor systems with high accuracy and low false alarm rate.
- The detection rate is insensitive to the designated false alarm rate. As long as the detection rate is maintained at a high level, the designated false alarm rate should be set as low as possible, empirically less than 5%.
- In those two three-sensor systems, any sensor or any combination of sensors being attacked can cause the disparity error beyond the threshold.

Based on these findings, we further develop the identification approach for the proposed framework.

V. ATTACK IDENTIFICATION

In this section, we present the second procedure of our framework which identifies the compromised sensors in a system with one LiDAR S_0 and n cameras, namely, S_1 to S_n from the right to the left, where $n \geq 3$, based on the empirical findings from the detection method. This method is inspired by Error Correction Codes (ECC) and can achieve the identification as long as there are no more than $n - 2$ sensors being attacked simultaneously. In addition, we also demonstrate the proof of the correctness of our identification method, as well as show its effectiveness and accuracy via experiments.

We now introduce a few definitions that are used in the rest of this section. For every sensor S_i , its state s_i can switch between normal state and attack state

$$s_i = \begin{cases} 1, & \text{if } S_i \text{ is attacked,} \\ 0, & \text{otherwise,} \end{cases} \quad (4)$$

where $i \in \{0, 1, \dots, n\}$, $n \geq 3$. The sensor state vector in the system is defined as:

$$\mathbf{s} := [s_0, s_1, \dots, s_n], \quad (5)$$

which is the hidden ground truth that we try to identify.

For disparity error $E_{i,j,k}$ among sensors S_i , S_j , and S_k , we use $e_{i,j,k}$ to indicate whether $E_{i,j,k}$ exceeds the corresponding threshold $\theta_{i,j,k}$ in which

$$e_{i,j,k} = \begin{cases} 1, & \text{if } E_{i,j,k} > \theta_{i,j,k}, \\ 0, & \text{otherwise,} \end{cases} \quad (6)$$

where $i, j, k \in \{0, 1, \dots, n\}$, $i < j < k$. And similarly, we use the disparity error state vector \mathbf{e} to represent the states of disparity errors of three-sensor combinations in the system.

Since the system consists of one LiDAR and n cameras, the combination of any three sensors from it must be either one LiDAR with two cameras or three cameras. According to the empirical finding drawn from the experiments in the previous section, any sensor or any combination of sensors from such three-sensor sets being compromised leads to the corresponding disparity error higher than the threshold. Therefore, based on the definitions in Eqn. (4) and Eqn. (6), we have

$$e_{i,j,k} = (s_i \vee s_j \vee s_k), \quad (7)$$

where $i, j, k \in \{0, 1, \dots, n\}$, $i < j < k$, and \vee is logical OR operation.

A. Calculation of Disparity Error State Vector

Given $n + 1$ sensors, we use the leftmost sensor S_n as the reference camera and calculate disparity error state vector

$$\mathbf{e} := [e_{0,1,n}, e_{0,2,n}, \dots, e_{n-2,n-1,n}], \quad (8)$$

where $n \geq 3$. In the calculation, the disparity maps generated by S_n with every remaining sensor are compared with each other using the same standard described by Eqn. (2). Then, by the definition in Eqn. (6), \mathbf{e} is obtained via thresholding the resulted disparity errors from comparison. We also show this calculation process in the form of pseudocode in Algorithm 1 that takes the data of one LiDAR and n cameras and a list of thresholds as inputs and outputs the disparity error state vector \mathbf{e} . Specifically, it first generates n disparity maps with the sensor data and compares each two of them to obtain disparity errors, and then calculates \mathbf{e} by encoding disparity errors with thresholds.

Note that the thresholds for calculating \mathbf{e} are also determined offline using one designated false alarm rate r . The approach is similar as in the detection procedure. For every disparity error, we collect sufficient samples when the system is safe and consider $r \times 100\%$ of samples as virtual outliers. The thresholds are then set to the maximal values of inliers. Hence, Eqn. (3) can be rewritten as:

$$\frac{\# \text{ samples of } E_{i,j,n} > \theta_{i,j,n}}{\# \text{ samples of } E_{i,j,n}} = r, \quad (9)$$

where $r \in [0, 1]$, $i, j \in \{0, 1, \dots, n - 1\}$, $i < j$. Unlike the detection rate which is insensitive to r , our subsequent experiments indicate that the identification rate drops linearly with r increasing and the best identifying performance is achieved when $r = 1\%$.

Algorithm 1: Calculation of disparity error state vector \mathbf{e}

Input: \mathbf{D} : a list of sensor data with length equal to $|\mathbf{D}|$, where D_0 is point cloud and the rest are images; θ : a list of thresholds.

Output: \mathbf{e} : the disparity error state vector.

- 1 select D_n as reference image, where $n = |\mathbf{D}| - 1$;
- 2 get disparity map $DM_{0,n}$ using D_0 and D_n (Scenario 1 in Section IV);
- 3 **for** $i \leftarrow 1$ to $n - 1$ **do**
- 4 get disparity map $DM_{i,n}$ using D_i and D_n (Scenario 2 in Section IV);
- 5 **for** $i \leftarrow 0$ to $n - 2$ **do**
- 6 **for** $j \leftarrow i + 1$ to $n - 1$ **do**
- 7 get disparity error $E_{i,j,n}$ by comparing $DM_{i,n}$ with $DM_{j,n}$ (Section IV);
- 8 **if** $E_{i,j,n}$ exceeds threshold $\theta_{i,j,n}$ **then**
- 9 assign 1 to disparity error state $e_{i,j,n}$;
- 10 **else**
- 11 assign 0 to $e_{i,j,n}$;
- 12 **return** \mathbf{e} ;

B. Identification of Sensor State Vector

We now elaborate on how to infer \mathbf{s} according to \mathbf{e} . We consider the following three cases of \mathbf{e} :

- If all elements in \mathbf{e} are 0s, according to Eqn. (7) and Eqn. (8), $s_i = 0$ for $0 \leq i \leq n$. In other words, no sensor is attacked.
- If only some elements in \mathbf{e} are 0s, we have the following Lemma 1 to identify all attacked sensors.
- If all elements in \mathbf{e} are 1s and no more than $n - 2$ sensors are attacked simultaneously, we can repeatedly use Lemma 1 and Lemma 2 to identify all attacked sensors.

Lemma 1: In a system with $n + 1$ sensors, if there exist i_0, j_0 such that $e_{i_0,j_0,n} = 0$, then $s_i = e_{i,i_0,n}$ for $0 \leq i < i_0$, and $s_i = e_{i_0,i-1}$ for $i_0 < i < n$.

Proof: If there exist i_0, j_0 such that $e_{i_0,j_0,n} = 0$, according to Eqn. (7), we can have:

$$e_{i_0,j_0,n} = (s_{i_0} \vee s_{j_0} \vee s_n) = 0 \quad (10)$$

which implies

$$s_{i_0} = s_{j_0} = s_n = 0 \quad (11)$$

Now we only need to focus on the state of s_0, s_1, \dots, s_{n-1} .

For any $i \in \{0, 1, \dots, i_0 - 1\}$,

$$e_{i,i_0-1} = (s_i \vee s_{i_0} \vee s_n) = (s_i \vee 0 \vee 0) = s_i, \quad (12)$$

and for any $i \in \{i_0 + 1, i_0 + 2, \dots, n - 1\}$,

$$e_{i_0,i,n} = (s_{i_0} \vee s_i \vee s_n) = (0 \vee s_i \vee 0) = s_i. \quad (13)$$

Algorithm 2: Identification of sensor state vector \mathbf{s}

Input: \mathbf{D} : a list of sensor data with length equal to $|\mathbf{D}|$, where D_0 is point cloud and the rest are images; θ : a list of thresholds.

Output: \mathbf{A} : the list of compromised sensors.

- 1 select D_n as reference image, where $n = |\mathbf{D}| - 1$;
- 2 calculate \mathbf{e} using Algorithm 1 with \mathbf{D} and θ ;
- 3 **if** there exists 1 in \mathbf{e} **then**
- 4 **if** there exists 0 in \mathbf{e} **then**
- 5 // use Lemma 1 to infer \mathbf{s}
- 6 find i_0, j_0 which satisfy $e_{i_0,j_0,n} = 0$;
- 7 **for** $i \leftarrow 0$ to $n - 1$ **do**
- 8 assign $\min(i, i_0)$ to k_1 , $\max(i, i_0)$ to k_2 ;
- 9 **if** $e_{k_1,k_2,n} = 1$ and $k_1 \neq k_2$ **then**
- 10 sensor state s_i is 1;
- 11 push attacked sensor S_i into \mathbf{A} ;
- 12 **else**
- 13 // use Lemma 2 to infer \mathbf{s}
- 14 sensor state s_n is 1;
- 15 push attacked sensor S_n into \mathbf{A} ;
- 16 // infer rest sensors recursively
- 17 **if** $n > 3$ **then**
- 18 remove D_n from \mathbf{D} ;
- 19 go to line 1 to rerun with updated \mathbf{D} ;
- 20 **return** \mathbf{A} ;

Lemma 2: In a system with $n + 1$ sensors, when there are no more than $n - 2$ sensors being compromised simultaneously, if the elements of \mathbf{e} are all 1s, then $s_n = 1$.

Proof: Since there are no more than $n - 2$ sensors being attacked, the states of at least three sensors are 0s. If S_n is normal, then there exist i^* and j^* , where $i^* < j^*$, such that sensors S_{i^*} and S_{j^*} are normal, i.e., $s_{i^*} = s_{j^*} = s_n = 0$. In this case, $e_{i^*,j^*,n} = 0$, which contradicts the fact that all the elements of \mathbf{e} are 1s. Therefore, $s_n = 1$. ■

When all elements of \mathbf{e} in a system with $n + 1$ sensors are 1, though we cannot directly find out the states of all sensors, Lemma 2 can identify the last sensor's state. After that, we can virtually remove the last sensor and consider a system with n sensors. We recalculate the \mathbf{e} by Algorithm 1 for such n sensors, then determine \mathbf{s} according to Lemma 1 and Lemma 2. We repeat this process until the states of all sensors are identified. We also present this identification algorithm as pseudocode in Algorithm 2 that takes the same inputs as Algorithm 1 and outputs a list of attacked sensors. Specifically, it begins with calculating \mathbf{e} using Algorithm 1 with the inputs, and then infers the sensor state vector \mathbf{s} using Lemma 1 when some elements in \mathbf{e} are 0s. When all elements in \mathbf{e} are 1s, it first infers s_n , the state of the reference camera, using Lemma 2, and then excludes the data of the reference camera from the inputs and infers the rest of sensor states by rerunning Algorithm 2 with the updated inputs.

C. Experiments

- 1) **Setup:** To verify the effectiveness and evaluate the precision of our identification scheme, we conduct substantial

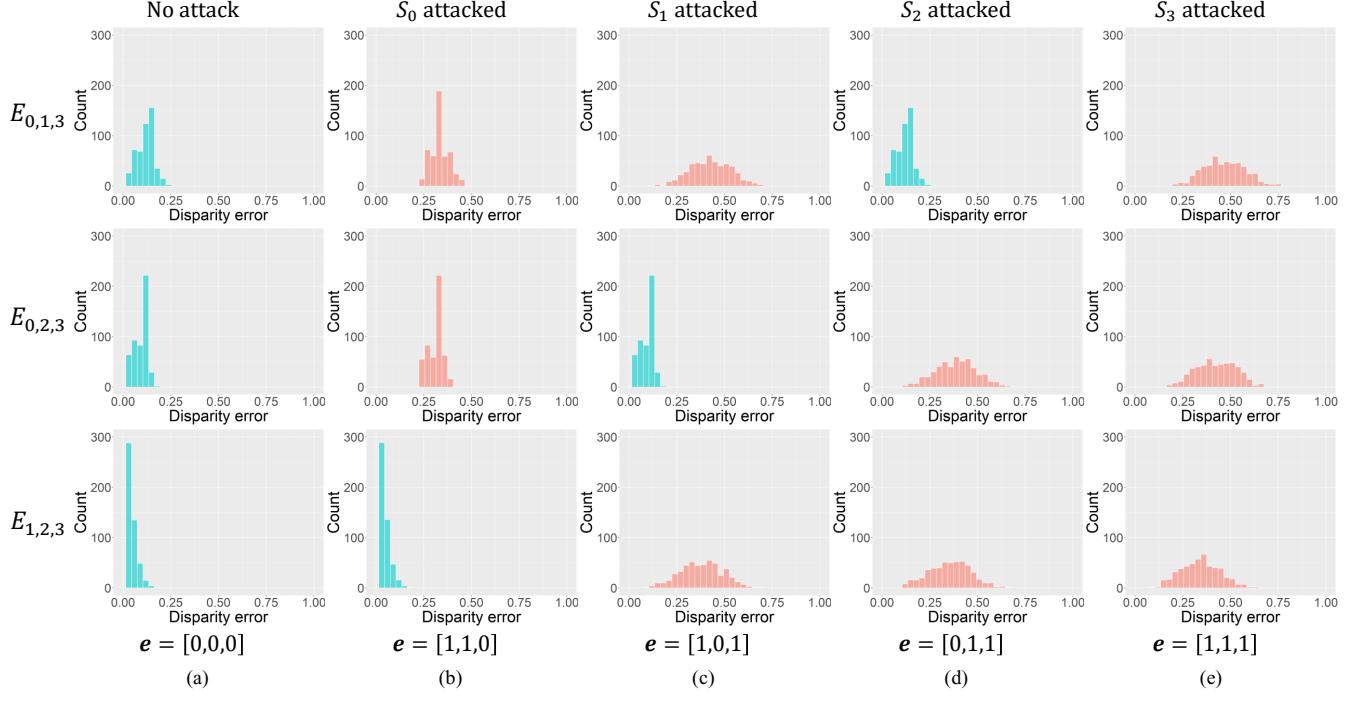


Fig. 9. Distributions of disparity error $E_{0,1,3}$ (1st row), disparity error $E_{0,2,3}$ (2nd row) and disparity error $E_{1,2,3}$ (3rd row) in normal scenario (1st column) and in attack scenarios (2nd–5th columns). Cyan bars mean that the disparity errors in a scenario involve no attacked sensor, while red bars indicate that the disparity errors in a scenario involve attacked sensor. (a) No disparity errors exceed thresholds, so $e = [0, 0, 0]$ indicating no optical attack; (b) $E_{0,1,3}$ and $E_{0,2,3}$ exceed thresholds, so $e = [1, 1, 0]$ indicating that S_0 is attacked; (c) $E_{0,1,3}$ and $E_{1,2,3}$ exceed thresholds, so $e = [1, 0, 1]$ indicating that S_1 is attacked; (d) $E_{0,2,3}$ and $E_{1,2,3}$ exceed thresholds, so $e = [0, 1, 1]$ indicating that S_2 is attacked; (e) All three disparity errors exceed thresholds, so $e = [1, 1, 1]$ indicating that S_3 is attacked.

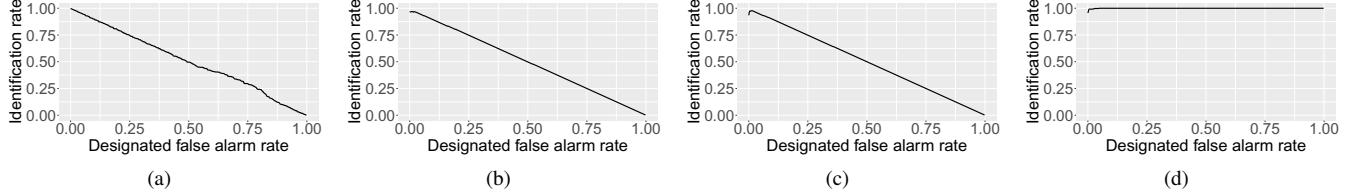


Fig. 10. Identification rate varies with the designated false alarm rate r in each attack scenario. (a) S_0 is attacked; (b) S_1 is attacked; (c) S_2 is attacked; (d) S_3 is attacked.

TABLE II
IDENTIFICATION RATE AT PARTICULAR VALUES OF THE DESIGNATED
FALSE ALARM RATE IN ATTACK SCENARIOS

r	Attacked Sensor			Average	
	S_0	S_1	S_2		
0%	99.80%	96.40%	93.40%	95.60%	96.30%
1%	98.80%	97.00%	97.60%	99.20%	98.15%
2%	98.00%	96.80%	97.20%	99.20%	97.80%
3%	96.60%	96.00%	96.20%	99.80%	97.15%
5%	94.80%	94.00%	94.20%	100%	95.75%

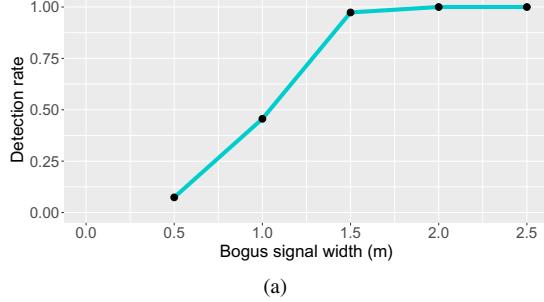
experiments. Since the identification scheme is designed to be in a form of recursion, the experiments here are for the base case where a system consists of one LiDAR (S_0) and $n = 3$ cameras (S_1 , S_2 , and S_3). And according to the constraint, there is at most one attacked sensor in the system.

In the experiments, we consider the normal scenario and all possible attack scenarios where each one of the four sensors gets attacked. We use the data of one LiDAR and three cameras from the customized KITTI raw dataset [11]. To generate the compromised sensor data for each attack scenario, we do the same as in Section III-C1. The depth estimation model is

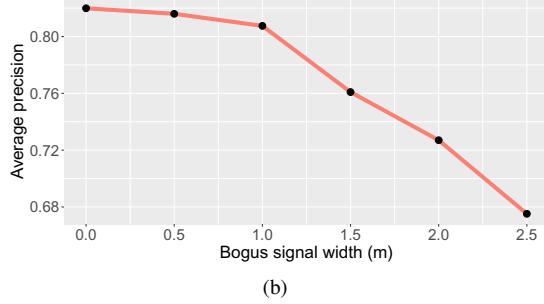
provided by [24]. As for metrics, we measure the disparity error distribution and the rate of correct identification for each attack scenario.

2) *Results:* We present the experimental results in Fig. 9, Fig. 10, and Table II. In Fig. 9, the three rows of sub-figures represent the distributions of disparity error $E_{0,1,3}$ (1st row), disparity error $E_{0,2,3}$ (2nd row), and disparity error $E_{1,2,3}$ (3rd row), respectively. And the columns in Fig. 9 represent five scenarios.

In Fig. 9, we can first compare the disparity errors in each row. Similar to the results in the last section, the results clearly illustrate that the distribution of disparity errors can help to detect whether there is an attack. For example, in the first row for $E_{0,1,3}$, when there is no attack on S_0 , S_1 , and S_3 , the disparity errors (cyan bars) are mostly less than 23%. By comparison, the disparity errors (red bars) are larger than 23% when any of the sensors is attacked. Such results affirm the feasibility and correctness of defining attacks according to the disparity error state. Three sub-figures in each of the five attack scenarios clearly show that there is



(a)



(b)

Fig. 11. Sensitivity of our framework and AP of PointRCNN when the attack is against LiDAR. (a) Detection rate for attacks against LiDAR varies with the width of bogus signal; (b) Average precision of PointRCNN varies with the width of bogus signal.

a unique pattern of the combination of $E_{0,1,3}$, $E_{0,2,3}$, and $E_{1,2,3}$ for each attack scenario. For instance, when there is no attack launched, the three disparity errors are all within certain boundaries, representing the disparity error state vector $e = [0, 0, 0]$. And if any one of the sensors is attacked, the disparity errors involving that sensor will exceed boundaries, leading to another unique e .

In Fig. 10, we vary the thresholds by adjusting the designated false alarm rate r and compute the corresponding identification rate in four attack scenarios. It is obvious that, if the attacked sensor is not the reference, then the identification rate drops linearly when r increases from 0 to 1. On the other hand, when the reference sensor S_3 is attacked, then the identification rate remains very close to 100%.

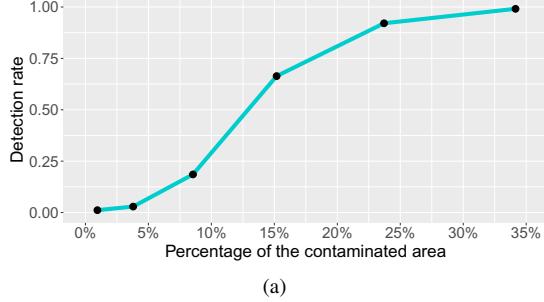
Such observations suggest that choosing a small r may lead to the best performance overall. To identify the best r , we conduct some experiments to investigate the impact of r , when it is within the range of 0 to 5%. The numerical results are shown in Table II. As we can see, the best average identification rate for the four attack scenarios occurs at $r = 1\%$.

D. Discussion

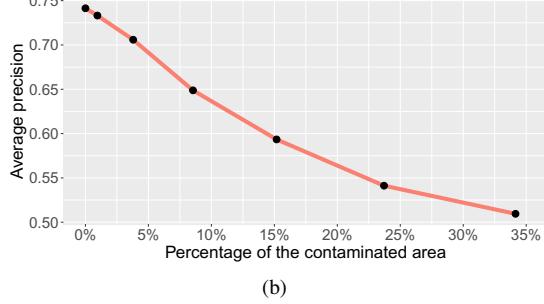
Though the identification method of our framework can accurately identify attacked sensors, it is limited to the condition where no more than $n - 2$ sensors are attacked simultaneously in a system with $n + 1$ sensors. We plan to address this limitation through cross-vehicle sensing data validation in our future studies.

VI. FRAMEWORK SENSITIVITY

With the best designated false alarm rate r determined, we conduct further experiments to investigate how sensitive our



(a)



(b)

Fig. 12. Sensitivity of our framework and AP of YOLOv3 when the attack is against camera. (a) Detection rate for attacks against camera varies with the percentage of the area contaminated by facula; (b) Average precision of YOLOv3 varies with the percentage of the area contaminated by facula.

framework is, namely, for optical attacks with what range of settings (width of bogus signal, size of facula) our framework works effectively. Empirically, the milder optical attacks are, the more difficult they get detected. Meanwhile, we also measure how much the perception function of AVs is influenced by the optical attacks with different settings using state-of-the-art object detection algorithms. Those algorithms usually possess a certain degree of resistance to minor optical attacks, so our framework does not have to be universally sensitive.

In this section, we use experiments to demonstrate that our framework has excellent sensitivity to the attacks on LiDAR with settings that object detection model cannot overcome. As for the attacks on camera, our framework is also sensitive in most cases, but shows limit when the attack is too mild. The experiments consist of two parts: the first part is for the sensitivity to the attacks on LiDAR, and the second part for the sensitivity to the attacks on camera.

A. Metrics

To measure the sensitivity, we use the detection rate of our framework with $r = 1\%$, since detection rate can also reflect the performance of identification procedure. As described in Section V, the identification procedure of our framework is directly developed upon multiple detection processes, so the identification rate is highly correlated with the detection rate.

As for evaluating the corresponding performance of object detection algorithms used in AVs, we follow the KITTI object detection benchmark [10] and calculate the average precision of vehicle detection with IoU threshold equal to 70%.

B. Experimental Setup

We conduct our experiments on the dataset provided in the KITTI object detection benchmark [10] which contains sensor data of one LiDAR and two cameras. As described in

Section III-C1, we divide the labeled part of the dataset into training set and test set. The training set is used to train the object detection models, while the test set is for generating compromised sensor data.

To find out the sensitivity of our framework to the optical attacks on LiDAR, we produce five affected point clouds for every set of sensor data in the test set. The five affected point clouds contain a bogus signal with a width of 0.5 meter, 1.0 meter, 1.5 meters, 2.0 meters, and 2.5 meters, respectively. And the object detection algorithm chosen for this part of experiments is PointRCNN [26], a state-of-the-art 3D object detection algorithm that takes a point cloud as input.

In term of the experiment setup for evaluating sensitivity to the attacks on camera, we generate six pairs of compromised stereo pictures for each set of sensor data in the test set. The left pictures of the six pairs are overlaid with a Gaussian facula with radius of 37.5 pixels, 75 pixels, 112.5 pixels, 150 pixels, 187.5 pixels, and 225 pixels, respectively. And the corresponding percentages of the contaminated area in images are 0.95%, 3.79%, 8.54%, 15.18%, 23.71%, and 34.15%. The object detection algorithm for this part of experiments is YOLOv3 [27], which is one of the most popular real-time object detection algorithms using images as input.

In the experiments, we feed the compromised sensor data to our framework and the selected object detection model, and then evaluate them via the aforementioned metrics. The PSMNet model used in the framework is provided by [24].

C. Experiment Results

1) *Sensitivity to the Attacks on LiDAR*: As shown in Fig. 11, with the increase of the width of the bogus signal, the detection rate of our framework surges, while the average precision of PointRCNN declines. The AP of PointRCNN decreases very slightly when the width of the bogus signal is within 1.0 meter, which implies that PointRCNN exhibits some resistance to minor disturbing signals. On the other hand, when the size of the bogus signal is larger, the average precision starts dropping rapidly. However, it should be noted that when the width is at 1.5 meters, the total decline of AP is only 0.06, while the detection rate of our framework already reaches nearly 100%. These results show clearly that our framework is highly sensitive to the attacks on LiDAR that cannot be mitigated by object detection algorithms.

2) *Sensitivity to the Attacks on Camera*: The experiment results for this part are illustrated in Fig. 12. The tendencies of the detection rate of our framework and the AP of object detection model are similar as those in the first part of the experiments. Particularly, when the percentage of the contaminated area is within the range of 5% to 15%, although our framework has a small detection rate, the AP of YOLOv3 maintains at a high level, which means that the perception has not been compromised due to small attacks. When the percentage of the contaminated area in images increases to 23.71%, the detection rate of our framework surpasses 90%. In the meantime, the AP of YOLOv3 drops significantly to 54.12%. These results suggest that our framework has a strong sensitivity to the attacks on camera when the contaminated area in images is greater than 20%. Less than that, the framework may show some limitations.

VII. CONCLUSION

In this paper, we have systematically investigated the mitigation of attacks on optical devices (LiDAR and camera) that are essential to perform accurate object and event detection and response (OEDR) tasks in autonomous driving systems. Specifically, we proposed a framework to detect and identify sensors that are under attack. For the attack detection, we considered two common three-sensor systems, (1) one LiDAR and two cameras, and (2) three cameras, and we developed effective procedures to detect any attack on each of them. Using real datasets and the state-of-the-art machine learning model, we conducted extensive experiments confirming that our detection scheme can detect attacks with high accuracy and a low false alarm rate. Based on the detection models, we further developed an identification model that is capable of identifying up to $n - 2$ attacked sensors in a system with one LiDAR and n cameras. For the identification procedure, we proved its correctness and used experiments to validate its performance. At last, we investigated the sensitivity of our framework and showed its excellence in this aspect.

In the future, we plan to study methods to further identify the damaged portion in image and point cloud and perform data recovery for the damaged portion using intact data from other sensors of the ego-vehicle, nearby infrastructure, or vehicles in vicinity. We also plan to investigate the identification task for autonomous driving system based on Multi-Access Edge Computing (MEC) and 5G in which sensor data of *multiple* vehicles can be exploited to overcome the limitation of our proposed identification solution.

REFERENCES

- [1] SAE J3016, “Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles,” 2018. [Online]. Available: https://www.sae.org/standards/content/j3016_201806/
- [2] A. Davies, “Waymo’s So-Called Robo-Taxi Launch Reveals a Brutal Truth,” 2018. [Online]. Available: <https://www.wired.com/story/waymo-self-driving-taxi-service-launch-chandler-arizona/>
- [3] A. Hawkins, “Teslas full self-driving feature may get early-access release by the end of 2019,” 2019. [Online]. Available: <https://www.theverge.com/2019/10/23/20929529/tesla-full-self-driving-release-2019-beta>
- [4] Apollo Open Platform, “Apollo 5.0 Perception Module,” 2019. [Online]. Available: <http://apollo.auto/platform/perception.html>
- [5] A. M. Wyglinski, X. Huang, T. Padir, L. Lai, T. R. Eisenbarth, and K. Venkatasubramanian, “Security of Autonomous Systems Employing Embedded Computing and Sensors,” *IEEE Micro*, vol. 33, no. 1, pp. 80–86, Jan. 2013.
- [6] K. Ren, Q. Wang, C. Wang, Z. Qin, and X. Lin, “The Security of Autonomous Driving: Threats, Defenses, and Future Directions,” *Proceedings of the IEEE*, vol. 108, no. 2, pp. 357–372, Feb. 2019.
- [7] J. Petit, B. Stottelaar, M. Feiri, and F. Kargl, “Remote attacks on automated vehicles sensors: Experiments on camera and LiDAR,” *Black Hat Europe*, vol. 11, Nov. 2015.
- [8] C. Yan, W. Xu, and J. Liu, “Can You Trust Autonomous Vehicles: Contactless Attacks against Sensors of Self-driving Vehicle,” in *DEFCON 24*, Aug. 2016.
- [9] H. Shin, D. Kim, Y. Kwon, and Y. Kim, “Illusion and Dazzle: Adversarial Optical Channel Exploits Against Lidars for Automotive Applications,” in *International Conference on Cryptographic Hardware and Embedded Systems*. Springer, Cham, Aug. 2017, pp. 445–467.
- [10] A. Geiger, P. Lenz, and R. Urtasun, “Are we ready for autonomous driving? The KITTI vision benchmark suite,” in *Proc. IEEE CVPR 2012*. IEEE, Jun. 2012, pp. 3354–3361.
- [11] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, “Vision meets robotics: The KITTI dataset,” *The International Journal of Robotics Research*, vol. 32, no. 11, pp. 1231–1237, Sep. 2013.
- [12] J.-R. Chang and Y.-S. Chen, “Pyramid Stereo Matching Network,” in *Proc. IEEE CVPR 2018*, Jun. 2018, pp. 5410–5418.

- [13] Y. Cao, C. Xiao, B. Cyr, Y. Zhou, W. Park, S. Rampazzi, Q. A. Chen, K. Fu, and Z. M. Mao, "Adversarial sensor attack on lidar-based perception in autonomous driving," in *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*. ACM, Nov. 2019, pp. 2267–2281.
- [14] M. Rofail, A. Alsafty, M. Matousek, and F. Kargl, "Multi-Modal Deep Learning for Vehicle Sensor Data Abstraction and Attack Detection," in *2019 IEEE International Conference of Vehicular Electronics and Safety (ICVES)*. IEEE, Sep. 2019, pp. 1–7.
- [15] R. Changalvala and H. Malik, "LiDAR Data Integrity Verification for Autonomous Vehicle," *IEEE Access*, vol. 7, pp. 138 018–138 031, Sep. 2019.
- [16] A. Bhoi, "Monocular Depth Estimation: A Survey," *arXiv preprint arXiv:1901.09402*, 2019.
- [17] H. Fu, M. Gong, C. Wang, K. Batmanghelich, and D. Tao, "Deep Ordinal Regression Network for Monocular Depth Estimation," in *Proc. IEEE CVPR 2018*, Jun. 2018, pp. 2002–2011.
- [18] X. Chen, K. Kundu, Y. Zhu, A. G. Berneshawi, H. Ma, S. Fidler, and R. Urtasun, "3D Object Proposals for Accurate Object Class Detection," in *Advances in Neural Information Processing Systems*. Curran Associates, Inc., 2015, pp. 424–432.
- [19] B. Yang, W. Luo, and R. Urtasun, "PIXOR: Real-Time 3D Object Detection From Point Clouds," in *Proc. IEEE CVPR 2018*, Jun. 2018, pp. 7652–7660.
- [20] J. Huang, V. Rathod, C. Sun, M. Zhu, A. Korattikara, A. Fathi, I. Fischer, Z. Wojna, Y. Song, S. Guadarrama, and K. Murphy, "Speed/Accuracy Trade-Offs for Modern Convolutional Object Detectors," in *Proc. IEEE CVPR 2017*, Jul. 2017, pp. 7310–7311.
- [21] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks," in *Proc. Advances in Neural Information Processing Systems 28 (NIPS 2015)*, 2015, pp. 91–99.
- [22] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in *Proc. IEEE CVPR 2016*, Jun. 2016, pp. 770–778.
- [23] M. Menze and A. Geiger, "Object Scene Flow for Autonomous Vehicles," in *Proc. IEEE CVPR 2015*, 2015, pp. 3061–3070.
- [24] Y. Wang, W.-L. Chao, D. Garg, B. Hariharan, M. Campbell, and K. Q. Weinberger, "Pseudo-LiDAR From Visual Depth Estimation: Bridging the Gap in 3D Object Detection for Autonomous Driving," in *Proc. IEEE CVPR 2019*, Jun. 2019, pp. 8445–8453.
- [25] N. Mayer, E. Ilg, P. Hausser, P. Fischer, D. Cremers, A. Dosovitskiy, and T. Brox, "A Large Dataset to Train Convolutional Networks for Disparity, Optical Flow, and Scene Flow Estimation," in *Proc. IEEE CVPR 2016*, Jun. 2016, pp. 4040–4048.
- [26] S. Shi, X. Wang, and H. Li, "PointRCNN: 3D Object Proposal Generation and Detection from Point Cloud," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 770–779.
- [27] J. Redmon and A. Farhadi, "Yolov3: An incremental improvement," *arXiv preprint arXiv:1804.02767*, 2018.



Jindi Zhang received the B.Eng. degree in automation from Tongji University, Shanghai, China, in 2016, and the M.Sc. degree in computer science from the Chinese University of Hong Kong, in 2017. He is currently a Ph.D. student in the Department of Computer Science at City University of Hong Kong. His research interests include multi-sensor application and security in autonomous driving.



Yifan Zhang received the B.Eng. degree in telecommunication from Xidian University, Xian, China, in 2015, and the M.Sc. degree in computer science from ENSIIE, Evry, France, in 2017. She is currently a Ph.D. student in the Department of Computer Science at City University of Hong Kong. Her research interests include intelligent transportation systems and motion planning in autonomous driving.



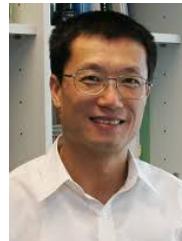
Kejie Lu (S'01-M'04-SM'07) received the B.Sc. and M.Sc. degrees from Beijing University of Posts and Telecommunications, Beijing, China, in 1994 and 1997, respectively. He received the Ph.D. degree in electrical engineering from the University of Texas at Dallas, Richardson, TX, USA, in 2003. In July 2005, he joined the University of Puerto Rico at Mayagüez, Mayagüez, Puerto Rico, where he is currently a professor in the Department of Computer Science and Engineering. His research interests include computer and communication networks, cyber-physical system, and network-based computing.



Jianping Wang received the B.S. and the M.S. degrees in computer science from Nankai University, Tianjin, China, in 1996 and 1999, respectively. She received the Ph.D. degree in computer science from the University of Texas at Dallas, in 2003. She is currently a professor in the Department of Computer Science at City University of Hong Kong. Her research interests include security and motion planning in autonomous driving, dependable networking, wireless networking, optical networking, cloud computing, service oriented networking, data center networks, integration of Fiber-Wireless networks, and network coding.



Kui Wu received the B.Sc. and M.Sc. degrees in computer science from Wuhan University, Wuhan, China, in 1990 and 1993, respectively, and the Ph.D. degree in computing science from the University of Alberta, Edmonton, AB, Canada, in 2002. He joined the Department of Computer Science, University of Victoria, Victoria, BC, Canada, in 2002, where he is currently a Professor. His current research interests include network performance analysis, online social networks, Internet of Things, and parallel and distributed algorithms.



Xiaohua Jia (F'13) received the BSc and MEng degrees in 1984 and 1987, respectively, from the University of Science and Technology of China, and the DSc degree in 1991 in information science from the University of Tokyo. He is currently the chair professor with the Department of Computer Science at the City University of Hong Kong. His research interests include cloud computing and distributed systems, computer networks, wireless sensor networks and mobile wireless networks. He is an editor of the *IEEE Transactions on Parallel and Distributed Systems* (2006–2009), *Wireless Networks*, *Journal of World Wide Web*, *Journal of Combinatorial Optimization*, etc. He is a fellow of IEEE (Computer Society).



Bin Liu received his B.S., M.S. and Ph.D. degrees in computer science and engineering from Northwestern Polytechnical University, Xian, China, in 1985, 1988 and 1993, respectively. From 1993 to 1995, he was a Postdoctoral Research Fellow in the National Key Laboratory of SPC and Switching Technologies, Beijing University of Post and Telecommunications, Beijing, China. In 1995, he transferred to the Department of Computer Science and Technology, Tsinghua University, where he is currently a Professor. His current research areas include high performance switches/routers, network processors, traffic measurement and management, high speed network security and greening Internet.