

CooPre: Cooperative Pretraining for V2X Cooperative Perception

Seth Z. Zhao¹ Hao Xiang¹ Chenfeng Xu² Xin Xia¹ Bolei Zhou¹ Jiaqi Ma¹
¹UCLA ²UC Berkeley

Abstract

Existing Vehicle-to-Everything (V2X) cooperative perception methods rely on accurate multi-agent 3D annotations. Nevertheless, it is time-consuming and expensive to collect and annotate real-world data, especially for V2X systems. In this paper, we present a self-supervised learning framework for V2X cooperative perception, which utilizes the vast amount of unlabeled 3D V2X data to enhance the perception performance. Specifically, multi-agent sensing information is aggregated to form a holistic view and a novel proxy task is formulated to reconstruct the LiDAR point clouds across multiple connected agents to better reason multi-agent spatial correlations. Besides, we develop a V2X bird-eye-view (BEV) guided masking strategy which effectively allows the model to pay attention to 3D features across heterogeneous V2X agents (i.e., vehicles and infrastructure) in the BEV space. Noticeably, such a masking strategy effectively pretrains the 3D encoder with a multi-agent LiDAR point cloud reconstruction objective and is compatible with mainstream cooperative perception backbones. Our approach, validated through extensive experiments on representative datasets (i.e., V2X-Real, V2V4Real, and OPV2V) and multiple state-of-the-art cooperative perception methods (i.e., AttFuse, F-Cooper, and V2X-ViT), leads to a performance boost across all V2X settings. Notably, CooPre achieves a 4% mAP improvement on V2X-Real dataset and surpasses baseline performance using only 50% of the training data, highlighting its data efficiency. Additionally, we demonstrate the framework’s powerful performance in cross-domain transferability and robustness under challenging scenarios. The code will be made publicly available at <https://github.com/sethzhao506/CooPre>.

1. Introduction

Achieving autonomy in complex and open traffic environments poses significant challenges for single-vehicle vision systems. These systems often suffer from occlusions and a limited perception range due to each vehicle’s singular viewpoint, limiting the capacity of current deep learning ap-




Figure 1. **Illustration of our CooPre framework.** Given multi-agent LiDAR input, CooPre pretrains the 3D encoder via LiDAR point cloud reconstruction. During the finetuning stage, the initialized 3D encoder will be used for 3D detection.

proaches to develop a holistic 3D representation in the interacting environment. Vehicle-to-Everything (V2X) Cooperative Perception [19, 37, 40, 47] emerges as a promising solution by providing each ego agent with a comprehensive understanding of the surrounding environment. Through collaboration among connected agents (vehicles or infrastructure), V2X facilitates the sharing of critical sensing information, thereby extending the perception range and mitigating occlusions. However, this paradigm introduces additional geometrical and topological information that the model must handle, necessitating the exploration of a robust representation that accounts for these elements. Obtaining such a representation often requires a large amount of annotated data, but the scale of current real-world V2X datasets [9, 21, 37, 42, 47, 53] is still limited compared to single-vehicle counterparts [1, 2, 8, 29]. Therefore, investigating the representation learning problem within the constraints of current V2X dataset scales is crucial.

In multi-agent perception systems, each ego agent must learn a representation that manages complex agent interaction information to achieve effective cooperation. Specifically, in V2X scenarios, the ego agent must handle different sensor configurations among various agents, which operate at different ranges and placement positions. Capturing this underlying distribution is challenging solely through the supervision of hand-labeled 3D bounding boxes. This complexity poses a challenge for the “train-from-scratch” paradigm, as its learned representation heavily depends on

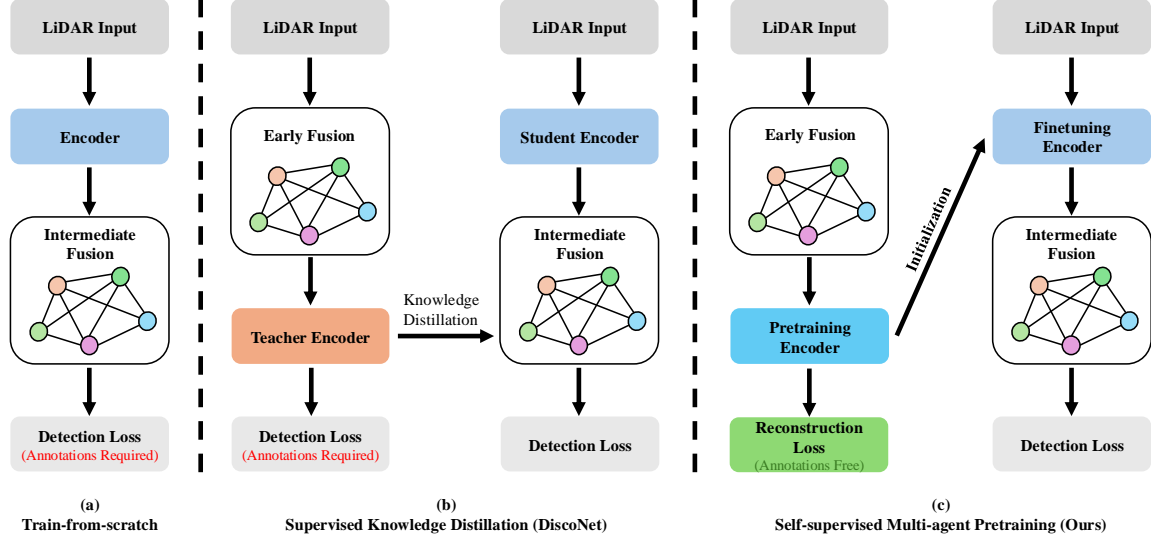


Figure 2. **Illustration of different training paradigms.** Compared with previous training paradigms, our paradigm uses multi-agent collaboration while being annotation-free.

the random initialization of model parameters and the quality and quantity of annotated 3D bounding boxes, thereby affecting overall performance. For example, as discussed in [24, 33, 36, 42, 50], such approach could be easily perturbed by synchronization or localization errors in real-world scenarios. DiscoNet [18] offers an alternative paradigm by employing a teacher model to guide the student model during training. However, a significant amount of annotations is still required to train an effective teacher model that can facilitate the learning of the student model.

To this end, we propose an effective self-supervised multi-agent **Cooperative Pretraining** framework (**CooPre**) that enables the model to learn meaningful *prior representation* of the holistic 3D environment before the perception task, as illustrated in Fig. 1. Our framework leverages the benefits of unlabeled LiDAR point cloud data transmitted from different agents, allowing the model to reconstruct the point cloud location and learn essential prior knowledge of scenarios (e.g., intersections or corridors) and LiDAR sensor distributions (e.g., range, placement position, and sparsity) of each agent from a bird-eye-view (BEV) perspective. In the enlarged perception field, this heterogeneous reconstruction task also helps mitigate issues related to sparse feature points in far-range and occlusion scenarios. In addition, this method can be seamlessly integrated with state-of-the-art cooperative perception methods to improve perception capability. Compared to the other two training paradigms in Fig. 2, our pretraining framework is annotation-free, which is the main contribution to representation learning in V2X cooperative perception.

Through extensive experiments and ablation studies across three representative V2X datasets (i.e., V2X-Real

[37], V2V4Real [42], and OPV2V [41]) and three different V2X fusion methods [4, 40, 41], we demonstrate the effectiveness of **CooPre** in V2X cooperative perception. Our findings show that: 1) **CooPre** boosts perception capabilities in scenarios involving occlusions and long-range perception, outperforming train-from-scratch and single-agent pretraining methods [22, 28]; 2) **CooPre** enables the model to learn important geometric and topological representations under heterogeneous sensor configurations, allowing better recognition of rigid-body objects like cars and trucks; 3) the learned representations from **CooPre** enhance data efficiency and demonstrate strong generalizability to unseen distributions, making them particularly advantageous for domain adaptation, especially in data scarcity scenarios. Moreover, our study underscores the critical role of well-learned 3D representations as a promising complement to task-specific design optimizations in V2X cooperative perception.

2. Related Works

2.1. Cooperative Perception

Single-vehicle systems struggle with occlusions and long-distance perception in complex traffic environments due to the limitation of the perception range of LiDAR devices. Cooperative systems, on the other hand, enhance detection performance by sharing raw data (Early Fusion), detection outputs (Late Fusion), or intermediate bird-eye-view (BEV) representations (Intermediate Fusion) among connected agents [11, 32, 41]. With the recent advancement of a variety of simulation and real-world dataset curations [9, 19, 37, 41, 42, 46, 47], many literature [4, 11, 32, 34, 39–

[41, 44] have been discussing the algorithmic designs of collaborative modes from vehicle-to-vehicle (V2V) collaboration to vehicle-to-everything (V2X) collaboration in different traffic scenarios. Noticeably, the intermediate fusion strategy has been the primary direction since it achieves the best trade-off between accuracy and bandwidth requirements. After applying a 3D encoder [14, 43, 51] to the input LiDAR feature, the intermediate fusion strategy involves projecting 3D features to BEV features where agents will perform interactions before final detection results. For example, AttFuse [41] utilizes a simple agent-wise single-head attention to fuse all features, whereas V2X-ViT [40] presents a unified vision transformer for robust multi-agent multi-scale perception. Despite their targeted designs, these methods all follow "train-from-scratch" paradigm and thus exhibit unstable performance when faced with V2X collaboration challenges, particularly due to sensor data heterogeneity issues [25, 35]. In this paper, we propose a model-agnostic pretraining framework to enhance the perception capability of these cooperative perception methods.

2.2. LiDAR-based Self-supervised Learning

Representation learning in autonomous driving [3, 5, 15, 16, 20, 22, 26, 28, 30, 38] has been prevalently investigated in single-vehicle systems. Stemming from the recent advancement of image reconstruction pretraining methods [6, 10], point cloud pretraining reconstruction methods [7, 12, 22, 26, 27, 30, 45, 48, 49, 52] are also proven effective in improving backbone model's robustness and generalizability. Recently, this approach has been applied to 3D representation learning of outdoor point clouds. Occupancy-MAE [26] applies a voxel-wise masking strategy to reconstruct masked voxels and predict occupancy. GD-MAE [45] proposes a multi-level transformer architecture and a multi-scale masking strategy with a lightweight generative decoder to recover masked patches. GeoMAE [30] formulates the pretraining target to be geometric feature predictions, such as pyramid centroid, occupancy, surface normal, and curvature of point clouds. BEV-MAE [22] uses a BEV-guided masking strategy to learn BEV feature representations. Notably, while such pretraining methods have shown great potential in developing general feature representations, their improvement is limited due to the restricted perception field caused by perception range and occlusion in single-vehicle systems. Additionally, there is limited discussion on pretraining in multi-agent systems, with the most similar work being CORE [31], which proposes a BEV map reconstruction as an auxiliary task alongside the original perception target. Our work, on the other hand, provides the well-learned BEV representations that facilitate perception performance without additional computation efforts during the finetuning stage.

3. Method

3.1. Framework Overview

Our CooPre framework, as illustrated in Fig. 3, is designed to facilitate 3D representation learning across heterogeneous V2X agents. To ensure robust cross-agent feature learning and mitigate feature sparsity issues, we propose three key design elements. First, we leverage early fusion to integrate LiDAR point cloud from collaborative agents, effectively expanding the ego agent's perception field. This design mitigates the sparse feature issue and enables the subsequent LiDAR point cloud reconstruction objective to be applied to both the ego and collaborative agents. Second, instead of reconstructing the point cloud in voxel space, we reconstruct masked information on the BEV plane. This design enhances the model's feature learning process from a BEV perspective. Third, we require the model to reconstruct masked LiDAR point clouds from both the ego and collaborative agents, allowing the 3D encoder to learn the multi-agent LiDAR sensor distribution regardless of the ego agent's perspective. Notably, our pretraining operates in BEV representation space, making it compatible with mainstream cooperative fusion methods [4, 40, 41] and widely used 3D encoder backbones [14, 43] during finetuning.

3.2. Multi-agent Reconstruction Pretraining

Early Fusion of Multi-agent LiDAR point cloud. We formulate the process of early fusion of raw LiDAR point cloud data across V2X agents as follows. Suppose we have ego agent A_{ego} and a set of N cooperative agents A_i for $i \in \{1 \dots N\}$ within the communication range, where each agent could be either Connected Autonomous Vehicle (CAV) or infrastructure. During the pretraining stage, each cooperative agent A_i shares LiDAR point clouds and metadata information, such as poses, extrinsics, and agent type, to A_{ego} . We assume the transmission of LiDAR point clouds and metadata is well-synchronized. Consequently, after projecting the point clouds of each cooperative agent to the ego agent's coordinate, the perception point cloud field of the ego agent A_{ego} includes its own LiDAR point clouds P_{ego} as well as the LiDAR point clouds P_i from each cooperative agent A_i . We refer to the collection of projected LiDAR point clouds from cooperative agents as $P_{coop} = \bigcup_i P_i$. In order for the encoder to learn invariant features from raw point cloud data from each agent, we perform data augmentations, including scaling, rotation, flip and downsampling [14]. These augmentations enhance the model's generalization capabilities during pretraining, allowing it to effectively accomplish the pretraining objectives under varying conditions and viewpoints. Thus, the collection of original and augmented multi-agent LiDAR point cloud becomes our pretraining data corpus.

V2X BEV-guided Masking Strategy. We design a V2X




Figure 3. **Detailed pipeline of our proposed CooPre framework.** During the pretraining stage, the 3D encoder would be asked to perform a cooperative reconstruction task utilizing the vast amount of unlabeled LiDAR point cloud data transmitted from different agents. The masking strategy will be applied in a BEV-guided manner to take into account the reconstruction of far-range point clouds.

BEV-guided masking strategy to enhance feature learning from a BEV perspective. For each ego agent A_{ego} , we suppose the shape of obtained BEV features from raw LiDAR input is $X \times Y \times C$. According to this shape, we define a BEV plane of size $X \times Y$ to divide up the plane and obtain BEV grids. We ensure that the size of masked BEV grids matches the resolution of BEV features obtained from 3D encoder. Each BEV grid $g_{i,j}$ would be used to determine whether the points inside the grid would be masked or not. After that, for each LiDAR point $p_k \in P_{ego} \cup P_{coop}$, we project it onto a corresponding BEV grid $g_{i,j}$ based on p_k 's x, y coordinates. With the help of projected point clouds from cooperative agents, the number of empty BEV grids largely decreases. We then randomly apply a high masking ratio towards non-empty BEV grids, where the point clouds inside the grid will be masked. Note that for a single-agent approach [22], this strategy often leads to numerous BEV grids with sparse point clouds, making self-supervision both difficult and imprecise. In contrast, multi-agent collaboration proves advantageous, as it enriches the sparsely populated BEV grids with point cloud data from other collaborative agents.

Decoder design. The decoder design is flexible and independent from downstream task since only the 3D encoder would be taken for finetuning. Unlike [45], we utilize a lightweight decoder with one convolution layer for training efficiency. The decoder would take unmasked point clouds as input and output a fixed number of LiDAR point clouds for the multi-agent reconstruction objective.

Multi-agent LiDAR point cloud Reconstruction Objective. As the ground-truth number of point clouds varies in each masked BEV grid, we utilize the Chamfer distance loss as the learning objective to measure the difference between predicted point clouds \hat{P} and ground-truth point clouds P . Specifically, suppose in the masked BEV grid $g_{i,j}$, the reconstruction loss is defined as the following equation:

$$L_{rec}(\hat{P}, P) = \frac{1}{|\hat{P}|} \sum_{\hat{p}_i \in \hat{P}} \min_{p_j \in P} \|\hat{p}_i - p_j\|_2^2 + \frac{1}{|P|} \sum_{p_i \in P} \min_{\hat{p}_j \in \hat{P}} \|p_i - \hat{p}_j\|_2^2 \quad (1)$$

We enforce the model to reconstruct the LiDAR point cloud from both ego and cooperative agent P_{ego} and P_{coop} . We

will show in the ablation study that such design is crucial for far-range feature learning across heterogeneous agents.

4. Experiments

4.1. Experiment Setting

Datasets. We evaluate our method on three datasets, namely V2X-Real [37], V2V4Real [42], and OPV2V [41]. V2X-Real is a large-scale, real-world V2X dataset that encompasses all V2X collaboration modes, including vehicle-centric (VC), infrastructure-centric (IC), vehicle-to-vehicle (V2V), and infra-to-infra (I2I). LiDAR sensors for this dataset were embodied with different sensor configurations and deployed in both intersection and corridor scenarios. Note that except for the V2V subset in the V2X-Real dataset, the rest of the subsets all exhibit such LiDAR sensor heterogeneity property. We also examine the effectiveness of our pretraining method on V2V4Real and OPV2V, two well-established benchmarks for real-world and simulated V2V cooperative perception.

Evaluation Metrics. Following previous evaluation protocols [37, 41, 42], we adopt the Average Precisions (AP) at the specified Intersection-over-Union (IoU) threshold as the metric to evaluate the detection performance. For V2X-Real [37] dataset, we evaluate the performance of different subclasses (i.e. car, pedestrian, and truck) and a final mean average precision (mAP) at the threshold of 0.3 and 0.5. For V2V4Real [42] dataset, we evaluate the performance of the vehicle class at the threshold of 0.5 and 0.7 in both synchronized and asynchronous modes. For OPV2V [41] dataset, we evaluate the performance of the vehicle class at the threshold of 0.5 and 0.7 at their two separate test sets. **Green** denotes the improvements over the corresponding backbone for all experiments.

Implementation Details. We evaluate our pretraining method using SECOND backbone [43] since it obtains higher performance and faster training speed in the real-world dataset [37, 42]. Note that our pretraining method could also be generalized to other 3D encoders such as PointPillar [14], as shown in Table 7. All experiments are conducted in one Nvidia A6000 GPU. We employ AdamW

Table 1. Performance on the test set of V2X-Real dataset.

Dataset	Method	Car		Pedestrian		Truck		Mean	
		AP0.3	AP0.5	AP0.3	AP0.5	AP0.3	AP0.5	AP0.3	AP0.5
V2X-Real VC	No Fusion	51.0	48.0	35.8	20.4	38.9	36.3	41.9	34.9
	Early Fusion	60.2	58.2	36.2	21.0	45.3	43.3	47.2	40.8
	Late Fusion	59.0	56.5	36.7	17.6	44.0	39.7	46.5	37.9
	AttFuse [41]	67.7	65.8	45.0	27.6	55.5	52.0	56.1	48.5
	F-Cooper [4]	64.9	62.7	45.5	26.7	39.8	35.1	50.1	41.5
	V2X-ViT [40]	63.3	61.2	42.2	24.8	38.4	36.0	48.0	40.0
	CooPre (AttFuse)	71.5 ^{+3.8}	70.2 ^{+4.4}	46.9 ^{+1.9}	28.0 ^{+0.4}	61.9 ^{+6.4}	58.3 ^{+6.3}	60.1 ^{+4.0}	52.2 ^{+3.7}
	CooPre (F-Cooper)	67.0 ^{+2.1}	65.7 ^{+3.0}	43.6 ^{-1.9}	25.4 ^{-1.3}	49.3 ^{+9.5}	44.9 ^{+9.8}	53.3 ^{+3.2}	45.3 ^{+3.8}
	CooPre (V2X-ViT)	66.1 ^{+2.8}	64.5 ^{+3.3}	45.3 ^{+3.1}	27.7 ^{+2.9}	43.5 ^{+5.1}	40.1 ^{+4.1}	51.7 ^{+3.7}	44.1 ^{+4.1}
V2X-Real V2V	No Fusion	49.8	47.7	43.6	28.7	38.3	37.0	43.9	37.8
	Early Fusion	62.5	60.0	43.3	28.1	47.3	44.7	51.0	44.3
	Late Fusion	51.6	50.0	44.6	29.1	37.1	35.8	44.4	38.3
	AttFuse [41]	66.9	65.3	46.0	30.4	53.8	50.2	55.5	48.6
	F-Cooper [4]	64.6	63.0	41.9	26.9	48.9	44.8	51.8	44.9
	V2X-ViT [40]	62.6	58.7	47.5	30.5	49.3	47.0	53.1	45.4
	CooPre (AttFuse)	71.4 ^{+4.5}	70.1 ^{+4.8}	49.7 ^{+3.7}	33.7 ^{+3.3}	57.5 ^{+3.7}	55.5 ^{+5.3}	59.5 ^{+4.0}	53.1 ^{+4.5}
	CooPre (F-Cooper)	68.6 ^{+4.0}	67.0 ^{+4.0}	42.0 ^{+0.1}	26.4 ^{-0.5}	51.0 ^{+2.1}	48.1 ^{+3.3}	53.9 ^{+2.1}	47.2 ^{+2.3}
	CooPre (V2X-ViT)	66.9 ^{+4.3}	64.5 ^{+5.8}	46.7 ^{-0.8}	29.2 ^{-1.3}	54.1 ^{+4.8}	52.3 ^{+5.3}	55.9 ^{+2.8}	48.7 ^{+3.3}
V2X-Real IC	No Fusion	55.7	48.6	35.2	20.4	46.3	45.5	45.8	38.2
	Early Fusion	73.6	66.2	45.0	25.7	49.9	47.4	56.2	46.4
	Late Fusion	74.5	72.5	47.8	27.4	66.1	57.8	62.8	52.6
	AttFuse [41]	84.5	82.1	61.0	40.6	59.9	59.1	68.5	60.6
	F-Cooper [4]	79.1	76.2	57.0	35.8	52.0	45.7	62.7	52.6
	V2X-ViT [40]	81.0	76.9	44.9	29.2	37.3	36.8	54.4	47.6
	CooPre (AttFuse)	86.2 ^{+1.7}	84.1 ^{+2.0}	60.5 ^{-0.5}	39.3 ^{-1.3}	61.5 ^{+1.6}	61.0 ^{+1.9}	69.4 ^{+0.9}	61.4 ^{+0.8}
	CooPre (F-Cooper)	84.1 ^{+5.0}	80.7 ^{+4.5}	58.4 ^{+1.4}	33.9 ^{-1.9}	51.8 ^{-0.2}	50.7 ^{+5.0}	64.8 ^{+2.1}	55.1 ^{+2.5}
	CooPre (V2X-ViT)	81.4 ^{+0.4}	77.3 ^{+0.4}	45.8 ^{+0.9}	29.1 ^{-0.1}	37.4 ^{+0.1}	37.2 ^{+0.4}	54.8 ^{+0.4}	47.9 ^{+0.3}
V2X-Real I2I	No Fusion	67.3	61.5	49.2	32.3	53.2	49.0	56.6	47.6
	Early Fusion	69.3	63.4	51.9	32.5	55.5	52.1	58.9	49.4
	Late Fusion	79.6	78.0	67.3	46.7	64.4	55.8	70.4	60.2
	AttFuse [41]	83.4	81.9	67.0	44.8	63.6	62.8	71.3	63.2
	F-Cooper [4]	85.1	82.0	66.1	44.3	51.3	50.8	67.5	59.0
	V2X-ViT [40]	82.0	79.5	68.9	45.4	59.0	57.0	69.9	60.6
	CooPre (AttFuse)	84.6 ^{+1.2}	82.7 ^{+0.8}	66.5 ^{-0.5}	45.2 ^{+0.4}	65.2 ^{+1.6}	62.5 ^{-0.3}	72.1 ^{+0.8}	63.5 ^{+0.3}
	CooPre (F-Cooper)	85.4 ^{+0.3}	82.1 ^{+0.1}	66.6 ^{+0.5}	45.1 ^{+0.8}	51.9 ^{+0.6}	51.1 ^{+0.3}	68.0 ^{+0.5}	59.4 ^{+0.4}
	CooPre (V2X-ViT)	84.0 ^{+2.0}	82.4 ^{+2.9}	66.9 ^{-2.0}	44.4 ^{-1.0}	59.1 ^{+0.1}	58.7 ^{+1.7}	70.0 ^{+0.1}	61.9 ^{+1.3}

Table 2. Performance on the test set of V2V4Real dataset.

Method	Sync (AP0.5/0.7)				Async (AP0.5/0.7)			
	Overall	0-30m	30m-50m	50m-100m	Overall	0-30m	30m-50m	50m-100m
No Fusion	52.6/36.2	76.6/58.7	45.9/28.7	10.0/5.0	52.6/36.2	76.6/58.7	45.9/28.7	10.0/5.0
Early Fusion	66.7/34.5	83.3/48.9	51.0/25.3	48.4/19.9	59.5/28.7	82.5/46.2	43.5/21.1	29.8/9.0
Late Fusion	69.9/40.2	81.4/44.0	60.7/38.5	56.9/36.2	63.9/32.5	78.9/40.3	56.8/32.9	39.4/16.8
AttFuse [41]	71.2/44.2	87.3/55.5	55.6/39.8	51.1/25.7	63.8/34.9	85.7/51.4	47.8/31.0	33.4/10.5
F-Cooper [4]	70.6/41.8	84.4/52.8	59.9/37.8	51.5/26.2	62.2/34.1	82.5/50.0	51.7/30.9	31.9/10.7
V2X-ViT [40]	71.7/43.3	87.8/59.5	58.6/34.1	50.4/23.0	60.0/32.9	83.3/51.3	47.4/25.7	26.3/9.0
CooPre (AttFuse)	74.3/49.3 ^{+3.1/+5.1}	87.4/60.2 ^{+0.1/+4.7}	63.4/45.2 ^{+7.8/+5.4}	55.6/32.0 ^{+4.5/+6.3}	66.1/39.9 ^{+2.3/+5.0}	85.9/57.0 ^{+0.2/+5.6}	54.3/34.9 ^{+6.5/+3.9}	36.1/14.0 ^{+2.7/+3.5}
CooPre (F-Cooper)	71.4/43.4 ^{+0.8/+1.6}	86.1/53.2 ^{+1.7/+0.4}	63.0/39.6 ^{+3.1/+1.8}	48.2/29.5 ^{-3.3/+3.3}	63.7/34.8 ^{+1.5/+0.7}	84.7/50.3 ^{+2.2/+0.3}	55.0/31.1 ^{+3.3/+0.2}	29.7/11.9 ^{-2.2/+0.8}
CooPre (V2X-ViT)	73.1/43.9 ^{+1.4/+0.6}	88.9/58.1 ^{+1.1/-1.4}	61.6/36.7 ^{+3.0/+2.6}	50.7/26.0 ^{+0.3/+3.0}	63.8/38.2 ^{+3.8/+5.3}	86.7/58.3 ^{-3.4/+7.0}	51.1/30.7 ^{+3.7/+5.0}	29.8/11.3 ^{+3.5/+2.3}

[13] optimizer with a weight decay of 1×10^{-2} to optimize our models. During the pretraining stage, we train the model with a batch size of 4 for 15 epochs using a learning

rate of 0.002, and we decay the learning rate with a cosine annealing [23]. We use a masking ratio of 0.7 in our main experiments and a fixed predicted point cloud number of

Table 3. Performance on the test set OPV2V dataset.

Method	Default		Culver	
	AP0.5	AP0.7	AP0.5	AP0.7
No Fusion	71.3	60.4	64.6	51.7
Early Fusion	87.7	81.3	82.1	73.8
Late Fusion	84.6	77.5	80.8	68.2
AttFuse [41]	89.3	82.6	87.5	76.0
F-Cooper [4]	87.8	81.7	89.0	79.7
V2X-ViT [40]	89.8	83.7	88.7	80.2
CooPre (AttFuse)	91.7 ^{+2.4}	86.6 ^{+4.0}	88.7 ^{+1.2}	80.0 ^{+4.0}
CooPre (F-Cooper)	89.0 ^{+1.2}	83.4 ^{+1.7}	90.2 ^{+1.2}	82.0 ^{+2.3}
CooPre (V2X-ViT)	91.9 ^{+2.1}	86.5 ^{+2.8}	89.9 ^{+1.2}	82.4 ^{+2.2}

20. During the fine-tuning stage, the optimization process is identical to the train-from-scratch baselines. We fine-tune the model for 40 epochs in OPV2V dataset, 60 epochs in V2V4Real dataset, and 20 epochs in V2X-Real dataset. We also add normal point cloud data augmentations for all finetuning experiments, including scaling, rotation, and flip [14]. Note that both pretraining and finetuning stages are conducted within the same dataset.




Figure 4. **Illustration of enhanced detection performance in long-range and occlusion scenarios.** Our multi-agent pretraining framework yields more robust perception capabilities in scenarios involving occlusions and long-range perception. Particularly, the occurrence of false perceptions significantly decreases in far-range scenarios where point cloud features are sparse. **Blue** and **red** 3D bounding boxes correspond to the ground-truth and detection outputs, respectively. **Orange** boxes denote the zoomed-in view of detection results.

Table 4. Results of cross-domain transferability of our method (Metric: AP0.5 for Car Category) evaluated under the test set of corresponding finetuning dataset (V2V4Real or V2X-Real V2V).

Pretrain	Finetune	V2V4Real	V2X-Real V2V
No Pretrain		71.2	65.3
V2V4Real		74.3 ^{+3.1}	68.2 ^{+2.9}
V2X-Real V2V		72.2 ^{+1.0}	70.1 ^{+4.8}
Combined		74.5 ^{+3.3}	70.5 ^{+5.2}

Table 5. Comparison with prior works in different perception ranges (Metric: mAP0.5).

Dataset	Pretraining Method	Overall	0-30m	30-50m	50-100m
V2X-Real VC	No Pretrain	48.5	61.0	53.5	36.9
	BEVContrast [28]	48.8	61.2	52.5	39.2
	Random Masking	50.9	63.9	55.4	40.3
	CORE [31]	51.6	64.7	55.4	40.7
	STAR [17]	50.0	63.9	55.6	38.4
	BEV-MAE [22]	51.1	64.5	55.7	39.4
	Ours	52.2	64.9	55.9	42.1
V2X-Real V2V	No Pretrain	48.6	64.2	52.3	36.0
	BEVContrast [28]	46.3	63.1	54.1	32.9
	Random Masking	50.9	65.5	54.7	40.4
	CORE [31]	51.9	65.9	55.0	41.8
	STAR [17]	51.2	66.0	54.4	40.6
	BEV-MAE [22]	50.5	65.4	54.6	38.7
	Ours	53.1	66.1	55.4	42.7

4.2. Main Results

In Table 1, we present the experimental results across all collaboration modes in V2X-Real [37] dataset. In the VC and V2V settings, CooPre substantially outperforms the train-from-scratch baselines across all classification categories. Take AttFuse [41] as example, the improvements are particularly notable in the Car and Truck categories, with a 4.4 mAP0.5 increase in the Car category and a 6.3 mAP0.5 increase in the Truck category in the VC setting, and a 4.8 mAP0.5 increase in the Car category and a 5.3 mAP0.5 increase in the Truck category in the V2V setting. This aligns with our pretraining design, which enables the model to learn more 3D geometrical and topological information beforehand, resulting in more accurate bounding box detection results under higher thresholds. On the other hand, our pretraining method shows an incremental effect on detecting pedestrians. We attribute this to two factors: 1) pedestrians are small-scale in nature and thus receive fewer LiDAR features than larger objects, and 2) pedestrians are non-rigid-body objects, making it more challenging for the model to learn their 3D features compared to rigid-body counterparts such as cars and trucks. On the infrastructure side, our pretraining method shows improvements for the Car and Truck categories, but we observe a subtle decrease in the Pedestrian category. We attribute this to the small-scale and non-rigid-body nature of pedestrians, which might




Figure 5. **Data efficiency experiment.** We evaluate the performance boost using CooPre under different ratios of finetuning data on V2X-Real dataset with AttFuse [41] as backbone.




Figure 6. **Comparison of CooPre across different cooperative fusion strategies.** We assess the performance enhancement provided by CooPre under various fusion strategies on the V2X-Real VC and V2V datasets with AttFuse [41] as backbone.




Figure 7. **Robustness assessment on localization errors and time delay.** Our method shows strong robustness against localization errors and time delay compared to the train-from-scratch baseline on V2V4Real dataset with AttFuse [41] as backbone.

be negatively affected by asynchronized results transmitted from connected vehicles, such as pose and localization errors [24, 33]. This asynchronization can alter the perception accuracy of static infrastructure observers. Such observation is also shown in the results of other cooperative methods [4, 40].

We show the generalizability of our method in another real-world dataset V2V4Real [42] and a simulation dataset OPV2V [41], as shown in Table 2 and Table 3, respectively. For the V2V4Real dataset, after we pretrain and finetune the model, we test it under synchronized and asynchronized modes for a fair comparison. The improvement is substan-

tial in both testing settings. For the OPV2V dataset, we outperform the baseline by a large margin across all test sets. These results demonstrate the generalizability of our pretraining method across different domains with different fusion methods.

4.3. Data Efficiency

In this section, we investigate the benefits of our pretraining method in scenarios with limited labeled data using AttFuse [41] as backbone. Specifically, we randomly sample 20%, 50%, and 80% of the training dataset and train the models with these annotated subsets. For our method, we pretrain the model on the entire training set and then finetune it on each sampled subset. As shown in Fig. 5, our method outperforms train-from-scratch baselines across all settings. Notably, the performance gain of CooPre increases as the percentage of data decreases. Additionally, CooPre provides crucial guidance when collaboration involves different sensor configurations. For instance, in the V2X-Real VC dataset, the baseline method struggles to learn a good representation with limited annotations and the performance drops dramatically due to data heterogeneity issues. With our pretraining method, the model learns meaningful prior knowledge of different sensor distributions, leading to substantial improvements even when finetuning with less labeled data. These findings demonstrate the effectiveness of our method in data scarcity scenarios.

4.4. Cross-domain Transferability

We evaluate the cross-dataset transferability of our pre-trained 3D encoder by finetuning it on another dataset using AttFuse [41] backbone. Since V2V4Real [42] does not support multi-class classification, we only examine the perception results on car category which is shared by both domains. To isolate the effects of data heterogeneity, we investigate the performance across two real-world V2V datasets (V2X-Real V2V and V2V4Real), which have similar sensor configurations. As shown in Table 4, pretraining on a different domain improves performance compared to the train-from-scratch baseline. However, it performs worse

than pretraining on the source domain due to the domain gap issue. We also conduct an experiment where we combine the V2X-Real V2V and V2V4Real datasets to create a large pretraining corpus. While this approach improves performance over pretraining on the source domain, the gains are less significant. This difference could be attributed to the scenario differences between the V2X-Real V2V and V2V4Real datasets. The former primarily focuses on intersection scenarios, whereas the latter includes a broader range of corridor scenarios in its training set.

4.5. Ablation Studies and Discussion

In this section, we conduct extensive experiments to explain how, where, and why our pretraining framework benefits current cooperative perception backbones.

Comparison with other pretraining methods. The comparison of pretraining methods is carried out in three aspects using AttFuse [41] as backbone. First, to evaluate reconstruction targets in a multi-agent pretraining setting, we adapt the reconstruction loss function from [17, 31] to align with our pretraining framework. Our approach demonstrates superior performance. Second, to analyze the impact of collaboration during pretraining, we compare our multi-agent pretraining method with the ego-agent pretraining approach from [22, 28]. As shown in Table 5, multi-agent pretraining results in further performance gains. Additionally, both single-agent and multi-agent pretraining outperform the train-from-scratch approach. Lastly, we note variability in [28], a contrastive learning based single-agent pretraining method, which may struggle to establish robust priors for multi-agent BEV representation compared to reconstruction-based pretraining methods.

Extent of improvements in different perception ranges. We also explore the enhancements across various perception ranges. As shown in Table 5 and Table 2, the most substantial improvements are observed in middle and long ranges compared to the baselines. Since our multi-agent pretraining method enables a significantly larger perception field and provides supervision signals for feature learning on cooperative agents, it demonstrates greater robustness in handling long-range perception and occlusions compared to single-agent pretraining methods, as illustrated in Fig. 4.

Improvements with different cooperative fusion strategies. While our primary focus is on the intermediate fusion strategy, our method also applies to other fusion strategies, as depicted in Fig. 6. We observe that our pretraining strategy benefits other cooperative fusion strategies as well, with the intermediate fusion strategy consistently demonstrating the best performance.

Robustness assessment. Following [40], we evaluate our method’s robustness towards localization error and time delay compared to the baseline method on the V2V4Real dataset, as illustrated in Fig. 7. Our method is less sen-

sitive to localization errors and shows its strong robustness against time delay.

Masking Ratio. Table 6 demonstrates the effect of the masking ratio within the range of 0.6 to 0.8. With a masking ratio of 0.7, the pretraining framework achieves the best performance.

Effectiveness on different 3D encoders. Table 7 shows our pretraining method consistently improves the accuracy of different 3D encoders [14, 43].

Table 6. Ablation on masking ratio.

Dataset	Method	Mask Ratio	mAP0.3	mAP0.5
V2X-Real VC	AttFuse [41]	0.6/0.7/0.8	59.1/ 60.1 /59.0	51.7/ 52.2 /51.0

Table 7. Ablation on different 3D encoders in terms of mAP0.3/0.5.

Dataset	Method	Second		PointPillar	
		Original	W/ CoopPre	Original	W/ CoopPre
V2X-Real VC	AttFuse [41]	56.1/48.5	60.1/52.2	42.5/33.8	44.0/35.2

5. Conclusion

We introduce **CooPre**, a multi-agent pretraining framework that prompts the representation to learn a holistic prior knowledge of the 3D environment before performing the perception task. The framework explores the intrinsic geometrical and topological information of scenarios and sensor distributions. Extensive experiments on representative datasets demonstrate the efficacy of the method as it outperforms the previous state-of-the-art methods in all V2X settings. Furthermore, we demonstrate this framework’s strong generalizability, cross-domain adaptability, and data efficiency in cooperative perception.

Limitations and Future work. Despite its strengths, CooPre has several limitations: (1) its reliance on raw point clouds during pretraining incurs substantial computational overhead, posing challenges for large-scale real-time applications; (2) the framework has not yet been benchmarked with recent cooperative fusion strategies such as Pyramid Fusion [25]; and (3) though well-suited for rigid-body object detection, performance degrades on deformable and small-scale targets like pedestrians. Future work will include: (1) extending this self-supervised pretraining paradigm to encompass cooperative prediction tasks; (2) developing communication-efficient pretraining strategies to alleviate the computational cost; and (3) integrating post-training quantization techniques to build a unified pretrain/post-train framework for efficient V2X cooperative perception.

References

- [1] Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multi-modal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020. 1
- [2] Ming-Fang Chang, John Lambert, Patsorn Sangkloy, Jagjeet Singh, Slawomir Bak, Andrew Hartnett, De Wang, Peter Carr, Simon Lucey, Deva Ramanan, and James Hays. Argoverse: 3D Tracking and Forecasting With Rich Maps. pages 8748–8757, 2019. 1
- [3] Hao Chen, Jiase Wang, Kun Shao, Furui Liu, Jianye Hao, Chenyong Guan, Guangyong Chen, and Pheng-Ann Heng. Traj-mae: Masked autoencoders for trajectory prediction. *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 8317–8328, 2023. 3
- [4] 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 the 4th ACM/IEEE Symposium on Edge Computing*, pages 88–100, 2019. 2, 3, 5, 6, 7
- [5] Jie Cheng, Xiaodong Mei, and Ming Liu. Forecast-MAE: Self-supervised pre-training for motion forecasting with masked autoencoders. *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023. 3
- [6] Christoph Feichtenhofer, Haoqi Fan, Yanghao Li, and Kaiming He. Masked autoencoders as spatiotemporal learners. *ArXiv*, abs/2205.09113, 2022. 3
- [7] Peng Gao, Teli Ma, Hongsheng Li, Jifeng Dai, and Yu Qiao. Convmae: Masked convolution meets masked autoencoders. *arXiv preprint arXiv:2205.03892*, 2022. 3
- [8] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3354–3361, 2012. 1
- [9] Ruiyang Hao, Siqi Fan, Yingru Dai, Zhenlin Zhang, Chenxi Li, Yuntian Wang, Haibao Yu, Wenxian Yang, Yuan Jirui, and Zaiqing Nie. Rcooper: A real-world large-scale dataset for roadside cooperative perception. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 1, 2
- [10] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross B. Girshick. Masked autoencoders are scalable vision learners. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15979–15988, 2021. 3
- [11] Yue Hu, Shaoheng Fang, Zixing Lei, Yiqi Zhong, and Siheng Chen. Where2comm: Communication-efficient collaborative perception via spatial confidence maps. *Advances in neural information processing systems*, 35:4874–4886, 2022. 2
- [12] Di Huang, Sida Peng, Tong He, Honghui Yang, Xiaowei Zhou, and Wanli Ouyang. Ponder: Point cloud pre-training via neural rendering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16089–16098, 2023. 3
- [13] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 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*, pages 12697–12705, 2019. 3, 4, 6, 8
- [15] Quanyi Li, Zhenghao Peng, Lan Feng, Qihang Zhang, Zhenghai Xue, and Bolei Zhou. Metadrive: Composing diverse driving scenarios for generalizable reinforcement learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. 3
- [16] Quanyi Li, Zhenghao Peng, Lan Feng, Zhizheng Liu, Chenda Duan, Wenjie Mo, and Bolei Zhou. Scenarionet: Open-source platform for large-scale traffic scenario simulation and modeling. *Advances in Neural Information Processing Systems*, 2023. 3
- [17] Yiming Li, Juexiao Zhang, Dekun Ma, Yue Wang, and Chen Feng. Multi-robot scene completion: Towards task-agnostic collaborative perception. In *6th Annual Conference on Robot Learning*. 6, 8
- [18] Yiming Li, Shunli Ren, Pengxiang Wu, Siheng Chen, Chen Feng, and Wenjun Zhang. Learning distilled collaboration graph for multi-agent perception. In *Thirty-fifth Conference on Neural Information Processing Systems (NeurIPS 2021)*, 2021. 2
- [19] Yiming Li, Dekun Ma, Ziyang An, Zixun Wang, Yiqi Zhong, Siheng Chen, and Chen Feng. V2X-Sim: Multi-Agent Collaborative Perception Dataset and Benchmark for Autonomous Driving. *IEEE Robotics and Automation Letters*, 7(4):10914–10921, 2022. Conference Name: IEEE Robotics and Automation Letters. 1, 2
- [20] Yiheng Li, Seth Z. Zhao, Chenfeng Xu, Chen Tang, Chenran Li, Mingyu Ding, Masayoshi Tomizuka, and Wei Zhan. Pre-training on synthetic driving data for trajectory prediction, 2023. 3
- [21] Yiming Li, Zhiheng Li, Nuo Chen, Moonjun Gong, Zonglin Lyu, Zehong Wang, Peili Jiang, and Chen Feng. Multi-agent multitraversal multimodal self-driving: Open mars dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 22041–22051, 2024. 1
- [22] Zhiwei Lin, Yongtao Wang, Shengxiang Qi, Nan Dong, and Ming-Hsuan Yang. Bev-mae: Bird’s eye view masked autoencoders for point cloud pre-training in autonomous driving scenarios. In *Proceedings of the AAAI conference on artificial intelligence*, 2024. 2, 3, 4, 6, 8
- [23] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019. 5
- [24] Yifan Lu, Quanhao Li, Baoan Liu, Mehrdad Dianati, Chen Feng, Siheng Chen, and Yanfeng Wang. Robust collaborative 3d object detection in presence of pose errors. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4812–4818. IEEE, 2023. 2, 7
- [25] Yifan Lu, Yue Hu, Yiqi Zhong, Dequan Wang, Siheng Chen, and Yanfeng Wang. An extensible framework for open het-

- erogeneous collaborative perception. In *The Twelfth International Conference on Learning Representations*, 2024. 3, 8
- [26] Chen Min, Liang Xiao, Dawei Zhao, Yiming Nie, and Bin Dai. Occupancy-mae: Self-supervised pre-training large-scale lidar point clouds with masked occupancy autoencoders. *IEEE Transactions on Intelligent Vehicles*, 2023. 3
- [27] Yatian Pang, Wenxiao Wang, Francis EH Tay, Wei Liu, Yonghong Tian, and Li Yuan. Masked autoencoders for point cloud self-supervised learning. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part II*, pages 604–621. Springer, 2022. 3
- [28] Corentin Sautier, Gilles Puy, Alexandre Boulch, Renaud Marlet, and Vincent Lepetit. BEVContrast: Self-supervision in bev space for automotive lidar point clouds. In *International Conference on 3D Vision (3DV)*, 2024. 2, 3, 6, 8
- [29] Pei Sun, Henrik Kretschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, Vijay Vasudevan, Wei Han, Jiquan Ngiam, Hang Zhao, Aleksei Timofeev, Scott Ettinger, Maxim Krivokon, Amy Gao, Aditya Joshi, Yu Zhang, Jonathon Shlens, Zhifeng Chen, and Dragomir Anguelov. Scalability in Perception for Autonomous Driving: Waymo Open Dataset. pages 2446–2454, 2020. 1
- [30] Xiaoyu Tian, Haoxi Ran, Yue Wang, and Hang Zhao. Geomae: Masked geometric target prediction for self-supervised point cloud pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13570–13580, 2023. 3
- [31] Binglu Wang, Lei Zhang, Zhaozhong Wang, Yongqiang Zhao, and Tianfei Zhou. Core: Cooperative reconstruction for multi-agent perception. *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 8676–8686, 2023. 3, 6, 8
- [32] Tsun-Hsuan Wang, Sivabalan Manivasagam, Ming Liang, Bin Yang, Wenyuan Zeng, James Tu, and Raquel Urtasun. V2VNet: Vehicle-to-Vehicle Communication for Joint Perception and Prediction. In *arXiv:2008.07519 [cs]*, 2020. 00010 ECC arXiv: 2008.07519. 2
- [33] Sizhe Wei, Yuxi Wei, Yue Hu, Yifan Lu, Yiqi Zhong, Siheng Chen, and Ya Zhang. Asynchrony-robust collaborative perception via bird’s eye view flow. In *Advances in Neural Information Processing Systems*, 2023. 2, 7
- [34] Zhuoyuan Wu, Yuping Wang, Hengbo Ma, Zhaowei Li, Hang Qiu, and Jiachen Li. Cmp: Cooperative motion prediction with multi-agent communication. *arXiv preprint arXiv:2403.17916*, 2024. 2
- [35] Hao Xiang, Runsheng Xu, and Jiaqi Ma. HM-ViT: Heteromodal Vehicle-to-Vehicle Cooperative perception with vision transformer, 2023. arXiv:2304.10628 [cs]. 3
- [36] Hao Xiang, Runsheng Xu, Xin Xia, Zhaoliang Zheng, Bolei Zhou, and Jiaqi Ma. V2xp-asg: Generating adversarial scenes for vehicle-to-everything perception. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3584–3591, 2023. 2
- [37] Hao Xiang, Zhaoliang Zheng, Xin Xia, Runsheng Xu, Letian Gao, Zewei Zhou, Xu Han, Xinkai Ji, Mingxi Li, Zonglin Meng, et al. V2x-real: a large-scale dataset for vehicle-to-everything cooperative perception. *arXiv preprint arXiv:2403.16034*, 2024. 1, 2, 4, 6
- [38] Chenfeng Xu, Tian Li, Chen Tang, Lingfeng Sun, Kurt Keutzer, Masayoshi Tomizuka, Alireza Fathi, and Wei Zhan. Pretram: Self-supervised pre-training via connecting trajectory and map. *arXiv preprint arXiv:2204.10435*, 2022. 3
- [39] Runsheng Xu, Zhengzhong Tu, Hao Xiang, Wei Shao, Bolei Zhou, and Jiaqi Ma. CoBEVT: Cooperative Bird’s Eye View Semantic Segmentation with Sparse Transformers. arXiv, 2022. arXiv:2207.02202 [cs]. 2
- [40] 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 *ECCV 2022*, pages 107–124, Cham, 2022. Springer Nature Switzerland. 1, 2, 3, 5, 6, 7, 8
- [41] 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 *2022 International Conference on Robotics and Automation (ICRA)*, pages 2583–2589, 2022. 2, 3, 4, 5, 6, 7, 8
- [42] Runsheng Xu, Xin Xia, Jinlong Li, Hanzhao Li, Shuo Zhang, Zhengzhong Tu, Zonglin Meng, Hao Xiang, Xiaoyu Dong, Rui Song, Hongkai Yu, Bolei Zhou, and Jiaqi Ma. V2V4Real: A Real-World Large-Scale Dataset for Vehicle-to-Vehicle Cooperative Perception. pages 13712–13722, 2023. 1, 2, 4, 7
- [43] Yan Yan, Yuxing Mao, and Bo Li. Second: Sparsely embedded convolutional detection. *Sensors*, 18(10), 2018. 3, 4, 8
- [44] Dingkan Yang, Kun Yang, Yuzheng Wang, Jing Liu, Zhi Xu, Rongbin Yin, Peng Zhai, and Lihua Zhang. How2comm: Communication-efficient and collaboration-pragmatic multi-agent perception. In *Thirty-seventh Conference on Neural Information Processing Systems (NeurIPS)*, 2023. 3
- [45] Honghui Yang, Tong He, Jiaheng Liu, Hua Chen, Boxi Wu, Binbin Lin, Xiaofei He, and Wanli Ouyang. Gd-mae: Generative decoder for mae pre-training on lidar point clouds. In *CVPR*, 2023. 3, 4
- [46] Haibao Yu, Yizhen Luo, Mao Shu, Yiyi Huo, Zebang Yang, Yifeng Shi, Zhenglong Guo, Hanyu Li, Xing Hu, Jirui Yuan, and Zaiqing Nie. DAIR-V2X: A Large-Scale Dataset for Vehicle-Infrastructure Cooperative 3D Object Detection. pages 21361–21370, 2022. 2
- [47] Haibao Yu, Wenxian Yang, Hongzhi Ruan, Zhenwei Yang, Yingjuan Tang, Xu Gao, Xin Hao, Yifeng Shi, Yifeng Pan, Ning Sun, Juan Song, Jirui Yuan, Ping Luo, and Zaiqing Nie. V2X-Seq: A Large-Scale Sequential Dataset for Vehicle-Infrastructure Cooperative Perception and Forecasting. arXiv, 2023. arXiv:2305.05938 [cs]. 1, 2
- [48] Xumin Yu, Lulu Tang, Yongming Rao, Tiejun Huang, Jie Zhou, and Jiwen Lu. Point-bert: Pre-training 3d point cloud transformers with masked point modeling. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 3
- [49] Renrui Zhang, Ziyu Guo, Peng Gao, Rongyao Fang, Bin Zhao, Dong Wang, Yu Qiao, and Hongsheng Li. Point-

m2ae: Multi-scale masked autoencoders for hierarchical point cloud pre-training. *arXiv preprint arXiv:2205.14401*, 2022. [3](#)

- [50] Zhaoliang Zheng, Xin Xia, Letian Gao, Hao Xiang, and Jiaqi Ma. Cooperfuse: A real-time cooperative perception fusion framework. In *2024 IEEE Intelligent Vehicles Symposium (IV)*, pages 533–538, 2024. [2](#)
- [51] Yin Zhou and Oncel Tuzel. VoxelNet: End-to-End Learning for Point Cloud Based 3D Object Detection. pages 4490–4499, 2018. [3](#)
- [52] Haoyi Zhu, Honghui Yang, Xiaoyang Wu, Di Huang, Sha Zhang, Xianglong He, Tong He, Hengshuang Zhao, Chunhua Shen, Yu Qiao, and Wanli Ouyang. Ponderv2: Pave the way for 3d foundation model with a universal pre-training paradigm. *arXiv preprint arXiv:2310.08586*, 2023. [3](#)
- [53] Walter Zimmer, Gerhard Arya Wardana, Suren Sritharan, Xingcheng Zhou, Rui Song, and Alois C. Knoll. Tumtraf v2x cooperative perception dataset. In *2024 IEEE/CVF International Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE/CVF, 2024. [1](#)