

MADE: Malicious Agent Detection for Robust Multi-Agent Collaborative Perception

Yangheng Zhao^{*†}, Zhen Xiang^{*‡}, Sheng Yin^{*†}, Xianghe Pang[†], Yanfeng Wang[†], Siheng Chen^{¶†}

Abstract—Recently, multi-agent collaborative (MAC) perception has been proposed and outperformed the traditional single-agent perception in many applications, such as autonomous driving. However, MAC perception is more vulnerable to adversarial attacks than single-agent perception due to the information exchange. The attacker can easily degrade the performance of a victim agent by sending harmful information from a malicious agent nearby. In this paper, we propose Malicious Agent Detection (MADE), a reactive defense specific to MAC perception that can be deployed by an agent to accurately detect and then remove any potential malicious agent in its local collaboration network. In particular, MADE inspects each agent in the network independently using a semi-supervised anomaly detector based on a double-hypothesis test with the Benjamini-Hochberg procedure for false positive control. For the two hypothesis tests, we propose a *match loss* statistic and a *collaborative reconstruction loss* statistic, respectively, both based on the *consistency* between the agent to be inspected and the ego agent deployed with our detector. We comprehensively evaluate MADE on a benchmark 3D dataset, V2X-sim, and a real-road dataset, DAIR-V2X, comparing it to baseline defenses. Notably, with the protection of MADE, the drops in the average precision compared with the best-case ‘Oracle’ defender are merely 1.27% and 0.28%, respectively.

I. INTRODUCTION

Through decades of efforts in computer vision and machine learning, single-agent perception has achieved remarkable success in object detection, tracking, and segmentation [38], [12], [44], [45], [30]. However, the traditional single-agent perception still suffers from a number of inevitable limitations due to an individual perspective, such as occlusion and long-range issues [40], [22]. To this end, multi-agent collaborative (MAC) perception has been proposed recently, which allows multiple agents to share complementary perceptual information with each other, promoting a more holistic perception [18], [34], [14], [10].

While existing studies on MAC perception mainly focus on improving perception performance [24], the associated security protocols are severely underexplored. In particular, the information sharing between agents potentially creates more opportunities for adversaries to break the entire system. For example, adversarial attacks ([28]) have been extended to MAC object detection recently, where the performance of an arbitrary victim agent can be easily degraded by the adversary using an optimized feature map sent from a compromised agent in the collaboration network [32]. Since generic defenses such as adversarial training cannot

effectively mitigate this threat, the deployment of MAC perception tools remains risky in safety-critical scenarios due to potential economic losses and human fatalities [23], [21].

In this paper, we propose Malicious Agent Detection (MADE), a reactive defense specific to adversarial attacks against MAC object detectors. MADE leverages the nature of agent collaboration to accurately detect malicious agents in an ego agent’s collaboration network and then remove them. It is based on a multi-test using conformal p-values [33], [1], and with the Benjamini-Hochberg (BH) procedure for the false positive control [4]. The multi-test is composed of two hypothesis tests on the *consistency* between the ego agent and the agent to be inspected. Specifically, we propose two novel detection statistics, a *match loss* statistic and a *collaborative reconstruction loss* statistic, as the consistent metrics for the two hypothesis tests, respectively. An agent is claimed to be a malicious one if it is ‘inconsistent’ with the ego agent as inferred by our multi-test. Note that our MADE is different from generic proactive defenses since we not only protect agents from potential adversarial attacks but also help to catch the malicious entities launching the attack, which is socially impactful [26], [29].

Our main contributions are summarized as follows:

- We propose MADE, a reactive, detection-based defense to address adversarial attacks against MAC object detectors by detecting and removing malicious agents in the local collaboration network.
- We propose a multi-test detection framework for MADE, with two novel statistics, *match loss* and *collaborative reconstruction loss*, to assess the consistency between each putative malicious agent and the ego agent.
- We comprehensively evaluate MADE on V2X-SIM and DAIR-V2X. We show that MADE outperforms the previous state-of-the-art by recovering 22.13% and 22.60% average precision on the two datasets, respectively, merely 1.27% and 0.28% lower than the average precision achieved by the ‘Oracle’ defender with ground-truth knowledge about the malicious agents, respectively.

II. RELATED WORK

A. MAC Perception

A MAC perception framework involves a group of agents, each deployed with the same model to produce its own perception output via collaboration with the other agents. Different from other multi-agent learning paradigms, such as distributed learning focusing on training-time collaboration [5], federated learning focusing on training data

[†] Shanghai Jiao Tong University

[‡] University of Illinois, Urbana-Champaign

^{*} Equal contribution

[¶] Corresponding to sihengc@sjtu.edu.cn

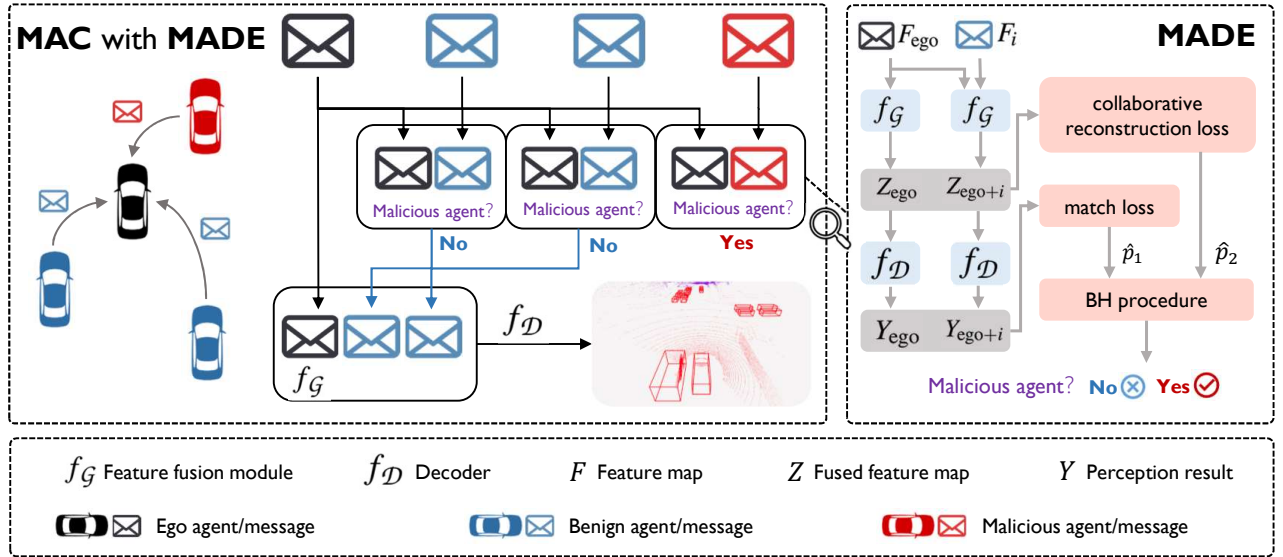


Fig. 1. Illustration of a standard MAC perception pipeline protected by our proposed MADE. Message (i.e. intermediate feature map) from each agent will be inspected by MADE before fusing – malicious agents will be removed once detected by MADE. Specifically, for any agent a_i (with a feature map F_i) to be inspected, we obtain Z_{ego} directly from the ego agent’s feature map F_{ego} (without fusion) and obtain a fused feature map Z_{ego+i} by fusing F_i with F_{ego} . We then obtain the bounding box proposals Y_{ego} and Y_{ego+i} from Z_{ego} and Z_{ego+i} , respectively. A *collaborative reconstruction loss* is computed on Z_{ego} and Z_{ego+i} , while a *match loss* statistic is computed on Y_{ego} and Y_{ego+i} . Finally, conformal p-values are computed for the two tests respectively, which are then used for a multi-test with the BH procedure, leading to the inference result for agent a_i .

privacy [20], or multi-modal learning focusing on the aggregation of different sensory inputs [3], MAC perception addresses the collaboration between agents *during inference* [2]. In this paper, we consider state-of-the-art MAC object detectors that conduct the collaboration in the intermediate feature space, which achieves the best trade-off between communication cost and perception performance [10], [34], [14], [10], [37].

B. Adversarial Attacks and Defenses for MAC Perception

Adversarial (evasion) attacks originated from the study of the worst-case robustness of deep neural network (DNN) classifiers [28], and have been extensively studied for single-agent object detection [35], [36], [9], [6], [15], [17], [31], [43]. Recently, adversarial attacks have been extended to MAC object detection [32]. The adversary applies an optimized perturbation to the feature map sent from a malicious agent to the victim agent to degrade its performance in bounding box prediction. Such an attack is shown evasive to generic defenses against adversarial attacks, particularly the adversarial training approach [19].

Prior to our work, a defense based on geometric consistency has been proposed to detect adversarial attacks against MAC object detectors [41]. However, this approach requires each agent to generate and share additional information only for the defense purpose (i.e., information not useful to MAC object detection); thus, it will not be further considered here since we do not allow the defender to manipulate the communication approach hardwired to the MAC object detector. Closest to our work is the detection-based defense in [13], where potential malicious agents are identified using a consensus loss. As will be shown experimentally, this defense is outperformed by our MADE, which leverages the

power of two hypothesis tests with two different detection statistics, respectively.

III. THREAT MODEL

A. MAC Object Detection Pipeline

We consider a collaboration network of N agents a_1, \dots, a_N . For each agent, the input $X \in \mathcal{X}$ is the local sensory data and the output $Y = \{y^{(1)}, \dots, y^{(L)}\} \in \mathcal{Y}$ is an ordered set of L bounding box proposals. Each bounding box proposal is a $k + d$ dimensional vector represented by $y = [p_1, \dots, p_k, z_1, \dots, z_d]^T$, where p_1, \dots, p_k are the posteriors for all $k > 1$ classes in the domain, respectively, and z_1, \dots, z_d jointly specify the spatial location of the bounding box. We consider the state-of-the-art MAC object detector pipeline with the following three stages: 1) **Encoding stage**. For each agent a_i , an encoder $f_{\mathcal{E}} : \mathcal{X} \rightarrow \mathcal{F}$ is deployed to encode the local input X_i captured by a_i into a feature map $F_i \in \mathcal{F}$. 2) **Collaboration stage**. Through a communication channel, each agent a_i receives the feature maps $\{F_j\}_{j \neq i}$ generated by all the other $(N - 1)$ agents (using the same $f_{\mathcal{E}}$). These received feature maps are aggregated with the local feature map F_i using an asymmetrical function $f_{\mathcal{G}} : \mathcal{F}^N \rightarrow \mathcal{Z}$, which produces a fused feature map $Z_i = f_{\mathcal{G}}(F_i; \{F_j\}_{j \neq i}) \in \mathcal{Z}$. 3) **Decoding stage**. The output $Y_i \in \mathcal{Y}$ of each agent a_i is obtained by applying a decoder $f_{\mathcal{D}} : \mathcal{Z} \rightarrow \mathcal{Y}$ to its fused feature map Z_i .

B. Adversarial Attack Against MAC Object Detector

We consider the adversarial attack in [32], where a_v denotes the victim agent targeted by the attacker, and a_m ($m \neq v$) is the malicious agent controlled by the attacker. For each agent a_i , the feature map obtained by applying the encoder $f_{\mathcal{E}}$ to its local input is denoted by

$F_i \in \mathcal{F}$. The attacker aims to degrade the performance of the victim agent a_v by applying an optimized perturbation to the feature map F_m sent from agent a_m to a_v . Let $Y'(\delta) = \{y'^{(1)}(\delta), \dots, y'^{(L)}(\delta)\} = f_{\mathcal{D}} \circ f_{\mathcal{G}}(F_v; (F_m + \delta) \cup \{F_i\}_{i \neq v, i \neq m})$ be the L bounding box proposals when F_m is perturbed by δ , the attacker aims to optimize δ by solving:

$$\underset{\|\delta\|_p \leq \epsilon}{\text{minimize}} \quad \sum_{l=1}^L \mathcal{L}_{\text{adv}}(y^{(l)}, y'^{(l)}(\delta), \varrho(y^{(l)})) \quad (1)$$

where $\|\cdot\|_p$ represents ℓ_p norm. The categorization function $\varrho(\cdot) : \mathbb{R}^{k+d} \rightarrow \{0, 1\}$ decides whether a bounding box proposal is associated with a ‘background’ box (‘0’) or a ‘foreground’ box associated with an object (‘1’). The ‘per-box loss’ in the sum of (1) is defined by:

$$\mathcal{L}_{\text{adv}}(y, y', \varrho(y)) = \begin{cases} 0 & \varrho(y) = 0 \\ -\log(1 - p'_u) \cdot \text{IoU}(y, y') & \varrho(y) = 1 \end{cases}$$

where $u = \arg \max_{c \in \mathcal{C} \setminus k} p_c$ is the largest confidence class (except the ‘background’ class k) associated with the proposal y . p'_c is the posterior associated with proposal y' for any class $c \in \mathcal{C}$. Following [32], the attack optimization problem is solved by projected gradient descent (PGD) [19].

While the procedure above corresponds to a single attacker with one malicious agent, in our experiments, we will consider both cases where there are multiple independent attackers and cases where one attacker controls multiple agents to launch a ‘collaborative’ attack.

While we focus on the detection of malicious agents based on their shared feature maps in this work, it’s important to note that the integrity and security of the communication protocols during the collaboration stage also affect the robustness of the system [42]. This important direction will be further explored in our future work.

IV. MALICIOUS AGENT DETECTION METHOD

A. Overview

Without loss of generality, we nominally set the N -th agent a_N to be a benign agent (dubbed the ‘ego’ agent) deployed with our detector. The goal here is to infer if any of the agents a_1, \dots, a_{N-1} is malicious, and then remove the detected agents from the local collaboration network. However, there is no prior knowledge about the presence of any malicious agents or their identities. The ground truth for the objects in the scene is also unavailable to the ego agent. Moreover, no additional communication with other agents is allowed, except for the exchange of the feature maps, which is the default for MAC object detection. The ego agent is only assumed with a set \mathcal{D}_{B} of (benign) feature maps (e.g.) collected locally from past instances. These constraints make the detection problem very challenging.

Our key idea is to inspect each agent in a_1, \dots, a_{N-1} by assessing its ‘consistency’ with the ego agent a_N in both outputs and intermediate features. Specifically, we propose a *match loss* statistic for the output-level consistency (see Sec. IV-B) and a *collaborative reconstruction loss* statistic for

the intermediate-level consistency (see Sec. IV-C). While the two statistics can be leveraged for malicious agent detection independently through hypothesis testing, in this paper, we construct a multi-test to jointly exploit the power of the two statistics, with the false detection rate controlled by the Benjamini-Hochberg (BH) procedure [4] (see Sec. IV-D). The procedure of MADE is summarized in Fig. 1.

B. Match Loss

The match loss statistic is an asymmetric metric designed to assess the consistency between two sets of bounding box proposals, say, $Y = \{y^{(1)}, \dots, y^{(L)}\}$ and $Y' = \{y'^{(1)}, \dots, y'^{(L)}\}$. Let \mathcal{I}_c and \mathcal{I}'_c be the two sets of indices of all the proposed bounding boxes in Y and Y' that are predicted to some class $c \in \mathcal{C}$, respectively. Let \mathfrak{S}_c be all *matches* between \mathcal{I}_c and \mathcal{I}'_c that maps *each* $l \in \mathcal{I}_c$ to a *unique* $l' \in \mathcal{I}'_c$. Note that to guarantee such uniqueness when $|\mathcal{I}'_c| < |\mathcal{I}_c|$, we pad \mathcal{I}'_c to the same size as \mathcal{I}_c with arbitrary ‘dummy’ indices not appeared in $\{1, \dots, L\}$ and associate each ‘dummy’ index with an empty box \emptyset . Then, the *match loss* between Y and Y' is defined by:

$$\mathcal{L}_{\text{m}}(Y, Y') = \sum_{c \in \mathcal{C}} \min_{\sigma \in \mathfrak{S}_c} \frac{1}{|\mathcal{I}_c|} \sum_{l \in \mathcal{I}_c} \mathcal{L}_{\text{box}}(y^{(l)}, y'^{(\sigma(l))}; c)$$

where the optimal match for each class c is solved using the Hungarian match algorithm [11]. Here, for any class $c \in \mathcal{C}$ and any two arbitrary bounding box proposals y and y' ,

$$\mathcal{L}_{\text{box}}(y, y'; c) = \max(p_c - p'_c, 0) + \phi(1 - \text{IoU}(z, z')) \quad (2)$$

where p_c and p'_c represent the posteriors of class c associated with bounding boxes y and y' , respectively, and ϕ is the parameter balancing the two terms.

Focusing on the first term in Eq. (2), a *large* loss will be induced only if p'_c is much smaller than p_c , which is the usual case when the ego agent collaborates with a malicious agent. By contrast, collaboration with a benign agent will not create a big decrement in p'_c . Similarly, for the second term in Eq. (2), collaboration with a malicious (or benign) agent will likely cause a large (or small) shift in the position of the proposed bounding boxes. Thus, the match loss tends to be *large* if the agent to be inspected is *malicious*.

C. Collaborative Reconstruction Loss

We propose the collaborative reconstruction loss statistic to measure the consistency of the *fused* feature maps of the ego agent with and without the collaboration with the agent a_i to be inspected. For convenience, these two fused feature maps are denoted as $Z_{\text{ego}+i} = f_{\mathcal{G}}(F_{\text{ego}}; \{F_i, 0, \dots, 0\})$ and $Z_{\text{ego}} = f_{\mathcal{G}}(F_{\text{ego}}; \{0, \dots, 0\})$ respectively. If a_i is a benign agent ‘consistent’ with the ego agent, the *residual feature map* $R_{\text{ego}+i} = Z_{\text{ego}+i} - Z_{\text{ego}}$ will likely be less noisy than for a_i being a malicious agent. In other words, the residual feature maps associated with malicious agents and benign agents follow different distributions.

However, for our practical detection problem, only the feature maps associated with benign agents are available to

the defender. Thus, we propose to train an autoencoder on the residual feature maps for benign agents (simulated on \mathcal{D}_B), such that residual feature maps for malicious agents will exhibit abnormally large reconstruction loss when going through the autoencoder. Specifically, we denote the set of all ordered pairs of benign feature maps in \mathcal{D}_B as:

$$\mathcal{F}_{\text{pair}}(\mathcal{D}_B) = \{(F, F') | F \in \mathcal{D}_B, F' \in \mathcal{D}_B, F \neq F'\}$$

The training set is obtained from $\mathcal{F}_{\text{pair}}(\mathcal{D}_B)$, denoted by:

$$\mathcal{R} = \{f_{\mathcal{G}}(F; \{F', 0, \dots, 0\}) - f_{\mathcal{G}}(F; \{0, \dots, 0\}) | (F, F') \in \mathcal{F}_{\text{pair}}(\mathcal{D}_B)\}$$

Then the autoencoder $f_{\mathcal{AE}} : \mathcal{Z} \rightarrow \mathcal{Z}$ is trained by minimizing the average reconstruction loss on \mathcal{R} , i.e.:

$$\underset{\theta}{\text{minimize}} \quad \frac{1}{|\mathcal{R}|} \sum_{R \in \mathcal{R}} \mathcal{L}_r(R, f_{\mathcal{AE}}(R; \theta))$$

where θ represents the parameters of the autoencoder $f_{\mathcal{AE}}$ and will be dropped in the sequel for brevity. Our *collaborative reconstruction loss* statistic for any residual feature map $R \in \mathcal{Z}$ is defined by the following mean square error:

$$\mathcal{L}_r(R, f_{\mathcal{AE}}(R)) = \|R - f_{\mathcal{AE}}(R)\|_2,$$

which tends to be large for malicious agents.

D. Detection Inference Based on Multi-Test

Our detection inference is based on *multi-test* using *conformal p-values* computed on calibration sets. For each agent a_i to be inspected, we first obtain the match loss statistic $\mathcal{L}_m(Y_{\text{ego}}, Y_{\text{ego}+i})$ and the collaborative reconstruction loss statistic $\mathcal{L}_r(R_{\text{ego}+i}, f_{\mathcal{AE}}(R_{\text{ego}+i}))$, respectively. Then, our detection inference can be cast into the following multi-test:

$$\begin{aligned} H_{1,0} : (Y_{\text{ego}}, Y_{\text{ego}+i}) &\sim P_{Y,Y'}; & H_{1,1} : (Y_{\text{ego}}, Y_{\text{ego}+i}) &\not\sim P_{Y,Y'} \\ H_{2,0} : R_{\text{ego}+i} &\sim P_R; & H_{2,1} : R_{\text{ego}+i} &\not\sim P_R \end{aligned} \quad (3)$$

Here, $P_{Y,Y'}$ is the joint distribution for (output) random variables $Y = (f_{\mathcal{D}} \circ f_{\mathcal{G}})(F; \{0, \dots, 0\})$ and $Y' = (f_{\mathcal{D}} \circ f_{\mathcal{G}})(F; \{F', 0, \dots, 0\})$, and P_R is the distribution for (residual feature map) random variable $R = f_{\mathcal{G}}(F; \{F', 0, \dots, 0\}) - f_{\mathcal{G}}(F; \{0, \dots, 0\})$, where F and F' are i.i.d. random variables for benign feature maps.

Since the ego agent is always benign, ideally, the rejection of any of these two hypotheses will be caused by a_i being malicious. Here, we employ the Benjamini-Hochberg (BH) procedure to control the overall false detection rate of the multi-test [8], [27], [4]. We first obtain a p-value for each statistic under its associated null hypothesis independently. Since both $P_{Y,Y'}$ and P_R in (3) are analytically intractable, for each statistic, we compute a conformal p-value on a *calibration set* obtained from \mathcal{D}_B by:

$$\begin{aligned} \hat{p}_1 &= \frac{1 + |\{(Y, Y') \in \mathcal{J}_1^{\text{cal}} | \mathcal{L}_m(Y, Y') \geq \mathcal{L}_m(Y_{\text{ego}}, Y_{\text{ego}+i})\}|}{1 + |\mathcal{J}_1^{\text{cal}}|} \\ \hat{p}_2 &= \frac{1 + |\{R \in \mathcal{J}_2^{\text{cal}} | \mathcal{L}_r(R, f_{\mathcal{AE}}(R)) \geq \mathcal{L}_r(R_{\text{ego}+i}, f_{\mathcal{AE}}(R_{\text{ego}+i}))\}|}{1 + |\mathcal{J}_2^{\text{cal}}|} \end{aligned}$$

where subscripts ‘1’ and ‘2’ correspond to the match loss and the collaborative reconstruction loss respectively and $\mathcal{J}_i^{\text{cal}}$ denotes the calibration set. Then, following the BH procedure, a_i is deemed a malicious agent if:

$$\min\{\hat{p}_1, \hat{p}_2\} \leq \frac{\alpha}{2} \quad \text{or} \quad \max\{\hat{p}_1, \hat{p}_2\} \leq \alpha$$

where α is the desired FPR (e.g., the classical $\alpha = 0.05$ for statistical testing).

V. EXPERIMENTS

A. Experimental Setup

Dataset. We conduct experiments on two benchmarks, a vehicle-to-everything dataset V2X-SIM and a real-road vehicle-to-infrastructure dataset DAIR-V2X [14], [39]. For V2X-SIM, we conduct an out-of-distribution evaluation by reserving all 8 scenes with 2 or 3 agents for the defender for detection, while 3 scenes with 4 agents are used for performance evaluation. DAIR-V2X involves two agents (vehicle and infrastructure) in each scene. We use 4800 frames (approximately 8 minutes in real time) to train the MAC object detector, with the remaining 1800 frames (about 3 minutes in real time) for detection and evaluation, based on LiDAR’s 10Hz capture frequency.

MAC object detector. For V2X-SIM, we use a pre-trained DiscoNet [14], where the local input LiDAR point cloud for each agent is encoded to 32×32 bird eye view (BEV) feature maps with 256 channels before communication. For DAIR-V2X, we train a multi-scale Where2comm with temperature 20 [10], where the LiDAR point cloud for each agent is encoded to 100×252 , 50×126 , and 25×63 feature maps with 64, 128, and 256 channels, respectively.

Attack details. For V2X-SIM with four agents per frame, we generate four attack instances with the four agents taking turns as the ego agent; and the malicious agent is randomly chosen from the remaining three agents. For DAIR-V2X with only two agents, we generate one attack instance for each frame by setting the vehicle as the ego agent and the infrastructure as the malicious agent. In each attack instance, we manipulate the feature map sent from the malicious agent to the ego agent following our threat model in Sec. III-B. We consider maximum perturbation sizes (measured by the ℓ_∞ norm of the perturbation δ) $\epsilon \in \{0.05, 0.1, 0.2\}$ and $\epsilon \in \{1.0, 1.5, 2.0\}$ for V2X-SIM and DAIR-V2X, respectively. PGD is conducted for 10 steps with step size $\epsilon/5$ for V2X-SIM and 80 steps with step size $\epsilon/30$ for DAIR-V2X. In Fig. 2, we show the visualization of the feature maps being manipulated by the adversary for V2X-SIM and DAIR-V2X for attacks with $\epsilon = 0.2$ and $\epsilon = 2.0$ respectively. Our created attacks for evaluation are stealthy, with almost imperceptible artifacts on the manipulated feature maps, yet as shown in Tab. I, they effectively reduce the performance of the victim agent by 63.77% and 56.92%, respectively. Effective attacks provide a better assessment of defense methods.

Details of MADE. In each attack instance, we inspect all the agents other than the ego agent. The match loss is computed with the multiplier set to $\phi = 1.0$ in Eq. 2

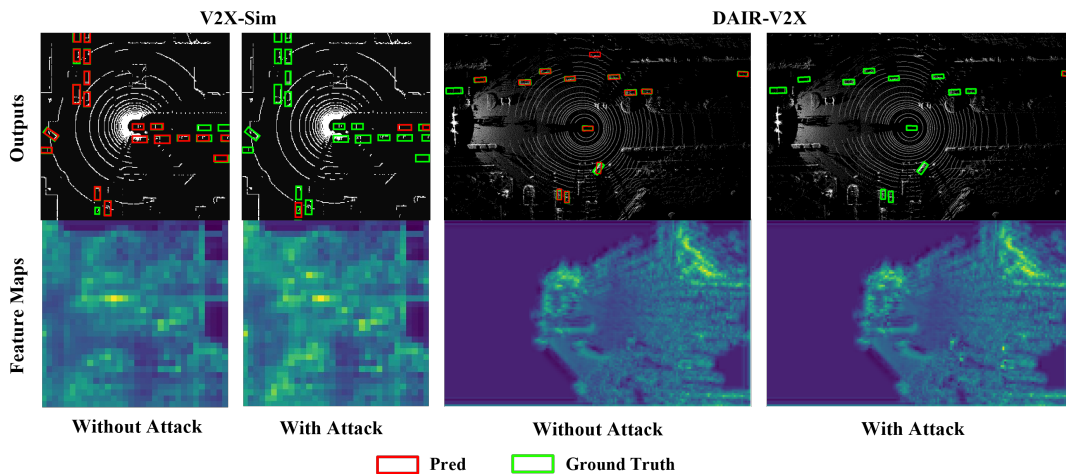


Fig. 2. Comparative visualization of intermediate feature maps and detection outputs with and without adversarial attack

TABLE I

END-TO-END DEFENSE PERFORMANCE OF MADE COMPARED WITH THE BASELINE.

		V2X-SIM				DAIR-V2X			
		no attack	$\epsilon = 0.05$	$\epsilon = 0.1$	$\epsilon = 0.2$	no attack	$\epsilon = 1.0$	$\epsilon = 1.5$	$\epsilon = 2.0$
(reference)	no defense	61.82	53.30	36.36	22.40	64.97	47.40	36.55	27.99
	oracle	61.82	59.76	59.76	59.76	64.97	59.43	59.43	59.43
(baseline)	ROBOSAC [13]	59.54	54.13	52.35	53.97	64.97	47.75	42.19	42.84
	ML only	61.91	56.56	57.72	57.79	64.87	57.11	58.48	58.95
(ours)	CRL only	60.38	54.77	54.47	50.35	60.16	57.99	57.47	58.04
	MADE (ML+CRL)	61.10	56.91	58.49	58.40	60.58	58.45	59.15	59.28

and the calibration sets containing 2800 and 1000 pairs of (benign) output bounding box proposals for V2X-SIM and DAIR-V2X, respectively. For the collaborative reconstruction loss, the training sets of the autoencoders for V2X-SIM and DAIR-V2X contain 23700 and 3811 instances, respectively. For V2X-SIM, we use U-Net [25] as the architecture of the autoencoder. Training is performed using Adam optimizer for 200 epochs with batch size 256 and learning rate 0.0001 (with a 0.1 decay per 50 epochs). For DAIR-V2X, we use the same architecture in [16] for the autoencoder. Training is performed using Adam optimizer for 200 epochs with batch size 64 and learning rate 0.001 (with a 0.1 decay per 50 iterations). Finally, we set the desired FPR to the classical $\alpha = 0.05$ for the multi-test inference procedure.

Evaluation metrics. We evaluate the effectiveness of defenses based on the end-to-end performance of the object detector measured by the average precision (AP) at 0.5 IoU (denoted by AP@0.5 and in percentage).

B. Defense Performance Evaluation

In Tab. I, we show the end-to-end AP@0.5 of our MADE against the attacks we created for evaluation, compared with the previous state-of-the-art approach ROBOSAC [13]. We also consider two single-test variants of MADE with the match loss (ML) statistic and the collaborative reconstruction loss (CRL) statistic respectively. Moreover, the worst-case performance when there is no defense and the best-case performance when the defender is an ‘Oracle’ with ground truth knowledge that can accurately identify and remove all malicious agents are shown for reference. As shown in the table, both defenses with a single test on our proposed

TABLE II

MALICIOUS AGENT DETECTION ACCURACY OF MADE ON V2X-SIM.

	$\epsilon = 0.05$		$\epsilon = 0.1$		$\epsilon = 0.2$	
	TPR	FPR	TPR	FPR	TPR	FPR
ML only	0.667	0.150	0.963	0.138	0.965	0.124
CRL only	0.255	0.025	0.733	0.018	0.825	0.015
MADE (ML+CRL)	0.593	0.076	0.958	0.069	0.947	0.055

ML or CRL statistic perform well with higher restored APs compared with the baseline for almost all attacks. Notably, the defense that only uses our ML statistic has already outperformed the baseline with uniformly higher restored APs. Furthermore, our MADE defense performs even better than the two single-test defenses against the two attacks on the two datasets by leveraging the detection power of both ML and CRL statistics, with the recovered AP@0.5 close to the best-case ‘Oracle’ defender for both datasets.

In addition, we demonstrate the true positive rate (TPR) and the false positive rate (FPR) of our MADE. Specifically, TPR is defined as the proportion of malicious agents being successfully detected, while FPR is defined as the proportion of benign agents being incorrectly detected. As shown in Tab II, both ML and CRL can be used to effectively distinguish malicious agents from benign ones independently, while our MADE combining the two statistics is the only method achieving both high TPR and controlled FPR close to the prescribed significance level $\alpha = 0.05$. The impact of removing these deficient agents on the end-to-end performance is minimal.

While the current experimental results demonstrate promising outcomes, we acknowledge that the range and

TABLE III
ABLATION STUDY ON PARAMETER ϕ IN MATCH LOSS.

		$\phi = 0.01$	$\phi = 0.5$	$\phi = 1$	$\phi = 2$	$\phi = 100$
V2X-SIM	attack	58.19	58.25	57.72	57.13	54.49
	no attack	61.73	60.75	61.91	60.10	60.01
DAIR-V2X	attack	59.41	59.12	58.48	57.84	54.47
	no attack	65.04	64.93	64.87	64.85	64.81

TABLE IV
PERFORMANCE OF ML, CRL, AND MADE AGAINST ATTACKS WITH MULTIPLE INDEPENDENT ATTACKERS AND ATTACKS WITH MULTIPLE COLLABORATIVE MALICIOUS AGENTS ON V2X-SIM.

No. malicious	Independent		Collaborative	
	2 agents	3 agents	2 agents	3 agents
no defense	21.55	15.76	20.20	14.35
oracle	57.40	49.56	57.40	49.56
ML only	55.33	48.68	53.57	45.47
CRL only	50.96	42.69	41.09	27.48
MADE	55.98	48.51	53.79	43.97

depth of the data could be expanded and made more rigorous. This will be an important direction for our future work.

C. Ablation Study

In this section, we conduct an ablation study on the parameter ϕ in match loss, which balances the classification difference (the first term in Eq.(2)) and the localization difference (the second term in Eq.(2)). As shown in Tab. III, for both datasets, the performance of MADE is generally insensitive to the choice of ϕ as long as it is close to 1.0, which is our default setting.

D. Detection Performance against Adaptive Attacks

Here, we consider more challenging scenarios with multiple malicious agents. In particular, there may be more than one attacker in the scene, and each attacker controls a malicious agent to launch an attack against the (ego) victim agent *independently*. Moreover, there may exist one attacker who controls multiple malicious agents to launch an attack *collaboratively* by jointly optimizing the feature maps sent from these malicious agents to the victim agent. Such a collaborative attack involving more than one agent can be regarded as an adaptive attack to our proposed defense. This is because we inspect each agent other than the ego agent independently, while for multiple malicious agents in collaboration, the performance degradation at the ego agent induced by a single malicious agent may not be severe enough to be detected, e.g., using the match loss statistic.

In Tab. IV, we show the end-to-end object detection performance (measured by AP@0.5) of ML, CRL, and the proposed MADE, against two and three coexisting but independent attackers, respectively, and two and three collaborative malicious agents, respectively. Here, we consider the V2X-SIM dataset only since the DAIR-V2X dataset has only two agents in each scene. With more malicious agents (either independent or collaborative) and fewer benign agents, the object detection performance of the victim agent is clearly degraded. Our MADE achieves 55.98% and 48.51% AP@0.5, merely 1.42% and 1.05% lower than the best-case ‘Oracle’ defender, respectively. For the collaborative attacks,

TABLE V
PERFORMANCE OF THE UNSUPERVISED ADAPTION OF MADE.

		Independent		Collaborative	
No. mal.	1 agent	2 agents	3 agents	2 agents	3 agents
AP@0.5	45.89	34.90	23.85	35.79	21.85

there is a slight drop in the performance of our MADE, but it still achieves 53.79% and 43.97% AP@0.5, merely 3.61% and 5.59% lower than the best-case ‘Oracle’ defender, respectively.

E. Unsupervised Adaption

Here, we adapt our proposed MADE defense to address a more challenging *unsupervised* detection scenario where past information is not allowed. The key intuition is that the manipulation of a feature map sent from a malicious agent is specific to the victim agent, which will not clearly affect the outputs of the other agents. Thus, if the ego agent (with output Y_{ego}) is indeed a victim agent targeted by one or more attackers, there will exist some agent a_i with output Y_i , such that the match loss $\mathcal{L}_m(Y_{\text{ego}}, Y_i)$ is a large outlier to the match losses computed for non-ego agents. For an instance with four agents a_1, \dots, a_4 where a_1 is the ego agent, for example, we aim to assess the atypicality of $\mathcal{L}_m(Y_1, Y_2)$, $\mathcal{L}_m(Y_1, Y_3)$, and $\mathcal{L}_m(Y_1, Y_4)$, with reference to the other nine match losses (‘centered’ at the three non-ego agents). To achieve this in an unsupervised fashion, we apply a robust detection rule based on median absolute deviation (MAD) [7]. We first compute the MAD of the ML statistics for all ordered pairs of agents in the scene. Then, we compute a score for each match loss \mathcal{L} of the ego agent by:

$$S = (\mathcal{L} - \text{median}) / (1.4826 \cdot \text{mad})$$

where the median is taken over all the match loss statistics, and the constant 1.4826 is associated with a 0.95 significance under the Gaussian assumption. The agent associated with the match loss is deemed to be a malicious agent if the score S is greater than 2. In Tab. V, we show the performance of the unsupervised adaption of our method against attacks with one, two, and three malicious agents on V2X-SIM with and without collaboration. For this challenging defense task, our unsupervised adaption achieves a decent performance for attacks with one malicious agent and some effect when there are more malicious agents.

VI. CONCLUSIONS

This paper focuses on the security issues of the emergent MAC perception systems. We propose a reactive defense that detects and removes malicious agents in the collaboration network of the ego agent. Our defense is based on a multi-test that involves two novel statistics, match loss and collaborative reconstruction loss, with the Benjamini-Hochber for false detection rate control. Comprehensive evaluations on both simulation and real-world benchmarks show that our defense is more effective than the previous state-of-the-art. It achieves 58.49% and 59.15% AP@0.5 on V2X-SIM and

DAIR-V2X, respectively, merely 1.27% and 0.28% lower than AP@0.5 achieved by the ‘Oracle’ defender, respectively.

ACKNOWLEDGMENTS

This research is supported by the National Key R&D Program of China under Grant 2021ZD0112801, NSFC under Grant 62171276, and the Science and Technology Commission of Shanghai Municipal under Grant 21511100900 and 22DZ2229005.

REFERENCES

- [1] Anastasios N. Angelopoulos and Stephen Bates. A gentle introduction to conformal prediction and distribution-free uncertainty quantification. 2021.
- [2] Kshitij Bakliwal, Maharshi Harshadbhai Dhada, Adria Salvador Palau, Ajith Kumar Parlikad, and Bhupesh Kumar Lad. A multi agent system architecture to implement collaborative learning for social industrial assets. *IFAC-PapersOnLine*, 51(11):1237–1242, 2018.
- [3] Tadas Baltrušaitis, Chaitanya Ahuja, and Louis-Philippe Morency. Multimodal machine learning: A survey and taxonomy. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2019.
- [4] Yoav Benjamini and Yoel Hochberg. Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society Series B (Methodological)*, 1995.
- [5] Mingzhe Chen, Deniz Gündüz, Kaibin Huang, Walid Saad, Mehdi Bennis, Aneta Vulgarakis Feljan, and H Vincent Poor. Distributed learning in wireless networks: Recent progress and future challenges. *IEEE Journal on Selected Areas in Communications*, 2021.
- [6] Kevin Eykholt, Ivan Evtimov, Earlene Fernandes, Bo Li, Amir Rahmati, Florian Tramèr, Atul Prakash, Tadayoshi Kohno, and Dawn Song. Physical adversarial examples for object detectors. In *Proceedings of the 12th USENIX Conference on Offensive Technologies*, 2018.
- [7] F. R. Hampel. The influence curve and its role in robust estimation. *Journal of the American Statistical Association*, 69, 1974.
- [8] Y. Hochberg and A. C. Tamhane. *Multiple Comparison Procedures*. John Wiley & Sons, Inc., USA, 1987.
- [9] Yu-Chih-Tuan Hu, Jun-Cheng Chen, Bo-Han Kung, Kai-Lung Hua, and Daniel Stanley Tan. Naturalistic physical adversarial patch for object detectors. In *ICCV*, 2021.
- [10] Yue Hu, Shaoheng Fang, Zixing Lei, Yiqi Zhong, and Siheng Chen. Where2comm: Communication-efficient collaborative perception via spatial confidence maps. In *NeurIPS*, 2022.
- [11] Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955.
- [12] Alex H Lang, Sourabh Vora, Holger Caesar, Lubing Zhou, Jiong Yang, and Oscar Beijbom. Pointpillars: Fast encoders for object detection from point clouds. In *CVPR*, 2019.
- [13] Yiming Li, Qi Fang, Jiamu Bai, Siheng Chen, Felix Juefei-Xu, and Chen Feng. Among us: Adversarially robust collaborative perception by consensus. In *ICCV*, 2023.
- [14] Yiming Li, Shunli Ren, Pengxiang Wu, Siheng Chen, Chen Feng, and Wenjun Zhang. Learning distilled collaboration graph for multi-agent perception. In *NeurIPS*, 2021.
- [15] Yuezun Li, Xiao Bian, Ming-Ching Chang, and Siwei Lyu. Exploring the vulnerability of single shot module in object detectors via imperceptible background patches. In *BMVC*, 2018.
- [16] Phillip Lippe. Tutorial 9: Deep Autoencoders. <https://github.com/phlippe>. [Online; accessed 14-Oct-2023].
- [17] Xin Liu, Huanrui Yang, Linghao Song, Hai Li, and Yiran Chen. Dpatch: Attacking object detectors with adversarial patches. In *AAAI Workshop on Artificial Intelligence Safety*, 2018.
- [18] Yen-Cheng Liu, Junjiao Tian, Nathaniel Glaser, and Zsolt Kira. When2com: Multi-agent perception via communication graph grouping. In *CVPR*, 2020.
- [19] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu. Towards deep learning models resistant to adversarial attacks. In *ICLR*, 2018.
- [20] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *AISTATS*, 2017.
- [21] the White House OSTP. Blueprint for an AI bill of rights, 2022.
- [22] Christopher Vincent Poulton, Matthew J Byrd, Peter Russo, Erman Timurdogan, Murshed Khandaker, Diedrik Vermeulen, and Michael R Watts. Long-range lidar and free-space data communication with high-performance optical phased arrays. *IEEE Journal of Selected Topics in Quantum Electronics*, 25(5):1–8, 2019.
- [23] Federal Register. Occupant Protection for Vehicles With Automated Driving Systems. <https://www.govinfo.gov/content/pkg/FR-2022-03-30/pdf/2022-05426.pdf>, 2022.
- [24] Shunli Ren, Siheng Chen, and Wenjun Zhang. Collaborative perception for autonomous driving: Current status and future trend. In Zhang Ren, Mengyi Wang, and Yongzhao Hua, editors, *Proceedings of 2021 5th Chinese Conference on Swarm Intelligence and Cooperative Control*, 2023.
- [25] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, 2015.
- [26] Jangirala Srinivas, Ashok Kumar Das, and Neeraj Kumar. Government regulations in cyber security: Framework, standards and recommendations. *Future Generation Computer Systems*, 92:178–188, 2019.
- [27] John D. Storey. A direct approach to false discovery rates. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 64(3):479–498, 2002.
- [28] C. Szegedy, W. Zaremba, I Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus. Intriguing properties of neural networks. In *ICLR*, 2014.
- [29] Yoshiyasu Takefuji. Connected vehicle security vulnerabilities [commentary]. *IEEE Technology and Society Magazine*, 2018.
- [30] Haotian Tang, Zhijian Liu, Shengyu Zhao, Yujun Lin, Ji Lin, Hanrui Wang, and Song Han. Searching efficient 3d architectures with sparse point-voxel convolution. In *ECCV*, 2020.
- [31] Simen Thys, Wiebe Van Ranst, and Toon Goedeme. Fooling automated surveillance cameras: Adversarial patches to attack person detection. In *CVPR Workshops*, June 2019.
- [32] J. Tu, T. Wang, J. Wang, S. Manivasagam, M. Ren, and R. Urtasun. Adversarial attacks on multi-agent communication. In *ICCV*, 2021.
- [33] Vladimir Vovk. Conditional validity of inductive conformal predictors. In *Proceedings of the Asian Conference on Machine Learning*, 2012.
- [34] Tsun-Hsuan Wang, Sivabalan Manivasagam, Ming Liang, Bin Yang, Wenyan Zeng, and Raquel Urtasun. V2vnet: Vehicle-to-vehicle communication for joint perception and prediction. In *ECCV*, 2020.
- [35] Cihang Xie, Jianyu Wang, Zhishuai Zhang, Yuyin Zhou, Lingxi Xie, and Alan Yuille. Adversarial examples for semantic segmentation and object detection. In *ICCV*, 2017.
- [36] Kaidi Xu, Gaoyuan Zhang, Sijia Liu, Quanfu Fan, Mengshu Sun, Hongge Chen, Pin-Yu Chen, Yanzhi Wang, and Xue Lin. Adversarial t-shirt! evading person detectors in a physical world. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *ECCV*, 2020.
- [37] Runsheng Xu, Hao Xiang, Zhengzhong Tu, Xin Xia, Ming-Hsuan Yang, and Jiaqi Ma. V2x-vit: Vehicle-to-everything cooperative perception with vision transformer. *arXiv preprint arXiv:2203.10638*, 2022.
- [38] Tianwei Yin, Xingyi Zhou, and Philipp Krahenbuhl. Center-based 3d object detection and tracking. In *CVPR*, 2021.
- [39] Haibao Yu, Yizhen Luo, Mao Shu, Yiyi Huo, Zebang Yang, Yifeng Shi, Zhenglong Guo, Hanyu Li, Xing Hu, Jirui Yuan, et al. Dair-v2x: A large-scale dataset for vehicle-infrastructure cooperative 3d object detection. In *CVPR*, 2022.
- [40] Guanqi Zhan, Weidi Xie, and Andrew Zisserman. A tri-layer plugin to improve occluded detection. In *BMVC*, 2022.
- [41] Qingzhao Zhang, Shuwei Jin, Ruiyang Zhu, Jiachen Sun, Xumiao Zhang, Qi Alfred Chen, and Z. Morley Mao. On data fabrication in collaborative vehicular perception: Attacks and countermeasures. In *USENIX*, 2024.
- [42] H. Zhao, A. Pacheco, V. Strobel, A. Reina, X. Liu, G. Dudek, and M. Dorigo. A generic framework for byzantine-tolerant consensus achievement in robot swarms. In *IROS 2023*, 2023.
- [43] Yue Zhao, Hong Zhu, Ruigang Liang, Qintao Shen, Shengzhi Zhang, and Kai Chen. Seeing isn’t believing: Towards more robust adversarial attack against real world object detectors. In *ACM CCS*, 2019.
- [44] Xingyi Zhou, Vladlen Koltun, and Philipp Krähenbühl. Tracking objects as points. In *ECCV*. Springer, 2020.
- [45] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. In *ICLR*, 2021.