

Spatio-Temporal Domain Awareness for Multi-Agent Collaborative Perception

Kun Yang^{1†} Dingkang Yang^{1†} Jingyu Zhang¹ Mingcheng Li¹ Yang Liu¹
Jing Liu¹ Hanqi Wang¹ Peng Sun² Liang Song^{1*}

¹Academy for Engineering and Technology, Fudan University ²Duke Kunshan University

{kunyang20, dkyang20, songl}@fudan.edu.cn

Abstract

Multi-agent collaborative perception as a potential application for vehicle-to-everything communication could significantly improve the perception performance of autonomous vehicles over single-agent perception. However, several challenges remain in achieving pragmatic information sharing in this emerging research. In this paper, we propose SCOPE, a novel collaborative perception framework that aggregates the spatio-temporal awareness characteristics across on-road agents in an end-to-end manner. Specifically, SCOPE has three distinct strengths: i) it considers effective semantic cues of the temporal context to enhance current representations of the target agent; ii) it aggregates perceptually critical spatial information from heterogeneous agents and overcomes localization errors via multi-scale feature interactions; iii) it integrates multi-source representations of the target agent based on their complementary contributions by an adaptive fusion paradigm. To thoroughly evaluate SCOPE, we consider both real-world and simulated scenarios of collaborative 3D object detection tasks on three datasets. Extensive experiments show the superiority of our approach and the necessity of the proposed components. The project link is <https://ydk122024.github.io/SCOPE/>.

1. Introduction

Perceiving complex driving environments is essential for ensuring road safety [23] and traffic efficiency [21] in autonomous driving. Despite the rapid development in the perception ability of intelligent vehicles through deep learning technologies [5, 6, 14, 41, 43], the conventional single-agent perception paradigm remains challenging. Single-agent perception system is generally vulnerable and defective due to the inevitable practical difficulties on the road,

such as restricted detection range and occlusion [47, 48]. To alleviate the inadequate observation constraints of isolated agents, multi-agent collaborative perception has emerged as a promising solution in recent years [17, 29]. With Vehicle-to-Everything (V2X) oriented communication applications, collaborative perception effectively promotes perspective complementation and information sharing among agents.

Based on recent LiDAR-based 3D collaborative perception datasets [35, 36, 37, 46], several impressive works have achieved a trade-off between collaborative perception performance and communication bandwidth of Autonomous Vehicles (AVs) via well-designed mechanisms. These remarkable efforts include knowledge distillation-based intermediate collaboration [17], handshake communication mechanism [20], V2X-visual transformer [36], and spatial information selection strategy [10]. However, there is still considerable room for progress in achieving a pragmatic and effective collaborative perception. Concretely, existing methods invariably follow a single-frame static perception pattern, suffering from the data sparsity shortcoming of 3D point clouds [47] and ignoring meaningful semantic cues in the temporal context [49]. The insufficient representation of the moving objects in the current frame potentially restricts the perception performance of target vehicles.

Moreover, spatial information aggregation in collaborative perception systems has exposed several problems. For the collaborator-shared message fusion, the per-agent/location-based fusion strategies of previous attempts [16, 17, 20, 28, 36] fail to deal with feature map misalignment from heterogeneous agents due to localization errors. Consequently, misleading features from collaborators could lead to object position misjudgment of the ego vehicle (*i.e.*, receiver) and harm its detection precision. For the ego vehicle message refinement, existing approaches [10, 17, 20, 21, 28, 33] rely on fused representations with collaborator information to implement detection, abandoning the natural perception advantages of the ego vehicle and introducing the potential noises. The ego-centered characteristics may contain locally critical spatial information that is not disturbed by noise from collaborative agents

[†]Equal contributions in no particular order. The work was done when Dingkang Yang was at the Institute of Meta-Medical, IPASS. The two first authors thank Runsheng Xu for providing constructive suggestions.

*Corresponding author.

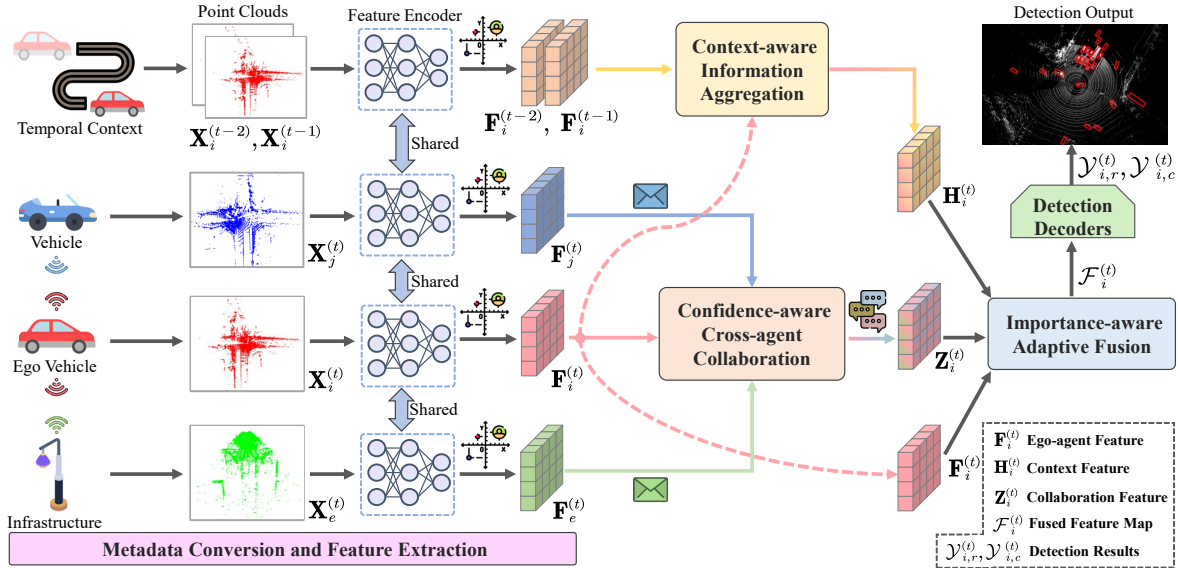


Figure 1. The overall architecture of the proposed SCOPE. The framework consists of five parts: metadata conversion and feature extraction, context-aware information aggregation, confidence-aware cross-agent collaboration, importance-aware adaptive fusion, and detection decoders. The details of each individual component are illustrated in Section 3.

with asynchronous sensor measurements. To this end, how to effectively break the aforementioned limitations becomes a priority for achieving robust collaborative perception.

Motivated by the above observations, we present SCOPE, a Spatio-temporal domain awareness multi-agent CollabORative PERception approach to tackle existing challenges jointly. From Figure 1, SCOPE achieves effective information collaboration among agents and precise perception of the ego agent via three proposed core components. Specifically, (i) for the data sparsity challenge in single-frame point clouds, the *context-aware information aggregation* is proposed to aggregate the ego agent’s context information from preceding frames. We employ selective information filtering and spatio-temporal feature integration modules to capture informative temporal cues and historical context semantics. (ii) For the collaborator-shared message fusion challenge, the *confidence-aware cross-agent collaboration* is introduced to guarantee that the ego agent aggregates complementary information from heterogeneous agents. Based on the confidence-aware multi-scale feature interaction, we facilitate holistic communication and mitigate feature map misalignment due to localization errors of collaborators. (iii) For the multi-feature integration challenge of the ego agent, the *importance-aware adaptive fusion* is designed to flexibly incorporate the advantages of diverse representations based on distinctive perception contributions. We evaluate the superiority of SCOPE on various collaborative 3D object detection datasets, including DAIR-V2X [46], OPV2V [37], and V2XSet [36]. Extensive experiments in real-world and simulated scenarios provide favorable evidence that our approach is competitive in

multi-agent collaborative perception tasks.

The main contributions can be summarized as follows:

- We present SCOPE, a novel framework for multi-agent collaborative perception. The framework facilitates information collaboration and feature fusion among agents, achieving a reasonable performance-bandwidth trade-off. Comprehensive experiments on collaborative detection tasks show that SCOPE outperforms previous state-of-the-art methods.
- To our best knowledge, we are the first to consider the temporal context of the ego agent in collaborative perception systems. Based on the proposed context-aware component, the target frame of point clouds efficiently integrates historical cues from preceding frames to capture valuable temporal information.
- We introduce two spatial information aggregation components to address the challenges of collaboration heterogeneity and fused representation unicity. These tailored components effectively perform multi-grained information interactions among agents and multi-source feature refinement of the ego agent.

2. Related Work

2.1. Multi-Agent Communication

Benefiting from rapid advances in learning-based technologies [18, 38, 39, 40, 42, 44, 45], communication has become an essential component in constructing robust multi-agent systems. Previous works [7, 25, 26, 27] typically focus on centralized paradigms or predefined protocols to

explore information-sharing mechanisms among multiple agents. However, these efforts are potentially domain-constrained and scenario-isolated. Recent methods [3, 9, 12] achieve effective communication strategies in diverse scenarios. For instance, Vain [9] determines which agents would share information by modeling the locality of interactions. ATOC [12] introduces an attentional communication strategy to learn when communicating and integrating useful information to execute collaborative decisions. In comparison, we focus on the more challenging 3D perception tasks in complex driving scenarios. Based on the proposed confidence-aware multi-scale information interaction, SCOPE achieves an information-critical and holistic communication pattern among heterogeneous agents.

2.2. Collaborative Perception

Collaborative perception investigates how to aggregate complementary perception semantics across connected agents to improve overall system performance. Current works are mainly categorized as early [2], intermediate [1, 10, 17, 20, 28, 34, 36, 37], and late [24, 31] fusion, where intermediate fusion is more favoured due to its good trade-off between performance and transmission bandwidth. In this case, several remarkable works explore efficient collaboration mechanisms. In DiscoNet [17], the teacher-student network is trained to aggregate the shared information. V2X-ViT [36] proposes heterogeneous multi-agent self-attention to handle device heterogeneity, and Xu *et al.* [34] design an axial attention module to capture sparse spatial semantics. Furthermore, Where2comm [10] filters irrelevant semantics by quantifying the spatial importance through confidence maps. Unlike these single-frame methods, we propose an information aggregation component to extract rich context information from local historical observations. Moreover, our cross-agent collaboration mechanism outperforms the previous per-agent/location message fusion models, significantly mitigating the impact of localization errors.

3. Methodology

In this paper, we focus on developing a robust collaborative perception system to improve the ego vehicle’s perception capability. Figure 1 illustrates the overall architecture of the proposed framework. It consists of five sequential parts, where the system input accommodates different types of connected agents (*i.e.*, infrastructures and AVs).

3.1. Metadata Conversion and Feature Extraction

During the collaboration, one of the connected agents is identified as the ego agent (*i*) to build a communication graph, where other agents within the communication range serve as collaborators. Upon receiving the metadata (*e.g.*, poses and extrinsics) broadcast by the ego agent, the collaborators project their local LiDAR point clouds

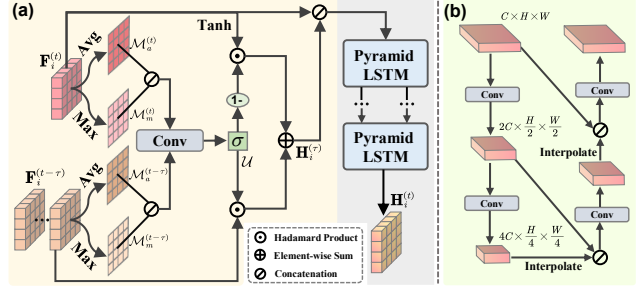


Figure 2. (a) The architecture of the proposed CIA component, including selective information filtering (*left*) and spatio-temporal feature integration (*middle*). (b) The multi-scale convolutional structure in the pyramid LSTM.

to the coordinate system of the ego agent. Similarly, the ego agent’s previous point cloud frames are synchronized to the current frame. Each agent $k \in \{1, \dots, K\}$ encodes the projected point clouds into Bird’s Eye View (BEV) features to extract local visual representations. Given the point cloud $\mathbf{X}_k^{(t)}$ of k -th agent at timestamp t , the extracted feature is $\mathbf{F}_k^{(t)} = f_{enc}(\mathbf{X}_k^{(t)}) \in \mathbb{R}^{C \times H \times W}$, where $f_{enc}(\cdot)$ is the PointPillar [14] encoder shared among all agents, and C, H, W stand for the channel, height, and width.

3.2. Context-aware Information Aggregation

Conventional 3D detection tasks [47, 49] have shown the potential of temporal information in improving perception performance. Unfortunately, current multi-agent collaborative efforts [10, 17, 20, 21, 28] invariably focus on exploring the current frame while ignoring context cues from previous frames. The single-frame solution fails to effectively detect fast-moving objects (*e.g.*, surrounding vehicles) due to the sparsity and insufficiency of point clouds. To this end, we present a Context-aware Information Aggregation (CIA) component to capture spatio-temporal representations of the ego agent’s previous frames to fuse valuable semantics. As shown in Figure 2(a), CIA consists of two core phases.

Selective Information Filtering. This phase aims to extract refined information by filtering redundant features from the target $\mathbf{F}_i^{(t)} \in \mathbb{R}^{C \times H \times W}$ and distilling meaningful characteristics from the previous $\mathbf{F}_i^{(t-\tau)} \in \mathbb{R}^{\tau \times C \times H \times W}$, where τ is the time shift. Concretely, the channel-wise average and max pooling operations $\Psi_{a/m}(\cdot)$ are first utilized to aggregate rich channel semantics, denoted as $\mathcal{M}_{a/m}^{(t)} = \Psi_{a/m}(\mathbf{F}_i^{(t)})$ and $\mathcal{M}_{a/m}^{(\tau)} = \Psi_{a/m}(\mathbf{F}_i^{(\tau)})$, where $\mathbf{F}_i^{(\tau)} = \sum_{n=1}^{\tau} \mathbf{F}_i^{(t-n)}$. These refined feature maps are combined to obtain a well-chosen spatial selection map. The procedure is expressed as follows:

$$\mathcal{U} = \sigma \cdot \aleph(\mathcal{M}_a^{(t)} \parallel \mathcal{M}_m^{(t)} + \mathcal{M}_a^{(\tau)} \parallel \mathcal{M}_m^{(\tau)}) \in \mathbb{R}^{H \times W}, \quad (1)$$

where σ is the sigmoid activation, \aleph denotes a convolutional layer for fusion, and \parallel denotes the concatenation operation.

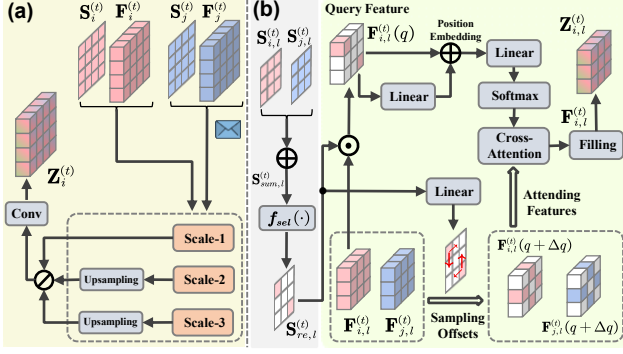


Figure 3. (a) The architecture of the proposed CCC component. (b) CCC consists of two phases at each scale: reference point proposal (middle) and deformable cross-attention module (right).

eration. Immediately, the information filtering strategy is implemented via pixel-wise attention weighting to enhance the salient historical features of $\mathbf{H}_i^{(\tau)} \in \mathbb{R}^{\tau \times C \times H \times W}$:

$$\mathbf{H}_i^{(\tau)} = (1 - \mathcal{U}) \odot \tanh(\mathbf{F}_i^{(t)}) + \mathcal{U} \odot \mathbf{F}_i^{(t-\tau)}. \quad (2)$$

Spatio-temporal Feature Integration. In order to integrate the historical prototype to improve the perception ability of the current representation, we introduce a pyramid LSTM [15] to learn the contextual dependencies of inter-frame features $\mathbf{F}_i^{(t)} \parallel \mathbf{H}_i^{(\tau)}$. The difference between the pyramid LSTM and vanilla LSTM [8] is that the initial matrix multiplication is replaced with a multi-scale convolutional structure, as shown in Figure 2(b). Note that the intrinsic long short-term memory mechanism in the LSTM is adopted, and please refer to [8] for details. In practice, multi-scale spatial features are extracted by two successive 2D convolutions among different scales, followed by batch normalization and ReLU activation. To achieve multi-level spatial semantic fusion, the downsampled features are progressively interpolated to deliver to the upsampled layers via lateral connections. The overall design facilitates learning spatio-temporal correlations in historical contexts, which benefits our multi-frame perception paradigm. Ultimately, the final hidden state in the pyramid LSTM is used as the refined context-aware feature $\mathbf{H}_i^{(t)} \in \mathbb{R}^{C \times H \times W}$.

3.3. Confidence-aware Cross-agent Collaboration

The cross-agent collaboration targets to enhance the ego agent’s visual representations by aggregating the complementary semantics from the collaborator-shared features. Existing works [10, 11, 36, 37] propose attention-based per-location feature fusion methods, which are vulnerable to localization errors and ignore the sparsity of point clouds. To tackle these issues, we implement a novel Confidence-aware Cross-agent Collaboration (CCC) component. We first obtain the confidence map $\mathbf{S}_k^{(t)}$ and binary selection matrix $\mathcal{C}_k^{(t)}$ of feature $\mathbf{F}_k^{(t)}$, $k \in \{1, \dots, K\}$ following [10].

The filtered feature shared by k -th agent to ego agent is $\mathbf{F}_k^{(t)} = \mathbf{F}_k^{(t)} \odot \mathcal{C}_k^{(t)}$. As Figure 3(a) shows, the CCC component encodes the features and confidence maps into three scales and performs feature fusion at each scale. $\mathbf{F}_{k,l}^{(t)}$ and $\mathbf{S}_{k,l}^{(t)}$ denote the feature and confidence map at l -th scale. We take two agents as an example to elaborate on the feature fusion in Figure 3(b), including the following two phases.

Reference Point Proposal. To predict the spatial locations containing meaningful object information, we obtain the element-wise summation of all confidence maps as $\mathbf{S}_{sum,l}^{(t)} = \sum_{k=1}^K \mathbf{S}_{k,l}^{(t)}$. Since the confidence map reflects the spatially critical levels, $\mathbf{S}_{sum,l}^{(t)}$ presents the potential positions of the targets within the detection range, called the reference points. Accordingly, we apply a threshold-based selection function $f_{sel}(\cdot)$ to extract the reference points $\mathbf{S}_{re,l}^{(t)}$. This design could actively guide the subsequent fusion network to concentrate on the essential spatial areas.

Deformable Cross-attention Module. We extract the ego agent’s feature $\mathbf{F}_{i,l}^{(t)}$ at the reference points as the initial query embedding and apply a linear layer to encode the locations of reference points into a position embedding. To resolve the misalignment of feature maps and acquire more robust representations against localization errors, the Deformable Cross-attention Module (DCM) aggregates the information from sampled keypoints through a deformable cross-attention layer [30]. Concretely, a linear layer is leveraged to learn the 2D spatial offset Δq of the reference point q for sampling attending keypoints at $q + \Delta q$. We simultaneously select M keypoints and extract their features as the attending features. The output of DCM at position q is:

$$\text{DCM}(q) = \sum_{a=1}^A \mathbf{W}_a \left[\sum_{k=1}^K \sum_{m=1}^M \phi(\mathbf{W}_b \mathbf{F}_{i,l}^{(t)}(q)) \mathbf{F}_{k,l}^{(t)}(q + \Delta q_m) \right], \quad (3)$$

where a indexes the attention head, $\mathbf{W}_{a/b}$ are the learnable weights, and $\phi(\cdot)$ is the softmax function. Ultimately, we derive a filling operation that fills $\text{DCM}(q)$ into the ego agent’s feature $\mathbf{F}_{i,l}^{(t)}$ based on the initial position q and outputs $\mathbf{Z}_{i,l}^{(t)}$. The output-enhanced features under the three scales are encoded into the same size and concatenated in the channel dimension. We utilize a 1×1 convolutional layer to fuse information of the three scales and get the final collaboration feature $\mathbf{Z}_i^{(t)} \in \mathbb{R}^{C \times H \times W}$.

3.4. Importance-aware Adaptive Fusion

Although previous works [17, 20, 21, 28, 36] have achieved impressive performance through ego features that aggregate collaborative information, they potentially suffer from noise interference introduced by collaborators with asynchronous measurements. A promising solution is to consider purely ego-centered characteristics, which contain the natural perception advantages of the target agent. To

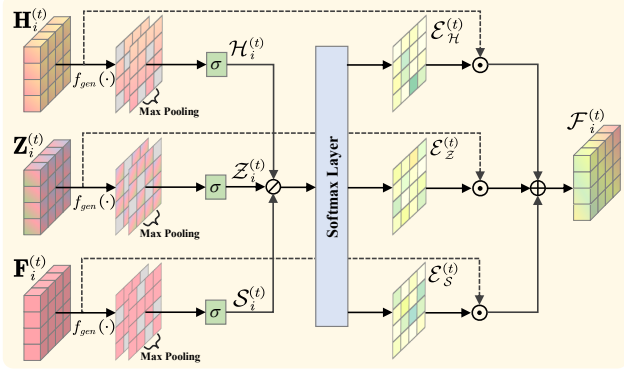


Figure 4. The architecture of the proposed IAF component.

this end, we propose an Importance-aware Adaptive Fusion (IAF) component to fuse multi-source features based on their complementary contributions. These informative features include $\mathbf{H}_i^{(t)}$, $\mathbf{Z}_i^{(t)}$, and $\mathbf{F}_i^{(t)}$. From Figure 4, the importance generator $f_{gen}(\cdot)$ is used to generate the spatial importance maps, which processes as follows:

$$\mathcal{H}_i^{(t)}, \mathcal{Z}_i^{(t)}, \mathcal{S}_i^{(t)} = \sigma \cdot \Psi_m(f_{gen}(\mathbf{H}_i^{(t)}, \mathbf{Z}_i^{(t)}, \mathbf{F}_i^{(t)})), \quad (4)$$

where $\mathcal{H}_i^{(t)}, \mathcal{Z}_i^{(t)}, \mathcal{S}_i^{(t)} \in [0, 1]^{H \times W}$. Then, these importance maps are concatenated to get the normalized attention maps across distinct features by the softmax function $\phi(\cdot)$:

$$\mathcal{E}_{\mathcal{H}}^{(t)} = \phi(\mathcal{H}_i^{(t)}) = \frac{\exp(\mathcal{H}_i^{(t)})}{\exp(\mathcal{H}_i^{(t)}) + \exp(\mathcal{Z}_i^{(t)}) + \exp(\mathcal{S}_i^{(t)})}. \quad (5)$$

A larger value in $\mathcal{E}_{\mathcal{H}}^{(t)} \in \mathbb{R}^{H \times W}$ implies that the corresponding feature position is more significant. Similarly, $\mathcal{E}_{\mathcal{Z}}^{(t)} = \phi(\mathcal{Z}_i^{(t)})$ and $\mathcal{E}_{\mathcal{S}}^{(t)} = \phi(\mathcal{S}_i^{(t)})$. Finally, we adaptively activate each feature map's perceptually critical spatial information by a weighted summation. The fused feature map $\mathcal{F}_i^{(t)} \in \mathbb{R}^{C \times H \times W}$ is obtained as follows:

$$\mathcal{F}_i^{(t)} = \mathcal{E}_{\mathcal{H}}^{(t)} \odot \mathbf{H}_i^{(t)} + \mathcal{E}_{\mathcal{Z}}^{(t)} \odot \mathbf{Z}_i^{(t)} + \mathcal{E}_{\mathcal{S}}^{(t)} \odot \mathbf{F}_i^{(t)}. \quad (6)$$

3.5. Detection Decoder and Objective Optimization

We leverage two detection decoders to decode the final fused feature $\mathcal{F}_i^{(t)}$ into the predicted outputs. The regression output is $\mathcal{Y}_{i,r}^{(t)} = f_{dec}^r(\mathcal{F}_i^{(t)}) \in \mathbb{R}^{7 \times H \times W}$, representing the position, size, and yaw angle of the bounding box at each location. The classification output reveals the confidence value of each predefined box to be a target or background, which is $\mathcal{Y}_{i,c}^{(t)} = f_{dec}^c(\mathcal{F}_i^{(t)}) \in \mathbb{R}^{2 \times H \times W}$. Following existing work [14], we adopt the smooth L_1 loss for regressing the bounding boxes and a focal loss [19] for classification.

4. Experiments

4.1. Datasets and Evaluation Metrics

Datasets. We conduct comprehensive experiments on three benchmark datasets for collaborative perception tasks, in-

cluding DAIR-V2X [46], V2XSet [36], and OPV2V [37]. These datasets provide annotated sensor data from multiple connected agents for each sample. Concretely, **DAIR-V2X** is a large-scale real-world dataset supporting collaborative 3D object detection, containing 71,254 samples and dividing the training/validation/testing sets as 5:2:3. Each sample contains the LiDAR point clouds from a vehicle and an infrastructure. During feature extraction, we transform the point cloud into the BEV feature map of the size (200, 504, 64). **V2XSet** is a public simulation dataset for V2X collaborative perception, including 73 representative scenarios and 11,447 frames of labeled point clouds generated by Carla [4]. The training/validation/testing sets are 6,694, 1,920, and 2,833 frames, respectively. The size of each extracted feature map is (100, 384, 64). **OPV2V** is a vehicle-to-vehicle collaborative perception dataset collected by Carla and OpenCDA [32]. It includes 11,464 frames of point clouds and RGB images with 3D annotation, splitting into training/validation/testing sets of 6,764, 1,981, and 2,719 frames, respectively. The 3D point cloud is extracted as a feature map of the size (100, 384, 64).

Evaluation Metrics. Following [36], we leverage the Average Precision (AP) at Intersection-over-Union (IoU) thresholds of 0.5 and 0.7 to evaluate the 3D object detection performance. The communication results count the message size by byte in the log scale with base 2 [10].

4.2. Implementation Details.

We build the proposed models on the Pytorch toolbox [22] and train them on four Nvidia Tesla V100 GPUs using the Adam optimizer [13]. The training settings of the DAIR-V2X, V2XSet, and OPV2V datasets follows: the batch sizes are {5, 3, 3} and the epochs are {30, 40, 40}. The initial learning rate is 2e-3 and decays every 15 epochs with a factor of 0.1. Our CIA module uses the previous 2 frames, the attention head in the CCC component is 8, and the number of keypoints is 15. To simulate the default localization and heading errors, we add Gaussian noise with a standard deviation of 0.2m and 0.2° to the initial position and orientation of the agents. The voxel resolution of the observation encoder $f_{enc}(\cdot)$ is set to 0.4m in height and width. The generator $f_{gen}(\cdot)$ follows a detection decoder from [14], and the summed confidence map $\mathcal{S}_{sum,l}^{(t)}$ is filtered using a threshold of 0.1 in the function $f_{sel}(\cdot)$. The decoders $\{f_{dec}^r(\cdot), f_{dec}^c(\cdot)\}$ are derived with two 1×1 convolutional layers. We provide results on the validation set of DAIR-V2X and the testing set of V2XSet and OPV2V.

4.3. Quantitative Evaluation

Detection Performance Comparison. Table 1 shows the 3D detection performance comparison results on the three datasets. We compare the following methods. No Fusion as a baseline considers only the point cloud data of ego vehi-




Figure 5. Collaborative perception performance comparison of SCOPE and Where2comm [10] on the DAIR-V2X, V2XSet and OPV2V datasets with varying communication volume.




Figure 6. Robustness to the localization error on the DAIR-V2X, V2XSet, and OPV2V datasets.

Table 1. Performance comparison on the DAIR-V2X, V2XSet, and OPV2V datasets. The results are reported in AP@0.5/0.7.

| Model | DAIR-V2X | V2XSet | OPV2V |
|---------------------|--------------------|--------------------|--------------------|
| | AP@0.5/0.7 | AP@0.5/0.7 | AP@0.5/0.7 |
| No Fusion | 50.03/43.57 | 60.60/40.20 | 68.71/48.66 |
| Late Fusion | 53.12/37.88 | 66.79/50.95 | 82.24/65.78 |
| Early Fusion | 61.74/46.53 | 77.39/50.45 | 81.30/68.69 |
| When2com [20] | 51.12/36.17 | 70.16/53.72 | 77.85/62.40 |
| V2VNet [28] | 56.01/42.25 | 81.80/61.35 | 82.79/70.31 |
| AttFuse [37] | 53.79/42.61 | 76.27/57.93 | 83.21/70.09 |
| V2X-ViT [36] | 54.26/43.35 | 85.13/68.67 | 86.72/74.94 |
| DiscoNet [17] | 54.29/44.88 | 82.18/63.73 | 87.38/73.19 |
| CoBEVT [34] | 54.82/43.95 | 83.01/62.67 | 87.40/74.35 |
| Where2comm [10] | 63.71/48.93 | 85.78/72.42 | 88.07/75.06 |
| SCOPE (Ours) | 65.18/49.89 | 87.52/75.05 | 89.71/80.62 |

cles without collaboration. Late Fusion allows agents to exchange detected results and utilizes non-maximum suppression to produce the final outputs. Early Fusion means the raw LiDAR point clouds are shared to complete the perception view. Meanwhile, the existing state-of-the-art (SOTA) models are fully considered, including When2com [20], V2VNet [28], AttFuse [37], V2X-ViT [36], DiscoNet [17], CoBEVT [34], and Where2comm [10]. Intuitively, the proposed SCOPE outperforms No Fusion, Late Fusion, and Early Fusion by large margins across all datasets, demonstrating the superiority of our perception paradigm. Moreover, our method significantly outperforms previous SOTA models in simulated and real-world (DAIR-V2X) scenarios. In particular, the SOTA performance is improved by 3.63% and 7.41% on the V2XSet and OPV2V datasets separately

in AP@0.7. Compared to single-frame and collaboration-only based efforts [10, 17, 20, 28, 36], SCOPE proves the effectiveness and rationality of simultaneously considering the temporal context, collaboration, and ego characteristics.

Comparison of Communication Volume. Evaluating performance at varying bandwidths plays an essential role in achieving pragmatic perception. To this purpose, Figure 5 shows the perception performance of our SCOPE and Where2comm [10] with various communication volumes. Note that we do not consider communication methods that use any extra model/data/feature compression for a fair comparison. We have the following observations: (i) SCOPE achieves better performance than previous models [17, 20, 28, 36] on all datasets with less communication volume. (ii) Compared to Where2comm [10], our method obtains a superior perception-communication trade-off across all the communication bandwidth choices.

Robustness to Localization and Heading Errors. Following the noise settings from [36], we perform experiments to evaluate the robustness against natural localization noises of existing methods. The noises are sampled from a Gaussian distribution with standard deviation $\sigma_{xyz} \in [0, 0.5]$ m, $\sigma_{heading} \in [0^\circ, 1.0^\circ]$. From Figures 6 and 7, the overall collaborative perception performance consistently deteriorates as errors continue to increase. For instance, When2com and V2VNet fail even worse than No Fusion on the DAIR-V2X when the localization and heading errors are over 0.2 m and 0.4° , respectively. Noticeably, our method consistently surpasses the previous SOTA models under all noise levels, clearly displaying the robustness of SCOPE against




Figure 7. Robustness to the heading error on the DAIR-V2X, V2XSet, and OPV2V datasets.

Table 2. Ablation study results of the proposed core components on all the datasets. **PL**: pyramid LSTM; **SIF**: selective information filtering; **DCM**: deformable cross-attention module; **RPP**: reference point proposal; **IAF**: importance-aware adaptive fusion.

| PL | SIF | DCM | RPP | IAF | DAIR-V2X AP@0.5/0.7 | V2XSet AP@0.5/0.7 | OPV2V AP@0.5/0.7 |
|----|-----|-----|-----|-----|------------------------|----------------------|---------------------|
| | | | | | 60.31/40.35 | 76.04/62.41 | 77.22/65.51 |
| ✓ | | | | | 61.75/44.58 | 80.36/65.74 | 81.19/69.02 |
| ✓ | ✓ | | | | 62.48/45.62 | 81.88/67.42 | 82.34/71.23 |
| ✓ | ✓ | ✓ | | | 64.06/47.87 | 85.47/72.29 | 86.96/76.84 |
| ✓ | ✓ | ✓ | ✓ | | 64.74/48.45 | 86.33/73.27 | 88.01/78.57 |
| ✓ | ✓ | ✓ | ✓ | ✓ | 65.18/49.89 | 87.52/75.05 | 89.71/80.62 |

pose errors. The reasonable explanations are: (i) the deformable cross-attention module facilitates capturing perceptually critical information across agents; (ii) the proposed multi-scale feature interaction alleviates the feature map misalignment about collaborative agents due to noises.

4.4. Ablation Studies

We perform thorough ablation studies on all datasets to understand the necessity of the different designs. Tables 2 and 3 show the following observations.

Importance of Core Components. We investigate the contribution of each core component in SCOPE. A base version (first row of Table 2) is provided as the gain reference. In the base model, we use the same fusion pattern to aggregate context information, collaborator-shared features, and multi-source features, where features are concatenated on the channel dimension and fused by a 1×1 convolutional network. As Table 2 demonstrates, we progressively add (1) Pyramid LSTM, (2) SIF, (3) DCM, (4) RPP, and (5) IAF and present the corresponding detection precision. The consistently increased detection results over the three datasets reveal the effectiveness of each proposed component, which all provide valuable performance gains.

Effect of Historical Frame Number. As shown in the top half of Table 3, we evaluate the effect of integrating different degrees of the temporal context from the ego agent. The interesting finding is that choosing a suitable historical frame count (*i.e.*, 2 frames) is beneficial in achieving better performance gains. The potential reason is that adding excessive historical frames may lead to accumulated errors under the noise setting and degrade the performance.

Table 3. Ablation study results of candidate designs and strategies on all the datasets. “w/” and “w/o” mean the with and without.

| Designs/Strategies | DAIR-V2X AP@0.5/0.7 | V2XSet AP@0.5/0.7 | OPV2V AP@0.5/0.7 |
|-----------------------------------|------------------------|----------------------|---------------------|
| Full Model | 65.18/49.89 | 87.52/75.05 | 89.71/80.62 |
| Effect of Historical Frame Number | | | |
| 1 Frame | 64.28/49.16 | 86.16/74.21 | 88.85/78.49 |
| 2 Frames (Default) | 65.18/49.89 | 87.52/75.05 | 89.71/80.62 |
| 3 Frames | 65.03/49.24 | 86.99/74.65 | 89.67/80.45 |
| Effect of Keypoint Number | | | |
| 9 Keypoints | 64.36/49.22 | 86.34/74.07 | 88.32/77.94 |
| 15 Keypoints (Default) | 65.18/49.89 | 87.52/75.05 | 89.71/80.62 |
| 21 Keypoints | 64.72/49.40 | 87.12/74.63 | 89.08/78.38 |
| Importance of Distinct Features | | | |
| w/o Ego-agent Feature | 64.96/49.34 | 86.47/74.24 | 88.63/79.38 |
| w/o Context Feature | 64.14/48.81 | 85.93/73.88 | 88.26/77.18 |
| w/o Collaboration Feature | 52.35/44.46 | 63.91/45.34 | 71.36/50.94 |
| Effect of Fusion Strategies | | | |
| w/ Summation Fusion | 63.65/48.91 | 85.81/72.64 | 88.05/76.13 |
| w/ Maximum Fusion | 64.92/49.27 | 86.28/73.61 | 87.95/78.17 |
| w/ Average Fusion | 64.85/48.96 | 86.55/73.46 | 88.24/78.73 |

Effect of Keypoint Number. We investigate the rationality of the keypoint number in the deformable cross-attention. Noticeably, the DCM component with 15 keypoints achieves better detection performance than other settings. This result shows that SCOPE mitigates the impact of localization errors and fully extracts essential spatial semantics by setting a reasonable keypoint number.

Importance of Distinct Features. We explore the importance of distinct representations by removing the ego-agent, context, and collaboration features (*i.e.*, $F_i^{(t)}$, $H_i^{(t)}$, $Z_i^{(t)}$). There is a significant drop in the model’s performance when the collaboration feature is removed, indicating that multi-agent information sharing dominates various collaborative perception tasks. Furthermore, the consistently decreased performance suggests that ego-centered characteristics (*i.e.*, pure ego-agent feature and corresponding context information) provide an indispensable contribution.

Effect of Fusion Strategies. Finally, we demonstrate the effectiveness of the proposed adaptive fusion strategy by providing candidate fusion mechanisms. Summation Fusion is a direct pixel-wise addition operation. Maximum and Average Fusion select elements at corresponding po-

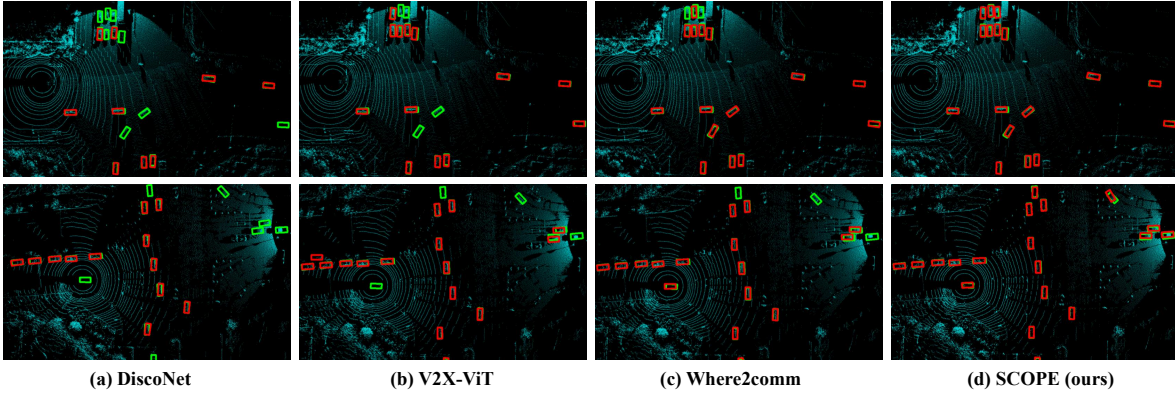


Figure 8. Qualitative comparison results in real-world scenarios from the DAIR-V2X dataset. Green and red boxes denote ground truths and detection results, respectively. Compared to the previous SOTA models, our method achieves more accurate detection results.




Figure 9. Visualization of learned multi-scale deformable cross-attention. Green cross marker denotes the proposed reference point. The sampled keypoints are represented by the brown dots whose colors reflect the corresponding attention weights.

sitions among features to retain the maximum or average, respectively, and then reshape them into a fused representation. The poor results in Table 3 show that the above mechanisms fail to tackle the complex multi-source feature fusion challenge. In contrast, the superiority of our IAF module is reflected in considering the complementary contributions of each feature at different positions, adaptively assigning higher importance to spatially beneficial semantics.

4.5. Qualitative Evaluation

Visualization of Detection Results. To intuitively illustrate the perception capability of different models, we provide a detection visualization comparison in Figure 8 under two challenging scenarios. SCOPE achieves more accurate and robust detection results than previous SOTA methods, where the predicted bounding boxes are well aligned to ground truths. In particular, DiscoNet [17], V2X-ViT [36], and Where2comm [10] fail to identify fast-moving objects (*i.e.*, more ground truth bounding boxes are not matched). In contrast, our method achieves holistic perception across heterogeneous agents by effectively aggregating the salient information in the spatio-temporal domains.

Visualization of Attention Map. We visualize the multi-scale deformable cross-attention in the CCC component. As Figure 9 shows, our tailored module can dynamically aggregate the semantics of multiple sampled keypoints at different scales to enhance the visual representation of the

reference point. Concretely, the deformable cross-attention learns to extract object-related valuable perception cues at lower scales by sampling keypoints inside or around targets. In comparison, it reasonably samples long-range features at higher scales to capture spatially critical information.

5. Conclusion

In this paper, we present SCOPE, a learning-based framework to address existing multi-agent collaborative perception challenges in an end-to-end manner. Our approach is the first time considering the temporal semantics from ego agents to capture valuable context cues through the context-aware component. Moreover, two spatial information aggregation components are introduced to achieve comprehensive information interaction among agents and refined feature fusion of ego agents. Extensive experiments on diverse 3D detection datasets prove the superiority of SCOPE and the effectiveness of our components. The future work focuses on extending SCOPE into more complex perception scenarios.

Acknowledgements

This work is supported by the Shanghai Key Research Laboratory of NSAI, the National Natural Science Foundation of China (Grant No.62250410368), and the Nanjing First Automobile Works Grant.

References

- [1] Qi Chen, Xu Ma, Sihai Tang, Jingda Guo, Qing Yang, and Song Fu. F-cooper: Feature based cooperative perception for autonomous vehicle edge computing system using 3D point clouds. In *Proceedings of ACM/IEEE Symposium on Edge Computing*, pages 88–100, Arlington Virginia, Nov. 2019. 3
- [2] Qi Chen, Sihai Tang, Qing Yang, and Song Fu. Cooper: Cooperative perception for connected autonomous vehicles based on 3d point clouds. In *2019 IEEE 39th International Conference on Distributed Computing Systems (ICDCS)*, pages 514–524, 2019. 3
- [3] Abhishek Das, Théophile Gervet, Joshua Romoff, Dhruv Batra, Devi Parikh, Mike Rabbat, and Joelle Pineau. Tarmac: Targeted multi-agent communication. In *International Conference on Machine Learning (ICML)*, pages 1538–1546, 2019. 3
- [4] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. Carla: An open urban driving simulator. In *Conference on Robot Learning (CoRL)*, pages 1–16, 2017. 5
- [5] Yangtao Du, Dingkan Yang, Peng Zhai, Mingchen Li, and Lihua Zhang. Learning associative representation for facial expression recognition. In *Proceedings of the IEEE International Conference on Image Processing (ICIP)*, pages 889–893, 2021. 1
- [6] Khaled El Madawi, Hazem Rashed, Ahmad El Sallab, Omar Nasr, Hanan Kamel, and Senthil Yogamani. Rgb and lidar fusion based 3d semantic segmentation for autonomous driving. In *Proceedings of the IEEE Intelligent Transportation Systems Conference (ITSC)*, pages 7–12, 2019. 1
- [7] Jakob Foerster, Ioannis Alexandros Assael, Nando De Freitas, and Shimon Whiteson. Learning to communicate with deep multi-agent reinforcement learning. *Advances in Neural Information Processing Systems (NIPS)*, 29, 2016. 2
- [8] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997. 4
- [9] Yedid Hoshen. Vain: Attentional multi-agent predictive modeling. *Advances in Neural Information Processing Systems (NIPS)*, 30, 2017. 3
- [10] Yue Hu, Shaoheng Fang, Zixing Lei, Yiqi Zhong, and Siheng Chen. Where2comm: Communication-efficient collaborative perception via spatial confidence maps. In *Advances in Neural Information Processing Systems (NIPS)*, 2022. 1, 3, 4, 5, 6, 8
- [11] Yue Hu, Yifan Lu, Runsheng Xu, Weidi Xie, Siheng Chen, and Yanfeng Wang. Collaboration helps camera overtake lidar in 3d detection. *arXiv preprint arXiv:2303.13560*, 2023. 4
- [12] Jiechuan Jiang and Zongqing Lu. Learning attentional communication for multi-agent cooperation. *Advances in Neural Information Processing Systems (NIPS)*, 31, 2018. 3
- [13] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015. 5
- [14] Alex H Lang, Sourabh Vora, Holger Caesar, Lubing Zhou, Jiong Yang, and Oscar Beijbom. Pointpillars: Fast encoders for object detection from point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12697–12705, 2019. 1, 3, 5
- [15] Zixing Lei, Shunli Ren, Yue Hu, Wenjun Zhang, and Siheng Chen. Latency-aware collaborative perception. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 316–332, 2022. 4
- [16] Jinlong Li, Runsheng Xu, Xinyu Liu, Jin Ma, Zicheng Chi, Jiaqi Ma, and Hongkai Yu. Learning for vehicle-to-vehicle cooperative perception under lossy communication. *IEEE Transactions on Intelligent Vehicles*, 2023. 1
- [17] Yiming Li, Shunli Ren, Pengxiang Wu, Siheng Chen, Chen Feng, and Wenjun Zhang. Learning distilled collaboration graph for multi-agent perception. *Advances in Neural Information Processing Systems (NIPS)*, 34:29541–29552, 2021. 1, 3, 4, 6, 8
- [18] Ming Liang, Bin Yang, Shenlong Wang, and Raquel Urtasun. Deep continuous fusion for multi-sensor 3d object detection. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 641–656, 2018. 2
- [19] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2980–2988, 2017. 5
- [20] Yen-Cheng Liu, Junjiao Tian, Nathaniel Glaser, and Zsolt Kira. When2com: Multi-agent perception via communication graph grouping. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4106–4115, 2020. 1, 3, 4, 6
- [21] Guiyang Luo, Hui Zhang, Quan Yuan, and Jinglin Li. Complementarity-enhanced and redundancy-minimized collaboration network for multi-agent perception. In *Proceedings of the ACM International Conference on Multimedia (ACM MM)*, pages 3578–3586, 2022. 1, 3, 4
- [22] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in Neural Information Processing Systems (NIPS)*, 32, 2019. 5
- [23] Tika Ram and Khem Chand. Effect of drivers’ risk perception and perception of driving tasks on road safety attitude. *Transportation Research Part F: Traffic Psychology and Behaviour*, 42:162–176, 2016. 1
- [24] Andreas Rauch, Felix Klanner, Ralph Raschofer, and Klaus Dietmayer. Car2x-based perception in a high-level fusion architecture for cooperative perception systems. In *IEEE Intelligent Vehicles Symposium*, pages 270–275, 2012. 3
- [25] Amanpreet Singh, Tushar Jain, and Sainbayar Sukhbaatar. Learning when to communicate at scale in multi-agent cooperative and competitive tasks. *arXiv preprint arXiv:1812.09755*, 2018. 2
- [26] Sainbayar Sukhbaatar, Rob Fergus, et al. Learning multi-agent communication with backpropagation. *Advances in Neural Information Processing Systems (NIPS)*, 29, 2016. 2
- [27] Ming Tan. Multi-agent reinforcement learning: Independent vs. cooperative agents. In *International Conference on Machine Learning (ICML)*, pages 330–337, 1993. 2

- [28] Tsun-Hsuan Wang, Sivabalan Manivasagam, Ming Liang, Bin Yang, Wenyuan Zeng, and Raquel Urtasun. V2vnet: Vehicle-to-vehicle communication for joint perception and prediction. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 605–621, 2020. [1](#), [3](#), [4](#), [6](#)
- [29] Hao Xiang, Runsheng Xu, Xin Xia, Zhaoliang Zheng, Bolei Zhou, and Jiaqi Ma. V2xp-asg: Generating adversarial scenes for vehicle-to-everything perception. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 3584–3591. IEEE, 2023. [1](#)
- [30] Zhu Xizhou, Su Weijie, Lu Lewei, Li Bin, Wang Xiaogang, and Dai Jifeng. Deformable detr: Deformable transformers for end-to-end object detection. In *International Conference on Learning Representations (ICLR)*, 2021. [4](#)
- [31] Runsheng Xu, Weizhe Chen, Hao Xiang, Xin Xia, Lantao Liu, and Jiaqi Ma. Model-agnostic multi-agent perception framework. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1471–1478. IEEE, 2023. [3](#)
- [32] Runsheng Xu, Yi Guo, Xu Han, Xin Xia, Hao Xiang, and Jiaqi Ma. Opencda: an open cooperative driving automation framework integrated with co-simulation. In *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*, pages 1155–1162. IEEE, 2021. [5](#)
- [33] Runsheng Xu, Jinlong Li, Xiaoyu Dong, Hongkai Yu, and Jiaqi Ma. Bridging the domain gap for multi-agent perception. *arXiv preprint arXiv:2210.08451*, 2022. [1](#)
- [34] Runsheng Xu, Zhengzhong Tu, Hao Xiang, Wei Shao, Bolei Zhou, and Jiaqi Ma. Cobevt: Cooperative bird’s eye view semantic segmentation with sparse transformers. In *Conference on Robot Learning (CoRL)*, 2022. [3](#), [6](#)
- [35] Runsheng Xu, Xin Xia, Jinlong Li, Hanzhao Li, Shuo Zhang, Zhengzhong Tu, Zonglin Meng, Hao Xiang, Xiaoyu Dong, Rui Song, et al. V2v4real: A real-world large-scale dataset for vehicle-to-vehicle cooperative perception. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13712–13722, 2023. [1](#)
- [36] Runsheng Xu, Hao Xiang, Zhengzhong Tu, Xin Xia, Ming-Hsuan Yang, and Jiaqi Ma. V2x-vit: Vehicle-to-everything cooperative perception with vision transformer. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [8](#)
- [37] Runsheng Xu, Hao Xiang, Xin Xia, Xu Han, Jinlong Li, and Jiaqi Ma. Opv2v: An open benchmark dataset and fusion pipeline for perception with vehicle-to-vehicle communication. In *Proceedings of the International Conference on Robotics and Automation (ICRA)*, pages 2583–2589, 2022. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#)
- [38] Dingkang Yang, Zhaoyu Chen, Yuzheng Wang, Shunli Wang, Mingcheng Li, Siao Liu, Xiao Zhao, Shuai Huang, Zhiyan Dong, Peng Zhai, and Lihua Zhang. Context deconfounded emotion recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 19005–19015, June 2023. [2](#)
- [39] Dingkang Yang, Shuai Huang, Haopeng Kuang, Yangtao Du, and Lihua Zhang. Disentangled representation learning for multimodal emotion recognition. In *Proceedings of the 30th ACM International Conference on Multimedia (ACM MM)*, page 1642–1651, 2022. [2](#)
- [40] Dingkang Yang, Shuai Huang, Yang Liu, and Lihua Zhang. Contextual and cross-modal interaction for multi-modal speech emotion recognition. *IEEE Signal Processing Letters*, 29:2093–2097, 2022. [2](#)
- [41] Dingkang Yang, Shuai Huang, Shunli Wang, Yang Liu, Peng Zhai, Liuzhen Su, Mingcheng Li, and Lihua Zhang. Emotion recognition for multiple context awareness. In *Proceedings of the European Conference on Computer Vision (ECCV)*, volume 13697, pages 144–162. Springer, 2022. [1](#)
- [42] Dingkang Yang, Haopeng Kuang, Shuai Huang, and Lihua Zhang. Learning modality-specific and -agnostic representations for asynchronous multimodal language sequences. In *Proceedings of the 30th ACM International Conference on Multimedia (ACM MM)*, page 1708–1717, 2022. [2](#)
- [43] Dingkang Yang, Yang Liu, Can Huang, Mingcheng Li, Xiao Zhao, Yuzheng Wang, Kun Yang, Yan Wang, Peng Zhai, and Lihua Zhang. Target and source modality co-reinforcement for emotion understanding from asynchronous multimodal sequences. *Knowledge-Based Systems*, page 110370, 2023. [1](#)
- [44] Kun Yang, Jing Liu, Dingkang Yang, Hanqi Wang, Peng Sun, Yanni Zhang, Yan Liu, and Liang Song. A novel efficient multi-view traffic-related object detection framework. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5, 2023. [2](#)
- [45] Kun Yang, Peng Sun, Jieyu Lin, Azzedine Boukerche, and Liang Song. A novel distributed task scheduling framework for supporting vehicular edge intelligence. In *IEEE International Conference on Distributed Computing Systems (ICDCS)*, pages 972–982, 2022. [2](#)
- [46] Haibao Yu, Yizhen Luo, Mao Shu, Yiyi Huo, Zebang Yang, Yifeng Shi, Zhenglong Guo, Hanyu Li, Xing Hu, Jirui Yuan, 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 (CVPR)*, pages 21361–21370, 2022. [1](#), [2](#), [5](#)
- [47] Zhenxun Yuan, Xiao Song, Lei Bai, Zhe Wang, and Wanli Ouyang. Temporal-channel transformer for 3d lidar-based video object detection for autonomous driving. *IEEE Transactions on Circuits and Systems for Video Technology*, 32(4):2068–2078, 2021. [1](#), [3](#)
- [48] Zixu Zhang and Jaime F Fisac. Safe occlusion-aware autonomous driving via game-theoretic active perception. *arXiv preprint arXiv:2105.08169*, 2021. [1](#)
- [49] Zixiang Zhou, Xiangchen Zhao, Yu Wang, Panqu Wang, and Hassan Foroosh. Centerformer: Center-based transformer for 3d object detection. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 496–513, 2022. [1](#), [3](#)