

# CoBEVMoE: Heterogeneity-aware Feature Fusion with Dynamic Mixture-of-Experts for Collaborative Perception

Lingzhao Kong<sup>1</sup>, Jiacheng Lin<sup>1</sup>, Siyu Li<sup>2,3</sup>, Kai Luo<sup>2,3</sup>, Zhiyong Li<sup>1,2,3</sup>, and Kailun Yang<sup>2,3</sup>

**Abstract**—Collaborative perception aims to extend sensing coverage and improve perception accuracy by sharing information among multiple agents. However, due to differences in viewpoints and spatial positions, agents often acquire heterogeneous observations. Existing intermediate fusion methods primarily focus on aligning similar features, often overlooking the perceptual diversity among agents. To address this limitation, we propose CoBEVMoE, a novel collaborative perception framework that operates in the Bird’s Eye View (BEV) space and incorporates a Dynamic Mixture-of-Experts (DMoE) architecture. In DMoE, each expert is dynamically generated based on the input features of a specific agent, enabling it to extract distinctive and reliable cues while attending to shared semantics. This design allows the fusion process to explicitly model both feature similarity and heterogeneity across agents. Furthermore, we introduce a Dynamic Expert Metric Loss (DEML) to enhance inter-expert diversity and improve the discriminability of the fused representation. Extensive experiments on the OPV2V and DAIR-V2X-C datasets demonstrate that CoBEVMoE achieves state-of-the-art performance. Specifically, it improves the IoU for Camera-based BEV segmentation by +1.5% on OPV2V and the AP@0.5 for LiDAR-based 3D object detection by +3.0% on DAIR-V2X-C, verifying the effectiveness of expert-based heterogeneous feature modeling in multi-agent collaborative perception. The source code will be made publicly available at <https://github.com/godk0509/CoBEVMoE>.

## I. INTRODUCTION

The perception module serves as the cornerstone of safe autonomous driving operations, providing reliable environmental priors for downstream tasks such as path planning [1]. Due to limitations in sensor coverage and viewpoint constraints inherent in single-agent perception, multi-agent collaborative perception has emerged as a key paradigm for enhancing the environmental awareness of autonomous driving [2]–[6]. Multiple intelligent agents, such as vehicles, form collaborative perception systems by sharing sensory information, enabling safe operation beyond visual range.

Due to the complex kinematic interactions among multiple vehicles, the effective fusion of perceptual information from diverse agents constitutes a core technical challenge

This work was supported in part by the National Natural Science Foundation of China (Grant No. U21A20518, No. 61976086, and No. 62473139), in part by the Hunan Provincial Research and Development Project (Grant No. 2025QK3019), and in part by the State Key Laboratory of Autonomous Intelligent Unmanned Systems (the opening project number ZZKF2025-2-10). (Corresponding authors: Zhiyong Li and Kailun Yang.)

<sup>1</sup>L. Kong, J. Lin, and Z. Li are with the School of Computer Science and Electronic Engineering, Hunan University, Changsha 410082, China.

<sup>2</sup>S. Li, K. Luo, Z. Li, and K. Yang are with the School of Artificial Intelligence and Robotics, Hunan University, Changsha 410012, China (email: zhiyong.li@hnu.edu.cn; kailun.yang@hnu.edu.cn).

<sup>3</sup>S. Li, K. Luo, Z. Li, and K. Yang are also with the National Engineering Research Center of Robot Visual Perception and Control Technology, Hunan University, Changsha 410082, China.

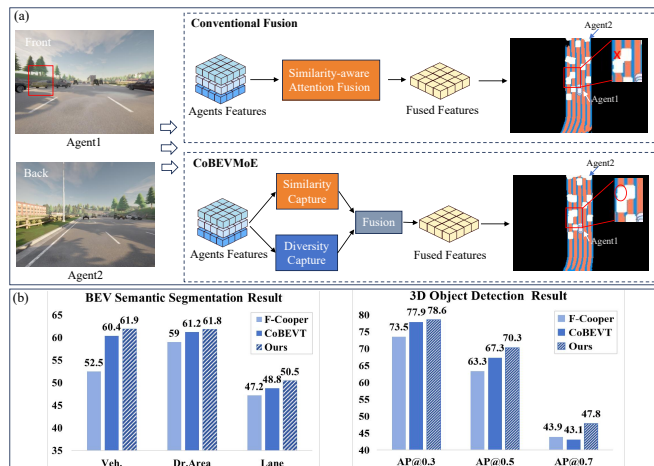


Fig. 1. (a) Conventional attention-based fusion suppresses agent-specific cues, causing missed perception (red box). Our DMoE-based fusion preserves heterogeneous features and correctly segments the target. (b) Quantitative comparison on OPV2V BEV semantic segmentation and DAIR-V2X-C 3D object detection, showing consistent IoU and AP improvements.

in collaborative perception systems. Bird’s-Eye View (BEV) represents 3D environmental information in a 2D planar space [7]. The BEV spaces of different vehicles can be aligned solely based on their relative poses, providing an efficient unified spatial framework for collaborative perception systems. V2VNet [8] introduced spatial transformation to aggregate BEV features from multiple agents. CoBEVT [9] utilized a Swin Transformer-based fusion backbone with alternating window-grid attention to improve the spatial granularity of multi-agent BEV feature fusion. V2X-ViT [10] further incorporates Transformer encoders to model long-range inter-agent dependencies. Although these methods significantly enhance collaborative perception performance, they primarily focus on shared semantic features among agents while overlooking their individualized characteristics.

Agents with significant positional and orientational disparities often observe unique or complementary features that may be occluded or outside the field of view of other agents. These viewpoint-specific cues can be highly informative for improving perception, especially under partial occlusions or sensor failures. As illustrated in Figure 1, when a target is observed by one agent but neglected by another, current similarity-guided fusion methods fail to detect this target. Motivated by this, a successful multi-agent collaboration method should account for both feature commonality and heterogeneity, allowing agents to contribute uniquely based on their observations.

In this work, we propose CoBEVMoE, a novel collabora-

orative perception framework designed to explicitly model both the similarity and diversity of multi-agent observations in the BEV space. The similar features across different agents can be effectively integrated using well-established attention mechanisms. Simultaneously, this paper aims to achieve parallel learning of differentiated features across various agents using independent modules. Interestingly, this concept is closely aligned with the Mixture of Experts (MoE) paradigm [11] where each expert independently acquires knowledge within its specialized domain, while a gating network dynamically controls the contribution of each expert. However, in traditional MoE systems, experts are initialized randomly. This paper aims to specialize each expert in learning the personalized information of individual vehicles. To achieve this, we propose a Dynamic Mixture-of-Experts (DMoE) module. It leverages the MoE principle to assign each agent a dynamically generated expert. These experts are conditioned on the local features of each agent and capture distinct, agent-specific patterns. All expert outputs are then adaptively fused through a gated aggregation mechanism that accounts for reliability and complementarity. This formulation allows our framework to go beyond uniform fusion, enabling it to retain unique perceptual cues from different viewpoints while still integrating shared semantics. To further encourage expert diversity and prevent representational collapse, we introduce a Dynamic Expert Metric Loss (DEML), which promotes proximity between each expert and the fused feature while pushing different experts apart.

We verify the effectiveness of our framework on two challenging collaborative perception benchmarks. CoBEV-MoE achieves notable gains across both camera and LiDAR modalities, with up to +1.5% **IoU** improvement on BEV segmentation and +3.0% **AP@0.5** gain on 3D detection, outperforming strong baselines such as CoBEVT [9] and AttFuse [12]. These results demonstrate the advantages of our expert-based heterogeneous fusion design.

Our main contributions are summarized as follows:

- We rethink the precise fusion of multi-source information in collaborative perception from the perspective of agent-specific diversity and propose a novel dynamic mixture-of-experts-based framework, CoBEVMoE, to improve collaborative perception accuracy.
- We design an expert kernel dynamic initialization Dynamic Mixture-of-Experts module, which guides the learning direction of experts using the personalized features of agents. Simultaneously, a Dynamic Expert Metric Loss is designed to ensure the differentiation among experts, thereby guaranteeing the effective learning of personalized features.
- We verify the proposed CoBEVMoE method on two large-scale collaborative perception datasets. CoBEV-MoE consistently outperforms strong baselines in 3D object detection and semantic segmentation.

## II. RELATED WORK

### A. Collaborative Perception

Most autonomous driving perception systems are built upon single-agent sensor data. While multi-sensor fusion improves perception accuracy, it still suffers from blind spots and a limited sensing range. Collaborative perception addresses these issues by enabling information sharing among multiple agents, leading to more comprehensive environmental awareness [13]–[16]. According to the fusion stage, collaborative perception methods are generally categorized into early, intermediate, and late fusion. Early fusion aggregates raw sensor data [17], but incurs prohibitive communication costs. Late fusion merges individual perception outputs [18], which is often sensitive to noise and error accumulation. Intermediate fusion performs collaboration at the feature level, offering a favorable trade-off between performance and communication efficiency, and has therefore become the dominant paradigm in recent studies [8], [19]–[21]. V2VNet [8] and DiscoNet [22] introduced intermediate fusion frameworks where BEV features were fused through a graph neural network. Recently, HEAL [20] and STAMP [21] addressed scalability and heterogeneity by pruning or learning dynamic communication graphs.

Subsequent research explored various feature fusion mechanisms to better handle spatial misalignment and inter-agent redundancy. AttFusion [12] and V2X-ViT [10] incorporated attention-based modules to weigh features from different agents adaptively. CoBEVT [9] proposed a Swin Transformer-based fusion backbone with window-grid alternating attention, enhancing local-global spatial interactions. Additionally, other works focused on robustness issues such as communication delays [23]–[27], positioning errors [28], [29], or adverse weather conditions [30]–[32] that may occur during the collaborative perception process. In practice, agents may observe the same scene from distinct perspectives, yielding non-overlapping but informative cues. Existing fusion strategies [8], [9] often suppress these subtle differences through averaging or global attention, which may degrade the final representation in complex, occluded scenes. To address this, our method introduces dynamically generated expert kernels that are initialized from unique features of each agent, allowing the model to extract and preserve unique perceptual cues during the fusion process.

### B. Mixture-of-Experts Models

Mixture-of-Experts (MoE) [11] is a conditional computation architecture that assigns input samples to a sparse subset of specialized expert networks via a gating mechanism. This approach has proven effective in scaling deep networks while maintaining diversity and efficiency. In NLP, models such as GShard [33] and Switch Transformer [34] leverage MoE to distribute computation and capture linguistic variation. In vision, MoE has been applied to tasks like image classification [35] and video understanding [36], showing strong performance gains on heterogeneous data.

Recent studies have extended MoE to multi-modal and multi-view perception. MetaBEV [37] incorporated

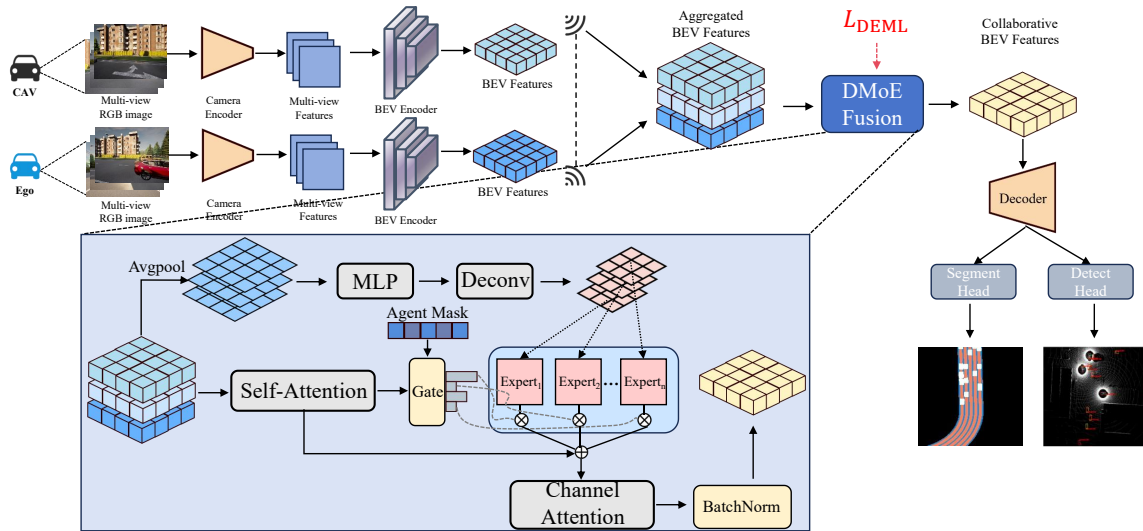


Fig. 2. Illustration of the proposed **CoBEVMoE** framework. Each agent first extracts BEV features from its raw sensory inputs, which are then spatially aligned and transmitted to a central aggregator. A Dynamic Mixture-of-Experts (DMoE) module generates agent-specific expert kernels and performs adaptive feature aggregation through a learned gating mechanism. The resulting fused representation is subsequently decoded for downstream perception tasks, including semantic segmentation and 3D object detection. To encourage inter-expert diversity while maintaining consistency with the fused feature, we introduce a Dynamic Expert Metric Loss (DEML), depicted in  $\mathcal{L}_{\text{DEML}}$  in the figure. This loss promotes diverse yet coherent expert representations, ultimately enhancing the quality of the final fusion.

an attention-based mixture of experts mechanism to learn modality-specific weights, enabling the adaptive fusion of reliable multi-modal features. MapExpert [38] used MoE to decode various non-cubic map elements, improving map accuracy and efficiency. MoE-Fusion [39] applied MoE to infrared and visible image fusion, achieving superior results.

These approaches directly apply the classical MoE paradigm to perception tasks, where each expert is learned implicitly within the model without a predefined or specialized learning objective. In contrast, this paper aims to anchor the learning objectives of the experts. To this end, we build upon the MoE foundation and propose a dynamic expert kernel, which leverages features from different agents to guide and specialize the learning direction of each expert.

### III. METHODOLOGY

#### A. Problem Formulation

Let  $\mathcal{A} = \{1, 2, \dots, N\}$  denote a set of  $N$  collaborating agents. Each agent  $i \in \mathcal{A}$  is equipped with multiple cameras (or other sensors) to observe its surroundings and generate local multi-view image features. These features are transformed into a bird's-eye-view (BEV) representation using a shared image-to-BEV encoder, yielding a local BEV feature map  $\mathbf{x}_i \in \mathbb{R}^{C \times H \times W}$ , where  $C$  is the feature dimension and  $H, W$  are spatial dimensions.

Each agent then transmits its BEV features to a central fusion module (or to neighboring agents), which aligns the features via spatial transformation using ego-pose information. The aligned features from all agents form a set:

$$\mathcal{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\}, \quad \mathbf{x}_i \in \mathbb{R}^{C \times H \times W}. \quad (1)$$

The goal of the collaborative perception system is to fuse the set  $\mathcal{X}$  into a unified representation  $\mathbf{F}_{\text{fused}} \in \mathbb{R}^{C \times H \times W}$  that accurately describes the shared environment. This fused

representation is then used for downstream tasks such as 3D object detection or BEV semantic segmentation.

The key challenge lies in effectively leveraging both the *redundancy* (*i.e.*, overlapping observations of the same object) and the *heterogeneity* (*i.e.*, complementary views and occlusion differences) among agents. To address this, we adopt an *intermediate fusion* paradigm, which provides a balanced compromise between early fusion, where raw data are directly shared, and late fusion, where only decision-level information is exchanged, by transmitting intermediate features among agents.

#### B. Overview

We revisit the collaborative perception problem from the perspective of agent-specific diversity and propose a new framework, CoBEVMoE. The framework consists of an image encoder, a BEV encoder, an intermediate collaboration module, a decoder, and a task head. Specifically, the image encoder (*i.e.*, ResNet34) extracts multi-view features from raw images, and the BEV encoder (*i.e.*, SinBEV Transformer) transforms them into BEV features. The intermediate collaboration module integrates our proposed Dynamic Mixture of Experts (DMoE) with Dynamic Expert Metric Loss (DEML) to fuse BEV features from all agents into a unified representation that improves perception. The decoder then processes the fused features, while the task head produces either BEV semantic segmentation or 3D object detection results. Our focus is on leveraging the diversity of agent perception features during collaboration to improve accuracy. To this end, we design the DMoE module with a DEML optimisation constraint.

#### C. Dynamic Mixture-of-Experts Fusion

Although prior work has demonstrated that MoE architecture can enhance the integration of multi-source features,

conventional MoE architectures are ill-suited for multi-agent collaborative perception due to the divergent observational perspectives of each agent. To tackle the challenge of fusing heterogeneous features in multi-agent collaboration, we propose a novel intermediate collaboration module dubbed DMoE. It is designed to explicitly model both inter-agent relationships and agent-specific feature variations, thereby further advancing performance. Motivated by these limitations, we design a Dynamic Mixture of Experts (DMoE) module tailored for multi-agent collaboration, which adaptively learns and emphasizes the key contributions from each agent and explicitly models both the shared and agent-specific features in collaborative perception.

*a) Expert Kernel Generation:* We first describe the construction of each expert in the proposed DMoE module. Unlike conventional designs where parameters are solely optimized via gradient descent, the parameters of each expert are dynamically generated conditioned on the input BEV features. Specifically, following the principle of dynamic convolution, each expert is instantiated as a dynamic convolutional layer whose kernels are adaptively generated from the input. After applying spatial transformation and feature aggregation across agents, we obtain the aligned BEV representations from the  $N$  collaborating agents, denoted as:

$$\mathbf{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\} \in \mathbb{R}^{B \times N \times C \times H \times W}, \quad (2)$$

where  $B$  is the batch size,  $N$  is the number of agents, and  $C, H, W$  are the channel and spatial dimensions.

For each agent  $k$ , the BEV feature  $\mathbf{x}_k$  is first subjected to spatial global average pooling:

$$\mathbf{z}_k = \text{AvgPool}(\mathbf{x}_k) \in \mathbb{R}^{B \times C'}, \quad (3)$$

where typically  $C' < C \cdot H \cdot W$ .

Then, the pooled vector  $\mathbf{z}_k$  is processed by an MLP:

$$\mathbf{h}_k = \text{MLP}(\mathbf{z}_k) \in \mathbb{R}^{B \times d}, \quad (4)$$

and then decoded via a deconvolution block to generate a convolutional kernel:

$$\mathbf{W}_k = \text{Deconv}(\mathbf{h}_k) \in \mathbb{R}^{B \times C \times C \times 3 \times 3}. \quad (5)$$

*b) Gating Mechanism:* After the expert kernels are generated, their contributions to the output are determined through a gating unit. The gating unit is designed to generate gating functions based on the input features and serves as a fundamental component of the proposed DMoE. It is responsible for coordinating the participation of each expert and combining their outputs. Specifically, we apply global average pooling along the channel dimension of the input collaborative features  $\mathbf{F}_c \in \mathbb{R}^{B \times C \times H \times W}$ , and then use a softmax function to generate the weights  $\alpha$  for each expert:

$$\alpha = \text{softmax}(\text{Gate}(\text{AvgPool}(\mathbf{F}_c))) \in \mathbb{R}^{B \times N}. \quad (6)$$

To handle invalid agents, we apply a binary mask  $\mathbf{M} \in \{0, 1\}^{B \times L}$  to the gating weights:

$$\tilde{\alpha} = \text{softmax}(\alpha \odot \mathbf{M} + (1 - \mathbf{M}) \cdot (-\infty)), \quad (7)$$

where  $\odot$  denotes element-wise multiplication.

*c) Domain Expert Construction:* In the proposed DMoE module, we construct an expert for each agent, enabling it to extract key features from its unique perspective and organically fuse them for higher-quality collaborative perception. Each expert is designed as a dynamic convolutional layer. When it receives a feature input  $\mathbf{F}_c \in \mathbb{R}^{B \times C \times H \times W}$ , it utilizes the unique convolution kernel generated by the Expert Kernel Generation module to produce the corresponding feature  $\mathbf{E}_k$ :

$$\mathbf{E}_k = \mathbf{W}_k * \mathbf{F}_c. \quad (8)$$

Therefore, each expert acquires unique knowledge and contributes distinctly to the feature fusion.

*d) DMoE Fusion Module:* The DMoE fusion module is composed of an expert kernel generation unit, a self-attention fusion layer, a gating mechanism, and a set of experts. Specifically, the aligned BEV features  $\mathbf{X}$  are fed into the expert kernel generation module, and the generated expert kernels are assigned to each expert. Then, the  $\mathbf{X}$  is directly input into the self-attention fusion layer to produce the preliminary fused BEV feature:

$$\mathbf{F}_c = \text{SelfAttn}(\mathbf{X}) \in \mathbb{R}^{B \times C \times H \times W}. \quad (9)$$

Here, we use the original SwapFusion module from CoBEVT [9] as our baseline. Each expert applies its generated convolution kernel  $\mathbf{W}_k$  to the preliminary fused feature  $\mathbf{F}_c$  to produce the expert output  $\mathbf{E}_k$ . All expert output features are stacked as:

$$\mathbf{E} = [\mathbf{E}_1, \mathbf{E}_2, \dots, \mathbf{E}_N] \in \mathbb{R}^{B \times N \times C \times H \times W}. \quad (10)$$

Finally, the expert weights obtained from the gating unit are combined with a residual connection to perform a weighted summation over the experts' outputs, enhancing the stability and expressiveness of the MoE output feature. This residual structure preserves both the shared semantic information and the expert-specific enhancements in the fused representation. The overall process proceeds as follows:

$$\mathbf{F}_{\text{MoE}} = \sum_{k=1}^L \tilde{\alpha}_k \cdot \mathbf{E}_k. \quad (11)$$

$$\mathbf{F}_{\text{res}} = \mathbf{F}_c + \mathbf{F}_{\text{MoE}}. \quad (12)$$

Unlike conventional fusion mechanisms [37], [39], the proposed DMoE not only captures spatial-temporal relationships across agents but also preserves the diversity of individual perspectives, enabling a more comprehensive and discriminative feature representation.

#### D. Dynamic Expert Metric Loss

Although the DMoE module effectively captures heterogeneous agent features, the extracted features may still lack sufficient regularization, potentially resulting in suboptimal performance. To address this, we introduce a Dynamic Expert Metric Loss (DEML) that serves as an internal regularizer for the DMoE module, aiming to prevent expert collapse by promoting functional diversity among agent-conditioned experts while maintaining consistency with the fused feature.

It differs from contrastive learning in that DEML operates on intra-sample expert outputs rather than learning instance-level representations.

a) *Expert Triplet Construction*: Contrastive learning-based metric learning has shown strong performance in tasks that require fine-grained feature representations, as it directly optimizes the embedding space to preserve clear distinctions among embeddings. Motivated by this, we construct a unique triplet for each expert.

The initial fused feature  $\mathbf{F}_c$  output by the self-attention fusion layer contains relatively complete semantic information. We take  $\mathbf{F}_c$  as the anchor to ensure that the experts' outputs do not deviate from the overall task objective. For each expert, we select its own output  $\mathbf{E}_k$  as the positive sample, and choose the output of the most similar expert (based on mean squared error) from the remaining experts as the hardest negative sample  $\mathbf{E}_j$ . This ensures that the distance between positive and negative samples is sufficiently large, encouraging diversity among experts while learning the specific characteristics of each agent's features:

$$T_k = (\mathbf{F}_c, \mathbf{E}_k, \mathbf{E}_j), \quad (13)$$

$$\mathbf{E}_j = \arg \min_{i \neq k} \text{MSE}(\mathbf{E}_i, \mathbf{E}_k). \quad (14)$$

b) *Expert Metric Loss*: Based on the constructed triplets, we define an expert-specific loss to guide the optimization process. Specifically, the  $L_2$  distance is employed to measure the distances between the anchor, positive, and negative samples. The objective is to minimize the distance between the anchor and positive sample while ensuring that the distance between the positive and negative samples exceeds a predefined margin. This encourages the positive sample to remain close to the anchor and the negative sample to stay farther away. The triplet loss for expert  $k$  is then defined as follows:

$$d_{\text{pos}}^k = \frac{1}{D} \|\mathbf{E}_k - \mathbf{F}_c\|_2^2, \quad (15)$$

$$d_{\text{neg}}^k = \min_{j \neq k} \frac{1}{D} \|\mathbf{E}_k - \mathbf{E}_j\|_2^2. \quad (16)$$

$$\mathcal{L}_{\text{triplet}}^k = \max(0, d_{\text{pos}}^k - d_{\text{neg}}^k + m), \quad (17)$$

where  $m > 0$  is a margin hyperparameter. And the overall expert regularization loss is formulated as:

$$\mathcal{L}_{\text{DEML}} = \frac{1}{N_{\text{valid}}} \sum_k (d_{\text{pos}}^k + \beta \cdot \mathcal{L}_{\text{triplet}}^k), \quad (18)$$

where  $N_{\text{valid}}$  represents the number of valid experts, and  $\beta$  is a balancing coefficient that weights the two terms.

Finally, we integrate the expert-specific DEML, scaled by a balancing hyperparameter  $\lambda$ , with the original loss objective to obtain the total training loss:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{task}} + \lambda \cdot \mathcal{L}_{\text{DEML}}, \quad (19)$$

where  $\mathcal{L}_{\text{task}}$  denotes the task-specific loss corresponding to different tasks, such as segmentation or detection.

## IV. EXPERIMENTS

### A. Datasets

To evaluate the effectiveness of the proposed method, we conducted extensive experiments on both the simulated dataset OPV2V [12] and the real-world dataset DAIR-V2X-C [18] for comprehensive verification.

**OPV2V** is a large-scale vehicle-to-vehicle collaborative perception dataset collected in CARLA and the cooperative driving automation tool OpenCDA. It comprises 12,000 LiDAR point cloud frames and RGB images with 230,000 annotations of 3D bounding boxes, with an image resolution of  $600 \times 800$ , providing a diverse set of scenarios for evaluating perception algorithms.

**DAIR-V2X-C** is a widely-used real-world vehicle-to-infrastructure collaborative perception dataset, consisting of 38,845 frames of image data and point cloud data, with an image resolution of  $1080 \times 1920$ , covering diverse urban driving scenarios for evaluation.

### B. Evaluation Metrics

To ensure fair and comprehensive evaluation, we adopt task-specific metrics for both BEV segmentation and 3D object detection.

For the **BEV semantic segmentation** task on OPV2V, we report the Intersection over Union (IoU) for each semantic category, including *vehicle*, *drivable area*, and *lane*. The IoU metric is defined as the ratio of the intersection to the union between predicted and ground truth masks, which effectively captures the overlap quality of semantic regions. Higher IoU indicates better segmentation accuracy and robustness against occlusion and viewpoint changes.

For the **3D object detection** task on DAIR-V2X-C, we employ the Average Precision (AP) at different Intersection-over-Union thresholds, specifically AP@0.3 and AP@0.5, following standard practice in V2X perception. The AP metric measures the area under the precision-recall curve, reflecting the trade-off between detection precision and recall. The use of multiple IoU thresholds allows us to evaluate both coarse and fine-grained localization accuracy of detected objects in collaborative settings.

To quantify inter-expert diversity, we report *Pairwise Cosine Diversity (PCD)*. Given expert outputs  $\{E_k\}_{k=1}^{N_{\text{valid}}}$ , we obtain compact vectors by global average pooling:  $\mathbf{v}_k = \text{GAP}(E_k) \in \mathbb{R}^{B \times C}$ . We compute the average pairwise cosine similarity

$$\text{ES} = \frac{2}{N_{\text{valid}}(N_{\text{valid}} - 1)} \sum_{i < j} \mathbb{E}_b \left[ \frac{\mathbf{v}_i^{(b)} \cdot \mathbf{v}_j^{(b)}}{\|\mathbf{v}_i^{(b)}\|_2 \|\mathbf{v}_j^{(b)}\|_2 + \epsilon} \right], \quad (20)$$

and define  $\text{PCD} = 1 - \text{ES}$ , where higher PCD indicates stronger diversity among experts.

### C. Experimental Setup

**Implementation Details.** Following CoBEVT [9], we assume all the AVs have a 70m communication range, and all

TABLE I

IOU OF BEV SEGMENTATION ON OPV2V DATASET CAMERA-TRACK.

Methods	Veh.	Dr.Area	Lane
F-Cooper [19]	52.5	59.0	47.2
AttFuse [12]	51.9	60.5	47.9
DiscoNet [22]	52.9	57.6	43.4
V2VNet [8]	53.5	58.2	43.8
CoBEVT [9]	60.4	61.2	48.8
Ours (CoBEVMoE)	<b>61.9 (+1.5)</b>	<b>61.8 (+0.6)</b>	<b>50.5 (+1.7)</b>

the vehicles outside this broadcasting radius of the ego vehicle will not have any collaboration. For the OPV2V camera-track BEV semantic segmentation, we choose ResNet34 as the image feature extractor. The transmitted BEV intermediate representation has a resolution of  $32 \times 32 \times 128$ . For the DAIR-V2X LiDAR-track 3D object detection, we select PointPillar [40] as the point cloud feature extractor and set the voxel resolution as  $(0.4, 0.4, 4)$  on the x, y, and z-axis. The architecture settings are the same as [40]. The MLP used in the dynamic kernel generator consists of two fully connected layers with a hidden dimension of 128 and ReLU activation. The kernel generator employs a transposed convolution with kernel size  $3 \times 3$ , stride 1, and no padding, which maps a 128-dimensional latent embedding to dynamic convolutional kernels of size  $128 \times 128 \times 3 \times 3$ . For the Dynamic Expert Metric Loss (DEML), we set the margin  $m = 0.5$  and the weighting factor  $\beta = 1.0$  for all experiments.

We train the whole model end-to-end with the Adam optimizer and the cosine annealing learning rate scheduler. We train all models on a single NVIDIA RTX A6000 GPU with 150 epochs and a batch size of 2.

#### D. Quantitative and Qualitative Evaluation

**OPV2V camera-track results.** Table I reports the performance of our method and several strong baselines on the camera-based BEV segmentation task. Compared to existing methods such as AttFuse [12] and CoBEVT [9], our proposed CoBEVMoE consistently achieves higher IoU across all categories. Compared to CoBEVT, our method improves the IoU for vehicle, drivable area, and lane segmentation by +1.5%, +0.6%, and +1.7%, respectively. This performance gain can be attributed to the joint effect of our DMoE module and the DEML. The DMoE module enables each agent to dynamically capture both redundant and complementary information, while the DEML encourages feature diversity across agents, mitigating information collapse. As shown in Fig. 3, different experts focus on diverse spatial patterns, complementing each other and contributing to a more comprehensive fusion representation.

**DAIR-V2X LiDAR-track results.** We further evaluate the 3D object detection performance of CoBEVMoE on the real-world DAIR-V2X-C dataset, which provides LiDAR point clouds under an infrastructure-to-vehicle collaboration setting. As shown in Table II, our method outperforms CoBEVT by +0.7% on AP@0.3 and +3.0% on AP@0.5. These results confirm that our framework not only excels in camera-based BEV segmentation but also generalizes effectively to geometric perception tasks in realistic V2X environments. The

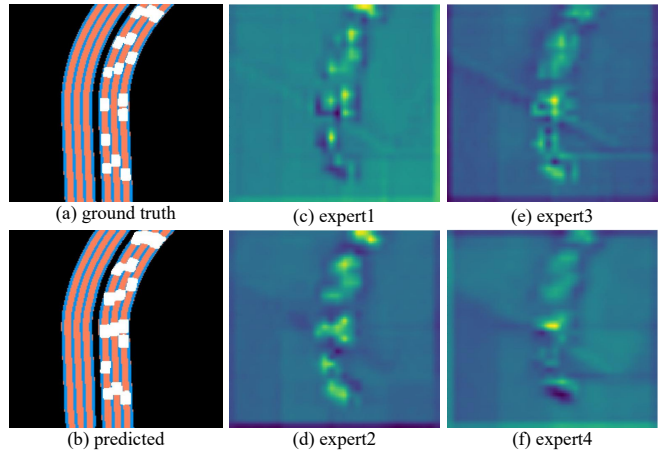


Fig. 3. Qualitative comparison among different expert feature maps. (a) shows the ground truth BEV semantic map. (b) shows the predicted BEV semantic map, while (c-f) depict the activation maps of four different experts in our DMoE module. The four experts focus on different spatial patterns, indicating the diversity in feature extraction. This complementary behavior contributes to a more comprehensive and robust feature fusion in the downstream task.

TABLE II

AVERAGE PRECISION (AP) OF 3D DETECTION ON THE DAIR-V2X LiDAR-TRACK.

Methods	AP@0.3	AP@0.5
F-Cooper [19]	72.3	62.0
AttFuse [12]	73.8	67.3
DiscoNet [22]	74.6	68.5
V2XViT [10]	77.6	52.1
CoBEVT [9]	77.9	67.3
Ours (CoBEVMoE)	<b>78.6 (+0.7)</b>	<b>70.3 (+3.0)</b>

improvements demonstrate that the proposed DMoE module can handle sensor noise, occlusion, and spatial misalignment, while the expert-based fusion strategy enhances adaptability to diverse real-world deployment scenarios.

Qualitative comparisons in Fig. 4 further demonstrate that CoBEVMoE generates more complete and accurate BEV segmentations than the baseline, particularly for distant or occluded targets.

Moreover, Fig. 5 visualizes BEV feature maps before and after DMoE fusion. The enhanced features exhibit clearer and more focused activations, indicating that our expert-based fusion improves feature quality and semantic alignment.

#### E. Ablation Studies

**Ablation studies on hyperparameters.** Table III reports the ablation study on the loss weight  $\lambda$  of the DEML, conducted on the OPV2V camera-based segmentation task. We vary  $\lambda$  from 0 to 1.0 with an interval of 0.2. As shown in the table, the performance remains relatively stable when  $\lambda$  ranges from 0.0 to 0.6, with the best results achieved at  $\lambda = 0.4$ , where vehicle, drivable area, and lane segmentation reach their highest IoUs. This indicates that a moderate loss weight provides sufficient regularization for the DMoE module, encouraging feature diversity without overwhelming the primary segmentation objective. In contrast, excessively small weights (e.g.,  $\lambda = 0.2$ ) fail to impose effective constraints, while overly large weights (e.g.,  $\lambda \geq 0.8$ ) dominate




Fig. 4. Qualitative results of collaborative BEV segmentation. From left to right: (a) ego vehicle’s front camera image, (b–c) front camera images from two collaborating agents (cav1 and cav2), (d) ground truth BEV segmentation, (e) result of the baseline fusion method, and (f) result of our proposed **CoBEVMoE**. Compared to the baseline, our method produces more complete and precise semantic segmentation, especially for distant or partially occluded targets, demonstrating its ability to effectively integrate complementary views and retain agent-specific information.

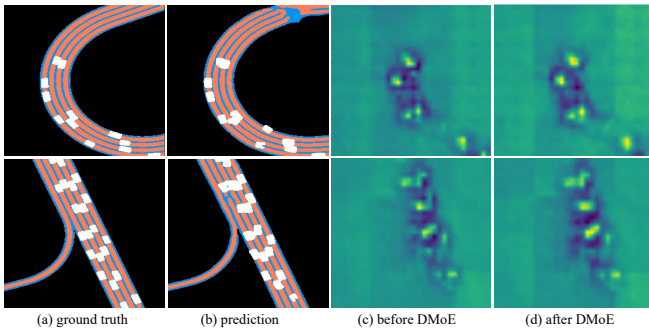


Fig. 5. Visualization of BEV feature maps before and after the DMoE fusion module. (a) shows the ground truth BEV semantic map. (b) shows the predicted BEV semantic map. (c) presents the feature activation of a representative channel before DMoE fusion, while (d) displays the output after fusion. The DMoE-enhanced feature map demonstrates clearer and more focused activations, indicating that our expert-based fusion improves feature quality and semantic alignment.

the optimization, leading to degraded semantic segmentation performance. These results highlight the importance of balancing the expert regularization and the main task loss for stable and effective training.

**Component analysis.** Table IV presents an ablation study on the key components of CoBEVMoE. From the results, we can see that each proposed component contributes incremental improvements. Adding a Vanilla MoE improves IoU by +0.4 on *vehicles*, while DMoE further boosts it by an additional +0.5. Incorporating the DEML yields a final gain of +0.6, leading to a total improvement of +1.5 in IoU over the CoBEVT baseline. We further quantify expert diversity using PCD. As shown in Table IV, enabling DEML significantly increases PCD from 0.10 to 1.53, which consistently correlates with the performance gains, indicating that DEML encourages agent-relevant expert diversity rather than arbitrary dispersion. Here, the *Vanilla MoE* baseline employs **five static experts**, which are **randomly initialized and jointly optimized during training**, while **sharing the**

TABLE III  
ABLATION RESULTS OF HYPERPARAMETERS  $\lambda$  ON THE OPV2V  
CAMERA-TRACK.

$\lambda$	Veh.	Dr.Area	Lane
0	61.3	61.6	50.1
0.2	61.4	61.8	50.2
0.4	<b>61.9</b>	<b>61.8</b>	<b>50.5</b>
0.6	61.4	61.7	50.2
0.8	60.7	61.3	49.2
1.0	59.8	61.0	48.7

**same gating network as DMoE.** We make the following observations:

- **Vanilla MoE vs. CoBEVT:** Simply using a standard mixture-of-experts module with static experts yields only marginal improvements. This indicates that without expert-specific design and regulation, the expert system fails to effectively disentangle useful agent-specific cues.
- **DMoE vs. Vanilla MoE:** Introducing dynamic kernel generation for each expert significantly boosts performance, as the learned kernels are conditioned on each agent’s local context. This dynamic adaptability enhances the model’s capacity to capture heterogeneity across agents.
- **DMoE + DEML:** The proposed Dynamic Expert Metric Loss further regularizes the expert outputs by explicitly encouraging diversity among experts while maintaining alignment with the fused feature. This alleviates feature redundancy and improves discriminability, leading to consistent gains across all segmentation categories.

Together, these results demonstrate that our proposed DMoE module and DEML are complementary and crucial for high-quality multi-agent feature fusion.

TABLE IV

ABLATION STUDY ON THE OPV2V DATASET. “DMoE” DENOTES OUR DYNAMIC MIXTURE-OF-EXPERTS MODULE. “DEML” INDICATES THE USE OF THE PROPOSED DYNAMIC EXPERT METRIC LOSS.

Method	Veh.	Dr.Area	Lane	PCD
Baseline (CoBEVT)	60.4	61.2	48.8	–
+ Vanilla MoE	60.8	61.2	49.1	0.10
+ DMoE (Ours)	61.3	61.6	50.1	1.07
+ DMoE + DEML (Ours)	<b>61.9</b>	<b>61.8</b>	<b>50.5</b>	1.53

## V. CONCLUSION

In this paper, we propose CoBEVMoE, a collaborative perception framework that leverages a dynamic mixture-of-experts module and a dynamic expert metric loss to capture both shared and agent-specific features during intermediate fusion. Experiments on OPV2V and DAIR-V2X demonstrate that our approach consistently outperforms strong baselines in semantic segmentation and 3D object detection, confirming its effectiveness across modalities and scenarios. In future work, we aim to extend CoBEVMoE to bandwidth-constrained settings and explore its integration with temporal modeling for multi-agent prediction and planning.

## REFERENCES

- [1] G. Velasco-Hernandez, D. J. Yeong, J. Barry, and J. Walsh, “Autonomous driving architectures, perception and data fusion: A review,” in *Proc. ICCP*, 2020, pp. 315–321.
- [2] X. Gao *et al.*, “A survey of collaborative perception in intelligent vehicles at intersections,” *IEEE Transactions on Intelligent Vehicles*, 2024.
- [3] S. Liu *et al.*, “Towards vehicle-to-everything autonomous driving: A survey on collaborative perception,” *arXiv preprint arXiv:2308.16714*, 2023.
- [4] Y. Li, Q. Fang, J. Bai, S. Chen, F. Juefei-Xu, and C. Feng, “Among us: Adversarially robust collaborative perception by consensus,” in *Proc. ICCV*, 2023, pp. 186–195.
- [5] B. Wang, L. Zhang, Z. Wang, Y. Zhao, and T. Zhou, “CORE: Cooperative reconstruction for multi-agent perception,” in *Proc. ICCV*, 2023, pp. 8676–8686.
- [6] Z. Chen, Y. Shi, and J. Jia, “TransIFF: An instance-level feature fusion framework for vehicle-infrastructure cooperative 3D detection with transformers,” in *Proc. ICCV*, 2023, pp. 18 159–18 168.
- [7] B. Zhou and P. Krähenbühl, “Cross-view transformers for real-time map-view semantic segmentation,” in *Proc. CVPR*, 2022, pp. 13 750–13 759.
- [8] 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 *Proc. ECCV*, 2020, pp. 605–621.
- [9] 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 *Proc. CoRL*, 2022, pp. 989–1000.
- [10] 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 *Proc. ECCV*, 2022, pp. 107–124.
- [11] N. Shazeer *et al.*, “Outrageously large neural networks: The sparsely-gated mixture-of-experts layer,” in *Proc. ICLR*, 2017.
- [12] 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 *Proc. ICRA*, 2022, pp. 2583–2589.
- [13] Y. Han, H. Zhang, H. Li, Y. Jin, C. Lang, and Y. Li, “Collaborative perception in autonomous driving: Methods, datasets, and challenges,” *IEEE Intelligent Transportation Systems Magazine*, vol. 15, no. 6, pp. 131–151, 2023.
- [14] Y. Zhou, J. Xiao, Y. Zhou, and G. Loianno, “Multi-robot collaborative perception with graph neural networks,” *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 2289–2296, 2022.
- [15] Y. Zhao, Z. Xiang, S. Yin, X. Pang, Y. Wang, and S. Chen, “MADE: Malicious agent detection for robust multi-agent collaborative perception,” in *Proc. IROS*, 2024, pp. 13 817–13 823.
- [16] K. Suresh, A. Rauniar, M. Corah, and S. Scherer, “Greedy perspectives: Multi-drone view planning for collaborative perception in cluttered environments,” in *Proc. IROS*, 2024, pp. 10 990–10 997.
- [17] Q. Chen, S. Tang, Q. Yang, and S. Fu, “Cooper: Cooperative perception for connected autonomous vehicles based on 3D point clouds,” in *Proc. ICDCS*, 2019, pp. 514–524.
- [18] H. Yu *et al.*, “DAIR-V2X: A large-scale dataset for vehicle-infrastructure cooperative 3D object detection,” in *Proc. CVPR*, 2022, pp. 21 329–21 338.
- [19] 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 *Proc. SEC*, 2019, pp. 88–100.
- [20] Y. Lu, Y. Hu, Y. Zhong, D. Wang, S. Chen, and Y. Wang, “An extensible framework for open heterogeneous collaborative perception,” in *Proc. ICLR*, 2024.
- [21] X. Gao, R. Xu, J. Li, Z. Wang, Z. Fan, and Z. Tu, “STAMP: Scalable task- and model-agnostic collaborative perception,” in *Proc. ICLR*, 2025.
- [22] Y. Li, S. Ren, P. Wu, S. Chen, C. Feng, and W. Zhang, “Learning distilled collaboration graph for multi-agent perception,” in *Proc. NeurIPS*, 2021, pp. 29 541–29 552.
- [23] Y. Lu *et al.*, “Robust collaborative 3D object detection in presence of pose errors,” in *Proc. ICRA*, 2023, pp. 4812–4818.
- [24] S. Wei *et al.*, “Asynchrony-robust collaborative perception via bird’s eye view flow,” in *Proc. NeurIPS*, 2023, pp. 28 462–28 477.
- [25] Z. Lei, S. Ren, Y. Hu, W. Zhang, and S. Chen, “Latency-aware collaborative perception,” in *Proc. ECCV*, 2022, pp. 316–332.
- [26] J. Shi, J. Zhao, L. Zhuo, X. Wang, X. Zhan, and H. Liu, “V2V cooperative perception with adaptive communication loss for autonomous driving,” *IEEE Transactions on Intelligent Transportation Systems*, 2025.
- [27] J. Zhang, K. Yang, Y. Wang, H. Wang, P. Sun, and L. Song, “ER-MVP: Communication-efficient and collaboration-robust multi-vehicle perception in challenging environments,” in *Proc. CVPR*, 2024, pp. 12 575–12 584.
- [28] Y. Yuan, H. Cheng, and M. Sester, “Keypoints-based deep feature fusion for cooperative vehicle detection of autonomous driving,” *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 3054–3061, 2022.
- [29] N. Vadivelu, M. Ren, J. Tu, J. Wang, and R. Urtasun, “Learning to communicate and correct pose errors,” in *Proc. CoRL*, 2021, pp. 1195–1210.
- [30] B. Li *et al.*, “V2X-DGW: Domain generalization for multi-agent perception under adverse weather conditions,” in *Proc. ICRA*, 2025.
- [31] J. Zhang *et al.*, “DSRC: Learning density-insensitive and semantic-aware collaborative representation against corruptions,” in *Proc. AAAI*, 2025, pp. 9942–9950.
- [32] Y. Ma *et al.*, “MACP: Efficient model adaptation for cooperative perception,” in *Proc. WACV*, 2024, pp. 3361–3370.
- [33] D. Lepikhin *et al.*, “GShard: Scaling giant models with conditional computation and automatic sharding,” in *Proc. ICLR*, 2021.
- [34] W. Fedus, B. Zoph, and N. Shazeer, “Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity,” *Journal of Machine Learning Research*, vol. 23, no. 120, pp. 1–39, 2022.
- [35] C. Riquelme *et al.*, “Scaling vision with sparse mixture of experts,” in *Proc. NeurIPS*, 2021, pp. 8583–8595.
- [36] Z. Yang, B. Wu, C. Feichtenhofer, H. Fan, X. Dai, and S. Xie, “MViTv2-MoE: Mixture-of-experts meets vision transformers,” *arXiv preprint arXiv:2203.16774*, 2022.
- [37] C. Ge *et al.*, “MetaBEV: Solving sensor failures for 3D detection and map segmentation,” in *Proc. ICCV*, 2023, pp. 8687–8697.
- [38] D. Zhang *et al.*, “MapExpert: Online HD map construction with simple and efficient sparse map element expert,” in *Proc. AAAI*, 2025, pp. 14 745–14 753.
- [39] B. Cao, Y. Sun, P. Zhu, and Q. Hu, “Multi-modal gated mixture of local-to-global experts for dynamic image fusion,” in *Proc. ICCV*, 2023, pp. 23 498–23 507.
- [40] A. H. Lang, S. Vora, H. Caesar, L. Zhou, J. Yang, and O. Beijbom, “PointPillars: Fast encoders for object detection from point clouds,” in *Proc. CVPR*, 2019, pp. 12 697–12 705.