

When Autonomous Vehicle Meets V2X Cooperative Perception: How Far Are We?

An Guo¹, Shuoxiao Zhang¹, Enyi Tang^{1*}, Xinyu Gao¹, Haomin Pang², Haoxiang Tian³,
Yanzhou Mu¹, Wu Wen², Chunrong Fang¹ and Zhenyu Chen^{4*}

¹State Key Laboratory for Novel Software Technology, Nanjing University, China

{guoan218, sx_zhang, xinyugao, muyanzhou}@smail.nju.edu.cn, {eytang, fangchunrong}@nju.edu.cn

²Guangzhou University, China, panghaomin@e.gzhu.edu.cn, wenwu@gzhu.edu.cn

³Nanyang Technological University, Singapore, haoxiang.tian@ntu.edu.sg

⁴Shenzhen Research Institute of Nanjing University, China, zyichen@nju.edu.cn

Abstract—Perceiving the complex driving environment precisely is crucial to the safe operation of autonomous vehicles. With the tremendous advancement of deep learning and communication technology, Vehicle-to-Everything (V2X) cooperative perception has the potential to address limitations in sensing distant objects and occlusion for a single-agent perception system. V2X cooperative perception systems are software systems characterized by diverse sensor types and cooperative agents, varying fusion schemes, and operation under different communication conditions. Therefore, their complex composition gives rise to numerous operational challenges. Furthermore, when cooperative perception systems produce erroneous predictions, the types of errors and their underlying causes remain insufficiently explored.

To bridge this gap, we take an initial step by conducting an empirical study of V2X cooperative perception. To systematically evaluate the impact of cooperative perception on the ego vehicle's perception performance, we identify and analyze six prevalent error patterns in cooperative perception systems. We further conduct a systematic evaluation of the critical components of these systems through our large-scale study and identify the following key findings: (1) The LiDAR-based cooperation configuration exhibits the highest perception performance; (2) Vehicle-to-infrastructure (V2I) and vehicle-to-vehicle (V2V) communication exhibit distinct cooperative perception performance under different fusion schemes; (3) Increased cooperative perception errors may result in a higher frequency of driving violations; (4) Cooperative perception systems are not robust against communication interference when running online. Our results reveal potential risks and vulnerabilities in critical components of cooperative perception systems. We hope that our findings can better promote the design and repair of cooperative perception systems.

Keywords—Autonomous Driving Systems, Cooperative Perception, Offline and Online Testing

I. INTRODUCTION

Autonomous driving systems (ADSs) have attracted significant attention from both industry and academia due to their great potential and long-term value [1]. Among the core components of ADS, the perception module plays a pivotal role, processing and interpreting multi-modal sensor data through deep neural networks (DNNs) to detect and localize obstacles. Recent advancements in deep learning and sensor fusion technologies have substantially improved the accuracy and robustness of perception systems in autonomous vehicles [2], [3]. Despite these remarkable developments, single-agent perception systems often produce inaccurate or

incomplete results due to inherent single-view limitations, which can significantly compromise the safety and performance of ADS. Unfortunately, despite extensive testing of ADS on public roads for tens of thousands of hours [4], these limitations can still result in incorrect system behavior or even severe accidents [5], [6].

Vehicle-to-everything (V2X) cooperative perception technology has emerged as a promising solution to address the inherent limitations of single-agent perception systems, particularly in detecting occluded objects and long-range obstacles [7], [8]. By enabling communication among road participants (e.g., connected vehicles and roadside units), V2X facilitates the sharing and fusion of perception information, thereby significantly improving overall perception accuracy. In particular, vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I) systems have received increasing attention. Major automotive manufacturers and technology leaders, including Ford, Volvo, and Baidu, are heavily investing in V2X platform development, aiming to empower autonomous vehicles with enhanced environmental perception capabilities for safe navigation in complex urban and highway scenarios [9], [10].

The rapid advancement of V2X cooperative perception systems has introduced a range of challenges and concerns. The primary concern lies in the reliability of the information and the effectiveness of its fusion. On the one hand, cooperative agents may introduce erroneous or misleading information into the perception pipeline. On the other hand, even when the information provided by cooperative agents is correct, improper fusion with the existing perception pipeline may still result in inaccurate perception outcomes. As a highly complex system, V2X cooperative perception integrates multiple critical components, including heterogeneous sensors (e.g., LiDAR and camera), diverse cooperating agents (e.g., V2V and V2I), environment understanding, and reliable wireless communication. While such integration can effectively address the performance limitations of standalone perception systems, it also introduces new vulnerabilities. Failures or inconsistencies in any individual component may undermine the robustness of the entire V2X pipeline, potentially leading to inaccurate perception and unsafe decision-making by the ego vehicle. Therefore, it is imperative to conduct comprehensive testing and analysis of each key component to identify potential sources of fragility and to evaluate their impact on the overall reliability of cooperative perception.

To bridge this gap, we take an initial step in this paper by

*Enyi Tang and Zhenyu Chen are corresponding authors.

conducting an empirical study of V2X cooperative perception systems. Figure 1 presents the high-level design and workflow of our study. We begin by performing a comprehensive analysis of V2X cooperative perception systems with different fusion schemes (detailed in Section II-A) and identifying six distinct cooperative perception error patterns. To investigate the performance of key components in the V2X cooperative perception process, we have formulated four research questions investigating the following aspects: (1) the effect of heterogeneous sensor configurations on perception performance (RQ1), (2) the performance differences between V2V and V2I cooperation modes (RQ2), (3) how cooperative perception errors lead to driving violations (RQ3), and (4) the impact of communication issues on system effectiveness (RQ4). These questions aim to identify critical challenges and highlight potential opportunities for advancing cooperative perception. Finally, we conduct empirical evaluations for each research question and identify several key findings, which we expect could provide feedback to enhance the quality of cooperative perception systems. The detailed findings are summarized as follows: (1) LiDAR-based cooperative configurations achieve the highest perception performance among various sensor setups; (2) V2I and V2V cooperation exhibit distinct performance across fusion schemes, with V2I performs better under the early fusion mechanism and V2V showing superior perception performance under intermediate and late fusion schemes; (3) increased cooperative perception errors could lead to a higher frequency of driving violations; and (4) cooperative perception systems are vulnerable to communication interference. Our findings reveal potential risks and vulnerabilities in critical components of cooperative perception systems, providing insights that could help improve their overall reliability.

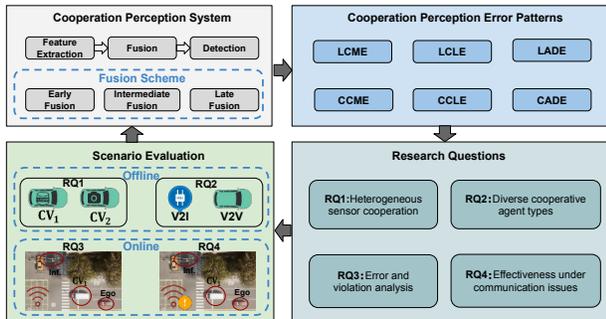


Fig. 1: Workflow of the high-level empirical study on cooperative perception.

The main contributions of this paper are summarized as follows:

- To the best of our knowledge, this work presents the first large-scale empirical study of a V2X cooperative perception system. We conduct a comprehensive analysis of error patterns to evaluate the performance of critical components within the cooperative perception systems.
- We further discuss existing V2X cooperative perception systems and outline potential future directions, including the advantages of various configurations in heterogeneous environments and opportunities for driving performance enhancement. Moreover, we come up with several novel findings and insights. These findings and insights can

benefit future research on cooperative perception and support the practical deployment of ADSs.

- To support the open science community, we make our source code publicly accessible¹ to facilitate the replication of our study and its application in extensive contexts.

This paper represents one of the earliest efforts to investigate cooperative perception systems. This study enables a quantitative investigation of these critical questions and facilitates further research in quality assurance for cooperative perception. In general, V2X cooperative perception plays a key role in enabling ADSs and potentially impacts autonomous driving applications and domains. Given the rapid shift toward a data-driven, intelligent networking era, we believe that early-stage investigation into cooperative perception systems will help practitioners better understand current limitations and develop more robust engineering techniques. Such efforts will pave the way toward designing safer and more reliable V2X-aided ADSs.

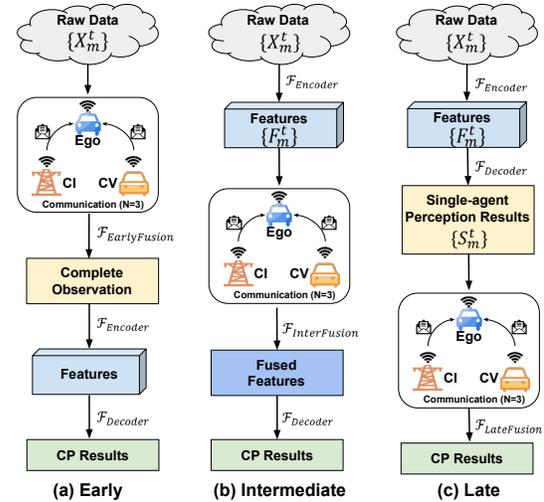


Fig. 2: Overview of cooperative perception schemes for autonomous driving.

II. BACKGROUND

A. Heterogeneous V2X Cooperative Perception

Owing to the inherent limitations of camera and LiDAR sensors, occlusion and long-range perception remain persistent challenges for autonomous systems in single-vehicle settings [11]. These limitations may result in critical failures in complex traffic scenarios. Conversely, cooperative perception systems enable multi-vehicle cooperative detection, overcoming the constraints of single-vehicle perception. By utilizing V2X communication, connected vehicles and infrastructure can share sensory data, yielding multiple viewpoints of obstacles and compensating for mutual blind spots [12], [13].

Multi-agent heterogeneous cooperative perception enables a group of agents to jointly perceive their environment and exchange complementary perception information. In real-world scenarios, various operationally and structurally different cooperative agents (e.g., vehicles and roadside units (RSUs)) are typically equipped with different sensors (e.g., LiDAR

¹<https://github.com/meng2180/V2X-Cooperative-Perception>

and cameras), sharing heterogeneous environmental perception information recorded in specific formats (e.g., images and point clouds). The ego vehicle then receives and fuses this heterogeneous information using a specific fusion scheme to accomplish the perception task (e.g., object detection). The following details the heterogeneous cooperative perception fusion scheme and the cooperative 3D object detection task.

Fusion Scheme. V2X cooperative perception methods can be divided into three categories based on the information fusion stage: *early fusion*, *late fusion*, and *intermediate fusion* [14]. A schematic overview of these cooperative perception schemes is provided in Figure 2. For mathematical clarity, consider a scenario with N connected agents including one ego vehicle (denoted as E) and $N - 1$ cooperative agents (denoted as C_i), where the local raw sensor data observed by the i -th agent at time step t is denoted as X_m^t ($m \in \{E, C_i\} | i \in 1, 2, \dots, N - 1$). As illustrated in Figure 2(a), early fusion involves aggregating raw sensor observations collected by all participating agents, fostering a comprehensive perspective [15]. The ego receives the raw observation $X_{C_i}^t$ ($i \in \{1, \dots, N - 1\}$). After that, it performs data fusion $\mathcal{F}_{\text{EarlyFusion}}$ on the received raw data with self-raw data and produces a fused complete observation. Then, the ego vehicle uses an encoder $\mathcal{F}_{\text{Encoder}}(\cdot)$ to extract high-level cooperative perception features from its fused complete observation. Finally, the ego vehicle uses a decoder $\mathcal{F}_{\text{Decoder}}(\cdot)$ to decode the cooperative perception feature and obtain the cooperative perception results. As illustrated in Figure 2(b), intermediate fusion shared intermediate features rather than raw data. Firstly, each agent extracts its own feature data F_m^t ($m \in \{E, C_i\} | i \in 1, 2, \dots, N - 1$) from its raw perception data X_m^t . Then, the extracted individual feature data of connected agents is shared with the ego vehicle via V2X communication. After receiving shared feature data $F_{C_i}^t$ from other $N - 1$ agents, the ego vehicle uses an intermediate fusion algorithm $\mathcal{F}_{\text{InterFusion}}(\cdot)$ to generate the fused cooperative perception feature data. Finally, the fused cooperative perception feature is fed into the decoder $\mathcal{F}_{\text{Decoder}}(\cdot)$ to obtain the final cooperative perception result. Late fusion involves exchanging individual perceptual outcomes, as depicted in Figure 2(c), requiring minimal data transmission. In this case, each agent independently processes its raw perception data X_m^t , generates its feature data F_m^t , and its final perception result S_m^t . After individual processing, the individual perception results from each agent are shared with the ego vehicle. Then, it uses a late fusion algorithm $\mathcal{F}_{\text{LateFusion}}(\cdot)$ to fuse the received individual perception result from all other $N - 1$ agents and its own perception result to generate the cooperative perception result.

The Preliminaries of Cooperative 3D Object Detection.

Cooperative 3D object detection serves as a core perception task for V2X communication systems, allowing ADS to achieve comprehensive environmental awareness. In cooperative 3D object detection, the cooperative perception system takes a multi-view scene data as input, which consists of an ego vehicle's data X_E^t and data $X_{\text{coop}}^t = \{X_{C_i}^t\}_{i=1}^{n-1}$ from the observation perspectives of $n - 1$ cooperative agents. The system then outputs a bounding box for each detected object based on the relative position of the ego vehicle, providing its 3d location $loc = [x, y, z]$, dimensions length l , width w , height h , and orientation through a heading angle yaw . The accuracy of cooperative 3D object detection is typically

evaluated using the Intersection over Union (IoU) metric [16]. IoU is computed as $IoU = \text{area}(b_p \cap b_{gt}) / \text{area}(b_p \cup b_{gt})$, where b_{gt} denotes the ground-truth 3D bounding box and b_p represents the predicted 3D bounding box, both projected in bird's-eye view. This metric quantifies the ratio of the overlapping area to the union area of the predicted and ground-truth boxes. An object is considered successfully detected if the IoU exceeds a predefined threshold τ .

B. Offline and Online Testing

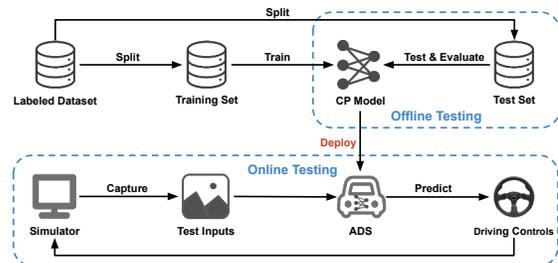


Fig. 3: Offline and online phases of cooperative perception system testing.

Testing has emerged as a well-established and effective approach for assessing the potential risks associated with deploying a software system in real-world scenarios. Unlike traditional software testing, deep learning-based cooperative perception system testing follows a structured workflow comprising two distinct phases: offline testing and online testing [17], [18], as illustrated in Figure 3. Offline testing constitutes a standard and essential step in the development of cooperative perception systems, typically conducted immediately after system training. Its primary objective is to ensure that the trained system demonstrates sufficient accuracy when applied to a test dataset. Following offline testing, online testing is performed by software testers by deploying the cooperative perception system on a specific ADS to examine its interaction with the operational environment under realistic conditions.

Specifically, during offline testing, cooperative perception systems are tested as a unit in open-loop mode. They receive test inputs that can be generated without the participation of the system under test, either manually or automatically. The system's performance is generally assessed by comparing the predicted output bounding boxes with the expected ground truth boxes. Conversely, in online testing, the cooperative perception system is evaluated in a closed-loop environment within a driving simulator. In this phase, test inputs are generated dynamically by the simulated environment, and the ADS's outputs are fed back into the simulator. Online testing evaluates cooperative perception systems by monitoring for safety-critical requirement violations (e.g., collisions or off-road) triggered by the ADS's output behavior.

III. COMPREHENSIVE ANALYSIS OF COOPERATIVE PERCEPTION SYSTEM BUGS

The purpose of V2X cooperative perception design is to help the ego vehicle have more comprehensive and accurate perception, but it may produce erroneous perception results during operation. Cooperative errors in complex cooperative perception systems can originate from multiple components,

including the ego vehicle, cooperative vehicles or infrastructures, and the information fusion process. This section presents a comprehensive analysis of the error patterns observed in cooperative perception systems and identifies six typical categories. Before giving a formal definition of cooperation error patterns, we give a formal representation of V2X-oriented cooperation perception.

A. Formal Representation

We first formalize the V2X-oriented cooperative perception system and the ego vehicle perception system. Consider $n - 1$ cooperative agents and an ego vehicle in the scene. Let X_E^t and $X_{C_i}^t$ represent the observation of the ego vehicle and the i -th cooperative agent at time step t , respectively. The cooperative perception systems \mathbb{CP} can be represented as:

$$\Phi_{CP}(X_E^t, \{X_{C_i}^t\}_{i=1}^{n-1}) = \{(b_i^{CP}) | i \in [n_{CP}]\}$$

Similarly, the ego vehicle perception system can be represented as:

$$\Phi_E(X_E^t) = \{(b_i^E) | i \in [n_E]\}$$

, where $x \in \{E, CP\}$ and Φ_x denotes the perception network of the corresponding perception system, b_i^x represents the i -th 3D bounding box in the perception results. Let n_x denote the number of 3D bounding boxes detected by the perception system. We define the index set of these bounding boxes as $[n_x] = \{1, 2, \dots, n_x\}$. Similarly, let b_i^T denote the i -th ground-truth bounding box and n_T the total number of ground-truth 3D bounding boxes. We use $B_x = \bigcup_{i=1}^{n_x} b_i^x$, $x \in \{T, E, CP\}$ to represent the set of all b_i^x .

B. Cooperative Perception Error Patterns

The primary purpose of the V2X cooperative perception system is to assist the autonomous driving ego vehicle in achieving more accurate perception. If V2X cooperation fails to improve or even degrades the ego vehicle's performance, it is considered a cooperative perception error. Based on this principle, we performed pattern mining on an existing cooperative perception dataset, systematically comparing ego perception results with cooperative perception results in identical scenes. This analysis revealed two major error categories: misleading cooperative errors (denoted as LE), where cooperation negatively affects otherwise correct ego perception results, and miscorrected cooperative errors (denoted as CE), where cooperation fails to resolve existing ego perception weaknesses such as occlusions or long-distance detections.

In this paper, we focus on the task of cooperative 3D object detection, which serves as a representative task for cooperative perception. Note that, in single-agent 3D detection tasks, error types are typically categorized as missing errors, localization errors, and additional detection errors, as established in prior research [19]. Missing errors occur when the perception system fails to detect an object; localization errors occur when producing an inaccurately localized perception bounding box of the object; and additional detection errors occur when the system predicts objects that do not exist in the scene. Hence, as shown in Figure 4, each cooperative perception error category (i.e. misleading cooperative errors (LE) and miscorrected cooperative errors (CE)), can be further subdivided according to these detection error types, resulting in six distinct cooperative perception error subcategories: misleading

cooperative missing error (LCME), misleading cooperative localization error (LCLE), misleading cooperative additional detection error (LADE), miscorrected cooperative missing error (CCME), miscorrected cooperative localization error (CCLE), and miscorrected cooperative additional detection error (CADE). For clarity, the first letter (i.e., L or C) in each error abbreviation indicates the cooperative perception error category (CE or LE), while the remaining letters specify the corresponding 3D object detection error type. In the following, we provide a detailed description of these error patterns.

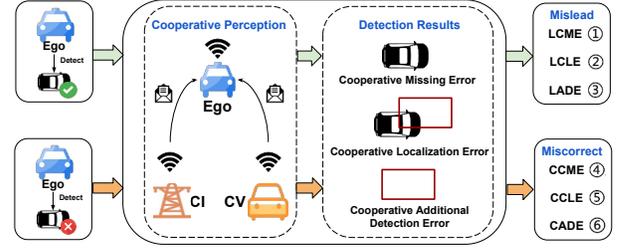


Fig. 4: Graphical representation of error patterns in cooperative perception systems.

1) *Misleading Cooperative Missing Error (LCME)*: A misleading cooperative missing error occurs (1) when an ego vehicle accurately perceives its objects during operation, and (2) the cooperative perception system may mislead the ego vehicle into losing the perception bounding box of the object due to inaccurate shared information or faulty fusion logic.

Definition. Consider a cooperative perception system \mathbb{CP} . \mathbb{CP} makes misleading cooperative missing error if

$$\exists b_i^T \in B_T, s.t. D_E(b_i^T) \wedge \forall b_j^{CP} \in B_{CP}, IoU(b_i^T, b_j^{CP}) = 0$$

, where D_E is the criteria to decide whether the ego vehicle system successfully matches the ground-truth bounding box b_i^T . $\forall b_j^{CP} \in B_{CP}, IoU(b_i^T, b_j^{CP}) = 0$ indicates that, for any predicted detection box b_j^{CP} of \mathbb{CP} , the IoU between the ground-truth bounding box b_i^T of the i -th object and b_j^{CP} is zero. This condition indicates that the cooperative perception system \mathbb{CP} has entirely missed the detection of the i -th object.

2) *Misleading Cooperative Localization Error (LCLE)*: A misleading cooperative localization error occurs (1) when an ego vehicle accurately perceives its objects during operation, and (2) the cooperative perception system may mislead the ego vehicle into producing an inaccurately localized perception bounding box of the object due to inaccurate shared information or faulty fusion logic.

Definition. Consider a cooperative perception system \mathbb{CP} . \mathbb{CP} makes misleading cooperative localization error if

$$\exists b_i^T \in B_T, s.t. D_E(b_i^T) \wedge \exists b_j^{CP} \in B_{CP}, 0 < IoU(b_i^T, b_j^{CP}) \leq \tau$$

, where $\exists b_j^{CP} \in B_{CP}, 0 < IoU(b_i^T, b_j^{CP}) \leq \tau$ indicates that there exists a predicted detection box b_j^{CP} from the cooperative perception system \mathbb{CP} such that the ground-truth bounding box b_i^T of the i -th object partially overlaps with b_j^{CP} , but the overlap falls below the threshold τ . This condition indicates that the detection box b_j^{CP} predicted by the cooperative perception system \mathbb{CP} has inaccurately localized the i -th object.

3) *Misleading Cooperative Additional Detection Error (LADE)*: Additional detection refers to the case that the system treats an arbitrary region without objects as an "object". A misleading cooperative additional detection error occurs (1)

when the ego vehicle makes the correct decision about an object that does not exist in the real world during operation, and (2) the cooperative perception system may mislead the ego vehicle into producing an additional detection of a bounding box due to inaccurate shared information or faulty fusion logic.

Definition. Consider a cooperative perception system $\mathbb{C}\mathbb{P}$. $\mathbb{C}\mathbb{P}$ makes a misleading cooperative additional detection error if

$$\begin{aligned} \exists b_j^{CP} \in B_{CP}, s.t. \forall b_m^E \in B_E, IoU(b_m^E, b_j^{CP}) = 0 \\ \wedge \forall b_i^T \in B_T, IoU(b_i^T, b_j^{CP}) = 0 \end{aligned}$$

, LADE is considered to occur when the predicted box b_j^{CP} generated by $\mathbb{C}\mathbb{P}$ neither overlaps with any predicted box b_m^E of the ego vehicle nor with any ground-truth bounding box b_i^T of the object to be detected. The first condition indicates that the ego vehicle does not produce an erroneous prediction similar to b_j^{CP} , while the second condition indicates that the cooperative perception system $\mathbb{C}\mathbb{P}$ makes an erroneous prediction by generating an additional detection box at a location where no obstacle actually exists.

4) *Miscorrected Cooperative Missing Error (CCME)*: A miscorrected cooperative missing error occurs (1) when an ego vehicle inaccurately perceives its objects during operation, and (2) the cooperative perception system incorrectly completes or corrects the ego vehicle's perception results to miss the detection of the object.

Definition. Consider a cooperative perception system $\mathbb{C}\mathbb{P}$. $\mathbb{C}\mathbb{P}$ makes miscorrected cooperative missing error if

$$\exists b_i^T \in B_T, s.t. \neg D_E(b_i^T) \wedge \forall b_j^{CP} \in B_{CP}, IoU(b_i^T, b_j^{CP}) = 0$$

, where $\neg D_E$ is the criteria to decide whether the ego vehicle system unsuccessfully matches the ground-truth bounding box b_i^T . The distinction between LCME and CCME lies solely in whether the ego vehicle's perception system successfully detects the object.

5) *Miscorrected Cooperative Localization Error (CCLE)*:

A miscorrected cooperation localization error occurs (1) when an ego vehicle inaccurately perceives its objects during operation, and (2) the cooperative perception system incorrectly completes or corrects the ego vehicle's perception results to produce an inaccurately localized perception bounding box of the object.

Definition. Consider a cooperative perception system $\mathbb{C}\mathbb{P}$. $\mathbb{C}\mathbb{P}$ makes miscorrected cooperative localization error if

$\exists b_i^T \in B_T, s.t. \neg D_E(b_i^T) \wedge \exists b_j^{CP} \in B_{CP}, 0 < IoU(b_i^T, b_j^{CP}) \leq \tau$, where $\neg D_E$ is the criteria to decide whether the ego vehicle system unsuccessfully matches the ground-truth bounding box b_i^T . The difference between LCLE and CCLE lies only in whether the ego vehicle's perception system successfully detects the object.

6) *Miscorrected Cooperative Additional Detection Error (CADE)*: A miscorrected cooperative additional detection error occurs (1) when the ego vehicle produces an additional detection of a bounding box during operation, and (2) the cooperative perception system fails to correct this erroneous detection, instead producing a false detection bounding box near the original additional detection location.

Definition. Consider a cooperative perception system $\mathbb{C}\mathbb{P}$. $\mathbb{C}\mathbb{P}$ makes a miscorrected cooperative additional detection error if

$$\begin{aligned} \exists b_j^{CP} \in B_{CP}, s.t. \exists b_m^E \in B_E, IoU(b_m^E, b_j^{CP}) > 0 \\ \wedge \forall b_i^T \in B_T, IoU(b_i^T, b_j^{CP}) = 0 \end{aligned}$$

Unlike LADE, CADE occurs when the prediction box b_m^E of the ego vehicle overlaps with the prediction box b_j^{CP} generated by $\mathbb{C}\mathbb{P}$. This indicates that the cooperative perception system $\mathbb{C}\mathbb{P}$ generates additional detection boxes without correcting the erroneous prediction of the ego vehicle.

IV. EMPIRICAL STUDY DESIGN

A. Research Questions

RQ1. How does equipping cooperative agents with heterogeneous sensors affect the performance of cooperative perception systems? Heterogeneous environments that integrate different sensor types may introduce new challenges for cooperative perception. RQ1 aims to evaluate the performance of agents equipped with diverse sensor configurations when cooperating with the ego vehicle. In this RQ, we focus on three cooperation scenarios: LiDAR-based cooperation, camera-based cooperation, and hybrid multimodal (LiDAR and camera) sensor cooperation. Specifically, in the first two cooperative scenarios, all agents are equipped with the same sensor type. In the final multi-modal scenario, participants comprise a heterogeneous mix of LiDAR-equipped and camera-equipped agents, with the ego vehicle randomly assigned one of the sensor configurations in each trial.

RQ2. What are the differences in cooperative perception performance between V2V and V2I cooperation modes? The cooperative perception system comprises multiple types of intelligent agents, such as vehicles and infrastructure, that contribute to environmental perception. Investigating the contribution of different types of collaborative agents to the ego vehicle's perception performance is of significant importance. In this research question, we focus on quantifying the differences in V2I and V2V cooperative perception performance under different fusion schemes.

RQ3. What is the relationship between cooperative perception errors and driving violations? To assess the potential impact of cooperative perception errors on driving violations, we first integrate the cooperative perception system into the ADS and run the integrated ADS in the simulator. This setup allows us to evaluate the system's environmental perception performance and to identify driving violations, such as collisions, during online testing. To investigate the correlation between cooperative perception system errors and ADS violations, we compare the number of cooperative perception errors observed in violation scenarios and non-violation scenarios. For ADS violation scenarios, we further investigate the distribution of cooperative perception errors over the entire duration from ADS startup to the point of violation.

RQ4. Do communication issues during online operations diminish the effectiveness of the cooperative perception system? During online operation, cooperative perception systems are inevitably affected by real-world communication interferences. Communication latency (CL) and pose error (PE) are two typical types of such interference. In practical scenarios, limited communication bandwidth [20] or transmission failures [21] may introduce delays in data exchange between participants, resulting in communication latency. Similarly, obstacles or GPS signal interference [22] in real-world environments can degrade the accuracy of participants' positioning information, leading to pose errors. In

this study, we investigate whether the performance of ADS equipped with cooperative perception capabilities is impacted by these two common forms of communication interference.

B. Datasets and Platform

To address RQ1, we employ the widely adopted OPV2V dataset [11], which has been extensively used in heterogeneous cooperative perception research. OPV2V is a large-scale, publicly available simulated dataset designed for V2V perception. It comprises over 70 diverse scenarios, 11,464 frames, and 232,913 annotated 3D vehicle bounding boxes, collected from eight towns in the CARLA simulator [23] as well as a digital reconstruction of Culver City, Los Angeles. For RQ2, we require a dataset that encompasses both cooperative vehicles and cooperative infrastructure within cooperative driving scenarios. We employ the V2XSet dataset [21] collected from CARLA, an open dataset for V2X cooperative perception that incorporates both cooperative vehicles and infrastructure across its training and testing scenarios. For RQ3 and RQ4, we use the CARLA-based testing platform V2Xverse [24], a comprehensive simulation environment for online collaborative autonomous driving. V2Xverse provides an end-to-end pipeline that includes a multi-agent driving dataset generation scheme and a codebase for deploying a full-stack cooperative driving system. Specifically, it supports constructing cooperative perception scenarios by adding connected vehicles and roadside units (RSUs), equipping them with sensors (e.g., cameras or LiDAR), and sharing their perception information with the ego vehicle. We evaluate the ADS equipped with cooperative perception on the default 32 CARLA Town05 routes, retaining only those routes where the ADS shows no violations under perfect perception to ensure that any observed failures are attributable to the cooperative perception system.

C. Cooperative Perception Systems

To evaluate system performance in heterogeneous environments, we conduct offline experiments (RQ1 & RQ2) on six cooperative perception systems employing different fusion schemes. These include one early fusion system, one late fusion system, and four state-of-the-art intermediate fusion systems (DiscoNet [25], V2X-ViT [21], AttFusion [11], F-Cooper [26]). All cooperative perception systems implemented LiDAR, camera, and multi-sensor (LiDAR and camera) cooperative driving sensor configurations based on HEAL [27]. For online performance evaluation (RQ3 & RQ4), we deploy both early and late fusion systems alongside two representative intermediate fusion architectures (V2X-ViT and F-Cooper) selected from our offline testing systems.

D. Implementation Details

In all experiments, we set the IoU threshold τ to 0.5, consistent with prior research [28]. For the online testing experiments, the cooperative agent used V2Xverse’s default test configuration, which enables cooperative perception with RSUs. Specifically, RSUs are strategically positioned along the roadside in proximity to the ego vehicle. To ensure that the ego vehicle consistently has at least one RSU available for valid collaboration, a new RSU is deployed every five seconds as the ego vehicle moves, positioned on the right side of the road, 12 meters ahead of the ego vehicle. For RQ4,

the transmission delay in V2X communication varies from 500 ms to 1000 ms. To simulate pose error, we introduce translational offsets along specific axes. To simulate pose errors, we introduce translational offsets along the x- and y-axes and rotational offsets around the z-axis to the cooperative agent. The translation offsets for both the x-axis and y-axis are set within the range of 0 to 0.6 meters, while the rotation angle around the z-axis is varied between 0 and 0.6 degrees. We randomly select parameters within the range for two abnormal communication conditions. For details on the experimental environment, please refer to the supplementary website [29].

E. Evaluation Metrics and Baseline

Evaluation Metrics. To evaluate the perception performance of the system in cooperative 3D object detection, we first apply the Intersection over Union (IoU) metric (Section II-A) to identify successful detections. We then assess detection performance using Average Precision (AP), a standard evaluation metric in V2X cooperative perception research [28], [30], which is defined as follows:

$$AP|_R = \frac{1}{|R|} \sum_{r \in R} \rho_{\text{interp}}(r)$$

For fair comparisons, we adopt the 11 recall positions from the Pascal VOC benchmark [31], denoted as $R_{11} = \{0, 0.1, 0.2, \dots, 1\}$. The interpolation function, $\rho_{\text{interp}}(r) = \max_{r', r' \geq r} \rho(r')$, is employed, where $\rho(r)$ represents the precision at recall r . A higher AP reflects better performance of cooperative perception systems.

To evaluate the driving performance of the ADS equipped with cooperative perception systems, we use the driving score defined in the CARLA Autonomous Driving Leaderboard 1.0 [32] as the evaluation metric. Specifically, the driving score (DS) is computed as $DS = R_k \times P_k$, where R_k is the percentage of the k -th route completed before the first violation of ADS, and P_k is the corresponding infraction penalty. The penalty rate P_k is defined as:

$$P_k = \prod_{j \in J} p_j^{n_j}$$

where p_j is the penalty rate for an incident type j from a given set of incident types J , and n_j is the number of occurrences of this type. Incident types refer to infractions, such as collisions with other vehicles, as defined in the CARLA Autonomous Driving Leaderboard [32].

To further evaluate the driving performance of the ADS equipped with cooperative perception systems in greater detail, we employ two additional metrics: collision rate (CR) and route completion rate (RCR). The CR refers to the ratio of the number of routes with collisions to the total number of test routes, while the RCR refers to the average percentage of route completion across all test routes.

Baseline. For each RQ, we introduce a baseline using a single-agent perception system. This baseline represents the ego vehicle performing environmental understanding independently, without incorporating information from cooperative agents. Specifically, in RQ1, the single-agent system employs either LiDAR or camera sensors for perception; in RQ2, we compare its perception performance with that of cooperative systems using V2V and V2I communication modes. For the online evaluations in RQ3 and RQ4, we compare the performance of the cooperative perception system under both

TABLE I: Cooperative error pattern analysis of cooperative perception systems under heterogeneous sensor configurations.

Systems	Sensor	LCME	LCLE	LADE	CCME	CCLE	CADE	AP
Early Fusion	L	0.0	0.0	0.0	0.3	0.0	1.2	0.96
	C	0.0	0.0	1.1	0.3	0.1	1.2	0.96
Late Fusion	C	0.0	2.0	2.5	0.6	4.0	2.0	0.73
	M	0.0	1.2	1.6	0.8	1.6	1.3	0.87
F-Cooper	L	0.0	1.0	0.3	0.3	0.4	0.9	0.87
	C	0.1	1.1	0.6	2.0	4.1	0.9	0.49
	M	0.0	1.1	0.8	0.8	2.5	1.1	0.68
DiscoNet	L	0.0	0.1	0.1	0.5	0.3	0.6	0.92
	C	0.1	0.8	0.2	2.5	2.8	1.2	0.51
	M	0.0	0.5	0.7	0.9	2.7	1.5	0.7
AttFusion	L	0.0	0.1	0.1	0.4	0.3	1.3	0.93
	C	0.1	0.8	0.1	2.4	3.1	0.9	0.53
	M	0.0	0.4	0.7	1.8	3.1	1.8	0.71
V2XViT	L	0.0	0.0	0.2	0.3	0.1	0.5	0.96
	C	0.0	0.7	0.6	2.0	3.1	0.9	0.6
	M	0.0	0.4	0.3	1.7	2.3	1.7	0.89

*Current lacks the early fusion design that leverages C and M configurations due to the absence of methods that directly accumulate images.

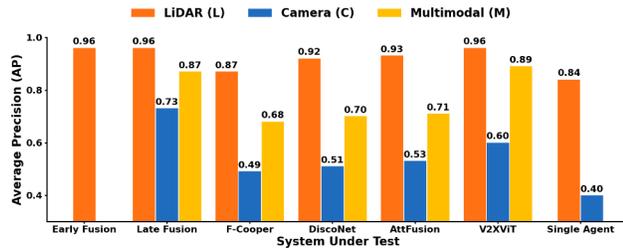


Fig. 5: Average precision of the system under test under heterogeneous sensor configurations.

normal and abnormal communication conditions against that of a single-agent perception system.

V. RESULTS AND FINDINGS

A. *RQ1. The LiDAR-based cooperation configuration exhibits the highest perception performance.*

1) *Experimental Setup:* To evaluate the performance of the cooperative perception system in a heterogeneous sensor environment, we conduct experiments using pre-trained systems under three cooperative configurations: LiDAR-based cooperation (L), camera-based cooperation (C), and multimodal sensor cooperation (M). Each pre-trained system is tested using the original test scenes and the default configuration of the HEAL heterogeneous framework. We assess the prediction performance of the three sensor configurations using the same set of test data, collect their detection outputs, and compute both the average number of cooperative perception errors per test frame and the overall AP.

2) *Experimental Results:* Table I presents an analysis of cooperative error patterns exhibited by cooperative perception systems under heterogeneous sensor configurations, detailing the average number of each error pattern per test frame. The perception performance achieved through LiDAR-based cooperation significantly exceeds that of the other sensor configurations. Specifically, for each tested system configuration,

it outperforms multimodal sensor cooperation by an average of 21.7% and camera-based cooperation by an average of 87.0%. Moreover, the total number of CCME and CCLE errors in camera-based cooperative perception is substantially higher than in LiDAR-based cooperative perception, leading to a pronounced decline in the AP of the camera-based configuration. Additionally, we observe that the number of cooperative localization errors is, on average, 122.0% and 38.8% greater than the number of cooperative missing errors and additional detection errors, respectively, across all configurations. Furthermore, in most configurations, the number of miscorrected cooperative errors exceeds that of misleading cooperative errors. This indicates that the majority of errors arise from the cooperative perception system's failure to correct inaccurate or incomplete information initially predicted by the ego vehicle. Figure 5 presents AP of the system under test under heterogeneous sensor configurations. When using the same type of sensors (LiDAR or camera), cooperative perception systems perform better than the single-agent system. Multimodal sensor cooperation outperforms camera-based single-agent perception, while its performance relative to LiDAR-based single-agent perception depends on the specific model architecture. To realize optimal cooperative perception, it is essential to appropriately configure sensor setups and select model architectures suited to heterogeneous environments.

Answer to RQ1: Compared to camera-based cooperation and multimodal sensor cooperation, using LiDAR for all cooperative agents yields superior performance. Additionally, we find that CCME and CCLE could be the primary factors contributing to the reduction in AP of the other two sensor configurations (C and M).

B. *RQ2. V2I and V2V communication exhibit distinct cooperative perception performance under different fusion schemes.*

1) *Experimental Setup:* To investigate differences in cooperative perception performance between V2V and V2I cooperation modes, we select cooperative vehicles and infrastructure to collaborate with the ego vehicle. To ensure a fair evaluation of the roles of cooperative vehicles and infrastructure, we randomly select a cooperative agent for each cooperation mode in each scenario. In addition, to eliminate the influence of sensor heterogeneity, all cooperative agents are equipped with the same type of sensor during the experiment. Specifically, we conduct experiments using configurations for LiDAR-based cooperation under the original test set. We further statistically analyze the effect of different types of cooperative agents on the cooperation of the ego vehicle at different distances. Specifically, we divide the detection targets into three ranges (refer to [30], [33]) according to the distance relative to the ego vehicle, including 0-30m (Short), 30-50m (Middle), and 50-100m (Long). Then we calculate the cooperative error numbers at each level and AP values.

2) *Experimental Results:* Figure 6 shows the AP under different cooperation modes. It demonstrates that V2V outperforms V2I under intermediate and late fusion, whereas V2I outperforms V2V under early fusion. Both V2V and V2I surpass single-agent perception, confirming their effectiveness in enhancing perception performance. Figure 7 illustrates the total cooperative perception error number of the system at varying distances under various cooperation modes. We find

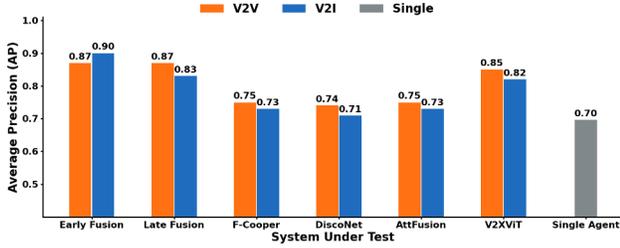


Fig. 6: Average precision of the system under test under different cooperation modes.

that the number of cooperative perception errors increases with distance, indicating that the greater the distance between the ego vehicle and the target obstacle, the lower the cooperative perception performance of both the cooperative vehicle and the cooperative infrastructure. For the complete analysis results of six cooperative perception error patterns, please refer to the supplementary website [29].

Answer to RQ2: Given the same number of cooperative agents, cooperative vehicles demonstrate superior perception performance under intermediate and late fusion schemes, whereas cooperative infrastructure performs better under the early fusion mechanism.

C. RQ3. Increased cooperative perception errors may result in a higher frequency of driving violations.

1) *Experimental Setup:* To evaluate the relationship between cooperative perception errors and driving violations during online operation, we connect the ADS equipped with the V2X cooperative perception to a simulator for testing. Each cooperative perception system is evaluated on the V2XVerse online testing platform using the selected routes described above in Section IV-B. Because autonomous driving accidents caused by perception systems are often precipitated by a sequence of erroneous perceptions within a short time window preceding a collision, and following the similar methodology of prior study [34], we focus our analysis on the final scene segment associated with each accident. During each route test, we record the following: (1) the number of six predefined errors of the cooperative perception system within the last 10 seconds of the scenario, and (2) whether a violation occurred (i.e., collision with a non-player character (NPC), off-road, or route incomplete). Then we calculate the driving score (DS), collision rate (CR), and average route completion rate (RCR) for all routes.

To further investigate the distribution of perception errors over the entire duration from the initiation of the ADS to the occurrence of a violation, we divide this period into consecutive 5-second intervals. Specifically, we record the number of errors occurring in the 0–5 seconds preceding the violation (denoted as t_3), the number of errors in the 5–10 seconds prior to the violation (denoted as t_2), and the average number of errors in all earlier intervals beyond 10 seconds before the violation (denoted as t_1).

Finally, to quantify the independent effects of different types of cooperative perception errors on driving violations, we employ a multivariable logistic regression model [35], which is commonly used in software engineering to evaluate the impact of predictors through p-values and odds ratios

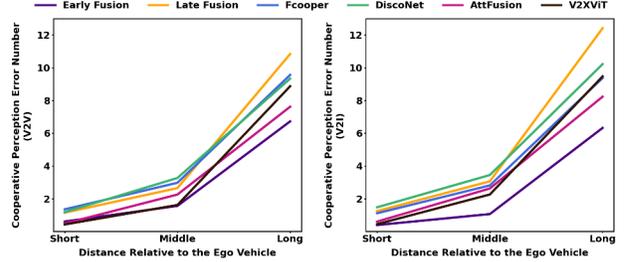


Fig. 7: The cooperative perception error number of the system at varying distances under V2V and V2I cooperation.

(ORs). Specifically, for each type of cooperative perception error, the null hypothesis assumed no association between the number of errors and the likelihood of a violation. Associations are considered statistically significant when the p-value is below the predefined threshold of 0.05. The OR represents the ratio of the probability of a violation to the probability of it not occurring, and it quantifies the change in violation odds associated with each additional error occurrence.

2) *Experimental Results:* Table II presents the driving performance of ADS equipped with different perception systems under normal communication conditions. The results indicate that all systems exhibit cooperative perception errors that lead to violations during online operation, with the V2XViT system achieving the highest overall driving score. Figure 8 shows the AP of the system under test along the driving route. As shown in Table II and Figure 8, under normal communication conditions, the cooperative perception system outperforms the single-agent perception system in terms of overall performance. Table III further provides an analysis of cooperative error patterns exhibited by the cooperative perception systems during route execution. The experimental results reveal that scenarios involving driving violations exhibit, on average, 15.4% lower perception performance compared to non-violation scenarios, and the number of cooperative perception errors increases. To further investigate the underlying causes of these violations, the distribution of cooperative perception errors over the entire period from ADS initialization to the occurrence of a violation is provided on the supplementary website [29]. The results show that the frequency of cooperative perception errors increases as the time of the violation approaches, with particularly notable increases in CCME and CCLE. Furthermore, statistical analysis reveals that only CCLE is statistically significant ($p = 0.047$) under normal driving conditions, with an OR of 1.019, indicating a 1.9% increase in the odds of violations for each additional CCLE error.

Answer to RQ3: Cooperative perception errors resulting from imperfect cooperative perception may lead to ADS violations. As the number of cooperative perception errors increases, the likelihood of violations also rises, with CCLE as a major contributor to driving violations.

D. RQ4. Cooperative perception systems are not robust against communication interference when running online.

1) *Experimental Setup:* To investigate whether the driving performance of ADS equipped with the cooperative perception system is affected by common communication interferences, we introduce communication latency (CL) and pose error (PE) to the cooperative agents. Specifically, we utilize the test

TABLE II: Driving performance of ADS equipped with different perception systems under normal and abnormal communication conditions.

Systems	Type	DS	CR	RCR
Early Fusion	Normal	77.1	26.3	79.4
	CL	72.9	31.6	76.2
	PE	64.2	42.1	69.0
Late Fusion	Normal	78.9	26.3	88.0
	CL	71.2	36.8	83.0
	PE	74.3	31.6	83.7
F-Cooper	Normal	75.9	31.6	81.7
	CL	71.2	31.6	76.8
	PE	62.5	42.1	70.6
V2XViT	Normal	81.0	26.3	87.9
	CL	65.1	42.1	76.0
	PE	77.2	26.3	84.8
Single Agent	Normal	69.1	36.8	74.7

*Single-agent perception remains unaffected by CL and PE.

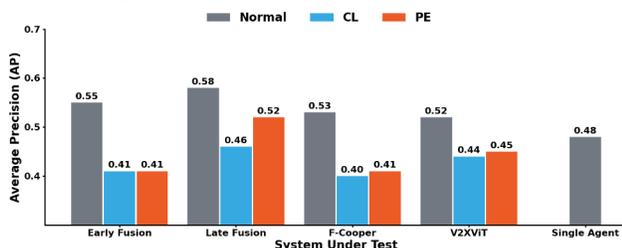


Fig. 8: Average precision of the system under test along the driving route.

routes from RQ3, introduce two types of communication issues into these routes, and simulate them randomly within a predefined parameter range. In the experiment, we record the same information as in RQ3 (i.e., the number of cooperative perception errors and whether a violation occurred) for each test route. Based on this information, we then calculate the AP, DS, CR, and RCR. Similar to RQ3, to quantify the impact of different types of cooperative perception errors on driving violations under abnormal communication conditions, we use the multivariate logistic regression model [35] for statistical analysis. Finally, to provide a more intuitive understanding, we select two representative collision cases attributable to communication latency and pose error for visualization.

2) *Experimental Results:* Table II presents the driving performance of ADS equipped with different cooperative perception systems under CL and PE. Compared to normal driving conditions, CL resulted in a 10.3% decrease in driving score and a 6.3% decrease in route completion rate, while the collision rate increased by 7.9%. Similarly, under PE conditions, the driving score and route completion rate decreased by 11.2% and 7.2%, respectively, and the collision rate increased by 7.9%. Figure 8 shows the average precision of the system under CL and PE. As shown in Table II and Figure 8, under abnormal communication conditions, such as CL and PE, the perception accuracy and driving performance of certain cooperative perception systems may even fall below that of a single agent. Table IV shows an analysis of online error patterns of cooperative perception systems under abnormal communication conditions. From the table, we can conclude that the perception performance of the cooperative perception system is signifi-

TABLE III: Online cooperative error pattern analysis of cooperative perception systems during route execution.

Systems	Type	LCME	LCLE	LADE	CCME	CCLE	CADE	AP
Early Fusion	Vio.	2.8	4.0	22.6	78.6	29.8	15.0	0.50
	No Vio.	2.3	2.3	14.4	51.9	26.6	6.5	0.56
Late Fusion	Vio.	0.3	0.5	13.5	53.3	37.2	24.5	0.51
	No Vio.	0.1	0.2	13.5	37.5	11.6	16.0	0.61
F-Cooper	Vio.	1.3	3.2	27.8	54.3	55.8	38.8	0.44
	No Vio.	0.7	0.8	12.1	51.4	27.1	25.8	0.56
V2XViT	Vio.	1.6	10.8	18.4	95.4	63.4	14.0	0.47
	No Vio.	0.6	3.9	13.1	63.9	20.1	10.1	0.54

cantly reduced under the conditions of communication delay and pose error, and the number of most cooperative perception errors increases significantly. Communication latency tends to result in more cooperative missing errors, while pose errors are more likely to cause increased cooperative localization errors. Furthermore, our statistical analysis reveals that under the CL condition, three types of errors exhibit significant effects: LCLE ($p = 0.024$, OR = 1.137), CCME ($p = 0.023$, OR = 1.015), and CADE ($p = 0.016$, OR = 1.037). Under the PE condition, LCLE emerges as the most critical predictor ($p = 0.014$, OR = 1.169), corresponding to a 16.9% increase in violation odds. These findings demonstrate that LCLE errors are the primary contributors to elevated violation rates under both CL and PE conditions. This can be attributed to the fact that abnormal communication conditions may transmit delayed or inconsistent information to the ego vehicle, thereby impairing its perception accuracy.

Figure 9 presents visualization samples illustrating representative collision cases caused by communication latency and pose errors. $time_0$ represents the initial timestamp of the scenario, while $time_1$ and $time_2$ denote timestamps corresponding to different moments during the scenario's execution under varying communication conditions. The first scenario (top row of Figure 9) involves a vehicle approaching from a distance and crossing an intersection. Under normal communication conditions, the ADS equipped with a cooperative perception system successfully detects the vehicle and decelerates accordingly. However, under communication delay, the system experiences a CCME, causing the ADS to fail to slow down and subsequently collide with the vehicle. The second scenario (second row of Figure 9) involves an occluded pedestrian crossing the road. Under normal communication conditions, the ADS equipped with a cooperative perception system detects the pedestrian in advance and stops to yield. However, in scenarios involving pose errors, the system encounters a LCLE and misjudges the pedestrian's position, causing the ADS to fail to decelerate in time and ultimately resulting in a collision. These two representative scenarios highlight the lack of robustness of V2X cooperative perception systems under abnormal communication conditions.

Answer to RQ4: Abnormal communication conditions (such as communication latency and pose error) can undermine the performance of the cooperative perception system and significantly increase the frequency of ADS violations. Under CL conditions, LCLE, CCME, and CADE are the primary contributors to driving violations, whereas under PE conditions, LCLE emerges as the dominant factor.

TABLE IV: Analysis of online error patterns of cooperative perception systems under abnormal communication conditions.

Systems	Type	Communication Latency							Pose Error						
		LCME	LCLE	LADE	CCME	CCLE	CADE	AP	LCME	LCLE	LADE	CCME	CCLE	CADE	AP
Early Fusion	Vio.	21.7	12.2	84.0	122.7	39.5	11.8	0.32	7.5	11.8	61.8	109.5	37.3	16.6	0.36
	No Vio.	6.6	1.5	42.7	69.1	26.3	5.9	0.45	4.9	5.3	40.0	61.0	36.0	7.6	0.45
Late Fusion	Vio.	0.1	1.3	47.8	92.0	30.8	34.4	0.38	0.0	3.1	29.1	67.7	57.1	19.6	0.45
	No Vio.	0.0	0.0	24.5	27.3	11.0	11.2	0.52	0.0	0.8	27.0	38.8	18.3	15.2	0.57
F-Cooper	Vio.	17.1	7.9	40.0	87.9	37.6	47.4	0.38	7.0	10.3	37.7	76.3	61.0	42.9	0.37
	No Vio.	10.7	2.6	33.3	59.8	45.9	31.8	0.40	7.0	3.9	17.4	45.0	56.6	28.9	0.43
V2XViT	Vio.	6.7	9.1	49.0	131.9	34.2	11.3	0.41	2.5	14.7	31.7	120.8	61.2	15.2	0.34
	No Vio.	1.6	4.7	36.9	72.7	17.3	8.5	0.46	2.5	4.0	21.5	62.1	18.4	9.7	0.49

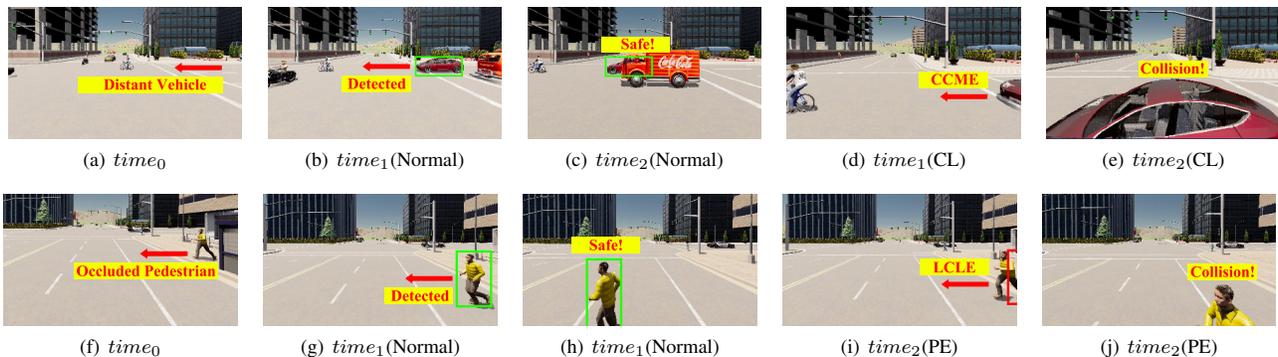


Fig. 9: Visualization samples showing representative collision cases attributable to communication latency and pose error.

VI. DISCUSSION

Based on the findings from RQ1–RQ4, existing cooperative perception systems require further improvement for practical use. First, such systems must operate in heterogeneous environments, where sensor configurations often prioritize LiDAR over cameras due to LiDAR’s superior depth perception capabilities (as point clouds inherently contain spatial information). Multimodal sensor fusion designs must balance cost-effectiveness and communication efficiency. In addition, V2V cooperation generally outperforms V2I under intermediate and late fusion, likely because infrastructure-mounted sensors capture a bird’s-eye view that is harder to fuse with ego vehicle features. Conversely, V2I cooperation shows superior performance in early fusion, potentially due to the infrastructure being equipped with higher resolution sensors. Considering that in the real world, the ego vehicle continuously interacts with diverse heterogeneous agents, cooperative perception systems should adapt to such conditions to ensure robust and reliable cooperative perception.

When an ADS equipped with a cooperative perception system operates online, communication delays and pose errors may further exacerbate ADS violations. In the real world, these two types of errors are prevalent. Even well-calibrated and synchronized sensors may still be inaccurate due to changes in the external environment. In order to deploy reliable cooperative perception systems, developers must carefully handle agent synchronization and calibration issues and design cooperative perception systems that can remain reliable in the face of abnormal communications.

We systematically summarize the error patterns of cooperative perception systems and find that a substantial number of cooperative perception errors occur during both offline and

online operation. The data-driven nature makes it challenging to train a robust cooperative perception system that satisfies safety and reliability requirements under all conditions. Consequently, cooperative perception system development must prioritize cooperative perception error mitigation in design and training processes.

Future Directions. Based on these insights, we summarize the following future directions:

- Cooperative perception systems should adapt to heterogeneous environments, where variations in varying sensor types and cooperative agent configurations could lead to inconsistent performance. Therefore, **the ongoing development of systems adapting to heterogeneous environments, coupled with the establishment of appropriate testing and optimization platforms, is essential.**
- There is an urgent need for techniques to continuously enhance the robustness and reliability of V2X cooperative perception systems. Given that V2X involves multiple critical components, **robustness assurance should encompass both offline and online aspects, including offline fault localization and repair, as well as monitoring under abnormal communications and perception during system online operation.**
- Cooperative perception systems should not perform worse than ego vehicle perception. However, our study systematically reveals the presence of cooperative perception error patterns, which can degrade system perception and even lead to driving accidents. **Designing efficient automated testing methodologies based on these patterns to detect and mitigate such cooperation errors represents a promising direction for future research.**

Threats to Validity. In terms of *construct validity*, a potential

threat arises from the classification of cooperative perception errors. Although this paper identifies six error patterns from the perspective of how cooperation influences the ego vehicle’s perception in the cooperative object detection task, alternative definitions of error patterns may exist under different viewpoints or cooperative perception tasks. Nevertheless, within the scope of this study, our classification is regarded as a complete and representative characterization of cooperative perception errors. In terms of *internal validity*, one of the primary threats is to conduct evaluations in a simulated environment. Simulated driving scenarios may differ both visually and physically from real-world environments due to the inherent gap between simulation-based testing and testing in the physical world. To alleviate this threat, we use the popular high-fidelity CARLA simulator to collect offline data or simulate online simulation scenarios, and we conduct multiple repeated experiments to ensure the reproducibility of our results. In terms of *external validity*, a potential threat is that our results may not generalize to other cooperative perception systems, datasets, or cooperation modes. To mitigate this threat, we try our best to collect a diverse set of cooperative perception systems with varying system structures and fusion schemes, and adopt two representative cooperation modes (V2V and V2I). We then evaluate them using widely adopted datasets with rich scenarios and diverse test routes across a variety of testing settings to ensure comprehensive assessment. Another potential threat lies in our task choice, as this study mainly focuses on cooperative object detection. Nevertheless, as one of the most representative tasks in current research on cooperative perception, cooperative object detection provides the most direct and critical input for downstream autonomous driving tasks, such as object tracking and motion planning. In future work, we plan to extend our classification approach to other cooperative perception tasks.

VII. RELATED WORKS

V2X Cooperative Perception. Cooperative perception systems facilitate the exchange of perception data among agents, thereby enabling each agent to achieve more comprehensive and accurate situational awareness. In this context, early fusion techniques involve the transmission of raw sensor data [36], whereas late fusion techniques transmit processed perception outputs [11]. To balance computational performance and bandwidth efficiency, recent studies have investigated intermediate feature transmission [37]. The first intermediate cooperative framework F-Cooper [26] balances bandwidth constraints and detection accuracy by using max pooling for V2V voxel feature fusion. To consider the potential relationships between multiple agents, several graph-based fusion approaches have been proposed [20], [25], [38]. For instance, V2VNet [20] leverages a spatial-aware graph neural network to model the communication among agents. In addition to graph learning, attention mechanisms have emerged as a powerful tool for exploring feature relationships [21], [39]–[42]. V2X-ViT [21] proposes a multi-scale window attention module to capture long-range spatial interaction on high-resolution detection. Recent studies have also focused on developing cooperative perception systems that remain robust under communication interruptions [43] and lossy transmission conditions [44]. However, current literature lacks investigation into how heterogeneous sensor configurations and varying cooperative agents affect the performance of cooperative perception systems, nor

have such systems been empirically tested through online deployment.

Autonomous Driving Perception System Testing. Safety is a critical priority in the development of ADSs, and various approaches, such as data-driven [45]–[47] and search-based [48]–[52] methods, have been proposed to test their safety. The perception system in an ADS serves as a key component in real-world applications, and rigorous testing is required to ensure the reliability of the overall system [3]. Various studies have focused on testing single-agent perception systems across various driving tasks [53]–[55]. To evaluate camera-based perception systems, a representative study [53] generates data by incorporating object instances into background images. Several studies also test LiDAR-based perception systems [54], [55]. Christian et al. [55] propose testing LiDAR systems by introducing semantic mutations to real-world data for the semantic segmentation task. Recently, due to the rapid development of V2X communication technology, some researchers have focused on the testing of cooperative perception systems [28], [56], [57]. Li et al. [56] focus on investigating the impact of variations and disturbances in the communication range on safety and potential hazards. Guo et al. pioneered CooTest [28], a framework for evaluating cooperative perception robustness under extreme conditions such as communication interference and adverse weather. CooTest [28], the most related work, lacks a comprehensive analysis of cooperative perception errors, is limited to LiDAR-based V2V cooperation, and operates offline. In contrast, this paper analyzes error patterns in depth, considers heterogeneous sensors (LiDAR and camera) and diverse cooperative agents (V2V and V2I), and is the first to evaluate cooperative perception performance in the online deployment.

VIII. CONCLUSION

This paper presents an empirical study of V2X cooperative perception aimed at understanding the limitations and potential risks of cooperative perception systems in real-world applications. To achieve this, we conduct a comprehensive analysis of error patterns and perform large-scale performance evaluations for multiple critical components of cooperative perception systems. First, we evaluate the performance differences of cooperative perception systems across diverse sensor types and cooperative agent configurations in heterogeneous environments. Additionally, we evaluate the cooperative perception systems online across 32 test routes, assessing their performance under both normal and abnormal communication conditions, including communication latency and pose errors. Our findings reveal potential vulnerabilities in key components of cooperative perception systems. Finally, we outline several promising directions for future research to develop more reliable and robust V2X cooperative perception systems.

ACKNOWLEDGEMENT

We would like to thank anonymous reviewers for their insightful comments. This work is supported partially by Shenzhen-Hong Kong-Macau Technology Research Programme (Type C) (Grant No.SGDX20230821091559018) and National Natural Science Foundation of China (Grant No. 62172210).

REFERENCES

- [1] M. A. Khan, H. El-Sayed, S. Malik, M. T. Zia, M. J. Khan, N. Alkaabi, and H. A. Ignatious, "Level-5 autonomous driving - are we there yet? A review of research literature," *ACM Comput. Surv.*, vol. 55, no. 2, pp. 27:1–27:38, 2023. [Online]. Available: <https://doi.org/10.1145/3485767>
- [2] L. Wen and K. Jo, "Deep learning-based perception systems for autonomous driving: A comprehensive survey," *Neurocomputing*, vol. 489, pp. 255–270, 2022. [Online]. Available: <https://doi.org/10.1016/j.neucom.2021.08.155>
- [3] Y. Li and J. Ibañez-Guzmán, "Lidar for autonomous driving: The principles, challenges, and trends for automotive lidar and perception systems," *IEEE Signal Process. Mag.*, vol. 37, no. 4, pp. 50–61, 2020. [Online]. Available: <https://doi.org/10.1109/MSP.2020.2973615>
- [4] "Waymo is way, way ahead on testing miles—that might not be a good thing," 2020. [Online]. Available: <https://arstechnica.com/cars/2020/01/waymo-is-way-way-ahead-on-testing-miles-that-might-not-be-a-good-thing/>
- [5] "U.s. probing fatal tesla crash that killed pedestrian," 2021. [Online]. Available: <https://www.reuters.com/business/autos-transportation/us-probing-fatal-tesla-crash-that-killed-pedestrian-2021-09-03/>
- [6] "U.s. agency probes tesla crashes that killed 2 motorcyclists," 2022. [Online]. Available: <https://www.cbsnews.com/news/tesla-crashes-kill-2-motorcyclists-autopilot-nhtsa/>
- [7] S. Hu, Z. Fang, Y. Deng, X. Chen, and Y. Fang, "Collaborative perception for connected and autonomous driving: Challenges, possible solutions and opportunities," *CoRR*, vol. abs/2401.01544, 2024. [Online]. Available: <https://doi.org/10.48550/arXiv.2401.01544>
- [8] V. Usinskis, M. Makulavicius, S. Petkevicius, A. Dzedzickis, and V. Bucinskis, "Towards autonomous driving: Technologies and data for vehicles-to-everything communication," *Sensors*, vol. 24, no. 11, p. 3411, 2024. [Online]. Available: <https://doi.org/10.3390/s24113411>
- [9] "Astri and baidu apollo collaborate to promote the implementation of c-v2x technology in hong kong," <https://www.astri.org/news-detail/astri-and-baidu-apollo-collaborate-to-promote-the-implementation-of-c-v2x-technology-in-hong-kong/>, 2023.
- [10] "V2v and v2x technology paves the way for autonomous driving," <https://www.assemblymag.com/articles/98013-v2v-and-v2x-technology-paves-the-way-for-autonomous-driving>, 2023.
- [11] R. Xu, H. Xiang, X. Xia, X. Han, J. Li, and J. Ma, "Opv2v: An open benchmark dataset and fusion pipeline for perception with vehicle-to-vehicle communication," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 2583–2589.
- [12] L. Hobert, A. Festag, I. Llatser, L. Altomare, F. Visintainer, and A. Kovács, "Enhancements of V2X communication in support of cooperative autonomous driving," *IEEE Commun. Mag.*, vol. 53, no. 12, pp. 64–70, 2015. [Online]. Available: <https://doi.org/10.1109/MCOM.2015.7355568>
- [13] C. Jung, D. Lee, S. Lee, and D. H. Shim, "V2x-communication-aided autonomous driving: System design and experimental validation," *Sensors*, vol. 20, no. 10, p. 2903, 2020. [Online]. Available: <https://doi.org/10.3390/s20102903>
- [14] T. Huang, J. Liu, X. Zhou, D. C. Nguyen, M. R. Azghadi, Y. Xia, Q. Han, and S. Sun, "V2X cooperative perception for autonomous driving: Recent advances and challenges," *CoRR*, vol. abs/2310.03525, 2023. [Online]. Available: <https://doi.org/10.48550/arXiv.2310.03525>
- [15] Z. Li, T. Yu, T. Suzuki, and K. Sakaguchi, "Building an SDVN framework for rsu-centric cooperative perception with heterogeneous V2X," in *20th IEEE Consumer Communications & Networking Conference, CCNC 2023, Las Vegas, NV, USA, January 8-11, 2023*. IEEE, 2023, pp. 1–7. [Online]. Available: <https://doi.org/10.1109/CCNC51644.2023.10060139>
- [16] R. Padilla, S. L. Netto, and E. A. B. da Silva, "A survey on performance metrics for object-detection algorithms," in *2020 International Conference on Systems, Signals and Image Processing, IWSSIP 2020, Niterói, Brazil, July 1-3, 2020*. IEEE, 2020, pp. 237–242. [Online]. Available: <https://doi.org/10.1109/IWSSIP48289.2020.9145130>
- [17] A. Stocco, B. Pulfer, and P. Tonella, "Model vs system level testing of autonomous driving systems: a replication and extension study," *Empir. Softw. Eng.*, vol. 28, no. 3, p. 73, 2023. [Online]. Available: <https://doi.org/10.1007/s10664-023-10306-x>
- [18] F. U. Haq, D. Shin, S. Nejati, and L. C. Briand, "Can offline testing of deep neural networks replace their online testing?" *Empir. Softw. Eng.*, vol. 26, no. 5, p. 90, 2021. [Online]. Available: <https://doi.org/10.1007/s10664-021-09982-4>
- [19] X. Gao, Z. Wang, Y. Feng, L. Ma, Z. Chen, and B. Xu, "Multitest: Physical-aware object insertion for testing multi-sensor fusion perception systems," in *Proceedings of the 46th IEEE/ACM International Conference on Software Engineering, ICSE 2024, Lisbon, Portugal, April 14-20, 2024*. ACM, 2024, pp. 139:1–139:13. [Online]. Available: <https://doi.org/10.1145/3597503.3639191>
- [20] T. Wang, S. Manivasagam, M. Liang, B. Yang, W. Zeng, and R. Urtaşun, "V2vnet: Vehicle-to-vehicle communication for joint perception and prediction," in *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part II*, ser. Lecture Notes in Computer Science, A. Vedaldi, H. Bischof, T. Brox, and J. Frahm, Eds., vol. 12347. Springer, 2020, pp. 605–621. [Online]. Available: https://doi.org/10.1007/978-3-030-58536-5_36
- [21] R. Xu, H. Xiang, Z. Tu, X. Xia, M. Yang, and J. Ma, "V2x-vit: Vehicle-to-everything cooperative perception with vision transformer," in *Computer Vision - ECCV 2022 - 17th European Conference, Tel Aviv, Israel, October 23-27, 2022, Proceedings, Part XXXIX*, ser. Lecture Notes in Computer Science, S. Avidan, G. J. Brostow, M. Cissé, G. M. Farinella, and T. Hassner, Eds., vol. 13699. Springer, 2022, pp. 107–124. [Online]. Available: https://doi.org/10.1007/978-3-031-19842-7_7
- [22] W. Liu, X. Xia, L. Xiong, Y. Lu, L. Gao, and Z. Yu, "Automated vehicle sideslip angle estimation considering signal measurement characteristic," *IEEE Sensors Journal*, vol. 21, no. 19, pp. 21 675–21 687, 2021.
- [23] A. Dosovitskiy, G. Ros, F. Codevilla, A. M. López, and V. Koltun, "CARLA: an open urban driving simulator," in *1st Annual Conference on Robot Learning, CoRL 2017, Mountain View, California, USA, November 13-15, 2017, Proceedings*, ser. Proceedings of Machine Learning Research, vol. 78. PMLR, 2017, pp. 1–16. [Online]. Available: <http://proceedings.mlr.press/v78/dosovitskiy17a.html>
- [24] G. Liu, Y. Hu, C. Xu, W. Mao, J. Ge, Z. Huang, Y. Lu, Y. Xu, J. Xia, Y. Wang *et al.*, "Towards collaborative autonomous driving: Simulation platform and end-to-end system," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025.
- [25] Y. Li, S. Ren, P. Wu, S. Chen, C. Feng, and W. Zhang, "Learning distilled collaboration graph for multi-agent perception," in *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, M. Ranzato, A. Beygelzimer, Y. N. Dauphin, P. Liang, and J. W. Vaughan, Eds., 2021, pp. 29 541–29 552. [Online]. Available: <https://proceedings.neurips.com/paper/2021/hash/f702defbc67edb455949f46babab0c18-Abstract.html>
- [26] Q. Chen, X. Ma, S. Tang, J. Guo, Q. Yang, and S. Fu, "F-cooper: feature based cooperative perception for autonomous vehicle edge computing system using 3d point clouds," in *Proceedings of the 4th ACM/IEEE Symposium on Edge Computing, SEC 2019, Arlington, Virginia, USA, November 7-9, 2019*, S. Chen, R. Onishi, G. Ananthanarayanan, and Q. Li, Eds. ACM, 2019, pp. 88–100. [Online]. Available: <https://doi.org/10.1145/3318216.3363300>
- [27] Y. Lu, Y. Hu, Y. Zhong, D. Wang, Y. Wang, and S. Chen, "An extensible framework for open heterogeneous collaborative perception," in *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024. [Online]. Available: <https://openreview.net/forum?id=KkrDUGIASK>
- [28] A. Guo, X. Gao, Z. Chen, Y. Xiao, J. Liu, X. Ge, W. Sun, and C. Fang, "Cootest: An automated testing approach for V2X communication systems," in *Proceedings of the 33rd ACM SIGSOFT International Symposium on Software Testing and Analysis, ISSTA 2024, Vienna, Austria, September 16-20, 2024*, M. Christakis and M. Pradel, Eds. ACM, 2024, pp. 1453–1465. [Online]. Available: <https://doi.org/10.1145/3650212.3680373>
- [29] "Supplementary website," <https://sites.google.com/view/v2x-empirical/>, 2025.
- [30] R. Xu, X. Xia, J. Li, H. Li, S. Zhang, Z. Tu, Z. Meng, H. Xiang, X. Dong, R. Song, H. Yu, B. Zhou, and J. Ma, "V2v4real: A real-world large-scale dataset for vehicle-to-vehicle cooperative perception," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June*

- 17-24, 2023. IEEE, 2023, pp. 13712–13722. [Online]. Available: <https://doi.org/10.1109/CVPR52729.2023.01318>
- [31] M. Everingham, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman, “The pascal visual object classes (voc) challenge,” *International journal of computer vision*, vol. 88, pp. 303–338, 2010. [Online]. Available: <https://doi.org/10.1007/s11263-009-0275-4>
- [32] “Evaluation criteria for the leaderboard 1.0.” [Online]. Available: https://leaderboard.carla.org/evaluation_v1_0/
- [33] S. Liu, C. Gao, Y. Chen, X. Peng, X. Kong, K. Wang, R. Xu, W. Jiang, H. Xiang, J. Ma *et al.*, “Towards vehicle-to-everything autonomous driving: A survey on collaborative perception,” *arXiv preprint arXiv:2308.16714*, 2023. [Online]. Available: <https://doi.org/10.48550/arXiv.2308.16714>
- [34] Z. Zhong, Z. Hu, S. Guo, X. Zhang, Z. Zhong, and B. Ray, “Detecting multi-sensor fusion errors in advanced driver-assistance systems,” in *ISSTA '22: 31st ACM SIGSOFT International Symposium on Software Testing and Analysis, Virtual Event, South Korea, July 18 - 22, 2022*, S. Ryu and Y. Smaragdakis, Eds. ACM, 2022, pp. 493–505. [Online]. Available: <https://doi.org/10.1145/3533767.3534223>
- [35] N. Cerpa, M. Bardeen, B. A. Kitchenham, and J. M. Verner, “Evaluating logistic regression models to estimate software project outcomes,” *Inf. Softw. Technol.*, vol. 52, no. 9, pp. 934–944, 2010. [Online]. Available: <https://doi.org/10.1016/j.infsof.2010.03.011>
- [36] Q. Chen, S. Tang, Q. Yang, and S. Fu, “Cooper: Cooperative perception for connected autonomous vehicles based on 3d point clouds,” in *39th IEEE International Conference on Distributed Computing Systems, ICDCS 2019, Dallas, TX, USA, July 7-10, 2019*. IEEE, 2019, pp. 514–524. [Online]. Available: <https://doi.org/10.1109/ICDCS.2019.00058>
- [37] Y. Han, H. Zhang, H. Li, Y. Jin, C. Lang, and Y. Li, “Collaborative perception in autonomous driving: Methods, datasets, and challenges,” *IEEE Intell. Transp. Syst. Mag.*, vol. 15, no. 6, pp. 131–151, 2023. [Online]. Available: <https://doi.org/10.1109/MITS.2023.3298534>
- [38] Y. Zhou, J. Xiao, Y. Zhou, and G. Loianno, “Multi-robot collaborative perception with graph neural networks,” *IEEE Robotics Autom. Lett.*, vol. 7, no. 2, pp. 2289–2296, 2022. [Online]. Available: <https://doi.org/10.1109/LRA.2022.3141661>
- [39] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention is all you need,” in *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, I. Guyon, U. von Luxburg, S. Bengio, H. M. Wallach, R. Fergus, S. V. N. Vishwanathan, and R. Garnett, Eds., 2017, pp. 5998–6008. [Online]. Available: <https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fdb053c1c4a845aa-Abstract.html>
- [40] R. Xu, Z. Tu, H. Xiang, W. Shao, B. Zhou, and J. Ma, “Cobevt: Cooperative bird’s eye view semantic segmentation with sparse transformers,” in *Conference on Robot Learning, CoRL 2022, 14-18 December 2022, Auckland, New Zealand*, ser. Proceedings of Machine Learning Research, K. Liu, D. Kulic, and J. Ichnowski, Eds., vol. 205. PMLR, 2022, pp. 989–1000. [Online]. Available: <https://proceedings.mlr.press/v205/xu23a.html>
- [41] J. Tan, F. Lyu, L. Li, F. Hu, T. Feng, F. Xu, Z. Zhang, R. Yao, and L. Wang, “Dynamic v2x perception from road-to-vehicle vision,” *IEEE Transactions on Intelligent Vehicles*, 2024.
- [42] C. Lin, D. Tian, X. Duan, J. Zhou, D. Zhao, and D. Cao, “V2vformer: Vehicle-to-vehicle cooperative perception with spatial-channel transformer,” *IEEE Trans. Intell. Veh.*, vol. 9, no. 2, pp. 3384–3395, 2024. [Online]. Available: <https://doi.org/10.1109/TIV.2024.3353254>
- [43] S. Ren, Z. Lei, Z. Wang, M. Dianati, Y. Wang, S. Chen, and W. Zhang, “Interruption-aware cooperative perception for V2X communication-aided autonomous driving,” *IEEE Trans. Intell. Veh.*, vol. 9, no. 4, pp. 4698–4714, 2024. [Online]. Available: <https://doi.org/10.1109/TIV.2024.3371974>
- [44] J. Li, R. Xu, X. Liu, J. Ma, Z. Chi, J. Ma, and H. Yu, “Learning for vehicle-to-vehicle cooperative perception under lossy communication,” *IEEE Trans. Intell. Veh.*, vol. 8, no. 4, pp. 2650–2660, 2023. [Online]. Available: <https://doi.org/10.1109/TIV.2023.3260040>
- [45] A. Guo, Y. Zhou, H. Tian, C. Fang, Y. Sun, W. Sun, X. Gao, A. T. Luu, Y. Liu, and Z. Chen, “Sovar: Build generalizable scenarios from accident reports for autonomous driving testing,” in *Proceedings of the 39th IEEE/ACM International Conference on Automated Software Engineering, ASE 2024, Sacramento, CA, USA, October 27 - November 1, 2024*, V. Filkov, B. Ray, and M. Zhou, Eds. ACM, 2024, pp. 268–280. [Online]. Available: <https://doi.org/10.1145/3691620.3695037>
- [46] A. Gambi, T. Huynh, and G. Fraser, “Generating effective test cases for self-driving cars from police reports,” in *Proceedings of the ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering, ESEC/SIGSOFT FSE 2019, Tallinn, Estonia, August 26-30, 2019*, M. Dumas, D. Pfahl, S. Apel, and A. Russo, Eds. ACM, 2019, pp. 257–267. [Online]. Available: <https://doi.org/10.1145/3338906.3338942>
- [47] X. Zhang and Y. Cai, “Building critical testing scenarios for autonomous driving from real accidents,” in *Proceedings of the 32nd ACM SIGSOFT International Symposium on Software Testing and Analysis, ISSTA 2023, Seattle, WA, USA, July 17-21, 2023*, R. Just and G. Fraser, Eds. ACM, 2023, pp. 462–474. [Online]. Available: <https://doi.org/10.1145/3597926.3598070>
- [48] Y. Huai, Y. Chen, S. Almanee, T. Ngo, X. Liao, Z. Wan, Q. A. Chen, and J. Garcia, “Doppelgänger test generation for revealing bugs in autonomous driving software,” in *45th IEEE/ACM International Conference on Software Engineering, ICSE 2023, Melbourne, Australia, May 14-20, 2023*. IEEE, 2023, pp. 2591–2603. [Online]. Available: <https://doi.org/10.1109/ICSE48619.2023.00216>
- [49] C. Hildebrandt, M. von Stein, and S. G. Elbaum, “Physcov: Physical test coverage for autonomous vehicles,” in *Proceedings of the 32nd ACM SIGSOFT International Symposium on Software Testing and Analysis, ISSTA 2023, Seattle, WA, USA, July 17-21, 2023*, R. Just and G. Fraser, Eds. ACM, 2023, pp. 449–461. [Online]. Available: <https://doi.org/10.1145/3597926.3598069>
- [50] Z. Zhong, G. E. Kaiser, and B. Ray, “Neural network guided evolutionary fuzzing for finding traffic violations of autonomous vehicles,” *IEEE Trans. Software Eng.*, vol. 49, no. 4, pp. 1860–1875, 2023. [Online]. Available: <https://doi.org/10.1109/TSE.2022.3195640>
- [51] H. Tian, X. Han, G. Wu, A. Guo, Y. Zhou, J. Zhang, S. Li, J. Wei, and T. Zhang, “An llm-empowered adaptive evolutionary algorithm for multi-component deep learning systems,” in *AAAI-25, Sponsored by the Association for the Advancement of Artificial Intelligence, February 25 - March 4, 2025, Philadelphia, PA, USA*, T. Walsh, J. Shah, and Z. Kolter, Eds. AAAI Press, 2025, pp. 20894–20902. [Online]. Available: <https://doi.org/10.1609/aaai.v39i19.34303>
- [52] A. Guo, Y. Feng, Y. Cheng, and Z. Chen, “Semantic-guided fuzzing for virtual testing of autonomous driving systems,” *J. Syst. Softw.*, vol. 212, p. 112017, 2024. [Online]. Available: <https://doi.org/10.1016/j.jss.2024.112017>
- [53] S. Wang and Z. Su, “Metamorphic object insertion for testing object detection systems,” in *35th IEEE/ACM International Conference on Automated Software Engineering, ASE 2020, Melbourne, Australia, September 21-25, 2020*. IEEE, 2020, pp. 1053–1065. [Online]. Available: <https://doi.org/10.1145/3324884.3416584>
- [54] A. Guo, Y. Feng, and Z. Chen, “Lirtest: augmenting lidar point clouds for automated testing of autonomous driving systems,” in *ISSTA '22: 31st ACM SIGSOFT International Symposium on Software Testing and Analysis, Virtual Event, South Korea, July 18 - 22, 2022*, S. Ryu and Y. Smaragdakis, Eds. ACM, 2022, pp. 480–492. [Online]. Available: <https://doi.org/10.1145/3533767.3534397>
- [55] G. Christian, T. Woodlief, and S. Elbaum, “Generating realistic and diverse tests for lidar-based perception systems,” in *2023 IEEE/ACM 45th International Conference on Software Engineering (ICSE)*. IEEE, 2023, pp. 2604–2616. [Online]. Available: <https://doi.org/10.1109/ICSE48619.2023.00217>
- [56] L. Li, X. Zhang, S. Tang, Z. Zhang, and J. Zhao, “Impact of V2V communication on robustness of autonomous driving systems,” in *35th IEEE International Symposium on Software Reliability Engineering, ISSRE 2024 - Workshops, Tsukuba, Japan, October 28-31, 2024*. IEEE, 2024, pp. 151–154. [Online]. Available: <https://doi.org/10.1109/ISSREW63542.2024.00073>
- [57] F. Hawlader, F. Robinet, and R. Frank, “Cooperative perception using V2X communications: An experimental study,” in *100th IEEE Vehicular Technology Conference, VTC Fall 2024, Washington, DC, USA, October 7-10, 2024*. IEEE, 2024, pp. 1–7. [Online]. Available: <https://doi.org/10.1109/VTC2024-Fall63153.2024.10757448>