

Breaking Data Silos: Cross-Domain Learning for Multi-Agent Perception from Independent Private Sources

Jinlong Li¹, Baolu Li¹, Xinyu Liu¹, Runsheng Xu², Jiaqi Ma², Hongkai Yu^{1*}

Abstract—The diverse agents in multi-agent perception systems may be from different companies. Each company might use the identical classic neural network architecture based encoder for feature extraction. However, the data source to train the various agents is independent and private in each company, leading to the Distribution Gap of different private data for training distinct agents in multi-agent perception system. The data silos by the above Distribution Gap could result in a significant performance decline in multi-agent perception. In this paper, we thoroughly examine the impact of the distribution gap on existing multi-agent perception systems. To break the data silos, we introduce the Feature Distribution-aware Aggregation (FDA) framework for cross-domain learning to mitigate the above Distribution Gap in multi-agent perception. FDA comprises two key components: Learnable Feature Compensation Module and Distribution-aware Statistical Consistency Module, both aimed at enhancing intermediate features to minimize the distribution gap among multi-agent features. Intensive experiments on the public OPV2V and V2XSet datasets underscore FDA’s effectiveness in point cloud-based 3D object detection, presenting it as an invaluable augmentation to existing multi-agent perception systems. The code is available at <https://github.com/jinlong17/BDS-V2V>.

I. INTRODUCTION

Although advancements in deep learning have greatly enhanced single-agent perception, the challenges of long-range detection, occlusion issues, and field of view still persist [1]–[3]. Multi-agent perception systems offer a solution by leveraging Vehicle-to-Everything (V2X) communication technology, allowing multiple nearby agents to share visual information, including detection results, raw sensor data, and intermediate features [4]–[7].

Existing multi-agent perception methods often operate under a strong assumption: all agents are trained using identical training data. However, this assumption might be not true. The diverse agents in multi-agent perception systems may be from different companies. Some companies whose perception systems are established on the same popular open platform might use the identical classic neural network architecture-based encoder for feature extraction. However, the data source to train the various agents is independent and private in each company as depicted in Fig. 1, leading to the Distribution Gap of different private data for training distinct agents in multi-agent perception systems. As an example of Distribution Gaps, let us compare the two widely-used large-scale simulated V2X cooperative perception datasets, namely OPV2V [5] and V2XSet [8]. We find key differences in the following two perspectives.

¹Cleveland State University, Cleveland Vision & AI Lab. ²University of California, Los Angeles, UCLA Mobility Lab.

*Corresponding Author: h.yu19@csuohio.edu

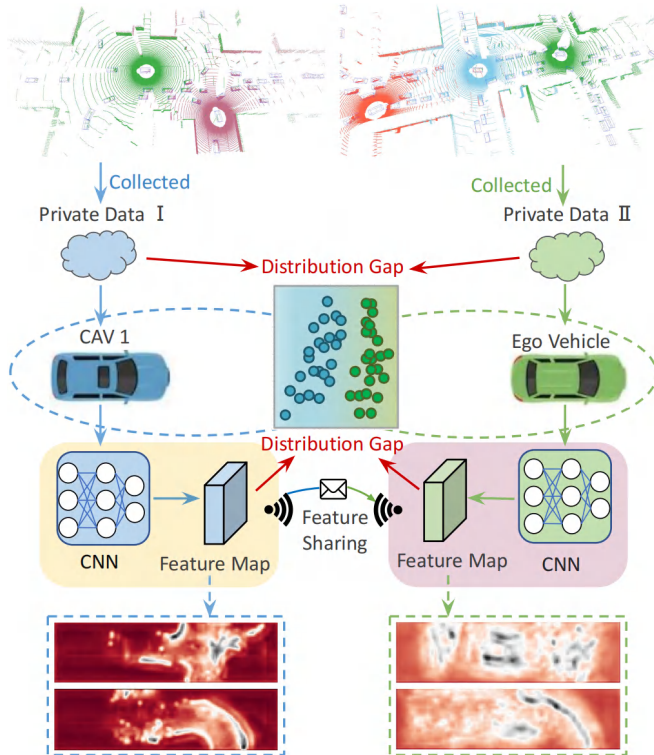


Fig. 1: Illustration of the Distribution Gap of different independent private data for training distinct agents in multi-agent perception. Here we use V2X cooperative perception in autonomous driving as an example.

- **Agent Number and Type:** The presence of diverse agent numbers within the same scenario can yield a different pool of shared visual information. Variations in the quantity and types of agents, *e.g.*, Connected and Automated Vehicles (CAVs), and Infrastructures can also contribute to disparities in the sharing of intermediate features.
- **Scene:** The density of LiDAR point cloud data can vary significantly based on the size and shape of the roadway environment in which the agents operate, leading to fluctuations in feature sharing.

The data silos by the above Distribution Gap could result in a significant performance decline in multi-agent perception. A naive solution for solving this problem might be using federated learning-based methods. Federated learning represents a privacy-preserving collaborative machine learning paradigm [9], [10], designed to facilitate the collaborative training of a shared global model among multiple clients




Fig. 2: **Architecture of the proposed Feature Distribution-aware Aggregation (FDA) framework.** FDA leverages LFCM to generate residual compensation map to enhance other-agent features, then utilizes DSCM to mitigate Distribution Gaps.

by exchanging information pertaining to model parameters. However, in the autonomous driving industry, *the data collection and annotation are extremely expensive and with business secrets, so the companies usually reject to share independent and private data sources due to business competition and privacy concerns.* As a result, breaking data silos for multi-agent perception is never studied before.

To address the Distribution Gap issue, we introduce the **Feature Distribution-aware Aggregation** framework for Multi-agent perception system on 3D object detection task, named as FDA. We choose the task of V2X cooperative perception for the point cloud-based 3D object detection as an investigation in this paper. Specifically, we proposed two key components: 1) Learnable Feature Compensation Module (LFCM) to generate a residual compensation map for CAV’s intermediate feature considering the large-scale feature information; 2) Distribution-aware Statistical Consistency Module (DSCM) to diminish the distribution gap between CAV feature and ego feature in terms of feature distributions. We conducted extensive experiments on two public V2X perception datasets, namely OPV2V [5] and V2XSet [8], to justify the effectiveness of our proposed method. Our contributions are summarized as follows:

- To the best of our knowledge, we propose a novel Feature Distribution-aware Aggregation framework, dubbed as FDA, which is the **first research** on multi-agent perception to address the Distribution Gap of different independent private data for training distinct agents.
- We propose a novel Learnable Feature Compensation Module (LFCM) to effectively mitigate discrepancies arising from various agents and a Distribution-aware Statistical Consistency Module (DSCM) to diminish the differences between intermediate features extracted

from the ego agent and other agents.

- We evaluate the proposed FDA framework on the large-scale simulated OPV2V and V2XSet datasets, whose experimental results demonstrate our superior performance in point cloud-based 3D object detection.

II. RELATED WORK

Perception on multi-agent system. Multi-vehicle perception systems aim to overcome the limitations of single-vehicle systems by leveraging multi-vehicle information. Collaboration modules are commonly developed by researchers to improve efficiency and performance. Typically, three approaches are used for aggregating multi-vehicle observations: raw data fusion, feature combination during processing, and output fusion. State-of-the-art methods often opt for sharing intermediate neural features to balance accuracy and bandwidth requirements [4], [5], [8]. Attfuse [5] harnesses self-attention models for the decompressed features amalgamation. SCOPE [11] introduces a learning-based framework that addresses multi-agent challenges with a first focus on the temporal context of the ego agent. V2X-ViT [8] presents a unified ViT architecture for V2X perception which is able to capture the heterogeneous nature of V2X systems. CoBEVT [12] is a pioneering multi-agent perception framework that collaboratively generates predictions using a ViT to enhance performance. Despite the impressive performance of these methods in V2V perception, their deep learning-based network is trained with homogeneous data from ego vehicles and other vehicles. There exists an inherent data gap in the real world. In this paper, we aim to tackle this challenge.

Vehicle-to-Everything Dataset. Numerous datasets, *e.g.*, Cityscapes [13] and KITTI [14], serve single-agent perception. Large-scale datasets are crucial for robust cooperative perception models in multi-agent systems. How-

ever, collecting real-world multi-agent data is challenging and costly. To address this, simulators like CARLA [15] and OpenCDA [16] are used to gather cooperative perception data. V2X-Sim [17], for instance, provides a simulated dataset for multi-agent perception in V2X-assisted autonomous driving, including records from various agents. For simulated V2V perception, OPV2V [5] offers an extensive open-source dataset with diverse scenes simulating real-world traffic dynamics, including a digital rendition of Culver City. In real-world data, DAIR-V2X [18] features multi-modal, multi-view data extracted from real-world Vehicle-Infrastructure Cooperative Autonomous Driving scenarios. V2V4Real is exclusively tailored for V2V cooperative autonomous perception. To address the shortage of real-world sequential data, V2X-Seq [19] captures essential elements from natural scenery. Current state-of-the-art cooperative perception methods assume uniform data distributions for ego vehicle and surrounding CAV encoders, using identical training datasets. This paper investigates cooperative perception performance in the context of diverse data distributions.

Deployment of Multi-Agent System. Compared to single-agent systems, multi-agent systems with V2V communication introduce challenges like communication latency, lossy communication, localization errors, and adversarial attacks, which can undermine collaboration benefits [7], [8]. Recent research has made progress in enhancing system robustness. V2X-ViT [8] uses a Vision Transformer (ViT) to handle GPS localization errors and sensing information delays. To address localization errors, Vadivelu et al. [20] proposed a pose regression module that learns a correction parameter to predict the true relative transformation from noisy data. The LC-aware Repair Network [21] is introduced to enhance collaborative perception robustness under lossy communication conditions, addressing packet loss problems. A model-agnostic framework [22] is proposed to address model heterogeneity in collaborative perception. [7] proposes MPDA framework to narrow the domain gap in spatial resolution, channel count, and patterns in multi-agent perception. [23] introduces ROBOSAC to enhance adversarially robust collaborative perception by promoting consensus among robots during collaboration. In this paper, we aim to address the distribution gap arising from differently distributed data to enhance the robustness of cooperative perception.

III. METHODOLOGY

A. Overview of the Feature Distribution-aware Aggregation

The Feature Distribution-aware Aggregation framework via the V2V cooperative perception pipeline is illustrated in Fig. 2. First, we select an ego vehicle among the CAVs to create a spatial graph that encompasses nearby CAVs within the communication range. Because of the similar sharing capabilities of CAVs and infrastructures, we consider each intelligent infrastructure as a CAV during methodology development in this paper. All the other nearby CAVs will project their own LiDAR data onto the ego vehicle’s coordinate frame, based on both the ego vehicle’s and their own GPS poses. The point clouds from the ego and CAVs are denoted

as $\mathbf{P}_{ego} \in \mathbb{R}^{4 \times m}$ and $\mathbf{P}_{cav} \in \mathbb{R}^{4 \times m}$, respectively. Each CAV has its own Encoder for LiDAR feature extraction. After feature extraction of each CAV, the ego vehicle receives the neighboring CAV visual features via V2V communication. The intermediate features aggregated from N surrounding CAVs are denoted as $\mathbf{F}^{cav} \in \mathbb{R}^{N \times H \times W \times C}$, and the ego intermediate features are denoted as $\mathbf{F}^{ego} \in \mathbb{R}^{1 \times H \times W \times C}$. These, along with features from other CAVs, are fed into our proposed LFCM and DSCM modules to reduce feature disparities. After processing by our proposed modules, the intermediate features are fused using the Feature Fusion Network (FFN). Finally, the resulting fused feature maps are fed into a prediction header for 3D bounding-box regression and classification. We formulate the original V2V cooperative perception system for LiDAR-based 3D object detection as

$$F(\mathbf{P}_{cav}, \mathbf{P}_{ego}) = \mathbf{H}(\text{FFN}(\mathbf{F}^{cav}, \mathbf{F}^{ego})), \quad (1)$$

$$\mathbf{F}^{cav} = \mathbf{E}_{cav}(\mathbf{P}_{cav}), \quad \mathbf{F}^{ego} = \mathbf{E}_{ego}(\mathbf{P}_{ego}), \quad (2)$$

where $\text{FFN}(\cdot)$ is the Feature Fusion Network responsible for fusing the features of CAVs and the ego vehicle, and \mathbf{H} is the prediction header for 3D object detection. Existing cooperative perception methods assume that the encoder weights (w) of \mathbf{E}_{ego} and \mathbf{E}_{cav} are shared, indicating they are trained with the same distribution data (*i.e.*, $\mathbf{E}_{ego}^w = \mathbf{E}_{cav}^w$). However, in the real-world deployment, $\mathbf{E}_{ego}^w \neq \mathbf{E}_{cav}^w$, leading to a distribution gap between CAVs and the ego vehicle for most existing models. To address this, we propose a feature distribution-aware aggregation framework, which can serve as a plug-in module in the V2V cooperative perception system. Then, the Eq. (1) can be modified as

$$F(\mathbf{P}_{cav}, \mathbf{P}_{ego}) = \mathbf{H}(\text{FFN}(\text{FDA}(\mathbf{F}_{cav}, \mathbf{F}_{ego}), \mathbf{F}_{ego})), \quad (3)$$

where $\text{FDA}(\cdot)$ refers to our proposed Feature Distribution-aware Aggregation framework, comprising the learnable feature compensation module and the distribution-aware statistical consistency module, as depicted in Fig. 2.

B. Learnable Feature Compensation Module

After encoding by \mathbf{E}_{ego} and \mathbf{E}_{cav} , we obtain intermediate neural features \mathbf{F}^{ego} and \mathbf{F}^{cav} from the ego vehicle and other CAVs, respectively. Due to the difference between \mathbf{E}_{ego}^w and \mathbf{E}_{cav}^w based on distinct training data, even when the same point cloud is fed into both encoders, different feature maps are obtained. To mitigate the impact of varying weights in the encoders, we introduce a Learnable Feature Compensation Module (LFCM). The LFCM is an encoder-decoder architecture with skip connections, inspired by the work of [24], and depicted in Fig. 2. The encoder and decoder components of our LFCM employ 5×5 convolutional layers to capture large-scale spatial features. The shared features \mathbf{F}^{cav} from CAVs are fed into the LFCM to predict residual compensation scores for the entire feature set, generating a spatial compensation map $\mathbf{M}_c \in \mathbb{R}^{N \times H \times W \times C}$ of the same size as \mathbf{F}^{cav} . Subsequently, \mathbf{M}_c is added to \mathbf{F}^{cav} , incorporating residual compensation levels into consideration, resulting in

the enhanced compensated feature $\widehat{\mathbf{F}}^{cav}$. The computation of the proposed LFCM is formulated as

$$\widehat{\mathbf{F}}^{cav} = \text{LFCM}(\mathbf{F}^{cav}) + \mathbf{F}^{cav}, \quad (4)$$

where $\text{LFCM}(\cdot)$ represents our proposed LFCM responsible for generating the residual compensation map.

C. Distribution-aware Statistical Consistency Module

After processing with our LFCM, we obtain enhanced compensated features for the CAV, denoted as $\widehat{\mathbf{F}}^{cav}$, and the ego intermediate features \mathbf{F}^{ego} . To address the distribution gap between $\widehat{\mathbf{F}}^{cav}$ and \mathbf{F}^{ego} , it is crucial to examine the relationship between statistical differences and feature distributions. Previous studies [25] have established a positive correlation between statistical differences and distribution disparities. In order to minimize the discrepancy in feature distributions between $\widehat{\mathbf{F}}^{cav}$ and \mathbf{F}^{ego} , we introduce the Maximum Mean Discrepancy (MMD) [26] distance as a metric. Let $\mathcal{F}^{ego} = \{\mathbf{F}_i^{ego}\}$ and $\mathcal{F}^{cav} = \{\widehat{\mathbf{F}}_i^{cav}\}$ represent the sets of ego and compensated CAV features, respectively. Our goal is to enhance the compensated CAV features to reduce the distribution gap with ego features in terms of feature distributions by tuning the parameters of LFCM. This process can be formulated as follow:

$$\arg \min_{\text{LFCM}} L_{mmd}(\mathcal{F}^{ego}, \mathcal{F}^{cav}), \quad (5)$$

where $L_{mmd}(\cdot)$ represents the MMD loss between the two sets of intermediate features [26].

D. Loss Function

For 3D object detection, as in [7], [8], we compute the smooth L1 loss for bounding box regression and apply focal loss [27] for classification. In the context of our FDA framework, we employ the Maximum Mean Discrepancy (MMD) loss to address distribution differences between ego and other CAV agents. The final loss function is a combination of these two losses:

$$L_{total} = \lambda L_{det} + \omega L_{mmd}, \quad (6)$$

where λ and ω are the balancing coefficients, both ranging within [0, 1].

IV. EXPERIMENTS

Dataset: Our experiments are conducted on two publicly available benchmark datasets for V2V/V2X cooperative perception tasks: OPV2V [5] and V2XSet [8]. **OPV2V** is a large-scale simulated dataset tailored for V2V cooperative perception tasks, collected using the CARLA [15] and OpenCDA [16] platforms. It comprises 73 diverse scenes featuring varying numbers of connected vehicles. The dataset is divided into training (6,764 frames), validation (1,981 frames), and testing (2,719 frames) sets. **V2XSet** is another extensive simulated dataset designed for V2X cooperative perception tasks, also collected using CARLA. This dataset provides LiDAR data from multiple autonomous vehicles and roadside intelligent infrastructure, all timestamped within the same scenarios. Its training/validation/testing set is split into 6,694, 1,920, and 2,833 frames, respectively

A. Experimental Setup

Evaluation Metrics: Following [5], [8], our performance evaluation centers around the final 3D vehicle detection accuracy. We set the evaluation range as $x \in [-140, 140]$ meters, $y \in [-40, 40]$ meters, where all CAVs are included in this spatial range in the experiment. We measure the accuracy with Average Precisions (AP) at Intersection-over-Union (IoU) thresholds of 0.5 and 0.7.

Experimental Details: In our investigation aimed at addressing the distribution gap in LiDAR-based 3D object detection, we make the assumption that all agents employ the same cooperative perception method. However, there is a distinction between the weight of encoders used for the ego vehicle and the CAVs, denoted as $\mathbf{E}_{ego}^w \neq \mathbf{E}_{cav}^w$. Unlike traditional training strategies where all encoders are trained on the same dataset, we introduce a novel approach. Specifically, we use the OPV2V and V2XSet training sets to train the ego encoder and CAV encoders, respectively. This configuration involves scenarios where the CAV encoders are trained on the OPV2V training set, while the ego encoder is trained on the V2XSet training set, defined as $OPV2V \rightarrow V2XSet$. The testing sets of these datasets (*i.e.*, OPV2V and V2XSet) are used to evaluate our cooperative methods. We assess our models under two key settings:

- 1) **V2XSet \rightarrow OPV2V:** All cooperative perception methods are evaluated on the OPV2V testing set, with CAV encoders trained on V2XSet training set while the ego encoder is trained on OPV2V training set.
- 2) **OPV2V \rightarrow V2XSet:** Here, all cooperative perception methods are evaluated on the V2XSet testing set, with CAV encoders trained on OPV2V training set while the ego encoder is trained on V2XSet training set.

Compared Methods: We utilize four state-of-the-art intermediate fusion methods to assess the distribution gap, comprising two attention-based fusion methods, Attfuse [5] and V2VAM [28], and two ViT-based fusion methods, V2X-ViT [8] and CoBEVT [12]. We establish a *No Fusion* baseline where no collaboration is involved in the system. To illustrate the significant impact of the distribution gap, we train and evaluate these four fusion methods under two distinct settings. Additionally, we finetune these models based on the specific dataset used for the ego encoder. Furthermore, to showcase the effectiveness of our proposed FDA framework in addressing the distribution gap, we apply our proposed LFCM and DSCM to finetune these fusion methods. This finetuning process serves to mitigate the distribution gap between the ego and CAV encoders.

Implementation Details: All cooperative perception models utilize PointPillar [29] as the backbone. We employ the Adam optimizer [30] with an initial learning rate of 10^{-3} , gradually decaying it every 10 epochs using a factor of 0.1. The hyperparameters align with those settings of [8]. To simulate real-world scenarios where CAV encoders originate from diverse sources that are challenging to update, we freeze the parameters of CAV encoders during finetuning. All models are trained using two RTX 3090 GPUs. The

TABLE I: **3D detection performance on OPV2V testing set under V2XSet \rightarrow OPV2V setting.** We show Average Precision (AP) at IoU=0.5, 0.7. *Dist. Gap* represents the distribution gap scenario, *No Dist* represents agents' encoders consistently trained on the same training set. *Dist* stands for Agents' encoders trained on different training sets, causing a distribution gap. Note that the training set type for original CAV encoders and ego encoder training is defined as S_{cav} , S_{ego} , respectively.

$S_{cav} \parallel S_{ego}$	<i>Dist. Gap</i>	NO Fusion		Attfuse [5]		V2VAM [28]		V2X-ViT [8]		CoBEVT [12]	
		AP@0.5	AP@0.7	AP@0.5	AP@0.7	AP@0.5	AP@0.7	AP@0.5	AP@0.7	AP@0.5	AP@0.7
OPV2V \parallel OPV2V	<i>No Dist</i>	67.9	60.2	90.7	81.0	91.6	84.9	90.2	79.2	92.5	85.2
V2XSet \parallel OPV2V	<i>Dist</i>	67.9	60.2	36.1	13.0	74.7	56.6	42.7	19.5	72.7	55.2
V2XSet \parallel OPV2V	<i>Dist (Finetuned)</i>	67.9	60.2	69.6	43.4	87.7	72.0	74.7	36.6	77.0	56.0
V2XSet \parallel OPV2V	<i>Dist + FDA</i>	67.9	60.2	85.7	69.0	88.3	79.4	89.7	75.2	89.0	74.7

TABLE II: **3D detection performance on V2XSet testing set under OPV2V \rightarrow V2XSet setting.** We show Average Precision (AP) at IoU=0.5, 0.7. Note that the training set type for original CAV encoders and ego encoder training are defined as S_{cav} , S_{ego} , respectively

$S_{cav} \parallel S_{ego}$	<i>Dist. Gap</i>	NO Fusion		Attfuse [5]		V2VAM [28]		V2X-ViT [8]		CoBEVT [12]	
		AP@0.5	AP@0.7	AP@0.5	AP@0.7	AP@0.5	AP@0.7	AP@0.5	AP@0.7	AP@0.5	AP@0.7
V2XSet \parallel V2XSet	<i>No Dist</i>	60.6	40.2	79.2	64.1	90.2	81.6	88.2	71.2	84.4	66.0
OPV2V \parallel V2XSet	<i>Dist</i>	60.6	40.2	45.1	32.7	70.6	53.9	50.0	36.4	55.5	40.7
OPV2V \parallel V2XSet	<i>Dist (Finetuned)</i>	60.6	40.2	64.6	46.8	84.6	70.4	68.0	47.9	74.9	54.7
OPV2V \parallel V2XSet	<i>Dist + FDA</i>	60.6	40.2	77.2	57.2	89.7	81.3	79.0	59.6	82.9	64.3

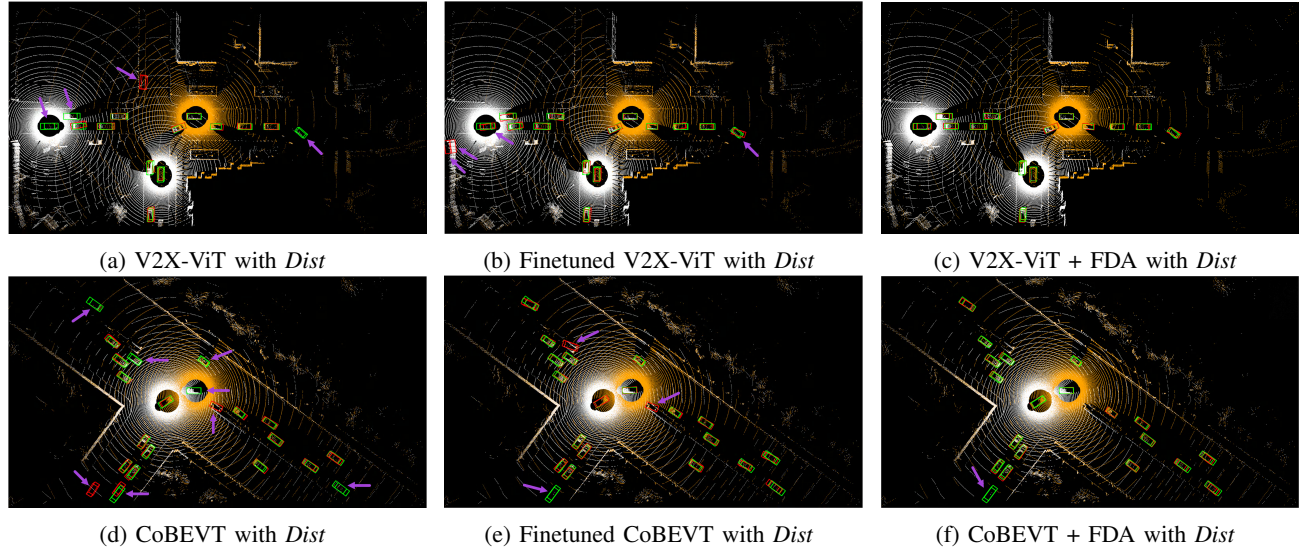


Fig. 3: **3D object detection visualization.** Orange point cloud is ego vehicle, and the white color of point clouds are CAVs. Green and red 3D bounding boxes represent the ground truth and prediction respectively. The detection results of the proposed FDA are clearly more accurate. False detection errors are highlighted using purple arrows.




Fig. 4: **Visualization of intermediate features before and after our FDA.** Two samples of point clouds are selected on OPV2V testing set to evaluate CoBEVT [12], where Orange point cloud is ego vehicle. It is evident that after applying the FDA, the intermediate features from CAVs exhibit more similar patterns to those of the ego. Bright pixels may tend to represent the objects.

TABLE III: Ablation study of the proposed FDA method (LFCM + DSCM) on OPV2V testing set under $V2XSet \rightarrow OPV2V$ setting. We show Average Precision (AP) at IoU=0.5.

LFCM	DSCM	V2X-ViT [8]	CoBEVT [12]
		42.7 (<i>Dist</i>)	72.7 (<i>Dist</i>)
✓		79.1 (+36.4)	80.3 (+7.6)
	✓	83.2 (+40.5)	83.9 (+11.2)
✓	✓	89.7 (+47.0)	89.0 (+16.3)

coefficients for the detection loss (λ) and our MMD loss (ω) are both set to 1.0.

B. Quantitative Evaluation

Performance Analysis of Distribution Gap: Table I provides the 3D object detection results on the OPV2V testing set. Under the $V2XSet \rightarrow OPV2V$ setting, we explore the impact of distribution disparities on the detection performance. Compared to the *No Dist*, where all agents are consistently trained with the OPV2V training set, the *Dist* for Attfuse [5] and V2X-ViT [8] exhibits a significant drop to 36.1%/13.0% and 42.7%/19.5% for AP@0.5/0.7, respectively. In the *Dist*, the agents of CAVs are trained on the V2XSet training set, leading to even lower performance than a single-agent perception system, *i.e.*, *NO Fusion*. These notable performance drops underscore the highly negative impact of the distribution gap. We perform finetuning on the four intermediate fusion methods using the OPV2V training set, specifically enhancing the ego encoder and feature fusion network while keeping the CAVs’ encoder parameters fixed. Finetuning improves performance in all fusion methods, but they still fall short of the *NO Fusion* baseline in AP@0.7. The key reason behind this limitation is that the CAVs’ encoder parameters remain fixed, potentially leading to sub-optimal feature extraction from the new and previously unseen distribution data. Our FDA method effectively addresses these limitations. When applied to all fusion methods, FDA achieves substantial improvements that nearly restore the original detection performance under a distribution gap. For instance, Attfuse and V2X-ViT with FDA exhibit impressive improvements of 49.6%/56.0% and 47.0%/55.7% for AP@0.5/0.7, respectively. Furthermore, the results presented in Table II also demonstrate the effectiveness of our proposed FDA method on the V2XSet dataset. Clearly, our proposed FDA method efficiently preserves the collaborative benefits under different distribution-based agents. It successfully mitigates the impact of distribution gaps, resulting in excellent cooperative perception performance.

Ablation Study: Table III presents the results of using two specially designed components in FDA framework on the detection performance of V2X-ViT [8] and CoBEVT [12]. The incorporation of LFCM and DSCM into V2X-ViT [8] leads to performance improvements of 36.4% and 40.5%, respectively, showing the effectiveness of our design.

3D Detection Visualization: We present visual comparisons of various methods in the OPV2V testing set scenario, illustrating their impact on V2X-ViT [8] and CoBEVT [12] in Fig. 3. Under the $V2XSet \rightarrow OPV2V$ setting, V2X-ViT and CoBEVT exhibit numerous false-negative and false-positive detection errors when no specially designed modules are applied. Even after direct finetuning of these two models on the OPV2V training set, detection performance improves, but noticeable missing detections and false positives persist, as evident in (b) and (e) of Fig. 3. However, when employing our proposed FDA, which takes into account the distribution gap, V2X-ViT and CoBEVT demonstrate more robust performance, successfully detecting most objects and mitigating false-negative and false-positive detection errors.

Feature Visualization: To analyze the impact of the distribution gap on agents, we visually present intermediate features in Fig. 4 using two point cloud samples. Under the $V2XSet \rightarrow OPV2V$ setting, observable disparities exist between ego and CAVs’ features. Our FDA effectively alleviates this gap, rendering the features more akin to those of the ego. In the case of CoBEVT, our FDA enhances local details, narrowing the gap with ego’s features, as seen in (b-d). This visual evidence underscores the effectiveness of our FDA.

V. CONCLUSIONS

In this paper, we present pioneering research on the Distribution Gap, stemming from disparate private training data for distinct agents in multi-agent perception systems, resulting in data silos. We analyze the impact of the Distribution Gap on existing cooperative perception methods. To mitigate data silos, we propose a Feature Distribution-aware Aggregation framework, comprising a Learnable Feature Compensation Module and a Distribution-aware Statistical Consistency Module. We empirically assess the FDA framework on two public and extensive simulated datasets, OPV2V and V2XSet, demonstrating its superior performance in point cloud-based 3D object detection. Our findings emphasize the crucial role of addressing the distribution gap in multi-agent perception systems, contributing to the advancement of cooperative perception strategies with potential implications for autonomous driving.

REFERENCES

- [1] 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 *ACM/IEEE Symposium on Edge Computing*, 2019, pp. 88–100.
- [2] R. Xu, X. Xia, J. Li, H. Li, S. Zhang, Z. Tu, Z. Meng, H. Xiang, X. Dong, R. Song *et al.*, “V2v4real: A real-world large-scale dataset for vehicle-to-vehicle cooperative perception,” in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 13 712–13 722.
- [3] B. Tang, Y. Zhong, C. Xu, W.-T. Wu, U. Neumann, Y. Zhang, S. Chen, and Y. Wang, “Collaborative uncertainty benefits multi-agent multimodal trajectory forecasting,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023.
- [4] T.-H. Wang, S. Manivasagam, M. Liang, B. Yang, W. Zeng, and R. Urtasun, “V2vnet: Vehicle-to-vehicle communication for joint perception and prediction,” in *European Conference on Computer Vision*. Springer, 2020, pp. 605–621.

- [5] 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 *International Conference on Robotics and Automation*. IEEE, 2022, pp. 2583–2589.
- [6] Y. Hu, S. Fang, Z. Lei, Y. Zhong, and S. Chen, "Where2comm: Communication-efficient collaborative perception via spatial confidence maps," in *Advances in Neural Information Processing Systems*, 2022.
- [7] R. Xu, J. Li, X. Dong, H. Yu, and J. Ma, "Bridging the domain gap for multi-agent perception," in *IEEE International Conference on Robotics and Automation*. IEEE, 2023, pp. 6035–6042.
- [8] R. Xu, H. Xiang, Z. Tu, X. Xia, M.-H. Yang, and J. Ma, "V2x-vit: Vehicle-to-everything cooperative perception with vision transformer," in *European Conference on Computer Vision*. Springer, 2022, pp. 107–124.
- [9] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Artificial intelligence and statistics*. PMLR, 2017, pp. 1273–1282.
- [10] Q. Yang, Y. Liu, T. Chen, and Y. Tong, "Federated machine learning: Concept and applications," *ACM Transactions on Intelligent Systems and Technology*, vol. 10, no. 2, pp. 1–19, 2019.
- [11] K. Yang, D. Yang, J. Zhang, M. Li, Y. Liu, J. Liu, H. Wang, P. Sun, and L. Song, "Spatio-temporal domain awareness for multi-agent collaborative perception," *IEEE/CVF International Conference on Computer Vision*, 2023.
- [12] 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*. PMLR, 2023, pp. 989–1000.
- [13] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele, "The cityscapes dataset for semantic urban scene understanding," in *IEEE conference on computer vision and pattern recognition*, 2016, pp. 3213–3223.
- [14] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? the kitti vision benchmark suite," in *IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2012, pp. 3354–3361.
- [15] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "Carla: An open urban driving simulator," in *Annual Conference on Robot Learning*. PMLR, 2017, pp. 1–16.
- [16] R. Xu, H. Xiang, X. Han, X. Xia, Z. Meng, C.-J. Chen, C. Correa-Jullian, and J. Ma, "The openca open-source ecosystem for cooperative driving automation research," *IEEE Transactions on Intelligent Vehicles*, pp. 1–13, 2023.
- [17] Y. Li, D. Ma, Z. An, Z. Wang, Y. Zhong, S. Chen, and C. Feng, "V2x-sim: Multi-agent collaborative perception dataset and benchmark for autonomous driving," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 10914–10921, 2022.
- [18] H. Yu, Y. Luo, M. Shu, Y. Huo, Z. Yang, Y. Shi, Z. Guo, H. Li, X. Hu, J. Yuan *et al.*, "Dair-v2x: A large-scale dataset for vehicle-infrastructure cooperative 3d object detection," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 21 361–21 370.
- [19] H. Yu, W. Yang, H. Ruan, Z. Yang, Y. Tang, X. Gao, X. Hao, Y. Shi, Y. Pan, N. Sun *et al.*, "V2x-seq: A large-scale sequential dataset for vehicle-infrastructure cooperative perception and forecasting," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 5486–5495.
- [20] N. Vaidvelu, M. Ren, J. Tu, J. Wang, and R. Urtasun, "Learning to communicate and correct pose errors," in *Conference on Robot Learning*. PMLR, 2021, pp. 1195–1210.
- [21] 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 Transactions on Intelligent Vehicles*, 2023.
- [22] R. Xu, W. Chen, H. Xiang, X. Xia, L. Liu, and J. Ma, "Model-agnostic multi-agent perception framework," in *IEEE International Conference on Robotics and Automation*. IEEE, 2023, pp. 1471–1478.
- [23] Y. Li, Q. Fang, J. Bai, S. Chen, F. Juefei-Xu, and C. Feng, "Among us: Adversarially robust collaborative perception by consensus," *IEEE/CVF International Conference on Computer Vision*, 2023.
- [24] B. Mildenhall, J. T. Barron, J. Chen, D. Sharlet, R. Ng, and R. Carroll, "Burst denoising with kernel prediction networks," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 2502–2510.
- [25] Y. Hou, Q. Guo, Y. Huang, X. Xie, L. Ma, and J. Zhao, "Evading deepfake detectors via adversarial statistical consistency," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 12 271–12 280.
- [26] A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola, "A kernel two-sample test," *Journal of Machine Learning Research*, vol. 13, no. 25, pp. 723–773, 2012.
- [27] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, "Focal loss for dense object detection," in *IEEE International Conference on Computer Vision*, 2017, pp. 2980–2988.
- [28] 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 Transactions on Intelligent Vehicles*, vol. 8, no. 4, pp. 2650–2660, 2023.
- [29] A. H. Lang, S. Vora, H. Caesar, L. Zhou, J. Yang, and O. Beijbom, "Pointpillars: Fast encoders for object detection from point clouds," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 12 697–12 705.
- [30] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *International Conference on Learning Representations.*, 2017.