

Among Us: Adversarially Robust Collaborative Perception by Consensus

Yiming Li^{1,*} Qi Fang^{1,*} Jiamu Bai¹ Siheng Chen^{2,3} Felix Juefei-Xu⁴ Chen Feng^{1,†}
¹New York University ²Shanghai Jiao Tong University ³Shanghai AI Laboratory ⁴Meta AI
 {yimingli, qifang, jb7082}@nyu.edu, sihengc@sjtu.edu.cn, felixxu@meta.com, cfeng@nyu.edu

<https://github.com/coperception/ROBOSAC>

Abstract

Multiple robots could perceive a scene (e.g., detect objects) collaboratively better than individuals, although easily suffer from adversarial attacks when using deep learning. This could be addressed by the adversarial defense, but its training requires the often-unknown attacking mechanism. Differently, we propose ROBOSAC, a novel sampling-based defense strategy generalizable to unseen attackers. Our key idea is that collaborative perception should lead to consensus rather than dissensus in results compared to individual perception. This leads to our hypothesize-and-verify framework: perception results with and without collaboration from a random subset of teammates are compared until reaching a consensus. In such a framework, more teammates in the sampled subset often entail better perception performance but require longer sampling time to reject potential attackers. Thus, we derive how many sampling trials are needed to ensure the desired size of an attacker-free subset, or equivalently, the maximum size of such a subset that we can successfully sample within a given number of trials. We validate our method on the task of collaborative 3D object detection in autonomous driving scenarios.

1. Introduction

Perception is a fundamental capability for autonomous robots to understand their surroundings [1–3]. Single-robot perception suffers from long-range or occlusion issues which stem from limited sensing capabilities and inadequate individual viewpoints [4]. Therefore, collaborative perception (co-perception) is proposed to provide more viewpoints for each robot via communication, so that robots can see further and better [5–7].

In literature, raw-data-level and decision-level fusion both demonstrate satisfactory performance in terms of robustness and precision [8, 9]. The recent development of deep learning has revolutionized many fields including

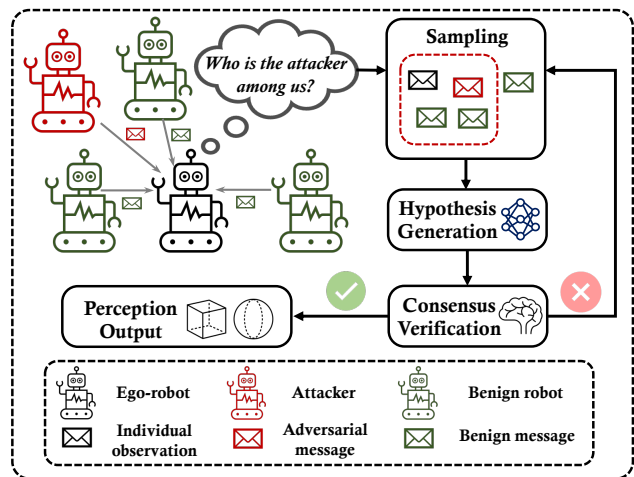


Figure 1: **Overview of ROBOSAC.** Ego-robot aims to find several benign collaborators with a hypothesize-and-verify procedure until reaching a consensus or using up the sampling budget. Consensus is checked between the results with and without the selected teammates.

robotic perception, and feature-level fusion has been proposed in which intermediate representations from a deep neural network (DNN) are shared amongst robots. Unlike raw-data-level and decision-level fusion approaches, feature-level fusion presents the advantages of good compressibility and the preservation of contextual information, further enhancing the performance-bandwidth trade-off in multi-robot perception [5, 10–12].

Although the original motivation for collaborative perception is to promote resilience and robustness via information sharing, the communication channel could potentially become a wide-open backdoor in DNN-based perception models due to the well-known adversarial vulnerability of DNNs [13]. Prior work has shown that a maliciously-crafted imperceptible perturbation added on the shared feature can drastically alter the perception output, jeopardizing the perception system [14]. To solve the safety concerns, adversarial training has been exploited [14], yet it introduces extra overhead during training and fails to generalize

*indicates equal contribution

†Corresponding author. The work is supported by NSF grants 2238968 and 2026479.

to unseen attackers [15]. Besides, adversarial training may lead to a small loss of accuracy [16]. In a word, it is still non-trivial to achieve *computationally-efficient* and *generalizable* adversarial defense in collaborative perception.

In this work, different from applying adversarial training after indiscriminately using all messages, we propose to enable the ego-robot to intelligently select benign collaborators from teammates, instead of naively trusting all the teammates. Inspired by random sample consensus (RANSAC) in robust estimation [17], we propose **ROBust cOllaborative SAmples Consensus (ROBOSAC)**, a general sampling-based framework for adversarially robust collaborative perception. Our key idea is that the robot is supposed to reach a consensus with its teammates after collaboration, rather than largely diverging from its individual perception. Specifically, ROBOSAC utilizes the hypothesize-and-verify workflow: the robot samples a subset of teammates and compares the results with and without the sampled teammates. After the consensus is verified, indicating no attackers *among us*, the robot can output the perceptual results generated in collaboration with teammates for further decision-making, as shown in Fig. 1. Different from the widely-used adversarial training, ROBOSAC is attacker-agnostic and thus can easily generalize to unseen adversarial learning algorithms.

Meanwhile, ROBOSAC can be customized for either strong performance or high efficiency: as more benign teammates leading to better performance require more computation to reject the attackers, there exists a performance-computation trade-off in ROBOSAC. Formally, under various attacker ratios, we can compute the *maximum number of attacker-free collaborators* that could be found given a specific sampling budget, and the *upper bound of the number of sampling* ensuring a desired number of benign teammates, all for achieving a guaranteed consensus probability. Additionally, we propose an adaptive probing approach to handle the scenario of unknown attacker ratios, starting from trusting all the teammates and then gradually becoming more vigilant. Our contributions are summarized as:

- We develop ROBOSAC, a scalable, generalizable, and generally-applicable adversarially robust collaborative perception framework via multi-robot consensus.
- We propose aggressive-to-conservative probing (A2CP) with retrospect to estimate the attacker ratio efficiently.
- We conduct experiments on collaborative 3D object detection in safety-critical autonomous driving to validate the effectiveness of ROBOSAC.

2. Related Works

Collaborative perception. To solve the fundamental issues of single-robot perception such as limited field-of-

view [18–21], multi-robot collaboration has been exploited to ameliorate the precision, robustness, and resilience of the perception system [22]. Previous works primarily investigate multi-robot perception in the aerial scenario [10, 23, 24] and autonomous driving [8, 25, 26], in different tasks such as object detection, semantic segmentation, and depth estimation. There are three kinds of communication strategies in multi-robot perception: (1) raw-data-level fusion, (2) feature-level fusion, and (3) output-level fusion. Amongst all three methods, feature-level fusion transmits learned intermediate representations of deep neural networks. Since the intermediate representations are easy to compress and are equipped with contextual knowledge of the environment, feature-level fusion demonstrates better performance-bandwidth trade-off, thus is widely applied in autonomous robots [5, 10, 11, 27]. Nevertheless, the adversarial robustness of feature-level fusion is underexplored.

Adversarial perception. Adversarial vulnerability in DNNs [13] can endanger the learning-based single-robot perception systems in safety-critical scenarios like autonomous driving [28–31]. For multi-robot perception, [14] reveals that an indistinguishable adversarial noise added on the shared intermediate representation can result in a lot of false detections. Adversarial attack can be classified into white-box attack and black-box attacks [32]. White-box attacker has full information about the DNNs [13, 33, 34], while black-box attacker is usually less effective than white-box attack since attackers have no access to the target models. The information about the model is obtained through query [35] or inferred through highly transferable surrogate models [36, 37]. For defense, adversarial training is proposed by incorporating adversarial samples into training stages, yet it requires the knowledge of attackers. Hence, to realize a generalizable adversarial defense against unseen attackers is still non-trivial and underexplored.

RANSAC. Random sample consensus (RANSAC) is a well-known robust estimation algorithm applicable to datasets containing a number of outliers. It was first proposed by Fischler and Bolles to solve Location Determination Problem (LDP) [38]. RANSAC employs a hypothesize-and-verify pipeline to iteratively select a sample of data points in the process of fitting the optimal model. At first, RANSAC was largely applied in works in the image domain, and gradually extension work of RANSAC is proposed, *e.g.*, MSAC (M-estimator SAmples and Consensus) [39] and MLESAC (Maximum Likelihood Estimation SAmples and Consensus) [40]. Currently, RANSAC is widely used in computer vision [41, 42] and robotics [43], *e.g.*, to estimate the fundamental matrix and remove outlier correspondences in Structure from Motion (SfM) [44]. In this work, we for the first time exploit the idea of sample consensus in the problem of robust collaborative perception.

3. ROBOSAC

In this section, we present the problem setup for collaborative perception under adversarial attacks in section 3.1, followed by the revisiting of RANSAC in section 3.2 and the illustration of a general defense framework termed ROBOSAC in section 3.3 (attacker ratio known) and section 3.4 (attacker ratio unknown).

3.1. Problem setup

Terminology. We first introduce our terminology in this work. Benign robots that share their truly-observed information are termed *collaborators*. Adversarial robots that share carefully-crafted harmful messages are termed *attackers*. The robot that tries to exploit the collaborators' information while protecting itself from attackers is termed *ego-robot*. All the robots other than the *ego-robot* are termed *teammates*, including both *collaborators* and *attackers*.

Assumption and setup. We consider two scenarios: (1) *static team*: the ego-robot communicates with the same teammates, and (2) *dynamic team*: the ego-robot meets and communicates with different teammates. We assume that the attacker ratio is fixed in both scenarios. *Regarding the attack*, attackers have access to the ego-robot's viewpoint and utilize adversarial learning to generate imperceptible noises added to its original messages, to significantly degrade the output of the ego-robot. *Regarding the defense*, the ego-robot is not aware of the specific attacking strategy, but it can identify whether it has been attacked based on the change in output space.

Objective and challenge. On one hand, the ego-robot cannot totally trust others and need to carefully use the messages to avoid being attacked. On the other hand, they also cannot fully ignore the messages to keep away from attackers, since they still need complementary information to enhance their own restricted perception. The objective for each ego-robot is: *given a certain computation budget for consensus verification, how to make full use of the messages shared by others while avoiding being attacked?* The challenges for this problem lie in two aspects: (1) *generalizability*: how to achieve generalizable defense given that there could be different and unseen attackers; and (2) *scalability*: how to realize computationally-efficient defense especially when there are a large number of teammates.

3.2. Revisit of RANSAC

RANSAC uses a hypothesize-and-verify workflow to robustly fit a model based on a set of data points in the presence of noise. The workflow mainly includes three steps: (1) produce a set of model hypotheses by sampling minimal data points for the fitted model; (2) evaluate hypotheses with some consensus metrics such as the number of inliers; (3) choose the best hypothesis. An optional fourth step is of-

Algorithm 1: Workflow of ROBOSAC

Input: The total number of teammates S , messages from teammates $\{\mathbf{M}_i\}_{i=1,\dots,S}$, message of an ego-robot \mathbf{M}_0 , perception model f_θ , difference measure d , consensus threshold ϵ , sampling budget N , attacker ratio η , probability of at least one successful sampling p

Output: The perception results for the ego-robot \mathbf{Y}_0 at current timestamp

- 1: Obtain the perception results of only using the ego-robot's message: $\hat{\mathbf{Y}}_0 = f_\theta(\mathbf{M}_0)$
- 2: Calculate the guaranteed maximum number of collaborators : $s = \lfloor \frac{\ln [1 - (1-p)^{\frac{1}{N}}]}{\ln(1-\eta)} \rfloor$
- 3: $n = 0$
- 4: **while** $n < N$ **do**
- 5: $n \leftarrow n + 1$
- 6: Sample s teammates randomly $\{\mathbf{M}_j\}_{j=1,\dots,s}$
- 7: Obtain the perception results $\hat{\mathbf{Y}}_s = f_\theta(\mathbf{M}_0, \{\mathbf{M}_j\}_{j=1,\dots,s})$
- 8: **if** $d(\hat{\mathbf{Y}}_s, \hat{\mathbf{Y}}_0) \leq \epsilon$ **then** ▷ Consensus
- 9: $\mathbf{Y}_0 = \hat{\mathbf{Y}}_s$
- 10: **break** ▷ Early stop
- 11: **else if** $n = N - 1$ **then** ▷ Dissensus
- 12: $\mathbf{Y}_0 = \hat{\mathbf{Y}}_0$ ▷ No collaboration
- 13: **break**
- 14: **end if**
- 15: **end while**

ten employed in practice, which refines the chosen hypothesis with all inliers. RANSAC can compute the required number of sampling to ensure a high probability of at least one successful sampling under different outlier ratios.

3.3. Workflow of ROBOSAC

One naive solution to the problem is iteratively verifying each teammate based on the consensus in the output space, yet it is not scalable. We instead propose to iteratively sample a subset of teammates until they reach a consensus, given a certain sampling budget. Assume that the attacker ratio is known as η , and the ego-robot plans to use the information from s teammates so that it keeps on sampling s teammates until reaching a consensus. The sampling budget is N , and a successful sampling is one that contains no attackers amongst the sampled s robots. The detailed procedures for an ego-robot are: (1) produce perception results using its individual observation; (2) sample s teammates and fuse their messages to generate perception results; (3) verify the consensus between the results in (1) and (2); (4) output the results in (2) if there are no attackers, otherwise, continue to implement (2). Formally, our objective is to maximize the probability of at least one successful sampling which is calculated by $p = 1 - [1 - (1 - \eta)^s]^N$. The workflow is shown in Algorithm 1.

Performance-computation trade-off. Different from




Figure 2: **A numerical example of ROBOSAC** (probability of at least one successful sampling is 0.99). (a) Guaranteed maximum number of collaborators given a certain sampling budget. (b) The maximum number of sampling given a desired number of collaborators.

RANSAC, ROBOSAC can be customized for the computation (budget N) or the performance related to the amount of beneficial information (number of collaborators s). Given a probability p to ensure that there is at least one successful sampling within a sampling budget N , the maximum number of attacker-free collaborators that could be found is:

$$s = \left\lfloor \frac{\ln[1 - (1-p)^{\frac{1}{N}}]}{\ln(1-\eta)} \right\rfloor, \quad (1)$$

which means that the ego-robot is able to have s collaborators to enhance its perception in a safe manner. In turn, given a probability p to ensure that there is at least one successful sampling and a desired number of collaborators s , we need to sample at most:

$$N = \left\lceil \frac{\ln(1-p)}{\ln[1 - (1-\eta)^s]} \right\rceil, \quad (2)$$

which means that the sampled teammates can reach a consensus before the N -th sampling. One numerical example is illustrated in Fig. 2. Here the probability of at least one successful sampling is 0.99. In (a), given a sampling budget 5, the ego-robot can make sure to have 4 collaborators to improve the perception when the attacker ratio is 10%; in (b), given a desired number of collaborators 5, the ego-robot can make final decisions before the 12-th sampling when the attacker ratio is 20%.

3.4. Attacker Ratio Estimation

The ego-robot might not be aware of the attacker ratio when it enters a novel environment. To estimate the attacker ratio as quickly as possible, we develop an **Aggressive-to-Conservative Probing (A2CP)** approach, which starts from trusting all teammates and gradually reduces the number of sampled teammates if the previous attempt fails. The basic idea is that once we find a consensus subset with s robots, the ratio of benign collaborators will not be less than $\frac{s}{S}$, where S is the total number of teammates. Take $S = 5$ as

Algorithm 2: Workflow of A2CP with Retrospect

Input: The number of teammates S , messages from teammates $\{\mathbf{M}_i\}_{i=1,\dots,S}$, message of an ego-robot \mathbf{M}_0 , perception model f_θ , difference measure d , consensus threshold ϵ , ascending-ordered array of discretized possible ratios $\{\mathbf{R}_k = h_k/S\}_{k=1,\dots,K}$ ($h_k \in [1..S]$), sampling budget $N > K$, probability of at least one successful sampling p

Output: The estimated attacker ratio $\hat{\eta}$

```

1:  $\hat{\eta} \leftarrow 1.0$ 
2:  $\{\mathbf{U}_k\}_{k=1,\dots,K} \leftarrow \mathbf{0}$   $\triangleright$  Upper-bound of attempts

3:  $\{\mathbf{T}_k\}_{k=1,\dots,K} \leftarrow \mathbf{0}$   $\triangleright$  Counter of attempts
4: for  $k$  in  $[1, \dots, K]$  do
5:    $\mathbf{U}_k = \left\lceil \frac{\ln(1-p)}{\ln[1 - (1-\mathbf{R}_k)^{S(1-\mathbf{R}_k)}]} \right\rceil$ 
6: end for
7: for each frame in the scene do
8:    $n = 0$ 
9:   Obtain the individual results:  $\hat{\mathbf{Y}}_0 = f_\theta(\mathbf{M}_0)$ 
10:  while  $n < N$  and  $\mathbf{T} < \mathbf{U}$  do
11:    for  $k \in [1, \dots, K]$  do
12:      if  $\mathbf{T}_k < \mathbf{U}_k$  then
13:         $n \leftarrow n + 1$ 
14:         $\eta \leftarrow \mathbf{R}_k$ 
15:        Sample  $S(1-\eta)$  teammates
16:         $\{\mathbf{M}_j\}_{j=1,\dots,S(1-\eta)}$ 
17:        Obtain the perception results
18:         $\hat{\mathbf{Y}}_s = f_\theta(\mathbf{M}_0, \{\mathbf{M}_j\}_{j=1,\dots,S(1-\eta)})$ 
19:        if  $d(\hat{\mathbf{Y}}_s, \hat{\mathbf{Y}}_0) \leq \epsilon$  then  $\triangleright$  Consensus
20:           $\hat{\eta} \leftarrow \eta$   $\triangleright$  Update
21:           $\mathbf{T}_k \leftarrow \mathbf{U}_k, \dots, \mathbf{T}_K \leftarrow \mathbf{U}_K$ 
22:           $\triangleright$  Stop probing higher ratios
23:          break
24:        else  $\triangleright$  Dissensus
25:           $\mathbf{T}_k \leftarrow \mathbf{T}_k + 1$ 
26:        end if
27:      end for
28:    end while
29:  end for

```

an example, we pre-define an array of discretized possible ratios in ascending order $\mathbf{R} = [0.0, 0.2, 0.4, 0.6, 0.8]$ and start from probing $\eta = 0.0$ ($s = 5$) which indicates no attackers. If the first attempt succeeds, there is no need to probe other ratios and the ego-robot can collaborate with all teammates. If not, we continue to probe a higher ratio with fewer collaborators (in this case $\eta = 0.2$ and $s = 4$) and stop testing the unprobed ratios once consensus is verified.

Note that we further propose a retrospect-based mechanism to avoid missing the probed lower ratios due to randomness. For example, if probing of $\eta = 0.2$, $s = 4$ fails yet probing of $\eta = 0.4$, $s = 3$ succeeds, we cannot conclude the ratio yet because it is more difficult for the previ-

ous more aggressive attempt to reach a consensus. Actually, there could be either 1 or 2 attackers. We continue to probe $\eta = 0.2, s = 4$ until enough attempts have been made. We use \mathbf{T}_k to count the number of probing for a ratio \mathbf{R}_k , and \mathbf{U}_k to denote the upper bound probing attempt of this ratio, which is derived from equation (2). See Algorithm 2 for the overall workflow of the proposed attacker ratio estimator.

4. Adversarially Robust Collaborative 3D Detection with ROBOSAC

In this section, we apply our ROBOSAC framework in collaborative 3D object detection under adversarial attacks in autonomous driving. Vehicle-to-vehicle (V2V) communication could improve the robustness, safety, and efficiency of autonomous driving systems in light of more viewpoints and computational resources. However, there are still safety concerns regarding the communication channel [14]. We consider that multiple vehicles located in the same geographical location are sharing information. An *ego-vehicle* attempts to maximize the usage of benign *collaborators* while defending against *attackers*.

Preliminaries. Each vehicle indexed by i ($i = 0, \dots, S$) is equipped with a 3D LiDAR sensor to generate a bird’s eye view (BEV) occupancy grid map $\mathbf{B}_i \in \{0, 1\}^{W \times L \times H}$ defined in its local coordinate. Here S is the total number of teammates, the dimensions W , L , and H respectively denote the width, length, and height of the BEV map.

Collaborative 3D detection. A collaborative detector shared amongst vehicles is composed of an encoder denoted by f_ψ , an aggregator and a decoder which are collectively denoted by f_θ for simplicity. f_ψ will take the BEV map as input and generate an intermediate feature map $\mathbf{M}_i \in \mathbb{R}^{\frac{W}{K} \times \frac{L}{K} \times C} = f_\psi(\mathbf{B}_i)$ as the transmitted message, where K indicates the downsampling scale in the neural network, and C denotes the feature dimension. After deciding to use the messages from s teammates $\{\mathbf{M}_j\}_{j=1, \dots, s}$ ($s \leq S$), the ego-vehicle indexed by 0 will use f_θ to produce a set of bounding boxes $\hat{\mathbf{Y}}_s = f_\theta(\mathbf{M}_0, \{\mathbf{M}_j\}_{j=1, \dots, s})$. During training, f_ψ and f_θ are jointly learned by minimizing the detection loss $\mathcal{L}_{det}(\hat{\mathbf{Y}}_s, \mathbf{Y}_{gt})$, where \mathbf{Y}_{gt} denotes ground-truth boxes. \mathcal{L}_{det} includes a classification loss and a regression loss following existing detectors [45].

Adversarial attack. We consider the white box attack where the attacker has full access to the target model because the detector model is shared amongst vehicles. The adversarial message is generated by gradient-based optimization to maximize the vehicle’s detection errors. Specifically, at inference, model parameters are frozen, and the objective for attacker v is to fool the ego-vehicle by sending an indistinguishable adversarial message $\mathbf{M}_v + \delta$:

$$\max_{\|\delta\| \leq \Delta} \mathcal{L}_{det}(f_\theta(\mathbf{M}_0, \mathbf{M}_v + \delta, \{\mathbf{M}_j\}_{j \neq v}), \mathbf{Y}_{gt}), \quad (3)$$

where $\delta \in \mathbb{R}^{\frac{W}{K} \times \frac{L}{K} \times C}$ with the same size as \mathbf{M}_v is the optimized perturbation, and is constrained by Δ to ensure its imperceptibility. In practice, the detection results of solely using the message of ego-vehicle $\hat{\mathbf{Y}}_0 = f_\theta(\mathbf{M}_0)$ could replace \mathbf{Y}_{gt} in case the ground-truth is not available. Such kind of carefully-crafted adversarial messages can create a lot of false positives/negatives for the ego-vehicle, which raises concerns for safety-critical autonomous vehicles.

Adversarial defense. To defend against such kind of adversarial messages, and maximize the usage of complementary messages, the ego-vehicle is supposed to intelligently select teammates to ensure that there are no adversarial messages in the exploited messages $\{\mathbf{M}_j\}_{j=1, \dots, s}$, while maximizing the number of benign collaborators s . To this end, we use the ROBOSAC framework shown in Algorithm 1. Specifically, the ego-vehicle obtains S messages $\{\mathbf{M}_i\}_{i=1, \dots, S}$ from all the teammates, and it samples s messages each time $\{\mathbf{M}_j\}_{j=1, \dots, s}$ and then verify the consensus between two detections $\hat{\mathbf{Y}}_s = f_\theta(\mathbf{M}_0, \{\mathbf{M}_j\}_{j=1, \dots, s})$ and $\hat{\mathbf{Y}}_0 = f_\theta(\mathbf{M}_0)$. Since the adversarial learning will result in a large number of false detections, $\hat{\mathbf{Y}}_s$ will significantly differ from $\hat{\mathbf{Y}}_0$ once there exist adversarial messages in $\{\mathbf{M}_j\}_{j=1, \dots, s}$. In 3D detection, the difference measure d is represented by the intersection-over-union (IoU) between two sets of boxes after Hungarian matching [46], and a consensus threshold ϵ is used to determine whether or not there is consensus. The ego-vehicle will keep on sampling until the sampled vehicles reach a consensus. In practice, we may want to give a certain sampling budget N to avoid too many computations for $\hat{\mathbf{Y}}_s = f_\theta(\mathbf{M}_0, \{\mathbf{M}_j\}_{j=1, \dots, s})$, and the ego-vehicle can compute the maximum number of benign collaborators which could be found, and use it as the sample size s based on Eq. 1, with the given sampling budget as well as the attacker ratio. In a word, the ego-vehicle can have s collaborators in consensus to aid in the promotion of its perception within N sampling steps.

5. Experiments

5.1. Experimental setup

Dataset and detector. We employ V2X-Sim [22] to verify our method. The dataset contains 5Hz LiDAR point clouds recorded by different vehicles at the same intersection. Regarding the multi-robot detector, we employ a simple average-based collaborative perception method which calculates the mean of the feature maps from different vehicles and feeds the aggregated features into the decoder to generate the final perception results. The detector backbone is a simple anchor-based method [45]. We follow the training procedures and evaluation protocols in [11].

Adversarial attack implementation. We use three attackers: projected gradient descent (PGD) [47], Carlini & Wagner Attack (C&W) [48], and basic iterative method

ROBOSAC			Actual Steps			Success
η	s	N	Avg	Min	Max	Rate
0.2	1	3	1.32	1	6	0.96
	2	5	1.76	1	6	0.97
	3	7	2.31	1	7	1.00
	4	9	4.89	1	19	0.89
0.4	1	6	1.80	1	8	0.98
	2	11	3.06	1	11	1.00
	3	19	10.36	1	39	0.86
0.6	1	10	2.46	1	8	1.00
	2	27	8.29	1	46	0.97
0.8	1	21	4.73	1	17	1.00

Table 1: **Validation of our derivation** under different attacker ratio η and desired number of collaborators s .

Setup	AP		Avg. Steps	FPS
	IoU=0.5	IoU=0.7		
Dynamic team	77.5	74.9	2.73	14.2
Static team	78.8	76.5	2.30	36.8

Table 2: **Comparison between scenarios of dynamic and static team**: averaged over 10 experiments. Agent 1 in scene #8, with $\eta = 0.2$, $s = 3$, $N = 7$, $\epsilon = 0.3$

(BIM) [34]. The number of iterations in PGD/BIM is set to 15 (iteration of C&W is set to 30), the step size is set to 0.1, and the magnitude is $\Delta = 0.5$. The consensus threshold is set to $\epsilon = 0.3$. We use average precision (AP) at IoU thresholds of 0.5 and 0.7 to evaluate the detection performance, and we utilize the performance of detection to evaluate the effectiveness of our defense strategies. We adopt scene #8 with 100 frames (each frame contains 6 vehicles).

Static and dynamic teammates. Regarding the static team, the ego-robot can establish a stable partnership with reliable collaborators. Once a requisite number of attacker-free teammates is identified, the ego-robot becomes capable of distinguishing potential attackers and subsequently disregards their messages in subsequent frames. In the scenario of a dynamic team, we need to deploy ROBOSAC at each frame, since now the ego-robot fails to rely on its past judgments of the teammates. For the first scenario, the computational burden introduced by ROBOSAC is negligible because only several initial frames need consensus verification. Addressing the more complex second scenario requires striking a balance between performance and computational demands under diverse conditions.

5.2. Quantitative results

Validation of the derivation. Given the attacker ratio η and the desired number of collaborators s , the maximum number of sampling N could be given by Eq. 2 based on

Method	AP		Success Rate
	IoU=0.5	IoU=0.7	
Upper-bound++	81.3	79.8	–
Upper-bound	78.0	76.0	–
Sampling budget: 7	77.3	74.7	1.00
Sampling budget: 5	76.6	74.1	0.96
Sampling budget: 3	74.8	73.2	0.77
Sampling budget: 1	69.8	67.4	0.45
Lower-bound	63.5	60.1	–
No Defense	39.7	39.0	–

Table 3: **Detection performance with different sampling budgets** when $\eta = 0.2$ and $s = 3$. Upper-bound++ denotes collaborative perception with all 5 benign robots. Upper-bound means collaborative perception with 3 out of 5 benign robots. Lower-bound denotes individual perception.

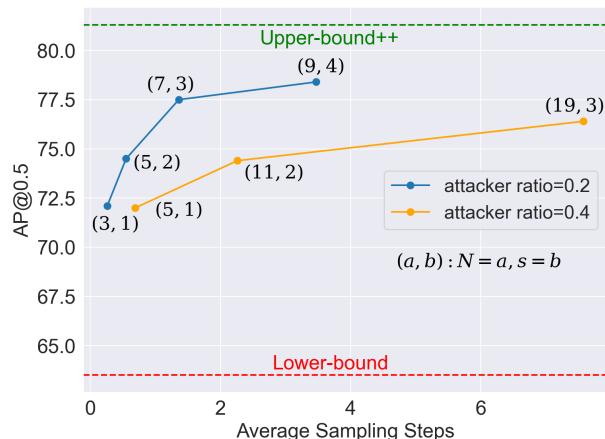


Figure 3: **Performance-computation trade-off plot.**

a certain probability $p = 0.99$. We validate this formulation by allowing the ego-vehicle to keep on sampling until reaching consensus, under different η and s , and the results are shown in Table 1. For each processed frame, it is considered a success if the actual steps taken to achieve consensus are not larger than the computed N , and the ratio of successful frames is referred to as the *success rate*. We see that a consensus can be mostly achieved within the theoretical upper-bound N . Meanwhile, the average sampling steps taken for consensus are usually less than 50% of N .

Static team. Since the teammates are not changed, so once the attackers are identified, the ego-robot could avoid being attacked in the following frames. The results are shown in Table 2. We see that the frames-per-second (FPS) is 36.8, which is satisfactory for real-time applications if the ROBOSAC is only needed at the initial stage.

Dynamic team. In the scenario of dynamic teammates, the ego-robot needs to identify the collaborators at each frame. Since the actual sampling steps are usually much

Method	AP		Success Rate	FPS
	IoU=0.5	IoU=0.7		
w/ Temporal Consistency	75.8	73.6	0.85	19.2
w/o Temporal Consistency	77.5	74.9	0.98	14.2

Table 4: **Comparison between w/ temporal consistency and w/o temporal consistency:** agent 1 in scene #8, with $\eta = 0.2$, $s = 3$, $N = 7$, $\epsilon = 0.3$

Method	AP	
	IoU=0.5	IoU=0.7
Upper-bound++	81.8	79.6
ROBOSAC (against PGD attack)	77.9	75.6
PGD Trained (White-box Defense)	75.6	73.0
ROBOSAC (against C&W attack)	74.5	71.1
C&W on PGD Trained (Black-box Defense)	43.2	40.8
Lower-bound	64.1	62.0
No Defense (PGD attack)	44.2	43.7

Table 5: **Quantitative results of generalizability test** on agent 1 in scene #8, with $\eta = 0.2$, $s = 3$, $\epsilon = 0.3$. The average sampling step and FPS are 2.67 and 15.5. Black-box defense is unaware of the attacker, while white-box defense knows the attacker and employs the adversarial training. Upper-bound++ denotes collaborative perception with all 5 benign robots. Lower-bound is individual perception.

less than the upper bound, we can lower the sampling budget in practice to pursue a balance between performance and computation. Assume we aim to find three attacker-free teammates when $\eta = 0.2$, the upper bound of the sampling budget is then $N = 7$. To save computations, we set the sampling budget as $N = 1, 3, 5, 7$ respectively and report the detection performance as well as the success rate (successfully find three attacker-free teammates within the specified sampling budget) in Table 3. We find that limiting the sampling budget to 5 still maintains a success rate of 0.96. Meanwhile, only sampling once in each frame (about 15% of N) can still achieve a success rate of 0.45 and an AP of 69.8/67.4 which are still better than solely using ego-vehicle’s information (63.5/60.1). In addition, we conduct precision-computation trade-off analyses under different $N - s$ pairs. As shown in Fig. 3, detection precision is higher when there are more collaborators, at the cost of more computations to reject attackers.

Temporal consistency. We further propose to use temporal consistency instead of the difference between collaborative and individual perception to save computations. Specifically, we compare the current output with the previous output for consensus verification. This can help to improve efficiency because we can reduce the times of model forward caused by individual perception. The results are shown in Table 4. We see that FPS is improved from 14.2

Ablation	AP		Success Rate
	IoU=0.5	IoU=0.7	
Threshold $\epsilon = 0.1$	77.5	75.4	0.96
Threshold $\epsilon = 0.2$	78.0	75.0	0.98
Threshold $\epsilon = 0.3$	77.8	76.1	1.00
Threshold $\epsilon = 0.4$	76.7	74.7	0.95
Threshold $\epsilon = 0.5$	76.1	73.9	0.87
PGD, 10 iterations	78.3	75.2	0.95
PGD, 15 iterations	77.8	76.1	1.00
BIM, 10 iterations	77.8	75.2	0.95
BIM, 15 iterations	77.5	74.5	0.97

Table 6: **Quantitative results of ablation studies.**

Metrics	Attacker Ratio					
	0.0	0.2	0.4	0.6	0.8	1.0
Frame # to reach final estimation	0	0.8	2.1	3.1	1.3	0
Final estimated ratio	0.0	0.22	0.42	0.60	0.80	1.0
Error of estimation	0.0	0.02	0.02	0.0	0.0	0.0
Total sampling steps	1	8.2	20.5	37.9	59	77
Success Rate	1.0	0.9	0.9	1.0	1.0	1.0

Table 7: **Quantitative results of attacker ratio estimation** on agent 1 in scene #8, with $N = 5$, $\epsilon = 0.3$, $\mathbf{R} = [0.0, 0.2, 0.4, 0.6, 0.8]$. All results are averaged over 10 repeated experiments.

to 19.2, yet the performance is still comparable to using individual perception results as a reference.

Generalizability. We compare the generalizability of ROBOSAC to adversarial training. We use PGD [47] which is the strongest one-stage gradient-based adversarial attack. The results are shown in Table 5: although adversarial training with PGD could effectively defend the PGD attack with a precision 75.6 (IoU=0.5), using a different attacker, Carlini & Wagner Attack [48], will largely degrade the precision to 43.2 (IoU=0.5). In contrast, our method can achieve comparable precision under both two attackers (77.9@IoU0.5 under PGD attack, and 74.5@IoU=0.5 under C&W attack). Our better generalizability stems from being attacker-agnostic: we do not rely on the knowledge of attackers while adversarial training does.

Attacker ratio estimation. In practice, using a budget of five samplings within a single frame yields acceptable results. As shown in Table 7, within the initial few frames, the actual attacker ratio can be efficiently probed. The estimated ratio can then be used to carry out ROBOSAC steps.

5.3. Ablation studies

We conduct ablation studies on consensus threshold as well as attack approaches (the other parameters are $N = 7$, $s = 3$, $\eta = 0.2$), and the results are shown in Table 6. Regarding the consensus threshold, a lower threshold indicates a lesser tolerance for outcomes that differ from indi-

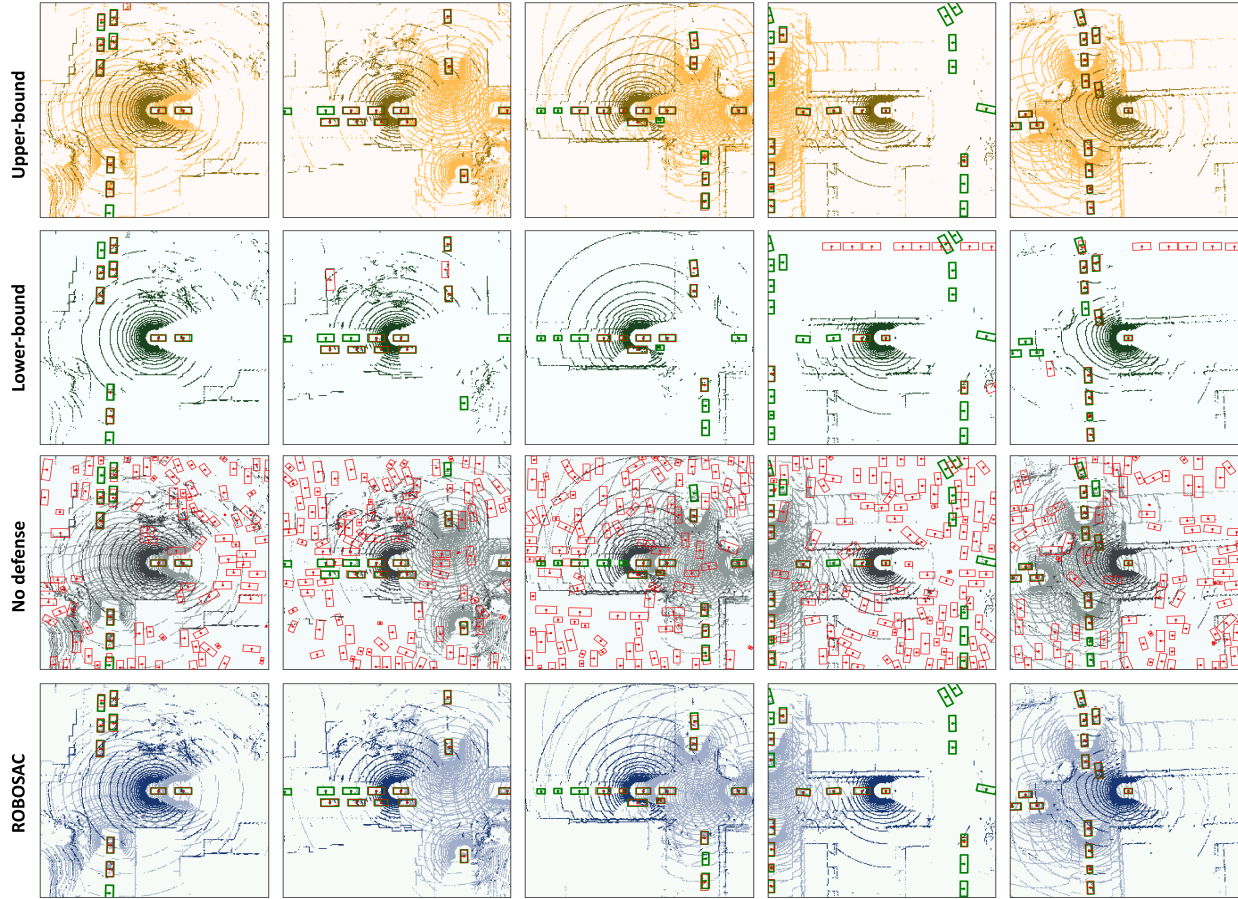


Figure 4: **Visualization of perception results** on V2X-Sim [22]. Red boxes denote predictions, and green boxes are ground truth. Upper-bound, lower-bound have the same meanings with that in Table 3.

vidual results, whereas a higher threshold indicates a higher likelihood of believing the variation is good. We find that our method achieves comparable performance with different consensus threshold. Meanwhile, our method remains unaffected by the type of attacker.

5.4. Computational Cost

Our method’s performance depends on the sampling steps, each requiring a forward propagation. Using an NVIDIA RTX 3090 GPU, the baseline detector averages 17ms for ego-only predictions and 27ms for collaborative predictions. With a 5Hz dataset, we allocate a 200ms time budget per frame, allowing up to 7 samplings. As actual sampling steps are lower, we can achieve a better performance-computation trade-off by sampling more attacker-free teammates, as shown in Fig. 3.

6. Limitations

We assume that the attacker ratio is fixed, yet the attacker ratio may vary in practice. Besides, we assume that although the input adversarial noise is imperceptible, its ef-

fect on the network output is significant (see Fig. 4), which has been observed in most existing attacking methods. Future attackers might develop dangerous yet subtle perturbations in both the input and output to bypass our “outlier-detection-based” defense mechanism, although currently, we are not aware of any such attacks.

7. Conclusion

In this work, we propose a novel adversarially robust collaborative perception framework termed ROBOSAC. It makes as much use of messages from benign collaborators as possible while resisting adversarial attackers within a certain computation budget. Moreover, we develop an aggressive-to-conservative probing method with retrospect for attacker ratio estimation in scenarios of unknown ratios. We validate our method in collaborative 3D detection for autonomous driving. Unlike adversarial training, our approach relies on the consistency in output space rather than the knowledge of a specific adversarial noise, thus is more generalizable. We believe our work will further improve the adversarial robustness of multi-robot systems.

References

- [1] Yuexin Ma, Tai Wang, Xuyang Bai, Huitong Yang, Yuenan Hou, Yaming Wang, Yu Qiao, Ruigang Yang, Dinesh Manocha, and Xinge Zhu. Vision-centric bev perception: A survey. *arXiv preprint arXiv:2208.02797*, 2022.
- [2] Hang Yin, Anastasia Varava, and Danica Kragic. Modeling, learning, perception, and control methods for deformable object manipulation. *Science Robotics*, 6(54):eabd8803, 2021.
- [3] You Li and Javier Ibanez-Guzman. Lidar for autonomous driving: The principles, challenges, and trends for automotive lidar and perception systems. *IEEE Signal Processing Magazine*, 37(4):50–61, 2020.
- [4] Luis Roldao, Raoul De Charette, and Anne Verroust-Blondet. 3d semantic scene completion: A survey. *International Journal of Computer Vision*, 130(8):1978–2005, 2022.
- [5] Tsun-Hsuan Wang, Sivabalan Manivasagam, Ming Liang, Bin Yang, Wenyuan Zeng, and Raquel Urtasun. V2vnet: Vehicle-to-vehicle communication for joint perception and prediction. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 605–621, 2020.
- [6] Yiming Li, Juexiao Zhang, Dekun Ma, Yue Wang, and Chen Feng. Multi-robot scene completion: Towards task-agnostic collaborative perception. In *6th Annual Conference on Robot Learning*, 2022.
- [7] Runsheng Xu, Zhengzhong Tu, Hao Xiang, Wei Shao, Bolei Zhou, and Jiaqi Ma. Cobevt: Cooperative bird’s eye view semantic segmentation with sparse transformers. In *6th Annual Conference on Robot Learning*, 2022.
- [8] Qi Chen, Sihai Tang, Q. Yang, and Song Fu. Cooper: Cooperative perception for connected autonomous vehicles based on 3d point clouds. In *IEEE 39th International Conference on Distributed Computing Systems (ICDCS)*, pages 514–524, 2019.
- [9] Eduardo Arnold, Mehrdad Dianati, Robert de Temple, and Saber Fallah. Cooperative perception for 3d object detection in driving scenarios using infrastructure sensors. *IEEE Transactions on Intelligent Transportation Systems*, 2020.
- [10] Yen-Cheng Liu, Junjiao Tian, Nathaniel Glaser, and Zsolt Kira. When2com: multi-agent perception via communication graph grouping. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4106–4115, 2020.
- [11] Yiming Li, Shunli Ren, Pengxiang Wu, Siheng Chen, Chen Feng, and Wenjun Zhang. Learning distilled collaboration graph for multi-agent perception. In *Advances in Neural Information Processing Systems*, volume 34, 2021.
- [12] Sanbao Su, Yiming Li, Sihong He, Songyang Han, Chen Feng, Caiwen Ding, and Fei Miao. Uncertainty quantification of collaborative detection for self-driving. In *IEEE International Conference on Robotics and Automation*, 2023.
- [13] Christian Szegedy, W. Zaremba, Ilya Sutskever, Joan Bruna, D. Erhan, Ian J. Goodfellow, and R. Fergus. Intriguing properties of neural networks. In *International Conference on Learning Representation*, 2014.
- [14] James Tu, Tsunhsuan Wang, Jingkang Wang, Sivabalan Manivasagam, Mengye Ren, and Raquel Urtasun. Adversarial attacks on multi-agent communication. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7768–7777, 2021.
- [15] Jiliang Zhang and Chen Li. Adversarial examples: Opportunities and challenges. *IEEE transactions on neural networks and learning systems*, 31(7):2578–2593, 2019.
- [16] Aditi Raghunathan, Sang Michael Xie, Fanny Yang, John Duchi, and Percy Liang. Understanding and mitigating the tradeoff between robustness and accuracy. In *Proceedings of the 37th International Conference on Machine Learning*, pages 7909–7919. PMLR, 2020.
- [17] José María Martínez-Otzeta, Itsaso Rodríguez-Moreno, Iñigo Mendiola, and Basilio Sierra. Ransac for robotic applications: A survey. *Sensors*, 23(1):327, 2022.
- [18] Anh-Quan Cao and Raoul de Charette. Monoscene: Monocular 3d semantic scene completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3991–4001, 2022.
- [19] Haobo Zuo, Changhong Fu, Sihang Li, Kunhan Lu, Yiming Li, and Chen Feng. Adversarial blur-deblur network for robust uav tracking. *IEEE Robotics and Automation Letters*, 8(2):1101–1108, 2023.
- [20] Yiming Li, Zhiding Yu, Christopher Choy, Chaowei Xiao, Jose M Alvarez, Sanja Fidler, Chen Feng, and Anima Anandkumar. Voxformer: Sparse voxel transformer for camera-based 3d semantic scene completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [21] Sihang Li, Changhong Fu, Kunhan Lu, Haobo Zuo, Yiming Li, and Chen Feng. Boosting uav tracking with voxel-based trajectory-aware pre-training. *IEEE Robotics and Automation Letters*, 8(2):1133–1140, 2023.
- [22] Yiming Li, Dekun Ma, Ziyang An, Zixun Wang, Yiqi Zhong, Siheng Chen, and Chen Feng. V2x-sim: Multi-agent collaborative perception dataset and benchmark for autonomous driving. *IEEE Robotics and Automation Letters*, 7(4):10914–10921, 2022.
- [23] Yen-Cheng Liu, Junjiao Tian, Chih-Yao Ma, Nathan Glaser, Chia-Wen Kuo, and Zsolt Kira. Who2com: Collaborative perception via learnable handshake communication. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6876–6883, 2020.
- [24] Yang Zhou, Jiahong Xiao, Yue Zhou, and Giuseppe Loianno. Multi-robot collaborative perception with graph neural networks. *IEEE Robotics and Automation Letters*, 2022.
- [25] Sanbao Su, Songyang Han, Yiming Li, Zhili Zhang, Chen Feng, Caiwen Ding, and Fei Miao. Collaborative multi-object tracking with conformal uncertainty propagation. *arXiv preprint arXiv:2303.14346*, 2023.

- [26] Seong-Woo Kim, B. Qin, Z. J. Chong, Xiaotong Shen, Wei Liu, M. Ang, Emilio Frazzoli, and D. Rus. Multivehicle cooperative driving using cooperative perception: Design and experimental validation. *IEEE Transactions on Intelligent Transportation Systems*, 16:663–680, 2015.
- [27] Nicholas Vadivelu, Mengye Ren, James Tu, Jingkang Wang, and Raquel Urtasun. Learning to communicate and correct pose errors. In *Conference on Robot Learning*, 2020.
- [28] J. Tu, Mengye Ren, Siva Manivasagam, Ming Liang, Bin Yang, R. Du, Frank C Z Cheng, and R. Urtasun. Physically realizable adversarial examples for lidar object detection. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 13713–13722, 2020.
- [29] Yiming Li, Congcong Wen, Felix Juefei-Xu, and Chen Feng. Fooling lidar perception via adversarial trajectory perturbation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 7898–7907, October 2021.
- [30] Yulong Cao, Chaowei Xiao, Benjamin Cyr, Yimeng Zhou, Won Park, Sara Rampazzi, Q. Chen, K. Fu, and Z. Morley Mao. Adversarial sensor attack on lidar-based perception in autonomous driving. In *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security*, 2019.
- [31] James Tu, Huichen Li, Xinchun Yan, Mengye Ren, Yun Chen, Ming Liang, Eilyan Bitar, Ersin Yumer, and Raquel Urtasun. Exploring adversarial robustness of multi-sensor perception systems in self driving. In *Conference on Robot Learning*, 2021.
- [32] Kui Ren, Tianhang Zheng, Zhan Qin, and Xue Liu. Adversarial attacks and defenses in deep learning. *Engineering*, 6(3):346–360, 2020.
- [33] Ian Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *International Conference on Learning Representations*, 2015.
- [34] Alexey Kurakin, Ian J Goodfellow, and Samy Bengio. Adversarial examples in the physical world. In *International Conference on Learning Representation*, 2017.
- [35] Wieland Brendel, Jonas Rauber, and Matthias Bethge. Decision-based adversarial attacks: Reliable attacks against black-box machine learning models. In *International Conference on Learning Representations*, 2018.
- [36] Shuyu Cheng, Yinpeng Dong, Tianyu Pang, Hang Su, and Jun Zhu. Improving black-box adversarial attacks with a transfer-based prior. *Advances in neural information processing systems*, 32, 2019.
- [37] Cihang Xie, Zhishuai Zhang, Yuyin Zhou, Song Bai, Jianyu Wang, Zhou Ren, and Alan L Yuille. Improving transferability of adversarial examples with input diversity. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2730–2739, 2019.
- [38] Martin A Fischler and Robert C Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981.
- [39] Phil Torr and Andrew Zisserman. Robust computation and parametrization of multiple view relations. In *Sixth International Conference on Computer Vision (IEEE Cat. No. 98CH36271)*, pages 727–732. IEEE, 1998.
- [40] Philip HS Torr and Andrew Zisserman. Mlesac: A new robust estimator with application to estimating image geometry. *Computer vision and image understanding*, 78(1):138–156, 2000.
- [41] Tomas Hodan, Daniel Barath, and Jiri Matas. Epos: Estimating 6d pose of objects with symmetries. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11703–11712, 2020.
- [42] Eugene Valassakis, Kamil Dreczkowski, and Edward Johns. Learning eye-in-hand camera calibration from a single image. In *Conference on Robot Learning*, pages 1336–1346. PMLR, 2022.
- [43] Kanji Tanaka and Eiji Kondo. Incremental ransac for online relocation in large dynamic environments. In *Proceedings 2006 IEEE International Conference on Robotics and Automation, 2006. ICRA 2006.*, pages 68–75. IEEE, 2006.
- [44] Shimon Ullman. The interpretation of structure from motion. *Proceedings of the Royal Society of London. Series B. Biological Sciences*, 203(1153):405–426, 1979.
- [45] Wenjie Luo, Bin Yang, and Raquel Urtasun. Fast and furious: Real time end-to-end 3d detection, tracking and motion forecasting with a single convolutional net. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 3569–3577, 2018.
- [46] Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955.
- [47] Aleksander Madry, Aleksandar Makelev, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representation*, 2018.
- [48] Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 39–57. IEEE, 2017.