# IFTR: An Instance-Level Fusion Transformer for Visual Collaborative Perception

Shaohong Wang<sup>1</sup>, Lu Bin<sup>1</sup>, Xinyu Xiao<sup>2</sup>, Zhiyu Xiang<sup>1</sup>, Hangguan Shan<sup>1</sup>, and Eryun Liu<sup>1\*</sup>

 <sup>1</sup> Zhejiang University, Hangzhou, China {wangsh0111,lubin2022, xiangzy, hshan, eryunliu}@zju.edu.cn
<sup>2</sup> Institute of Automation of Chinese Academy of Sciences, Beijing, China xinyu.xiao@nlpr.ia.ac.cn

**Abstract.** Multi-agent collaborative perception has emerged as a widely recognized technology in the field of autonomous driving in recent years. However, current collaborative perception predominantly relies on Li-DAR point clouds, with significantly less attention given to methods using camera images. This severely impedes the development of budgetconstrained collaborative systems and the exploitation of the advantages offered by the camera modality. This work proposes an instance-level fusion transformer for visual collaborative perception (IFTR), which enhances the detection performance of camera-only collaborative perception systems through the communication and sharing of visual features. To capture the visual information from multiple agents, we design an instance feature aggregation that interacts with the visual features of individual agents using predefined grid-shaped bird eye view (BEV) queries, generating more comprehensive and accurate BEV features. Additionally, we devise a cross-domain query adaptation as a heuristic to fuse 2D priors, implicitly encoding the candidate positions of targets. Furthermore, IFTR optimizes communication efficiency by sending instance-level features, achieving an optimal performance-bandwidth trade-off. We evaluate the proposed IFTR on a real dataset, DAIR-V2X, and two simulated datasets, OPV2V and V2XSet, achieving performance improvements of 57.96%, 9.23% and 12.99% in AP@70 metrics compared to the previous SOTAs, respectively. Extensive experiments demonstrate the superiority of IFTR and the effectiveness of its key components. The code is available at https://github.com/wangsh0111/IFTR.

**Keywords:** BEV Queries, Query Adaptation, Visual Collaborative Perception, Transformer

# 1 Introduction

3D object detection is a fundamental task in computer vision, aiming to locate objects in 3D physical space given sensor inputs. It is crucial in various

<sup>\*</sup> Corresponding author.

## 2 S. Wang et al.

applications such as autonomous driving [17, 27, 44, 48], robotics [14], and unmanned aerial vehicles [10]. Despite significant advancements in LiDAR-based methods [5, 13, 46, 50], camera-based approaches [26, 29, 35, 38, 46] have garnered widespread attention in recent years. In comparison to LiDAR, cameras have lower deployment costs, are more easily scalable, and provide valuable visual cues for detecting distant objects and identifying vision-based road elements (such as traffic lights and lane markings) [17]. However, due to the lack of depth information in RGB images, solely camera-based 3D object detection is often noticeably inferior to lidar-based detection in most scenarios.

The latest advancements in V2X communication technology and Intelligent Transportation Systems enable intelligent agents to share information between each other [1, 8, 28, 32, 47]. Through multi-agent interactions, the limitations of individual agent perception, such as occlusion and long-range issues, can be alleviated, thereby enhancing the performance of camera-based 3D object detection. Currently, the prevailing frameworks for collaborative perception [9, 36, 41–43] primarily involve the fusion at the Bird's Eye View (BEV) feature level. However, this paradigm is highly sensitive to the accuracy of BEV features for individual agents and faces challenges in spatial alignment, imposing stringent requirements on real-time poses between agents, leading to insufficient robustness in feature fusion. Nevertheless, due to the lack of depth information in RGB images, generating BEV features from a 2D plane is ill-posed. Therefore, the detection performance of visual collaborative perception methods based solely on BEV features level in camera settings is impacted by compounded errors, where inaccurate BEV features can severely impact the final performance.

Diverging from multi-camera 3D object detection, V2X collaborative perception must also account for the constraints of communication bandwidth. It is crucial to address how to reduce communication bandwidth while ensuring collaborative perception performance. Each agent must select visual features with the maximum information content to optimize communication efficiency.

In this paper, we propose IFTR, a robust and efficient transformer-based instance-level feature fusion framework. Our central idea is to extract spatial features from regions of interest in cross-agent camera views to generate BEV features, and implicitly encode candidate positions of objects in 3D space through a heuristic adaptive approach. Specifically, IFTR incorporates three key designs: i) Instance feature aggregation, aggregating spatial information from visual features of multiple agent camera views using predefined grid-shaped BEV queries. This is achieved through a transformer to implement robust and adaptive feature fusion, avoiding strict spatial alignment requirements in BEV feature-level fusion strategies. It also eliminates the dependency on high-accuracy single-agent BEV features and real-time pose accuracy. ii) Message selection and feature map reconstruction, reducing bandwidth consumption by sharing instance-level visual features rather than entire image features. iii) Cross-domain query adaptation, encoding instance-level features as object queries, combining them with the original learnable object queries to form the new hybrid object queries. This is then input into a deformable DETR head [2,51] for end-to-end 3D object detection.

To evaluate IFTR, we conduct extensive experiments on a real-world dataset, DAIR-V2X [47], and two simulated datasets, OPV2V [43] and V2XSet [42]. The experimental results indicate that IFTR significantly outperforms previous works in the performance-bandwidth trade-off across multiple datasets. To sum up, the main contributions are as follows:

- We introduce IFTR, a superior camera-only instance-level feature fusion framework based on transformers. This framework addresses occlusion and long-range issues through multi-agent collaboration, achieving more accurate and complete 3D object detection.
- We devise an instance feature aggregation that utilizes learnable BEV queries to aggregate spatial information from multiple agents, resulting in more accurate BEV features.
- We design a cross-domain query adaptation module by encoding instancelevel features as object queries, facilitating the implicit encoding of candidate positions for targets in 3D space, which alleviates the problem of target spatial distribution differences between the training set and the test set.
- We develope communication-efficient collaboration techniques by sharing the most informative visual features to reduce communication overhead. Our IFTR achieves state-of-the-art performance on a real world dataset, DAIR-V2X, and two simulated datasets, OPV2V and V2XSet, with AP@70 metrics improving by 57.96%, 9.23% and 12.99% compared to the previous SOTAs, respectively.

# 2 Related Work

#### 2.1 Camera-based 3D Object Detection

Given an RGB image and the corresponding camera parameters, the purpose of image-based 3D object detection is to classify and locate the objects of interest. Since images lack depth information, this problem is ill-posed and more challenging compared to 2D detection. FCOS3D [35] and SMOKE [23] extend 2D detection networks by utilizing fully convolutional networks to directly regress the depth for each object. Pseudo LiDAR [37] employs a depth estimation network to transform 2D images into 3D pseudo-point cloud signals, followed by the use of a lidar-based detection network for 3D detection. DETR3D [38] utilizes a set of sparse 3D queries to index 2D features extracted from multi-view images, samples corresponding features, enabling end-to-end 3D detection box prediction without the need for time-consuming post-processing like NMS (Non-Maximum Suppression).

Recently, there has been widespread attention on BEV representations due to their ability to clearly depict object positions and scales, coupled with efficient computational efficiency. OFT [30] and ImVoxelNet [31] project predefined voxels onto image features, generating voxel representations of the scene. LSS [26] and CaDDN [29] utilize view transformation modules to convert dense 2D image features into BEV space. CVT [49] leverages position encoding from camera perception and dense cross-attention to correlate perspective view and BEV view features. BEVDet [11, 12] and BEVDepth [16] use lift-splat operations for view transformation. PETR [21, 22] utilizes 3D position encoding and global crossattention for feature fusion, transforming 2D features into 3D position-aware representations. BEVFormer [17, 45] effectively aggregates spatio-temporal features from a surround-view camera and historical BEV features using deformable attention, generating unified BEV features. SparseBEV [19] adapts self-attention with an adaptive receptive field for feature aggregation in BEV space, enjoying the advantages of BEV space without explicitly constructing dense BEV features. In this work, we choose the single-scale and history-free BEVFormer-S [17] as the single-agent 3D object detector for computational efficiency.

#### 2.2 Collaborative Perception

Collaborative perception is an emerging application in multi-agent systems, aiming to enhance perceptual performance by facilitating communication among other agents for the shared exchange of information. Based on its message sharing strategy, it can be categorized into three types: early fusion (i.e. sharing raw sensor information), intermediate fusion (i.e. sharing intermediate layer features of neural networks), and late fusion (i.e. sharing 3D detection results, such as 3D bounding box position and confidence score). Early fusion usually requires large transmission bandwidth, while late fusion cannot provide valuable scenario context. Intermediate fusion achieves a better trade-off between performance and bandwidth. When 2 com [20] introduces a handshake mechanism to determine when to communicate and create sparse communication graph. OPV2V [43] utilizes a single-head self-attention module to fuse features, while F-Cooper [3] employs a maxout [7] fusion operation. V2VNet [36] utilizes a spatially aware graph neural network to aggregate shared feature representations among multiple vehicles. DiscoNet [15] adopts knowledge distillation to leverage the advantages of both early and intermediate collaboration. V2X-ViT [42] introduces a novel heterogeneous multi-agent attention module to fuse information across heterogeneous agents. Where 2 comm [9] introduces a spatial confidence map to reduce communication bandwidth consumption. TransIFF [4] further reduces communication bandwidth consumption by transmitting object queries. CoAlign [25] employs learnable or mathematical methods to correct pose errors for more accurate collaborative perception. HM-ViT [39] and HEAL [24] investigate multi-agent heterogeneous modality collaborative perception problems with different sensor modalities, expanding the scale of collaboration.

However, prior collaborative perception efforts [4,9,25,36,42,43] have primarily focused on lidar-based 3D object detection, yielding suboptimal performance in a camera-only setup. We speculate that the reason lies in the inferior quality of BEV features from a single-agent compared to lidar-based methods, resulting in suboptimal solutions at the BEV feature level. In this paper, we propose a novel camera-only instance-level feature fusion framework with lower communication bandwidth consumption. It generates more accurate BEV features and



Fig. 1: Overall architecture of IFTR. i) We employ the message selection and feature map reconstruction module to share instance-level features, reducing bandwidth consumption; ii) In instance feature aggregation (IFA), each BEV query interacts only with image features from regions of interest from multiple views; iii) In CDQA, we encode the feature map information and 3D positional information of each instance into 3D object query.

3D object queries tailored to each instance by leveraging instance-level features shared among multiple agents, thereby achieving more efficient and practical collaborative perception.

# 3 IFTR

This section presents IFTR, a robust and efficient camera-only instance-level feature fusion framework for collaborative perception. As depicted in Fig. 1, IFTR comprises an image encoder, message selection and feature map reconstruction, cascade instance feature aggregation, cross-domain query adaptation, and a DETR decoder. The image encoder is utilized to extract visual features from the input images. The proposed message selection and feature map reconstruction module are employed at the sender and receiver ends, respectively, for instance-level visual feature filtering and image feature reconstruction, effectively saving communication bandwidth. Based on visual features from multiple camera views, the proposed cascade instance feature aggregation generates comprehensive and accurate BEV features; the proposed cross-domain query adaptation module encodes instance-level visual features into object queries to encode candidate positions for the targets, combining them with the original learnable object queries to form the hybrid object queries. Ultimately, the DETR decoder

6 S. Wang et al.

takes the BEV features and the hybrid object queries as inputs to predict the target's category and location information.

## 3.1 Image Encoder

We leverage EfficientNet [33,34] to extract visual features from the input images because of its high performance and low computational complexity. Given Lagents and  $N_{cams}$  camera views for each agent,  $F_i^k \in \mathbb{R}^{fC \times fH \times fW}$  is used to represent the extracted feature map of the k-th camera view of the *i*-th agent, where fC, fH and fW are the channel, height and width of the feature map. Considering different agents may have different camera views, the feature maps of views that not exist are set to zero when the actual number of camera views less than  $N_{cams}$  and use a mask to indicate the valid feature maps.

#### 3.2 Message Selection and Feature Map Reconstruction

To achieve accurate and comprehensive 3D object detection, it is imperative for each agent to exchange visual features and leverage complementary information. We propose communication-efficient collaborative techniques by sharing the most informative visual features to reduce communication overhead. The intuition is that foreground regions contain more informative content than background regions. During collaboration, visual features of regions with objects can assist in recovering missed objects due to occlusion or long-range issues, while background regions can be omitted to reduce the system's communication load.

Given a feature map  $F_j^k$ , the message sent for sharing is selected by the following steps:

- 1. Extracting a set of candidate object instance by a 2D object detector. All the confidence values of instance are greater than a predefined threshold,  $c_{thre}$ .
- 2. Cropping the feature map of each instance by the predicted bounding box in inference stage or by the ground truth bounding box in the training stage.
- 3. The cropped feature map of the *n*-th selected instance,  $F_{j \to i}^{k,n}$ , and bounding box are sent to the *i*-th agent for sharing.

At the side of ego agent (or the *i*-th agent), the received object instance feature maps of the *k*-th camera view of the *j*-th agent are first combined together to reconstruct the feature map,  $\mathcal{F}_{j\to i}^k$ , of foreground regions of original feature map  $F_j^k$  by filling back the cropped features. The features in background region are set to zero.

## 3.3 Instance Feature Aggregation

Given a set of received foreground feature maps,  $\{\mathcal{F}_{j \to i}^k | j = 1, 2, \ldots, N_{agent}, k = 1, 2, \ldots, N_{cams}\}$ , where  $N_{agent}$  is the number of agents, the purpose of instance feature aggregation (IFA) module is to form a BEV feature map by aggregating

7



**Fig. 2:** (a) The architecture of our proposed Instance Feature Aggregation (IFA); (b) Multi-View Feature Aggregation (MVFA) illustrated in Sec. 3.3

features from all received foreground feature maps. Since all camera views are treated equally in IFA module, for simplicity, we represent all the foreground feature maps by  $\{\mathcal{F}_v | v = 1, 2, \ldots, N_v\}$ , where the total number of views is  $N_v = N_{agent} \times N_{cams}$ .

The BEV space around the ego agent is divided into a discrete grid space of size  $H \times W$ . The origin of ego coordinate is at the center of BEV grid. To construct a feature map on the BEV space of the ego agent by fusing reconstructed feature maps  $\{\mathcal{F}_i | i = 1, 2, \ldots, N_v\}$  from all collaborative views, we propose a deformable attention based cascade instance feature aggregation (IFA) module (see Fig. 2). The BEV feature map is obtained by the following steps:

- 1. A set of learnable parameters  $Q_{\text{bev}} \in \mathbb{R}^{C \times H \times W}$  is predefined and randomly initialized at the first IFA block.
- 2. For a given grid centered at (x, y), a set of points  $\{(x, y, z_h)|h = 1, 2, ..., N_{ref}\}$  is selected as reference points, where  $\{z_h|h = 1, 2, ..., N_{ref}\}$  are predefined heights for each grid. The reference points are used as feature indices to gather information from collaborative feature maps.
- 3. For a reference point  $(x, y, z_h)$  at grid (x, y), the corresponding feature vector is aggregated by

$$F_{bev}^{h}(x,y) = \frac{1}{|w|} \sum_{v=1}^{N_{v}} w_{v} f_{v}^{h},$$

where  $w_v \in \{0, 1\}$  indicates whether the reference point  $(x, y, z_h)$  is projected inside the image space of the v-th camera view, and  $|w| = \sum_v w_v$  is total number of views observing the reference point, and  $f_v^h$  is the output of a deformable attention module, which will be described later in this section.

- 8 S. Wang et al.
- 4. The aggregated BEV feature  $F_{bev}$  is the average feature over all  $N_{ref}$  reference points for each grid and followed by a feed forward network (FFN) and an Add&Norm module to get the output  $\hat{F}_{bev}$  of IFA module.
- 5. The above four steps assemble an IFA module. To sufficiently fuse all camera view features, we cascade 6 IFA modules sequentially, with the previous IFA output  $\hat{F}_{bev}$  as the input  $Q_{bev}$  of the next IFA.

Given a reference point  $(x, y, z_h)$  at grid (x, y) and the feature map  $\mathcal{F}_v$  of v-th camera view, if point  $(x, y, z_h)$  is observed by the v-th camera view, the deformable attention first takes a feature vector  $\hat{Q}_{\text{bev}}(x, y)$  from  $\hat{Q}_{\text{bev}}$  and the vector  $\hat{Q}_{\text{bev}}(x, y)$  is then fed into a MLP layer. The MLP layer output a list of offset vectors  $(o_1, o_2, \ldots, o_{N_{da}})$  and a weight vector  $\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_{N_{da}})$ , where  $N_{da}$  is the number of positions in attention, the output of deformable attention is

$$f_v^h = \sum_{i=1}^{N_{da}} \lambda_i \mathcal{F}_v(p+o_i),$$

where p is the projected image coordinate of  $(x, y, z_h)$  on v-th view.

#### 3.4 Cross-Domain Query Adaptation

The object decoder of BEVformer [17] is a DETR [2,51] decoder that receives a set of learned embeddings as object queries, aiming to learn potential locations of target objects during training. However, it takes a considerable amount of time for randomly initialized embeddings to converge to suitable positions. Additionally, during inference, object queries remain fixed for all images, while the spatial distribution of target objects may vary, leading to potentially inaccurate results. To address this issue, we employ a cross-domain query adaptation(CDQA) module to initialize object queries for decoder. This module implicitly encodes candidate positions for the target, making it easier for the detection head to capture the target object accurately.

The maximum number of queries for decoder is set to  $N_Q$ . The initial set of queries consists of two parts. The first part of queries is obtained from selected instances of all ego and collaborative camera views and the second part (the rest) of queries is randomly initialized.

Given a set of shared instance features  $\{(\mathcal{I}_n, B_n)|n = 1, 2, ..., N_I\}$  from all collaborative and ego views, where  $N_I$  is total number of instances,  $\mathcal{I}_n$  is the *n*-th cropped instance feature map, and  $B_n$  is the bounding box coordinates, the queries is obtained as following

- 1. Global average pooling (GAP) on instance feature maps and then followed by a MLP module to get a set of feature vector  $\{\mathcal{I}_n^p | n = 1, 2, ..., N_I\}$ .
- 2. Encode the instance position information into a vector by a MLP layer, obtaining a set of position embeddings  $\{\mathcal{V}_n | n = 1, 2, ..., N_I\}$ . The input to the MLP is the concatenate of three coordinates  $(p_1, p_2, c)$ , where  $p_1, p_2$  are the 3D coordinates of left-top and bottom-right point of instance bounding

box on the focal plane and c is the coordinate of corresponding camera origin. All coordinates are projected to the ego coordinate system. The three points indicate a cone where the object lies.

3. Combine the feature vectors and position embeddings by adding, obtaining the instance feature initialized set of queries  $\{\mathcal{I}_n^p + \mathcal{V}_n | n = 1, 2, ..., N_I\}$  for object decoder.

The union,  $Q_{\text{hybrid}}$ , of the instance feature initialized set of queries and the randomly initialized set of queries is then used as queries of object decoder.

## 3.5 Object Decoder

The object decoder of IFTR employs a deformable DETR detection head, taking the BEV features  $\hat{F}_{bev} \in \mathbb{R}^{C \times H \times W}$  and the hybrid object queries  $Q_{hybrid} \in \mathbb{R}^{N_Q \times C}$  as input, producing detection results  $\mathcal{O} = \Phi_{dec}(\hat{F}_{bev}, Q_{hybrid}) \in \mathbb{R}^{N_Q \times 8}$ . Each position of  $\mathcal{O}$  represents a rotated box with class  $(c, x, y, z, w, l, h, \theta)$ , denoting class confidence, position, size and angle. The objects are the final output of the proposed collaborative perception system. We utilize the  $\ell_1$  loss for regression and the focal loss [18] for classification.

# 4 Experiments

## 4.1 Experimental Setup

**Datasets.** To evaluate the performance of IFTR on the visual collaborative perception task, we conduct extensive experiments on three multi-agent datasets, including V2XSet [42], OPV2V [43], and DAIR-V2X [47]. V2XSet [42] is a simulated dataset supporting V2X perception, co-simulated by CARLA [6] and OpenCDA [40]. It comprises 73 representative scenes with 2 to 5 connected agents and 11,447 annotated 3D frames, with an image resolution of  $600 \times 800$ . OPV2V [43] is a large-scale vehicle-to-vehicle collaborative perception dataset, comprising 12,000 LiDAR point cloud frames and RGB images with 230,000 annotations of 3D bounding boxes, with an image resolution of  $600 \times 800$ . DAIR-V2X [47] is a real-world vehicle-to-infrastructure collaborative perception dataset, consists of 100 realistic scenes and 18,000 data samples, with an image resolution of  $1080 \times 1920$ .

**Implementation details.** We use the single-scale and history-free BEVFormer-S [17] as the 3D object detector for single agent. We employ EfficientNet [33,34] as the image backbone for better computation efficiency and utilize a smaller grid resolution (0.4 m) to retain fine-trained spatial details. The intermediate BEV feature map has a dimension of  $256 \times 256 \times 256$ , with a perception range defined as  $x \in [-51.2 \text{ m}, 51.2 \text{ m}]$  and  $y \in [-51.2 \text{ m}, 51.2 \text{ m}]$ . Following prior research [41–43], we only modified the fusion module of different intermediate fusion methods while keeping the other components such as camera feature extractor and detection head the same. As IFTR involves instance-level fusion of 10 S. Wang et al.

**Table 1:** 3D detection performance comparison on the V2XSet [42], OPV2V [43] and DAIR-V2X [47] datasets. The results are reported in Average Precision (AP) at IoU=0.30, 0.50 and 0.70 on perfect settings.

Mathad	V2XSet			OPV2V			DAIR-V2X		
Method	AP@30	AP@50	AP@70	AP@30	AP@50	AP@70	AP@30	AP@50	AP@70
No Collaboration	0.4343	0.3037	0.1379	0.5652	0.4594	0.2556	0.2102	0.0709	0.0080
Late Fusion	0.6803	0.5141	0.2559	0.8789	0.7762	0.5192	0.2952	0.1631	0.0421
V2VNet [36]	0.6886	0.5954	0.3900	0.8682	0.7906	0.5759	0.2581	0.1354	0.0276
V2X-ViT [42]	0.6690	0.5914	0.4123	0.8512	0.7841	0.5838	0.2544	0.1402	0.0277
Where2comm [9]	0.6969	0.6169	0.4396	0.8471	0.7714	0.5860	0.2613	0.1424	0.0296
CoBEVT [41]	0.6629	0.5884	0.4081	0.8709	0.8026	0.5934	0.2608	0.1453	0.0255
CoAlign [25]	0.7564	0.6479	0.3964	0.8689	0.8021	0.6046	0.2645	0.1470	0.0262
IFTR	0.8057	0.7173	0.4967	0.9233	0.8556	0.6604	0.3855	0.2058	0.0665

2D image features, we construct it following the method proposed in Sec. 3 and straightforwardly add a late fusion strategy to explore its combined effects with other fusion strategies.

**Evaluation metrics.** We evaluate the detection performance using Average Precision (AP) metrics with Intersection over Union (IoU) thresholds of 0.30, 0.50, and 0.70. Additionally, the communication volume follows the standard setting as [4,9,42,43] that counts the message size by byte in log scale with base 2.

#### 4.2 Quantitative Evaluation

**Benchmark comparison.** Tab. 1 compares the proposed IFTR with previous collaborative methods. We consider single-agent detection without collaboration (No collaboration), V2VNet [36], V2XViT [42], Where2comm [9], CoBEVT [41], CoAlign [25] and late fusion, where late fusion directly exchanges 3D detection boxes. We observe that the proposed IFTR outperforms previous methods on both real-world and simulated datasets, demonstrating the superiority of our model and its robustness to various realistic noises. Specifically, the SOTA performance of AP@70 on the V2XSet, OPV2V and DAIR-V2X datasets is improved by 12.99%, 9.23% and 57.96%, respectively. Compared to previous collaboration methods at the BEV level, IFTR significantly enhances the quality of BEV features through instance feature aggregation, and adapts the cross-domain query adaption to implicitly encode the positions of the target in 3D space, achieving more accurate perception.

Number of agents. In this experiment, we investigate the impact of the quantity of agents on the perceptual performance of IFTR on the OPV2V dataset. As depicted in Fig. 3a, the perceptual performance positively increases as more agents engage in cooperative perception. From Fig. 3b, it can be observed that, in



Fig. 3: IFTR steadily improves 3D detection performance as the number of agents grows. (a) The relationship between 3D detection performance and the maximum collaboration count on the OPV2V test set; (b) The relationship between 3D detection performance and collaboration count in a certain scene on the OPV2V test set.

a simple perception environment, the IFTR performance exceeds that of the single LiDAR-based detector [13] when the collaboration count is 2 for the AP@50 metric. Similarly, for the AP@70 metric, the IFTR performance surpasses the single LiDAR-based detector when the collaboration count is 5. It is noteworthy that the detection performance consistently improves with the increase in the number of collaborative agents. Even upon reaching the maximum collaboration quantity in the dataset, a rising trend in performance persists. Therefore, we advocate for active collaboration among more agents to achieve increasingly advanced perceptual performance.

**Robustness to localization noise.** We follow the localization noise settings in V2X-ViT [42] (Gaussian noise with a mean of 0m and a variance of 0m-0.6m) and validate the robustness of IFTR against realistic localization noise. Fig. 4 illustrates the detection performances as a function of localization noise level in V2XSet and DAIR-V2X datasets, respectively. We see: i) with the increasing localization noise, the performance of all intermediate collaborative methods deteriorates due to the misalignment of feature maps, while IFTR outperforms previous SOTAs at all the localization noise levels; ii) IFTR consistently outperforms the No Collaboration by a significant margin. The reason is the powerful transformer architecture in IFA can capture crucial perceptual cues and holistic information across agents.

#### 4.3 Ablation Studies



Fig. 4: Robustness to localization noise on the DAIR-V2X and V2XSet datasets. Gaussian noise with zero mean and a varying variance is introduced. IFTR consistently outperforms previous SOTAs.

**Table 2:** Ablation Study results of the proposed components on the OPV2V and DAIR-V2X datasets. IFA, CDQA respectively represent the incorporation of i) instance feature aggregation, and ii) cross-domain query adaptation. MASK denotes interacting exclusively with the foreground features of collaborators in the multi-view feature aggregation module.

IFA CDQA MASK		(	OPV2V		DAIR-V2X			
		MASK	AP@50	AP@70	$\operatorname{Comm}$	AP@50	AP@70	$\operatorname{Comm}$
			0.7762	0.5192	11.45	0.1631	0.0421	12.04
$\checkmark$			0.8480	0.6487	24.23	0.1998	0.0574	21.91
$\checkmark$	$\checkmark$		0.8556	0.6604	24.23	0.2058	0.0665	21.91
$\checkmark$	$\checkmark$	$\checkmark$	0.8458	0.6300	20.61	0.1978	0.0618	18.23

Effectiveness of different modules in IFTR. Here we investigate the effectiveness of various components in IFTR, with the base model being BEVFormer-S incorporating a late fusion strategy. We progressively adding i) IFA, ii) CDQA, and iii) MASK to assess the impact of each component on detection performance and communication volume, where MASK refers to interacting only with the foreground features of collaborators in the MVFA module. As Tab. 2 demonstrates i) IFA contributes the most significant performance gain, increasing AP@70 on OPV2V and DAIR-V2X by 24.94% and 36.34%, respectively. This is attributed to the MVFA design in IFA, which significantly improves the quality of BEV features for individual agents, enhancing BEV features and further improving the perceptual effectiveness of late fusion; ii) CDQA increases AP@70 on OPV2V and DAIR-V2X by 1.80% and 15.85%, respectively. This is because CDQA encodes 2D instances into 3D object queries, making the detection head less dependent on the initially learned object queries during training (i.e., gen-

13

**Table 3:** Comparison of different positional encoding strategies in the CDQA module on the OPV2V and DAIR-V2X datasets. None denotes not using positional encoding, Learned denotes the use of trainable positional encoding, and CE stands for employing the proposed cone encoding.

Strategy	OP	V2V	DAIR-V2X			
	AP@50	AP@70	AP@50	AP@70		
None	0.8493	0.6532	0.1950	0.0610		
Learned	0.8532	0.6588	0.2025	0.0628		
CE	0.8556	0.6604	0.2058	0.0665		

erating object queries adaptively based on 2D instances that are more suitable for 3D detection of the instance). It effectively alleviates the problem of differences in spatial distribution between the test set and the training set, and we speculate that this module will bring more significant performance improvement when there is a greater spatial distribution difference; iii) MASK reduces communication overhead by 12.30 times on the OPV2V dataset and 12.82 times on the DAIR-V2X dataset, the detection performance only decreased by 4.60% and 7.07% on the respective datasets for the AP@70 metric. This is highly useful for bandwidth-constrained application scenarios. It is worth noting that our proposed IFTR can also be combined with other fusion methods (such as CoBEVT [41] and CoAlign [25]), and we encourage researchers to explore their combined usage, believing it will lead to more advanced detection performance.

**Different positional encoding schemes in CDQA.** Due to the crucial importance of positional encoding in transformer-based architectures, in this subsection, we compare cone encoding with other commonly used positional encoding methods. As can be seen from Tab. 3, the performance using cone encoding is superior to other methods. This is attributed to cone encoding providing a potential cone space prior for instances, which can alleviate the difficulty of the detection head in capturing targets.

## 4.4 Qualitative Evaluation

Visualization of detection results. We conducted a qualitative analysis of the model's performance using typical samples from the OPV2V dataset, as illustrated in Fig. 5. The first column (a) shows the perception results of no collaboration, the second column (b) shows results with late fusion, and the third column (c) shows results using IFTR for collaborative perception. Overall, our approach generates more prediction bounding boxes that perfectly align with the ground truth. This improvement can be attributed to two main factors: i) the proposed IFA integrates meaningful information from nearby agents, significantly enhancing the quality of BEV features, and thereby improving detection performance; ii) the proposed CDQA provides effective spatial cues for objects in the scene, effectively compensating for the differences in target spatial distribution between the test set and the training set.



(a) No Collaboration

(c) IFTR

**Fig. 5:** Visualization of predictions from (a) No Collaboration, (b) Late Fusion and (c) IFTR on the OPV2V test set. Green and red 3D bounding boxes represent the ground truth and prediction respectively.

(b) Late Fusion

# 5 Conclusion

IFTR is an instance-level feature fusion framework based on transformer that aims to enhance the detection performance of camera-only collaborative perception systems through the communication and sharing of visual features. It comprises three key components: message selection and feature map reconstruction, instance feature aggregation, and cross-domain query adaptation. The message selection and feature map reconstruction module optimize communication efficiency by transmitting instance-level features. Instance feature aggregation interacts with the visual features of individual agents using predefined grid-shaped BEV queries, generating more comprehensive and accurate BEV features. The cross-domain query adaptation implicitly encodes candidate positions of targets, heuristically integrating 2D priors. Extensive experiments on several multi-agent datasets show the effectiveness of IFTR and the necessity of all its components.

# Acknowledgements

The work was supported by the Key Research & Development Plan of Zhejiang Province under Grant 2021C01196, in part by the Zhejiang Provincial Natural Science Foundation of China under Grant LR23F010006, in part by the Open Research Project Programme funded by the Science and Technology Development Fund (SKL-IOTSC(UM)-2024-2026) and the State Key Laboratory of Internet

15

of Things for Smart City (University of Macau) (Ref. No.: SKL-IoTSC(UM)-2024-2026/ORP/GA01/2023).

# References

- Betz, J., Zheng, H., Liniger, A., Rosolia, U., Karle, P., Behl, M., Krovi, V., Mangharam, R.: Autonomous vehicles on the edge: A survey on autonomous vehicle racing. IEEE Open Journal of Intelligent Transportation Systems 3, 458–488 (2022)
- Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: Endto-end object detection with transformers. In: European conference on computer vision. pp. 213–229. Springer (2020)
- Chen, Q., Ma, X., Tang, S., Guo, J., Yang, Q., Fu, S.: 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. pp. 88–100 (2019)
- Chen, Z., Shi, Y., Jia, J.: Transiff: An instance-level feature fusion framework for vehicle-infrastructure cooperative 3d detection with transformers. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 18205–18214 (2023)
- Deng, J., Shi, S., Li, P., Zhou, W., Zhang, Y., Li, H.: Voxel r-cnn: Towards high performance voxel-based 3d object detection. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 1201–1209 (2021)
- Dosovitskiy, A., Ros, G., Codevilla, F., Lopez, A., Koltun, V.: Carla: An open urban driving simulator. In: Conference on robot learning. pp. 1–16. PMLR (2017)
- Goodfellow, I., Warde-Farley, D., Mirza, M., Courville, A., Bengio, Y.: Maxout networks. In: International conference on machine learning. pp. 1319–1327. PMLR (2013)
- Guo, Y., Ma, J., Leslie, E., Huang, Z.: Evaluating the effectiveness of integrated connected automated vehicle applications applied to freeway managed lanes. IEEE Transactions on Intelligent Transportation Systems 23(1), 522–536 (2020)
- Hu, Y., Fang, S., Lei, Z., Zhong, Y., Chen, S.: Where2comm: Communicationefficient collaborative perception via spatial confidence maps. Advances in neural information processing systems 35, 4874–4886 (2022)
- Hu, Y., Fang, S., Xie, W., Chen, S.: Aerial monocular 3d object detection. IEEE Robotics and Automation Letters 8(4), 1959–1966 (2023)
- Huang, J., Huang, G.: Bevdet4d: Exploit temporal cues in multi-camera 3d object detection. arXiv preprint arXiv:2203.17054 (2022)
- Huang, J., Huang, G., Zhu, Z., Ye, Y., Du, D.: Bevdet: High-performance multicamera 3d object detection in bird-eye-view. arXiv preprint arXiv:2112.11790 (2021)
- Lang, A.H., Vora, S., Caesar, H., Zhou, L., Yang, J., Beijbom, O.: Pointpillars: Fast encoders for object detection from point clouds. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 12697–12705 (2019)
- Li, J., Li, R., Li, J., Wang, J., Wu, Q., Liu, X.: Dual-view 3d object recognition and detection via lidar point cloud and camera image. Robotics and Autonomous Systems 150, 103999 (2022)
- Li, Y., Ren, S., Wu, P., Chen, S., Feng, C., Zhang, W.: Learning distilled collaboration graph for multi-agent perception. Advances in Neural Information Processing Systems 34, 29541–29552 (2021)

- 16 S. Wang et al.
- Li, Y., Ge, Z., Yu, G., Yang, J., Wang, Z., Shi, Y., Sun, J., Li, Z.: Bevdepth: Acquisition of reliable depth for multi-view 3d object detection. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 37, pp. 1477–1485 (2023)
- Li, Z., Wang, W., Li, H., Xie, E., Sima, C., Lu, T., Qiao, Y., Dai, J.: Bevformer: Learning bird's-eye-view representation from multi-camera images via spatiotemporal transformers. In: European conference on computer vision. pp. 1–18. Springer (2022)
- Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: Proceedings of the IEEE international conference on computer vision. pp. 2980–2988 (2017)
- Liu, H., Teng, Y., Lu, T., Wang, H., Wang, L.: Sparsebev: High-performance sparse 3d object detection from multi-camera videos. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 18580–18590 (2023)
- Liu, Y.C., Tian, J., Glaser, N., Kira, Z.: When2com: Multi-agent perception via communication graph grouping. In: Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition. pp. 4106–4115 (2020)
- Liu, Y., Wang, T., Zhang, X., Sun, J.: Petr: Position embedding transformation for multi-view 3d object detection. In: European Conference on Computer Vision. pp. 531–548. Springer (2022)
- 22. Liu, Y., Yan, J., Jia, F., Li, S., Gao, A., Wang, T., Zhang, X.: Petrv2: A unified framework for 3d perception from multi-camera images. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3262–3272 (2023)
- Liu, Z., Wu, Z., Tóth, R.: Smoke: Single-stage monocular 3d object detection via keypoint estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. pp. 996–997 (2020)
- Lu, Y., Hu, Y., Zhong, Y., Wang, D., Chen, S., Wang, Y.: An extensible framework for open heterogeneous collaborative perception. arXiv preprint arXiv:2401.13964 (2024)
- Lu, Y., Li, Q., Liu, B., Dianati, M., Feng, C., Chen, S., Wang, Y.: Robust collaborative 3d object detection in presence of pose errors. In: 2023 IEEE International Conference on Robotics and Automation (ICRA). pp. 4812–4818. IEEE (2023)
- Philion, J., Fidler, S.: Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIV 16. pp. 194–210. Springer (2020)
- 27. Qi, C.R., Su, H., Mo, K., Guibas, L.J.: Pointnet: Deep learning on point sets for 3d classification and segmentation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 652–660 (2017)
- Raboy, K., Ma, J., Leslie, E., Zhou, F.: A proof-of-concept field experiment on cooperative lane change maneuvers using a prototype connected automated vehicle testing platform. Journal of Intelligent Transportation Systems 25(1), 77–92 (2021)
- Reading, C., Harakeh, A., Chae, J., Waslander, S.L.: Categorical depth distribution network for monocular 3d object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8555–8564 (2021)
- Roddick, T., Kendall, A., Cipolla, R.: Orthographic feature transform for monocular 3d object detection. arXiv preprint arXiv:1811.08188 (2018)
- Rukhovich, D., Vorontsova, A., Konushin, A.: Imvoxelnet: Image to voxels projection for monocular and multi-view general-purpose 3d object detection. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 2397–2406 (2022)

- 32. Shladover, S.E.: Opportunities and challenges in cooperative road vehicle automation. IEEE Open Journal of Intelligent Transportation Systems **2**, 216–224 (2021)
- Tan, M., Le, Q.: Efficientnet: Rethinking model scaling for convolutional neural networks. In: International conference on machine learning. pp. 6105–6114. PMLR (2019)
- 34. Tan, M., Le, Q.: Efficient netv2: Smaller models and faster training. In: International conference on machine learning. pp. 10096–10106. PMLR (2021)
- Wang, T., Zhu, X., Pang, J., Lin, D.: Fcos3d: Fully convolutional one-stage monocular 3d object detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 913–922 (2021)
- Wang, T.H., Manivasagam, S., Liang, M., Yang, B., Zeng, W., Urtasun, R.: 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 16. pp. 605–621. Springer (2020)
- 37. Wang, Y., Chao, W.L., Garg, D., Hariharan, B., Campbell, M., Weinberger, K.Q.: Pseudo-lidar from visual depth estimation: Bridging the gap in 3d object detection for autonomous driving. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8445–8453 (2019)
- Wang, Y., Guizilini, V.C., Zhang, T., Wang, Y., Zhao, H., Solomon, J.: Detr3d: 3d object detection from multi-view images via 3d-to-2d queries. In: Conference on Robot Learning. pp. 180–191. PMLR (2022)
- 39. Xiang, H., Xu, R., Ma, J.: Hm-vit: Hetero-modal vehicle-to-vehicle cooperative perception with vision transformer. arXiv preprint arXiv:2304.10628 (2023)
- Xu, R., Guo, Y., Han, X., Xia, X., Xiang, H., Ma, J.: Opencda: an open cooperative driving automation framework integrated with co-simulation. In: 2021 IEEE International Intelligent Transportation Systems Conference (ITSC). pp. 1155–1162. IEEE (2021)
- 41. Xu, R., Tu, Z., Xiang, H., Shao, W., Zhou, B., Ma, J.: Cobevt: Cooperative bird's eye view semantic segmentation with sparse transformers. arXiv preprint arXiv:2207.02202 (2022)
- 42. Xu, R., Xiang, H., Tu, Z., Xia, X., Yang, M.H., Ma, J.: V2x-vit: Vehicle-toeverything cooperative perception with vision transformer. In: European conference on computer vision. pp. 107–124. Springer (2022)
- Xu, R., Xiang, H., Xia, X., Han, X., Li, J., Ma, J.: Opv2v: An open benchmark dataset and fusion pipeline for perception with vehicle-to-vehicle communication. In: 2022 International Conference on Robotics and Automation (ICRA). pp. 2583– 2589. IEEE (2022)
- Yan, Y., Mao, Y., Li, B.: Second: Sparsely embedded convolutional detection. Sensors 18(10), 3337 (2018)
- 45. Yang, C., Chen, Y., Tian, H., Tao, C., Zhu, X., Zhang, Z., Huang, G., Li, H., Qiao, Y., Lu, L., et al.: Bevformer v2: Adapting modern image backbones to bird's-eyeview recognition via perspective supervision. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 17830–17839 (2023)
- Yin, T., Zhou, X., Krahenbuhl, P.: Center-based 3d object detection and tracking. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 11784–11793 (2021)
- 47. Yu, H., Luo, Y., Shu, M., Huo, Y., Yang, Z., Shi, Y., Guo, Z., Li, H., Hu, X., Yuan, J., et al.: Dair-v2x: A large-scale dataset for vehicle-infrastructure cooperative 3d object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 21361–21370 (2022)

- 18 S. Wang et al.
- Zhang, R., Qiu, H., Wang, T., Guo, Z., Cui, Z., Qiao, Y., Li, H., Gao, P.: Monodetr: Depth-guided transformer for monocular 3d object detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9155–9166 (2023)
- Zhou, B., Krähenbühl, P.: Cross-view transformers for real-time map-view semantic segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 13760–13769 (2022)
- Zhou, Y., Tuzel, O.: Voxelnet: End-to-end learning for point cloud based 3d object detection. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4490–4499 (2018)
- Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable detr: Deformable transformers for end-to-end object detection. arXiv preprint arXiv:2010.04159 (2020)