

V2X-ViT: Vehicle-to-Everything Cooperative Perception with Vision Transformer

Runsheng Xu^{1*}, Hao Xiang^{1*}, Zhengzhong Tu^{2*}, Xin Xia¹,
Ming-Hsuan Yang^{3,4}, and Jiaqi Ma¹[†]

¹ University of California, Los Angeles

² University of Texas at Austin

³ Google Research

⁴ University of California, Merced

Abstract. In this paper, we investigate the application of Vehicle-to-Everything (V2X) communication to improve the perception performance of autonomous vehicles. We present a robust cooperative perception framework with V2X communication using a novel vision Transformer. Specifically, we build a holistic attention model, namely V2X-ViT, to effectively fuse information across on-road agents (i.e., vehicles and infrastructure). V2X-ViT consists of alternating layers of heterogeneous multi-agent self-attention and multi-scale window self-attention, which captures inter-agent interaction and per-agent spatial relationships. These key modules are designed in a unified Transformer architecture to handle common V2X challenges, including asynchronous information sharing, pose errors, and heterogeneity of V2X components. To validate our approach, we create a large-scale V2X perception dataset using CARLA and OpenCDA. Extensive experimental results demonstrate that V2X-ViT sets new state-of-the-art performance for 3D object detection and achieves robust performance even under harsh, noisy environments. The code is available at <https://github.com/DerrickXuNu/v2x-vit>.

Keywords: V2X, Vehicle-to-Everything, Cooperative perception, Autonomous driving, Transformer

1 Introduction

Perceiving the complex driving environment precisely is crucial to the safety of autonomous vehicles (AVs). With recent advancements of deep learning, the robustness of single-vehicle perception systems has demonstrated significant improvement in several tasks such as semantic segmentation [32,10], depth estimation [54,11], and object detection and tracking [20,23,48,13]. Despite recent advancements, challenges remain. Single-agent perception system tends to suffer from occlusion and sparse sensor observation at a far distance, which can potentially cause catastrophic consequences [50]. The cause of such a problem is that

* Equal contribution. [†] Corresponding author: jiaqima@ucla.edu

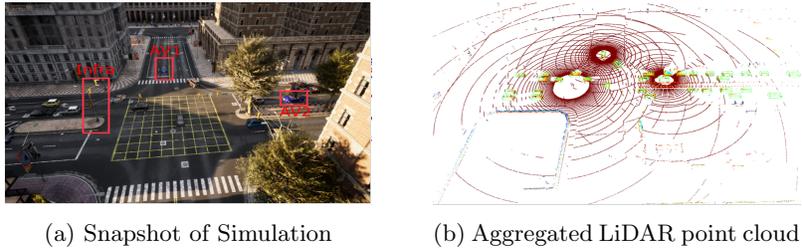


Fig. 1: **A data sample from the proposed V2XSet.** (a) A simulated scenario in CARLA where two AVs and infrastructure are located at different sides of a busy intersection. (b) The aggregated LiDAR point clouds of these three agents.

an individual vehicle can only perceive the environment from a single perspective with limited sight-of-view. To address these issues, recent studies [39, 5, 44, 4] leverage the advantages of multiple viewpoints of the same scene by investigating Vehicle-to-Vehicle (V2V) collaboration, where visual information (*e.g.*, detection outputs, raw sensory information, intermediate deep learning features, details see Sec. 2) from multiple nearby AVs are shared for a complete and accurate understanding of the environment.

Although V2V technologies have the prospect to revolutionize the mobility industry, it ignores a critical collaborator – roadside infrastructure. The presence of AVs is usually unpredictable, whereas the infrastructure can always provide supports once installed in key scenes such as intersections and crosswalks. Moreover, infrastructure equipped with sensors on an elevated position has a broader sight-of-view and potentially less occlusion. Despite these advantages, including infrastructure to deploy a robust V2X perception system is non-trivial. Unlike V2V collaboration, where all agents are homogeneous, V2X systems often involve a heterogeneous graph formed by infrastructure and AVs. The configuration discrepancies between infrastructure and vehicle sensors, such as types, noise levels, installation height, and even sensor attributes and modality, make the design of a V2X perception system challenging. Moreover, the GPS localization noises and the asynchronous sensor measurements of AVs and infrastructure can introduce inaccurate coordinate transformation and lagged sensing information. Failing to properly handle these challenges will make the system vulnerable.

In this paper, we introduce a unified fusion framework, namely V2X Vision Transformer or **V2X-ViT**, for V2X perception, that can jointly handle these challenges. Fig. 2 illustrates the entire system. AVs and infrastructure capture, encode, compress, and send intermediate visual features with each other, while the ego vehicle (*i.e.*, receiver) employs V2X-Transformer to perform information fusion for object detection. We propose two novel attention modules to accommodate V2X challenges: 1) a customized heterogeneous multi-agent self-attention module that explicitly considers agent types (vehicles and infrastructure) and their connections when performing attentive fusion; 2) a multi-scale window attention module that can handle localization errors by using multi-resolution

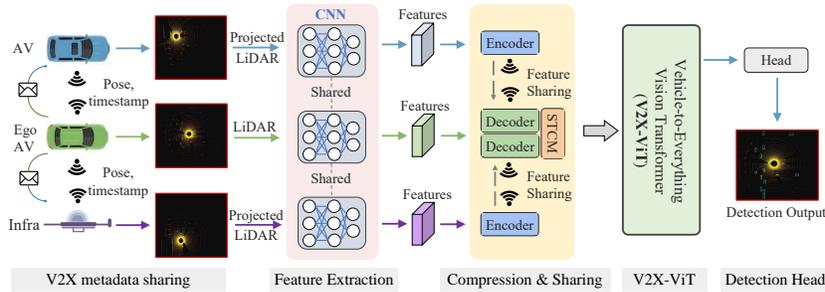


Fig. 2: **Overview of our proposed V2X perception system.** It consists of five sequential steps: V2X metadata sharing, feature extraction, compression & sharing, V2X-ViT, and the detection head. The details of each individual component are illustrated in Sec. 3.1.

windows in parallel. These two modules will adaptively fuse visual features in an iterative fashion to capture inter-agent interaction and per-agent spatial relationship, correcting the feature misalignment caused by localization error and time delay. Moreover, we also integrate a delay-aware positional encoding to further handle the time delay uncertainty. Notably, all these modules are incorporated in a single transformer that learns to address these challenges end-to-end.

To evaluate our approach, we collect a new large-scale open dataset, namely V2XSet, that explicitly considers real-world noises during V2X communication using the high-fidelity simulator CARLA [9], and a cooperative driving automation simulation tool OpenCDA. Fig. 1 shows a data sample in the collected dataset. Experiments show that our proposed V2X-ViT significantly advances the performance on V2X LiDAR-based 3D object detection, achieving a 21.2% gain of AP compared to single-agent baseline and performing favorably against leading intermediate fusion methods by at least 7.3%. Our contributions are:

- We present the first unified transformer architecture (V2X-ViT) for V2X perception, which can capture the heterogeneity nature of V2X systems with strong robustness against various noises. Moreover, the proposed model achieves state-of-the-art performance on the challenging cooperative detection task.
- We propose a novel heterogeneous multi-agent attention module (HMSA) tailored for adaptive information fusion between heterogeneous agents.
- We present a new multi-scale window attention module (MSWin) that simultaneously captures local and global spatial feature interactions in parallel.
- We construct V2XSet, a new large-scale open simulation dataset for V2X perception, which explicitly accounts for imperfect real-world conditions.

2 Related work

V2X perception. Cooperative perception studies how to efficiently fuse visual cues from neighboring agents. Based on its message sharing strategy, it can

be divided into 3 categories: 1) early fusion [5] where raw data is shared and gathered to form a holistic view, 2) intermediate fusion [39,44,36,4] where intermediate neural features are extracted based on each agent’s observation and then transmitted, and 3) late fusion [28,29] where detection outputs (*e.g.*, 3D bounding box position, confidence score) are circulated. As early fusion usually requires large transmission bandwidth and late fusion fails to provide valuable scenario context [39], intermediate fusion has attracted increasing attention because of its good balance between accuracy and transmission bandwidth. Several intermediate fusion methods have been proposed for V2V perception recently. OPV2V [44] implements a single-head self-attention module to fuse features, while F-Cooper employs *maxout* [14] fusion operation. V2VNet [39] proposes a spatial-aware message passing mechanism to jointly reason detection and prediction. To attenuate outlier messages, [36] regresses vehicles’ localization errors with consistent pose constraints. DiscoNet [21] leverages knowledge distillation to enhance training by constraining the corresponding features to the ones from the network for early fusion. However, intermediate fusion for V2X is still in its infancy. Most V2X methods explored late fusion strategies to aggregate information from infrastructure and vehicles. For example, a late fusion two-level Kalman filter is proposed by [26] for roadside infrastructure failure conditions. Xiangmo *et al.* [51] propose fusing the lane mark detection from infrastructure and vehicle sensors, leveraging Dempster-Shafer theory to model the uncertainty.

LiDAR-based 3D object detection. Numerous methods have been explored to extract features from raw points, voxels, bird-eye-view (BEV) images, and their mixtures. PointRCNN [31] proposes a two-stage strategy based on raw point clouds, which learns rough estimation in the first stage and then refines it with semantic attributes. The authors of [53,45] propose to split the space into voxels and produce features per voxel. Despite having high accuracy, their inference speed and memory consumption are difficult to optimize due to reliance on 3D convolutions. To avoid computationally expensive 3D convolutions, [20,46] propose an efficient BEV representation. To satisfy both computational and flexible receptive field requirements, [30,52,47] combine voxel-based and point-based approaches to detect 3D objects.

Transformers in vision. The Transformer [38] is first proposed for machine translation [38], where multi-head self-attention and feed-forward layers are stacked to capture long-range interactions between words. Dosovitskiy *et al.* [8] present a Vision Transformer (ViT) for image recognition by regarding image patches as visual words and directly applying self-attention. The full self-attention in ViT [38,8,12], despite having global interaction, suffers from heavy computational complexity and does not scale to long-range sequences or high-resolution images. To ameliorate this issue, numerous methods have introduced locality into self-attention, such as Swin [25], CSwin [7], Twins [6], window [40,34], and sparse attention [37,35,43]. A hierarchical architecture is usually adopted to progressively increase the receptive fields for capturing longer dependencies.

While these vision transformers have proven efficient in modeling homogeneous structured data, their efficacy to represent heterogeneous graphs has been

less studied. One of the developments related to our work is the heterogeneous graph transformer (HGT) [16]. HGT was originally designed for web-scale Open Academic Graph where the nodes are text and attributes. Inspired by HGT, we build a customized heterogeneous multi-head self-attention module tailored for graph attribute-aware multi-agent 3D visual feature fusion, which is able to capture the heterogeneity of V2X systems.

3 Methodology

In this paper, we consider V2X perception as a heterogeneous multi-agent perception system, where different types of agents (*i.e.*, smart infrastructure and AVs) perceive the surrounding environment and communicate with each other. To simulate real-world scenarios, we assume that all the agents have imperfect localization and time delay exists during feature transmission. Given this, our goal is to develop a robust fusion system to enhance the vehicle’s perception capability and handle these aforementioned challenges in a unified end-to-end fashion. The overall architecture of our framework is illustrated in Fig. 2, which includes five major components: 1) metadata sharing, 2) feature extraction, 3) compression and sharing, 4) V2X vision Transformer, and 5) a detection head.

3.1 Main architecture design

V2X metadata sharing. During the early stage of collaboration, every agent $i \in \{1 \dots N\}$ within the communication networks shares metadata such as poses, extrinsics, and agent type $c_i \in \{I, V\}$ (meaning infrastructure or vehicle) with each other. We select one of the connected AVs as the ego vehicle (e) to construct a V2X graph around it where the nodes are either AVs or infrastructure and the edges represent directional V2X communication channels. To be more specific, we assume the transmission of metadata is well-synchronized, which means each agent i can receive ego pose $x_e^{t_i}$ at the time t_i . Upon receiving the pose of the ego vehicle, all the other connected agents nearby will project their own LiDAR point clouds to the ego-vehicle’s coordinate frame before feature extraction.

Feature extraction. We leverage the anchor-based PointPillar method [20] to extract visual features from point clouds because of its low inference latency and optimized memory usage [44]. The raw point clouds will be converted to a stacked pillar tensor, then scattered to a 2D pseudo-image and fed to the PointPillar backbone. The backbone extracts informative feature maps $\mathbf{F}_i^{t_i} \in \mathbb{R}^{H \times W \times C}$, denoting agent i ’s feature at time t_i with height H , width W , and channels C .

Compression and sharing. To reduce the required transmission bandwidth, we utilize a series of 1×1 convolutions to progressively compress the feature maps along the channel dimension. The compressed features with the size (H, W, C') (where $C' \ll C$) are then transmitted to the ego vehicle (e), on which the features are projected back to (H, W, C) using 1×1 convolutions.

There exists an inevitable time gap between the time when the LiDAR data is captured by connected agents and when the extracted features are received by

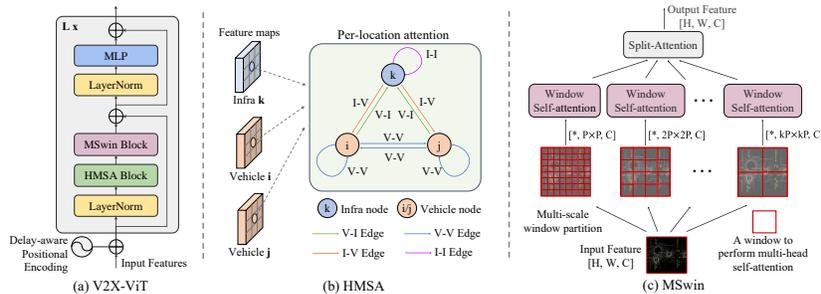


Fig. 3: **V2X-ViT architecture.** (a) The architecture of our proposed V2X-ViT model. (b) Heterogeneous multi-agent self-attention (HMSA) presented in Sec. 3.2. (c) Multi-scale window attention module (MSwin) illustrated in Sec. 3.2.

the ego vehicle (details see appendix). Thus, features collected from surrounding agents are often temporally misaligned with the features captured on the ego vehicle. To correct this delay-induced global spatial misalignment, we need to transform (*i.e.*, rotate and translate) the received features to the current ego-vehicle’s pose. Thus, we leverage a spatial-temporal correction module (STCM), which employs a differential transformation and sampling operator Γ_{ξ} to spatially warp the feature maps [17]. An ROI mask is also calculated to prevent the network from paying attention to the padded zeros caused by the spatial warp. **V2X-ViT.** The intermediate features $\mathbf{H}_i = \Gamma_{\xi}(\mathbf{F}_i^{t_i}) \in \mathbb{R}^{H \times W \times C}$ aggregated from connected agents are fed into the major component of our framework *i.e.*, V2X-ViT to conduct an iterative inter-agent and intra-agent feature fusion using self-attention mechanisms. We maintain the feature maps in the same level of high resolution throughout the entire Transformer as we have observed that the absence of high-definition features greatly harms the objection detection performance. The details of our proposed V2X-ViT will be unfolded in Sec. 3.2. **Detection head.** After receiving the final fused feature maps, we apply two 1×1 convolution layers for box regression and classification. The regression output is $(x, y, z, w, l, h, \theta)$, denoting the position, size, and yaw angle of the predefined anchor boxes, respectively. The classification output is the confidence score of being an object or background for each anchor box. We use the smooth ℓ_1 loss for regression and a focal loss [24] for classification.

3.2 V2X-Vision Transformer

Our goal is to design a customized vision Transformer that can jointly handle the common V2X challenges. Firstly, to effectively capture the heterogeneous graph representation between infrastructure and AVs, we build a heterogeneous multi-agent self-attention module that learns different relationships based on node and edge types. Moreover, we propose a novel spatial attention module, namely multi-scale window attention (MSwin), that captures long-range interactions at

various scales. MSwin uses multiple window sizes to aggregate spatial information, which greatly improves the detection robustness against localization errors. Lastly, these two attention modules are integrated into a single V2X-ViT block in a factorized manner (illustrated in Fig. 3a), enabling us to maintain high-resolution features throughout the entire process. We stack a series of V2X-ViT blocks to iteratively learn inter-agent interaction and per-agent spatial attention, leading to a robust aggregated feature representation for detection.

Heterogeneous multi-agent self-attention The sensor measurements captured by infrastructure and AVs possibly have distinct characteristics. The infrastructure’s LiDAR is often installed at a higher position with less occlusion and different view angles. In addition, the sensors may have different levels of sensor noise due to maintenance frequency, hardware quality *etc.* To encode this heterogeneity, we build a novel heterogeneous multi-agent self-attention (HSMA) where we attach types to both nodes and edges in the directed graph. To simplify the graph structure, we assume the sensor setups among the same category of agents are identical. As shown in Fig. 3b, we have two types of nodes and four types of edges, *i.e.*, node type $c_i \in \{I, V\}$ and edge type $\phi(e_{ij}) \in \{V-V, V-I, I-V, I-I\}$. Note that unlike traditional attention where the node features are treated as a vector, we only reason the interaction of features *in the same spatial position* from different agents to preserve spatial cues. Formally, HSMA is expressed as:

$$\mathbf{H}_i = \text{Dense}_{c_i} \left(\mathbf{ATT}(i, j) \cdot \mathbf{MSG}(i, j) \right)_{\forall j \in N(i)} \quad (1)$$

which contains 3 operators: a linear aggregator Dense_{c_i} , attention weights estimator \mathbf{ATT} , and message aggregator \mathbf{MSG} . The Dense is a set of linear projectors indexed by the node type c_i , aggregating multi-head information. \mathbf{ATT} calculates the importance weights between pairs of nodes conditioned on the associated node and edge types:

$$\mathbf{ATT}(i, j) = \text{softmax}_{\forall j \in N(i)} \left(\parallel_{m \in [1, h]} \text{head}_{\mathbf{ATT}}^m(i, j) \right) \quad (2)$$

$$\text{head}_{\mathbf{ATT}}^m(i, j) = \left(\mathbf{K}^m(j) \mathbf{W}_{\phi(e_{ij})}^{m, \mathbf{ATT}} \mathbf{Q}^m(i)^T \right) \frac{1}{\sqrt{C}} \quad (3)$$

$$\mathbf{K}^m(j) = \text{Dense}_{c_j}^m(\mathbf{H}_j) \quad (4)$$

$$\mathbf{Q}^m(i) = \text{Dense}_{c_i}^m(\mathbf{H}_i) \quad (5)$$

where \parallel denotes concatenation, m is the current head number and h is the total number of heads. Notice that Dense here is indexed by both node type $c_{i/j}$, and head number m . The linear layers in \mathbf{K} and \mathbf{Q} have distinct parameters. To incorporate the semantic meaning of edges, we calculate the dot product between Query and Key vectors weighted by a matrix $\mathbf{W}_{\phi(e_{ij})}^{m, \mathbf{ATT}} \in \mathbb{R}^{C \times C}$. Similarly, when parsing messages from the neighboring agent, we embed infrastructure and vehicle’s features separately via $\text{Dense}_{c_j}^m$. A matrix $\mathbf{W}_{\phi(e_{ij})}^{m, \mathbf{MSG}}$ is used to project the

features based on the edge type between source node and target node:

$$\mathbf{MSG}(i, j) = \parallel_{m \in [1, h]} \text{head}_{\text{MSG}}^m(i, j) \quad (6)$$

$$\text{head}_{\text{MSG}}^m(i, j) = \text{Dense}_{c_j}^m(\mathbf{H}_j) \mathbf{W}_{\phi(e_{ij})}^{m, \text{MSG}}. \quad (7)$$

Multi-scale window attention We present a new type of attention mechanism tailored for efficient long-range spatial interaction on high-resolution detection, called multi-scale window attention (MSwin). It uses a pyramid of windows, each of which caps a different attention range, as illustrated in Fig. 3c. The usage of variable window sizes can greatly improve the detection robustness of V2X-ViT against localization errors (see ablation study in Fig. 5b). Attention performed within larger windows can capture long-range visual cues to compensate for large localization errors, whereas smaller window branches perform attention at finer scales to preserve local context. Afterward, the split-attention module [49] is used to adaptively fuse information coming from multiple branches, empowering MSwin to handle a range of pose errors. Note that MSwin is applied on each agent independently without considering any inter-agent fusion; therefore we omit the agent subscript in this subsection for simplicity.

Formally, let $\mathbf{H} \in \mathbb{R}^{H \times W \times C}$ be an input feature map of a single agent. In branch j out of k parallel branches, \mathbf{H} is partitioned using window size $P_j \times P_j$, into a tensor of shape $(\frac{H}{P_j} \times \frac{W}{P_j}, P_j \times P_j, C)$, which represents a $\frac{H}{P_j} \times \frac{W}{P_j}$ grid of non-overlapping patches each with size $P_j \times P_j$. We use h_j number of heads to improve the attention power at j -th branch. More detailed formulation can be found in Appendix. Following [25, 15], we also consider an additional relative positional encoding \mathbf{B} that acts as a bias term added to the attention map. As the relative position along each axis lies in the range $[-P_j + 1, P_j - 1]$, we take \mathbf{B} from a parameterized matrix $\hat{\mathbf{B}} \in \mathbb{R}^{(2P_j-1) \times (2P_j-1)}$.

To attain per-agent multi-range spatial relationship, each branch partitions input tensor \mathbf{H} with different window sizes *i.e.* $\{P_j\}_{j=1}^k = \{P, 2P, \dots, kP\}$. We progressively decrease the number of heads when using a larger window size to save memory usage. Finally, we fuse the features from all the branches by a Split-Attention module [49], yielding the output feature \mathbf{Y} . The complexity of the proposed MSwin is *linear* to image size HW , while enjoying long-range multi-scale receptive fields and adaptively fuses both local and (sub)-global visual hints in parallel. Notably, unlike Swin Transformer [25], our multi-scale window approach requires no masking, padding, or cyclic-shifting, making it more efficient in implementations while having larger-scale spatial interactions.

Delay-aware positional encoding Although the global misalignment is captured by the spatial warping matrix Γ_ξ , another type of local misalignment, arising from object motions during the delay-induced time lag, also needs to be considered. To encode this temporal information, we leverage an adaptive delay-aware positional encoding (DPE), composed of a linear projection and a

learnable embedding. We initialize it with sinusoid functions conditioned on time delay Δt_i and channel $c \in [1, C]$:

$$\mathbf{p}_c(\Delta t_i) = \begin{cases} \sin\left(\Delta t_i/10000\frac{2c}{C}\right), & c = 2k \\ \cos\left(\Delta t_i/10000\frac{2c}{C}\right), & c = 2k + 1 \end{cases} \quad (8)$$

A linear projection $f: \mathbb{R}^C \rightarrow \mathbb{R}^C$ will further warp the learnable embedding so it can generalize better for unseen time delay [16]. We add this projected embedding to each agents' feature \mathbf{H}_i before feeding into the Transformer so that the features are temporally aligned beforehand.

$$\text{DPE}(\Delta t_i) = f(\mathbf{p}(\Delta t_i)) \quad (9)$$

$$\mathbf{H}_i = \mathbf{H}_i + \text{DPE}(\Delta t_i) \quad (10)$$

4 Experiments

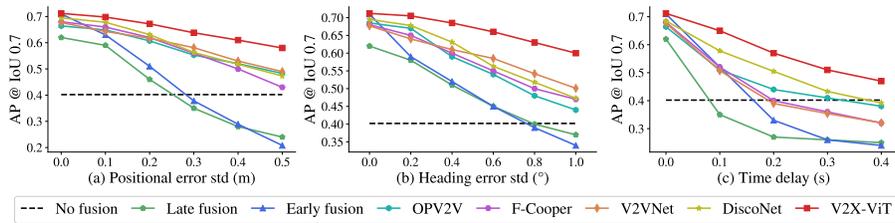
4.1 V2XSet: An open dataset for V2X cooperative perception

To the best of our knowledge, there exists no fully public V2X perception dataset suitable for investigating common V2X challenges such as localization error and transmission time delay. DAIR-V2X [2] is a large-scale real-world V2I dataset without V2V cooperation. V2X-Sim [21] is an open V2X simulated dataset but does not simulate noisy settings and only contains a single road type. OPV2V [44] contains more road types but are restricted to V2V cooperation. To this end, we collect a new large-scale dataset for V2X perception that explicitly considers these real-world noises during V2X collaboration using CARLA [9] and OpenCDA [42] together. In total, there are 11,447 frames in our dataset (33,081 samples if we count frames per agent in the same scene), and the train/validation/test splits are 6,694/1,920/2,833, respectively. Compared with existing datasets, V2XSet incorporates both V2X cooperation and realistic noise simulation. Please refer to the supplementary material for more details.

4.2 Experimental setup

The evaluation range in x and y direction are $[-140, 140]$ m and $[-40, 40]$ m respectively. We assess models under two settings: 1) *Perfect Setting*, where the pose is accurate, and everything is synchronized across agents; 2) *Noisy Setting*, where pose error and time delay are both considered. In the *Noisy Setting*, the positional and heading noises of the transmitter are drawn from a Gaussian distribution with a default standard deviation of 0.2 m and 0.2° respectively, following the real-world noise levels [41,22,1]. The time delay is set to 100 ms for all the evaluated models to have a fair comparison of their robustness against asynchronous message propagation.

Evaluation metrics. The detection performance is measured with Average Precisions (AP) at Intersection-over-Union (IoU) thresholds of 0.5 and 0.7. In

Fig. 4: **Robustness assessment** on positional and heading errors.Table 1: **3D detection performance comparison on V2XSet.** We show Average Precision (AP) at IoU=0.5, 0.7 on *Perfect* and *Noisy* settings, respectively.

Models	Perfect		Noisy	
	AP0.5	AP0.7	AP0.5	AP0.7
No Fusion	0.606	0.402	0.606	0.402
Late Fusion	0.727	0.620	0.549	0.307
Early Fusion	0.819	0.710	0.720	0.384
F-Cooper [4]	0.840	0.680	0.715	0.469
OPV2V [44]	0.807	0.664	0.709	0.487
V2VNet [39]	0.845	0.677	0.791	0.493
DiscoNet [21]	0.844	0.695	0.798	0.541
V2X-ViT (Ours)	0.882	0.712	0.836	0.614

this work, we focus on LiDAR-based vehicle detection. Vehicles hit by at least one LiDAR point from any connected agent will be included as evaluation targets.

Implementation details. During training, a random AV is selected as the ego vehicle, while during testing, we evaluate on a fixed ego vehicle for all the compared models. The communication range of each agent is set as 70 m based on [18], whereas all the agents out of this broadcasting radius of ego vehicle is ignored. For the PointPillar backbone, we set the voxel resolution to 0.4 m for both height and width. The default compression rate is 32 for all intermediate fusion methods. Our V2X-ViT has 3 encoder layers with 3 window sizes in MSwin: 4, 8, and 16. We first train the model under the *Perfect Setting*, then fine-tune it while fixing the backbone for *Noisy Setting*. We adopt Adam optimizer [19] with an initial learning rate of 10^{-3} and steadily decay it every 10 epochs using a factor of 0.1. All models are trained on Tesla V100 with 10^5 iterations.

Compared methods. We consider *No Fusion* as our baseline, which only uses ego-vehicle’s LiDAR point clouds. We also compare with *Late Fusion*, which gathers all detected outputs from agents and applies Non-maximum suppression to produce the final results, and *Early Fusion*, which directly aggregates raw LiDAR point clouds from nearby agents. For *intermediate fusion* strategy, we evaluate four state-of-the-art approaches: OPV2V [44], F-Cooper [4], V2VNet [39], and DiscoNet [21]. For a fair comparison, all the models use PointPillar as the backbone, and every compared V2V methods also receive infrastructure data, but they do not distinguish between infrastructure and vehicles.

4.3 Quantitative evaluation

Main performance comparison. Tab. 1 shows the performance comparisons on both *Perfect* and *Noisy Setting*. Under the *Perfect Setting*, all the cooperative methods significantly outperform *No Fusion* baseline. Our proposed V2X-ViT outperforms SOTA intermediate fusion methods by 3.8%/1.7% for AP@0.5/0.7. It is even higher than the ideal *Early fusion* by 0.2% AP@0.7, which receives complete raw information. Under noisy setting, when localization error and time delay are considered, the performance of *Early Fusion* and *Late Fusion* drastically drop to 38.4% and 30.7% in AP@0.7, even worse than single-agent baseline *No Fusion*. Although OPV2V [44], F-Cooper [4] V2VNet [39], and DiscoNet [21] are still higher than *No fusion*, their performance decrease by 17.7%, 21.1%, 18.4% and 15.4% in AP@0.7, respectively. In contrast, V2X-ViT performs favorably against the *No fusion* method by a large margin, *i.e.* 23% and 21.2% higher in AP@0.5 and AP@0.7. Moreover, when compared to the *Perfect Setting*, V2X-ViT only drops by less than 5% and 10% in AP@0.5 and AP@0.7 under *Noisy Setting*, demonstrating its robustness against normal V2X noises. The real-time performance of V2X-ViT is also shown in Tab. 4. The inference time of V2X-ViT is 57 ms, and by using only 1 encoder layer, V2X-ViT_S can still beat DiscoNet while reaching only 28 ms inference time, which achieves real-time performance.

Sensitivity to localization error. To assess the models’ sensitivity to pose error, we sample noises from Gaussian distribution with standard deviation $\sigma_{xyz} \in [0, 0.5]$ m, $\sigma_{heading} \in [0^\circ, 1.0^\circ]$. As Fig. 4 depicts, when the positional and heading errors stay within a normal range (*i.e.*, $\sigma_{xyz} \leq 0.2m, \sigma_{heading} \leq 0.4^\circ$ [1,22,41]), the performance of V2X-ViT only drops by less than 3%, whereas other intermediate fusion methods decrease at least 6%. Moreover, the accuracy of *Early Fusion* and *Late Fusion* degrade by nearly 20% in AP@0.7. When the noise is massive (*e.g.*, 0.5 m and 1°std), V2X-ViT can still stay around 60% detection accuracy while the performance of other methods significantly degrades, showing the robustness of V2X-ViT against pose errors.

Time delay analysis. We further investigate the impact of time delay with range [0, 400] ms. As Fig. 4c shows, the AP of *Late Fusion* drops dramatically below *No Fusion* with only 100 ms delay. *Early Fusion* and other intermediate fusion methods are relatively less sensitive, but they still drop rapidly when delay keeps increasing and are all below the baseline after 400 ms. Our V2X-ViT, in contrast, exceeds *No Fusion* by 6.8% in AP@0.7 even under 400 ms delay, which is much larger than usual transmission delay in real-world system[33]. This clearly demonstrates its great robustness against time delay.

Infrastructure vs. vehicles. To analyze the effect of infrastructure in the V2X system, we evaluate the performance between V2V, where only vehicles can share information, and V2X, where infrastructure can also transmit messages. We denote the number of agents as the total number of infrastructure and vehicles that can share information. As shown in Fig. 5a, both V2V and V2X have better performance when the number of agents increases. The V2X system has better APs compared with V2V in our collected scenes. We argue this is due to the

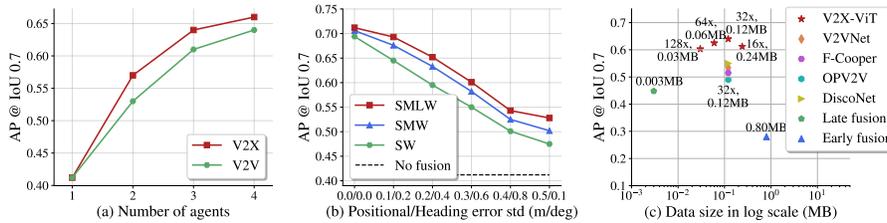


Fig. 5: **Ablation studies.** (a) AP *vs.* number of agents. (b) MSwin for localization error with window sizes: 4^2 (S), 8^2 (M), 16^2 (L). (c) AP *vs.* data size.

Table 2: **Component ablation study.** Table 3: **Effect of DPE** w.r.t. MSwin, SpAttn, HMSA, DPE represent time delay on AP@0.7.

adding i) multi-scale window attention, ii) split attention, iii) heterogeneous multi-agent self-attention, and iv) delay-aware positional encoding, respectively.

M5win	Sp-Attn	HMSA	DPE	AP0.5 / AP0.7
				0.719 / 0.478
✓				0.748 / 0.519
✓	✓			0.786 / 0.548
✓	✓	✓		0.823 / 0.601
✓	✓	✓	✓	0.836 / 0.614

Delay/Model	w/o DPE	w/ DPE
100 ms	0.639	0.650
200 ms	0.558	0.572
300 ms	0.496	0.514
400 ms	0.458	0.478

Table 4: **Inference time** measured on GPU Tesla V100.

Model	Time	AP0.7(prf/nsy)
V2X-ViT _S	28ms	0.696 / 0.591
V2X-ViT	57ms	0.712 / 0.614

better sight-of-view and less occlusion of infrastructure sensors, leading to more informative features for reasoning the environmental context.

Effects of transmission size. The size of the transmitted message can significantly affect the transmission delay, thereby affecting the detection performance. Here we study the model’s detection performance with respect to transmitted data size. The data transmission time is calculated by $t_c = f_s/v$, where f_s denotes the feature size and transmission rate v is set to 27 Mbps [3]. Following [27], we also include another system-wise asynchronous delay that follows a uniform distribution between 0 and 200 ms. See supplementary materials for more details. From Fig. 5c, we can observe: 1) Large bandwidth requirement can eliminate the advantages of cooperative perception quickly, *e.g.*, *Early Fusion* drops to 28%, indicating the necessity of compression; 2) With the default compression rate (32x), our V2X-ViT outperforms other intermediate fusion methods substantially; 3) V2X-ViT is insensitive to large compression rate. Even under a 128x compression rate, our model can still maintain high performance.

4.4 Qualitative evaluation

Detection visualization. Fig. 6 shows the detection visualization of OPV2V, V2VNet, DiscoNet, and V2X-ViT in two challenging scenarios under *Noisy set-*

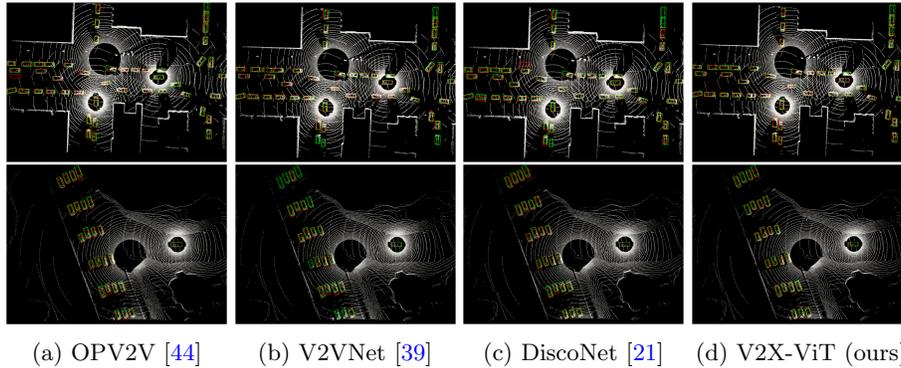


Fig. 6: **Qualitative comparison in a congested intersection and a highway entrance ramp.** Green and red 3D bounding boxes represent the ground truth and prediction respectively. Our method yields more accurate detection results. More visual examples are provided in the supplementary materials.

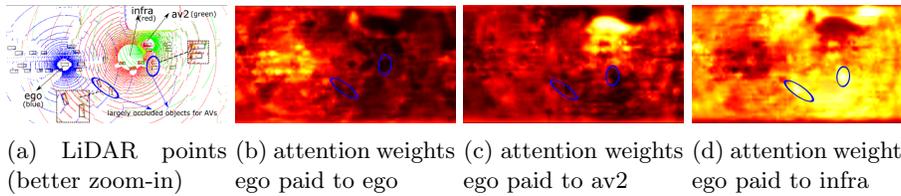


Fig. 7: **Aggregated LiDAR points and attention maps for ego.** Several objects are occluded (blue circle) from both AV’s perspectives, whereas infra can still capture rich point clouds. V2X-ViT learned to pay more attention to infra on occluded areas, shown in (d). We provide more visualizations in Appendix.

ting. Our model predicts highly accurate bounding boxes which are well-aligned with ground truths, while other approaches exhibit larger displacements. More importantly, V2X-ViT can identify more dynamic objects (more ground-truth bounding boxes have matches), which proves its capability of effectively fusing all sensing information from nearby agents. Please see Appendix for more results.

Attention map visualization. To understand the importance of infra, we also visualize the learned attention maps in Fig. 7, where brighter color means more attention ego pays. As shown in Fig. 7a, several objects are largely occluded (circled in blue) from both ego and AV2’s perspectives, whereas infrastructure can still capture rich point clouds. Therefore, V2X-ViT pays much more attention to infra on occluded areas (Fig. 7d) than other agents (Figs. 7b and 7c), demonstrating the critical role of infra on occlusions. Moreover, the attention map for infra is generally brighter than the vehicles, indicating more importance on infra seen by the trained V2X-ViT model.

4.5 Ablation studies

Contribution of major components in V2X-ViT. Now we investigate the effectiveness of individual components in V2X-ViT. Our base model is Point-Pillars with naive multi-agent self-attention fusion, which treats vehicles and infrastructure equally. We evaluate the impact of each component by progressively adding i) MSwin, ii) split attention, iii) HMSA, and iv) DPE on the *Noisy Setting*. As Tab. 2 demonstrates, all the modules are beneficial to the performance gains, while our proposed MSwin and HMSA have the most significant contributions by increasing the AP@0.7 4.1% and 6.6%, respectively.

MSwin for localization error. To validate the effectiveness of the multi-scale design in MSwin on localization error, we compare three different window configurations: i) using a single small window branch (SW), ii) using a small and a middle window (SMW), and iii) using all three window branches (SMLW). We simulate the localization error by combining different levels of positional and heading noises. From Fig. 5b, we can clearly observe that using a large and small window in parallel remarkably increased its robustness against localization error, which validates the design benefits of MSwin.

DPE Performance under delay. Tab. 3 shows that DPE can improve the performance under various time delays. The AP gain increases as delay increases.

5 Conclusion

In this paper, we propose a new vision transformer (V2X-ViT) for V2X perception. Its key components are two novel attention modules *i.e.* HMSA and MSwin, which can capture heterogeneous inter-agent interactions and multi-scale intra-agent spatial relationship. To evaluate our approach, we construct V2XSet, a new large-scale V2X perception dataset. Extensive experiments show that V2X-ViT can significantly boost cooperative 3D object detection under both perfect and noisy settings. This work focuses on LiDAR-based cooperative 3D vehicle detection, limited to single sensor modality and vehicle detection task. Our future work involves multi-sensor fusion for joint V2X perception and prediction.

Broader impacts and limitations. The proposed model can be deployed to improve the performance and robustness of autonomous driving systems by incorporating V2X communication using a novel vision Transformer. However, for models trained on simulated datasets, there are known issues on data bias and generalization ability to real-world scenarios. Furthermore, although the design choice of our communication approach (*i.e.*, project LiDAR to others at the beginning) has an advantage of accuracy (see supplementary for details), its scalability is limited. In addition, new concerns around privacy and adversarial robustness may arise during data capturing and sharing, which has not received much attention. This work facilitates future research on fairness, privacy, and robustness in visual learning systems for autonomous vehicles.

References

1. Rt3000. <https://www.oxts.com/products/rt3000-v3>, accessed: 2021-11-11 9, 11
2. Institute for AI Industry Research (AIR), T.U.: Vehicle-infrastructure cooperative autonomous driving: Dair-v2x dataset (2021) 9
3. Arena, F., Pau, G.: An overview of vehicular communications. *Future Internet* **11**(2), 27 (2019) 12
4. 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) 2, 4, 10, 11
5. Chen, Q., Tang, S., Yang, Q., Fu, S.: Cooper: Cooperative perception for connected autonomous vehicles based on 3d point clouds. In: *2019 IEEE 39th International Conference on Distributed Computing Systems (ICDCS)*. pp. 514–524. OPTorganization (2019) 2, 4
6. Chu, X., Tian, Z., Wang, Y., Zhang, B., Ren, H., Wei, X., Xia, H., Shen, C.: Twins: Revisiting the design of spatial attention in vision transformers. *arXiv preprint arXiv:2104.13840* **1**(2), 3 (2021) 4
7. Dong, X., Bao, J., Chen, D., Zhang, W., Yu, N., Yuan, L., Chen, D., Guo, B.: Cswin transformer: A general vision transformer backbone with cross-shaped windows. *arXiv preprint arXiv:2107.00652* (2021) 4
8. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929* (2020) 4
9. Dosovitskiy, A., Ros, G., Codevilla, F., Lopez, A., Koltun, V.: CARLA: An open urban driving simulator. In: *Proceedings of the 1st Annual Conference on Robot Learning*. pp. 1–16 (2017) 3, 9
10. El Madawi, K., Rashed, H., El Sallab, A., Nasr, O., Kamel, H., Yogamani, S.: Rgb and lidar fusion based 3d semantic segmentation for autonomous driving. In: *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*. pp. 7–12. OPTorganization (2019) 1
11. Fan, X., Zhou, Z., Shi, P., Xin, Y., Zhou, X.: Rafm: Recurrent atrous feature modulation for accurate monocular depth estimating. *IEEE Signal Processing Letters* pp. 1–5 (2022). <https://doi.org/10.1109/LSP.2022.3189597> 1
12. Fan, Z., Song, Z., Liu, H., Lu, Z., He, J., Du, X.: Svt-net: Super light-weight sparse voxel transformer for large scale place recognition. *AAAI* (2022) 4
13. Fan, Z., Zhu, Y., He, Y., Sun, Q., Liu, H., He, J.: Deep learning on monocular object pose detection and tracking: A comprehensive overview. *ACM Computing Surveys (CSUR)* (2021) 1
14. Goodfellow, I., Warde-Farley, D., Mirza, M., Courville, A., Bengio, Y.: Maxout networks. In: *International Conference on Machine Learning*. pp. 1319–1327. PMLR (2013) 4
15. Hu, H., Zhang, Z., Xie, Z., Lin, S.: Local relation networks for image recognition. In: *ICCV*. pp. 3464–3473 (2019) 8
16. Hu, Z., Dong, Y., Wang, K., Sun, Y.: Heterogeneous graph transformer. In: *Proceedings of The Web Conference 2020*. pp. 2704–2710 (2020) 5, 9
17. Jaderberg, M., Simonyan, K., Zisserman, A., et al.: Spatial transformer networks. In: *NeurIPS* (2015) 6

18. Kenney, J.B.: Dedicated short-range communications (dsrc) standards in the united states. *Proceedings of the IEEE* **99**(7), 1162–1182 (2011) [10](#)
19. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014) [10](#)
20. Lang, A.H., Vora, S., Caesar, H., Zhou, L., Yang, J., Beijbom, O.: Pointpillars: Fast encoders for object detection from point clouds. In: *CVPR*. pp. 12697–12705 (2019) [1](#), [4](#), [5](#)
21. Li, Y., Ren, S., Wu, P., Chen, S., Feng, C., Zhang, W.: Learning distilled collaboration graph for multi-agent perception. *NeurIPS* **34** (2021) [4](#), [9](#), [10](#), [11](#), [13](#)
22. Li, Y., Zhuang, Y., Hu, X., Gao, Z., Hu, J., Chen, L., He, Z., Pei, L., Chen, K., Wang, M., et al.: Toward location-enabled iot (le-iot): Iot positioning techniques, error sources, and error mitigation. *IEEE Internet of Things Journal* **8**(6), 4035–4062 (2020) [9](#), [11](#)
23. Liang, M., Yang, B., Wang, S., Urtasun, R.: Deep continuous fusion for multi-sensor 3d object detection. In: *ECCV*. pp. 641–656 (2018) [1](#)
24. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: *ICCV*. pp. 2980–2988 (2017) [6](#)
25. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. *arXiv preprint arXiv:2103.14030* (2021) [4](#), [8](#)
26. Mo, Y., Zhang, P., Chen, Z., Ran, B.: A method of vehicle-infrastructure cooperative perception based vehicle state information fusion using improved kalman filter. *Multimedia Tools and Applications* pp. 1–18 (2021) [4](#)
27. Rauch, A., Klanner, F., Dietmayer, K.: Analysis of v2x communication parameters for the development of a fusion architecture for cooperative perception systems. In: *2011 IEEE Intelligent Vehicles Symposium (IV)*. pp. 685–690. OPTorganization (2011) [12](#)
28. Rauch, A., Klanner, F., Rasshofer, R., Dietmayer, K.: Car2x-based perception in a high-level fusion architecture for cooperative perception systems. In: *2012 IEEE Intelligent Vehicles Symposium*. pp. 270–275. OPTorganization (2012) [4](#)
29. Rawashdeh, Z.Y., Wang, Z.: Collaborative automated driving: A machine learning-based method to enhance the accuracy of shared information. In: *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*. pp. 3961–3966. OPTorganization (2018) [4](#)
30. Shi, S., Guo, C., Jiang, L., Wang, Z., Shi, J., Wang, X., Li, H.: Pv-rcnn: Point-voxel feature set abstraction for 3d object detection. In: *CVPR*. pp. 10529–10538 (2020) [4](#)
31. Shi, S., Wang, X., Li, H.: Pointrcnn: 3d object proposal generation and detection from point cloud. In: *CVPR*. pp. 770–779 (2019) [4](#)
32. Treml, M., Arjona-Medina, J., Unterthiner, T., Durgesh, R., Friedmann, F., Schuberth, P., Mayr, A., Heusel, M., Hofmarcher, M., Widrich, M., et al.: Speeding up semantic segmentation for autonomous driving. In: *NeurIPS workshop MLITS* (2016) [1](#)
33. Tsukada, M., Oi, T., Ito, A., Hirata, M., Esaki, H.: Autoc2x: Open-source software to realize v2x cooperative perception among autonomous vehicles. In: *2020 IEEE 92nd Vehicular Technology Conference (VTC2020-Fall)*. pp. 1–6. OPTorganization (2020) [11](#)
34. Tu, Z., Talebi, H., Zhang, H., Yang, F., Milanfar, P., Bovik, A., Li, Y.: Maxim: Multi-axis mlp for image processing. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 5769–5780 (2022) [4](#)

35. Tu, Z., Talebi, H., Zhang, H., Yang, F., Milanfar, P., Bovik, A., Li, Y.: Maxvit: Multi-axis vision transformer. arXiv preprint arXiv:2204.01697 (2022) [4](#)
36. Vadivelu, N., Ren, M., Tu, J., Wang, J., Urtasun, R.: Learning to communicate and correct pose errors. arXiv preprint arXiv:2011.05289 (2020) [4](#)
37. Vaswani, A., Ramachandran, P., Srinivas, A., Parmar, N., Hechtman, B., Shlens, J.: Scaling local self-attention for parameter efficient visual backbones. In: CVPR. pp. 12894–12904 (2021) [4](#)
38. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: NeurIPS. pp. 5998–6008 (2017) [4](#)
39. Wang, T.H., Manivasagam, S., Liang, M., Yang, B., Zeng, W., Urtasun, R.: V2vnet: Vehicle-to-vehicle communication for joint perception and prediction. In: ECCV. pp. 605–621. Springer (2020) [2](#), [4](#), [10](#), [11](#), [13](#)
40. Wang, Z., Cun, X., Bao, J., Liu, J.: Uformer: A general u-shaped transformer for image restoration. arXiv preprint arXiv:2106.03106 (2021) [4](#)
41. Xia, X., Hang, P., Xu, N., Huang, Y., Xiong, L., Yu, Z.: Advancing estimation accuracy of sideslip angle by fusing vehicle kinematics and dynamics information with fuzzy logic. IEEE Transactions on Vehicular Technology (2021) [9](#), [11](#)
42. 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. OPTorganization (2021) [9](#)
43. 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) [4](#)
44. Xu, R., Xiang, H., Xia, X., Han, X., Liu, J., Ma, J.: Opv2v: An open benchmark dataset and fusion pipeline for perception with vehicle-to-vehicle communication. arXiv preprint arXiv:2109.07644 (2021) [2](#), [4](#), [5](#), [9](#), [10](#), [11](#), [13](#)
45. Yan, Y., Mao, Y., Li, B.: Second: Sparsely embedded convolutional detection. Sensors **18**(10), 3337 (2018) [4](#)
46. Yang, B., Luo, W., Urtasun, R.: Pixor: Real-time 3d object detection from point clouds. In: CVPR. pp. 7652–7660 (2018) [4](#)
47. Yang, Z., Sun, Y., Liu, S., Shen, X., Jia, J.: Std: Sparse-to-dense 3d object detector for point cloud. In: CVPR. pp. 1951–1960 (2019) [4](#)
48. Zelin, Z., Ze, W., Yueqing, Z., Boxun, L., Jiaya, J.: Tracking objects as pixel-wise distributions. arXiv preprint arXiv:2207.05518 (2022) [1](#)
49. Zhang, H., Wu, C., Zhang, Z., Zhu, Y., Lin, H., Zhang, Z., Sun, Y., He, T., Mueller, J., Manmatha, R., et al.: Resnest: Split-attention networks. arXiv preprint arXiv:2004.08955 (2020) [8](#)
50. Zhang, Z., Fisac, J.F.: Safe occlusion-aware autonomous driving via game-theoretic active perception. arXiv preprint arXiv:2105.08169 (2021) [1](#)
51. Zhao, X., Mu, K., Hui, F., Prehofer, C.: A cooperative vehicle-infrastructure based urban driving environment perception method using a ds theory-based credibility map. Optik **138**, 407–415 (2017) [4](#)
52. Zhong, Y., Zhu, M., Peng, H.: Vin: Voxel-based implicit network for joint 3d object detection and segmentation for lidars. arXiv preprint arXiv:2107.02980 (2021) [4](#)
53. Zhou, Y., Tuzel, O.: Voxnet: End-to-end learning for point cloud based 3d object detection. In: CVPR. pp. 4490–4499 (2018) [4](#)
54. Zhou, Z., Fan, X., Shi, P., Xin, Y.: R-msfm: Recurrent multi-scale feature modulation for monocular depth estimating. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 12777–12786 (2021) [1](#)