

# CoDTS: Enhancing Sparsely Supervised Collaborative Perception with a Dual Teacher-Student Framework

Yushan Han, Hui Zhang, Honglei Zhang, Jing Wang, Yidong Li\*

Key Laboratory of Big Data & Artificial Intelligence in Transportation (Beijing Jiaotong University), Ministry of Education  
School of Computer Science and Technology, Beijing Jiaotong University, Beijing 100044, China  
{21112034, huizhang1, honglei.zhang, jing\_w, ydli}@bjtu.edu.cn

## Abstract

Current collaborative perception methods often rely on fully annotated datasets, which can be expensive to obtain in practical situations. To reduce annotation costs, some works adopt sparsely supervised learning techniques and generate pseudo labels for the missing instances. However, these methods fail to achieve an optimal confidence threshold that harmonizes the quality and quantity of pseudo labels. To address this issue, we propose an end-to-end Collaborative perception Dual Teacher-Student framework (CoDTS), which employs adaptive complementary learning to produce both high-quality and high-quantity pseudo labels. Specifically, the Main Foreground Mining (MFM) module generates high-quality pseudo labels based on the prediction of the static teacher. Subsequently, the Supplement Foreground Mining (SFM) module ensures a balance between the quality and quantity of pseudo labels by adaptively identifying missing instances based on the prediction of the dynamic teacher. Additionally, the Neighbor Anchor Sampling (NAS) module is incorporated to enhance the representation of pseudo labels. To promote the adaptive complementary learning, we implement a staged training strategy that trains the student and dynamic teacher in a mutually beneficial manner. Extensive experiments demonstrate that the CoDTS effectively ensures an optimal balance of pseudo labels in both quality and quantity, establishing a new state-of-the-art in sparsely supervised collaborative perception.

## Introduction

Multi-agent collaborative perception (Han et al. 2023a) can enhance the perception capabilities of vehicles in autonomous driving scenarios, such as vehicle-to-vehicle (V2V), vehicle-to-infrastructure (V2I), and vehicle-to-everything (V2X) interactions. As an emerging research, recent works have contributed valuable datasets, efficient communication mechanisms, and effective fusion techniques. Additionally, several pioneer approaches aim to address practical challenges like localization errors, communication latency and model heterogeneity.

However, classic collaborative perception methods often depend on fully annotated datasets. Despite their considerable performance, they are challenged by being labor-

intensive and time-consuming. For example, in the case of LiDAR-based 3D object detection, it takes more than 100 seconds (Song, Lichtenberg, and Xiao 2015; Meng et al. 2020; Luo et al. 2022b) to annotate a precise 3D bounding box for individual perception datasets such as KITTI (Geiger, Lenz, and Urtasun 2012) and Waymo (Sun et al. 2020). Since collaborative perception involves multiple agents, the annotation cost becomes several times higher than that of individual perception. Moreover, the annotation of collaborative perception requires aligning shared objects across multiple agents, which increases annotation complexity. Therefore, the data annotation significantly impedes the practical application of collaborative perception.

To address the above high-cost issue, an alternative approach is sparsely supervised learning, where only one object per agent in each frame is randomly annotated, as shown in Figure 1(b). To mine missing instances, the current mainstream methods (Han et al. 2023b; Xia et al. 2024) primarily adopt a teacher-student framework. Specifically, SSC3OD (Han et al. 2023b) proposes a self-supervised learning approach to pre-train the teacher model. Then, it utilizes a static teacher-student framework to produce high-quality pseudo labels using high thresholds. HINTED (Xia et al. 2024) adopts a dynamic teacher-student framework, where the teacher mainly relies on high-level pseudo labels and a mixed-density student is incorporated to ensure that the model prioritizes hard instances. While these methods improve performance to a certain degree, the detection performance under sparse supervision still lags significantly behind that of full supervision. We argue that these methods primarily focus on the quality of pseudo labels, overlooking their quantity, thus resulting in suboptimal solutions.

Specifically, traditional teacher-student frameworks often employ higher thresholds to derive high-quality pseudo labels. However, varying threshold levels present unique benefits. As illustrated in Figure 1 (c)-(g), a lower threshold may capture the majority of foreground objects (low MPR) but at the expense of introducing substantial background noises (high FPR). In contrast, a higher threshold can ensure that the selection primarily includes foreground objects (low FPR) but may result in the loss of important foreground details (high MPR). Consequently, *how to find an optimal confidence threshold that balances the quality and quantity of pseudo labels remains a formidable challenge.*

\*Corresponding author (ydli@bjtu.edu.cn).

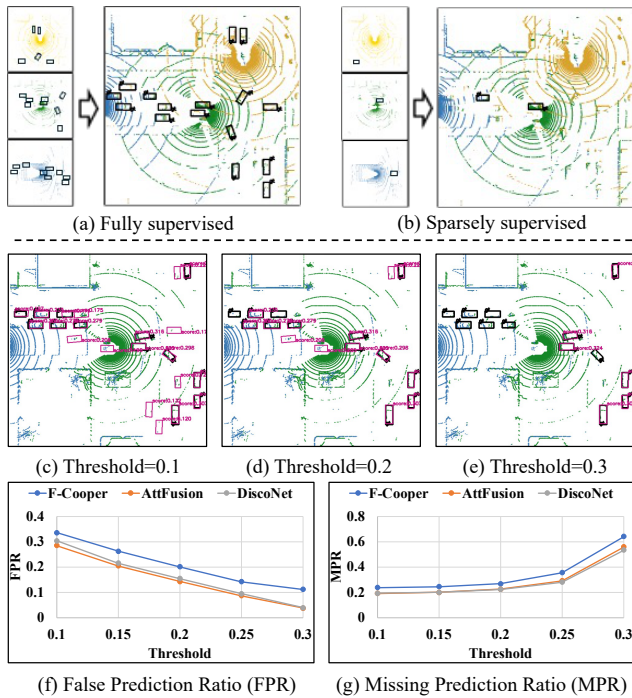


Figure 1: (a)-(b): Labels of fully and sparsely supervised collaborative perception. (c)-(g): The impact of different confidence thresholds on **pseudo labels**. The false prediction ratio (FPR) = the number of false predictions / the number of pseudo labels, and the missing prediction ratio (MPR) = the number of missed predictions / the number of ground truths.

To tackle this challenge, we propose an end-to-end collaborative perception dual teacher-student framework (CoDTS), which leverages adaptive complementary learning to generate pseudo labels that are both high in quality and quantity, ensuring optimal performance. This is achieved through a staged training strategy: i) In the warm-up stage, the Main Foreground Mining (MFM) module is introduced to ensure a sufficient quantity of pseudo labels. This mainly involves filtering the prediction of the static teacher using a low confidence threshold, which effectively pre-trains the student and dynamic teacher. ii) In the refinement stage, the MFM focuses on identifying high-quality positive instances by adjusting the confidence threshold of the static teacher to a higher value, which inevitably causes missing positive instances. To complement these instances, we introduce Supplement Foreground Mining (SFM) module, which identifies the missing ones by filtering the predictions of dynamic teachers with a dynamically adapted high confidence threshold. This adaptive complementary learning guarantees an optimal balance of both quality and quantity pseudo labels. Furthermore, during both stages, the Neighbor Anchor Sampling (NAS) module is used to select neighboring instances for the identified positives, enhancing the representations of pseudo labels. Besides, the framework trains the student and dynamic teacher through mutual learning, enabling them to make continuous improvements. Our contributions can be summarized as follows:

- We propose a novel dual teacher-student framework CoDTS, which effectively resolves the imbalance of pseudo labels between quality and quantity in traditional teacher-student frameworks. It enables sparsely supervised collaborative detectors to achieve comparable performance to fully supervised ones.
- The proposed MFM and SFM are used to identify both high-quality and high-quantity pseudo labels through adaptive complementary learning, with the threshold in SFM dynamically adapted according to current scenes. Besides, the NAS is employed to enhance the dense representation of pseudo labels.
- A staged training strategy is introduced to perform adaptive complementary learning in an end-to-end manner. By implementing this strategy, the student and dynamic teacher are optimized in a mutually beneficial manner, enabling them to evolve together and continuously enhance detection accuracy.
- Extensive experiments are conducted on V2X-Sim, OPV2V, DAIR-V2X and V2V4Real datasets to demonstrate the superior performance of CoDTS framework.

## Related Work

### Collaborative Perception

Collaborative perception enhances the individual perception (Zhang et al. 2022; Luo et al. 2023) by facilitating information exchange among multiple agents, enabling them to obtain more comprehensive views. Previous research has achieved notable advancements in effective communication (Hu et al. 2022; Yang et al. 2024) and adaptive fusion (Xu et al. 2022a; Luo et al. 2022a; Su et al. 2024) under ideal conditions and in addressing challenges like time delay (Yu et al. 2024) and localization errors (Zhang et al. 2024; Hong et al. 2024) in real-world applications. However, these methods heavily rely on complete annotations, which are labor-intensive and impede the practical deployment of collaborative perception.

### Sparsely Supervised 3D Object Detection

To mitigate the high annotation costs in LiDAR-based 3D object detection, researchers have turned to sparsely supervised approaches. SS3D (Liu et al. 2022) pioneers sparsely supervised 3D object detection, iteratively mining and storing reliable instances and backgrounds under sparse supervision to ultimately achieve full supervision. SSC3OD (Han et al. 2023b) first introduces sparsely supervised collaborative 3D object detection, utilizing a Pillar-based mask autoencoder to aid the detector in reasoning about high-level semantics. HINTED (Xia et al. 2024) employs a dynamic teacher-student framework and proposes a mixed-density student to boost performance in hard instances. Despite significant advances in single-vehicle 3D object detection, sparsely supervised collaborative 3D object detection has yet to achieve satisfactory performance. This work seeks to enhance the performance of collaborative 3D object detectors under sparse supervision.

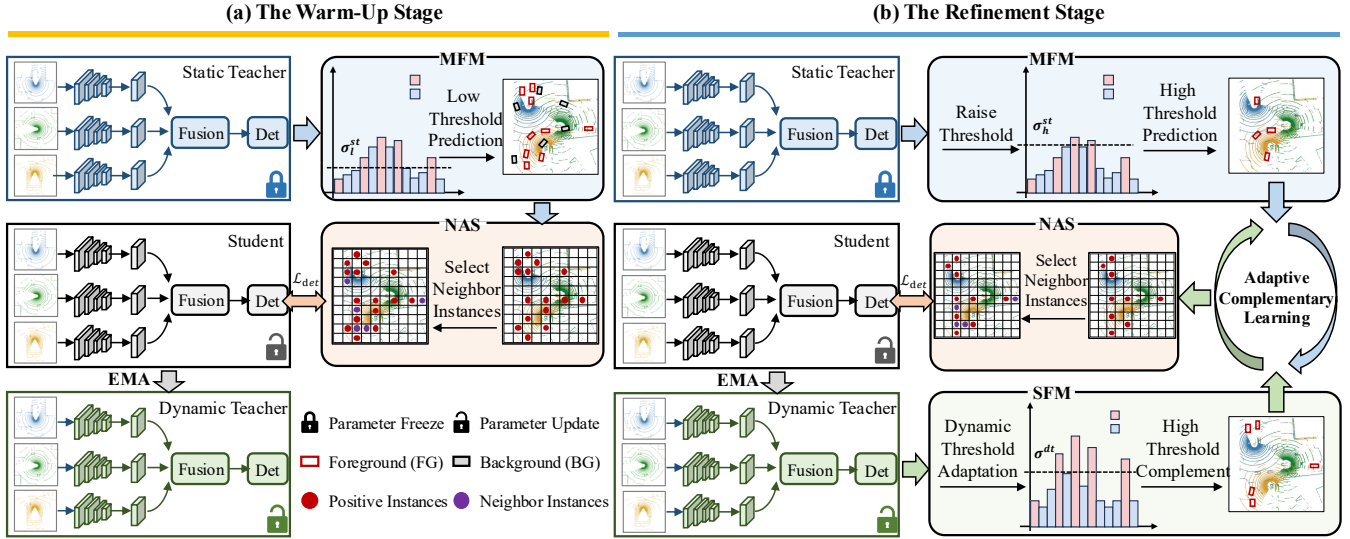


Figure 2: The CoDTS framework employs a staged training strategy. (a) In the warm-up stage, the MFM utilizes a low threshold to generate a sufficient amount of pseudo labels, which are used to guide the student and pre-train the dynamic teacher. (b) In the refinement stage, the MFM raises the threshold to obtain high-quality pseudo labels, while the SFM utilizes dynamic threshold adaptation to adjust the threshold and adaptively identify missing instances to complement the MFM. This Adaptive Complementary Learning ensures the generation of *both high-quality and high-quantity* pseudo labels, which are then merged to guide the student. Throughout both stages, the NAS is used to enhance the representations of pseudo labels.

## Methodology

### Overview Framework

To address the imbalance of pseudo labels between quality and quantity, we propose the CoDTS framework, which generates pseudo labels through adaptive complementary learning. Specifically, the MFM is proposed to generate high-quality pseudo labels based on the prediction of the static teacher, which inevitably misses some instances. To complement the MFM, the SFM is introduced to adaptively identify the missing instances based on the prediction of the dynamic teacher, ensuring the quality and quantity of pseudo labels. Additionally, the NAS is used to enhance the dense representations of pseudo labels during training by selecting neighboring instances. As shown in Figure 2, the proposed CoDTS framework employs a delicate staged training strategy. In the warm-up stage, the student and dynamic teacher are effectively pre-trained. In the refinement stage, we implement adaptive complementary learning to produce pseudo labels. Furthermore, the dynamic teacher and student are jointly trained in a mutually beneficial manner, enabling continuous optimization. Next, we will delve into the details of adaptive complementary learning, neighbor anchor sampling, and staged training strategy.

### Adaptive Complementary Learning

To generate both high-quality and high-quantity pseudo labels, we introduce the adaptive complementary learning regimen. Specifically, the MFM generates pseudo labels by using a high threshold to filter the prediction of the static teacher. Following this, the SFM adaptively complements the missing pseudo labels from MFM, with thresholds dynamically adjusting to the current scenario.

In a collaborative detector, each fused feature generates anchors  $R \in \mathbb{R}^{H \times W \times A}$ , where  $H$  and  $W$  represent the dimensions of the feature map, and  $A$  represents the number of anchors in each feature grid. The detection prediction consists of three types of information, i.e.,  $z = \{z_{cls}, z_{reg}, z_{dir}\}$ , where  $z_{cls}$ ,  $z_{reg}$  and  $z_{dir}$  represent the classification, regression and direction for each anchor, respectively. The details of MFM and SFM are as follows.

**Main Foreground Mining** Given the static teacher  $\theta^{st}$ , confidence threshold  $\sigma^{st}$  and overlap threshold  $\tau$ , the function of identifying positive instances  $R^{st}$  by MFM can be denoted as follows:

$$R^{st} = \text{MFM}(\theta^{st}, \sigma^{st}, \tau). \quad (1)$$

Specifically, the MFM first selects anchors  $R_0^{st}$  whose classification prediction  $z_{cls}^{st}$  is larger than confidence threshold  $\sigma^{st}$ . Subsequently, it utilizes the Non-Maximum Suppression (NMS) algorithm (Neubeck and Van Gool 2006) with an overlap threshold  $\tau$  to eliminate anchors with significant overlap. Finally, the remaining anchors are treated as positive instances  $R^{st}$ . This process can be formulated as follows:

$$R_0^{st} = \{R^{(i)} \mid z_{cls}^{st(i)} > \sigma^{st}, R^{(i)} \in R\},$$

$$R^{st} = \text{NMS}(R_0^{st}, \tau). \quad (2)$$

**Supplement Foreground Mining** The SFM aims to identify missing positive instances using the dynamic teacher  $\theta^{dt}$ . Similar to MFM, the SFM also utilizes a confidence threshold  $\sigma^{dt}$  and an overlap threshold  $\tau$  to identify instances  $R^{dt}$  with high confidence and low overlap, which can be denoted as:

$$R^{dt} = \text{SFM}(\theta^{dt}, \sigma^{dt}, \tau). \quad (3)$$




Figure 3: Visualization of **pseudo label** generation through adaptive complementary learning. (a)-(c) represent the pseudo labels generated by MFM, SFM, and their merged results, respectively.

The dynamic teacher  $\theta^{dt}$  exhibits varying detection performance across different scenes. For instance, it might predict high scores in some scenes but low scores in others, making it unreasonable to use a predefined static threshold. To ensure the dynamic teacher threshold  $\sigma^{dt}$  seamlessly adapts to diverse scenarios and effectively filters high-quality foregrounds, we employ a **dynamic threshold adaptation** method, eliminating the need for human intervention.

Specifically, we collect the predicted scores of the dynamic teacher for sparse labels in the current batch and apply the sigmoid function to generate classification scores. Since the sparse labels provide accurate supervision, these scores should generally be at a similar level, though some outliers may have lower scores. We use the K-Means clustering algorithm (MacQueen et al. 1967) to classify these scores into high-level and low-level clusters. The high-level cluster center, representing the average prediction level of the dynamic teacher  $\theta^{dt}$  in the current batch, is chosen as threshold  $\sigma^{dt}$ .

Since we pre-define  $A$  anchors for each feature grid, both the static and dynamic teachers generate predictions for each anchor. Consequently, different teachers may have different predictions for the same grid. To avoid conflicting predictions, we calculate the feature grid index for positive instances in  $R_1^{dt}$  and  $R^{st}$ , and only remain the instances in  $R_1^{dt}$  that are missed by  $R^{st}$ . We denote the index function of anchors as  $\text{Id}(\cdot)$ , and the overall process of SFM can be formulated as follows:

$$\begin{aligned} R_0^{dt} &= \{R^{(i)} \mid z_{cls}^{dt(i)} > \sigma^{dt}, R^{(i)} \in R\}, \\ R_1^{dt} &= \text{NMS}(R_0^{dt}, \tau), \\ R^{dt} &= \{R^{(i)} \mid \text{Id}(R^{(i)})/A \notin \text{Id}(R^{st})/A, R^{(i)} \in R_1^{dt}\}. \end{aligned} \quad (4)$$

Figure 3 illustrates the pseudo labels generated via adaptive complementary learning. The pseudo labels generated by MFM are all positive instances but miss many ground truths. Meanwhile, the pseudo labels generated by SFM effectively complement these missed instances. As a result, the merged pseudo labels from MFM and SFM encompass nearly all ground truths while maintaining high quality.

### Neighbor Anchor Sampling

While MFM and SFM effectively identify positive instances, the mined instances are too sparse compared to the entire feature grid, which makes it challenging to learn a




Figure 4: The process of NAS. First, we generate bounding boxes for positive instances. We then select neighbor instances that have a high overlap with these bounding boxes. Finally, the positive instances are combined with the neighbor instances to create dense positive instances.

foreground representation. Inspired by the fact that highly spatially overlapped instances should have the same label, we propose the NAS, which selects the neighbor instances based on these identified positive instances, thereby enhancing the dense representation of pseudo labels.

As shown in Figure 4, consider the positive instances  $R^{st}$  (represented as red dots) identified by MFM. We use the regression prediction  $z_{reg}^{st}$  of instances to calculate their corresponding bounding boxes  $B^{st}$  (represented as red rectangles). Subsequently, we compute the overlaps between all anchors  $R$  and the foreground bounding boxes  $B^{st}$  and set an overlap threshold  $\tau^{nei}$  to select a set of neighbor instances  $R^{nei}$  (represented as purple dots). The above procedure can be defined as:

$$\begin{aligned} R^{nei} &= \text{NAS}(R^{st}, \tau^{nei}) \\ &= \{R^{(i)} \mid \text{IoU}(R^{(i)}, B^{st}) > \tau^{nei}, R^{(i)} \in R\}. \end{aligned} \quad (5)$$

By utilizing NAS, we enhance the representation of pseudo labels, which assists collaborative detectors in learning these instances more effectively. Finally, we will combine the sparse labels with the pseudo labels to generate the ultimate pseudo label. If both the sparse and pseudo labels identify the same anchor as a positive instance, we retain only the sparse label to ensure supervision accuracy.

### Staged Training Strategy

To implement adaptive complementary learning, our CoDTS employs a delicate staged training strategy. As shown in Algorithm 1, the staged training strategy includes two phases: the warm-up stage and the refinement stage.

**The Warm-Up Stage** The purpose of the warm-up stage is to generate *high-quantity* pseudo labels to guide the student and pre-train the dynamic teacher, which is used for the generation of complementary pseudo labels in the subsequent stage. First, we use the MFM to identify positive instances by applying a low threshold  $\sigma_l^{st}$  to filter the predictions of the static teacher. This ensures that the pseudo labels include sufficient foreground information. Next, we use the NAS to select neighboring instances, enhancing the dense representation of the pseudo labels.

After each training iteration  $iter$ , the student  $\theta^s$  utilizes the exponential moving average (EMA) with a smoothing

coefficient  $\alpha$  to update the parameter weight of the dynamic teacher  $\theta^{dt}$ . By the end of the warm-up stage, both the student and the dynamic teacher can effectively identify most foregrounds. The process of EMA is defined as follows:

$$\theta_{iter}^{dt} = \begin{cases} \left(1 - \frac{1}{iter}\right) \times \theta_{iter-1}^{dt} + \frac{1}{iter} \times \theta_{iter}^s, & 1 - \frac{1}{iter} < \alpha \\ \alpha \theta_{iter-1}^{dt} + (1 - \alpha) \theta_{iter}^s. & \text{otherwise} \end{cases} \quad (6)$$

By employing the EMA strategy, we can view the slowly evolving dynamic teacher as an ensemble of the students at various iterations, which enhances the stability and robustness of the prediction capabilities of the dynamic teacher. Additionally, positive instance features remain consistent throughout training, while negative instance features exhibit random diversity. Consequently, the dynamic teacher will assign higher confidence to positive instances, which sets the foundation for utilizing the dynamic teacher to identify missing positive instances in the subsequent stage.

**The Refinement Stage** The refinement stage aims to generate *both high-quality and high-quantity* pseudo labels to improve the detection accuracy of both student and dynamic teacher. This is accomplished through adaptive complementary learning between MFM and SFM. Specifically, we first raise the threshold of MFM to a high level  $\sigma_h^{st}$  to identify positive instances, ensuring the quality of pseudo labels but inevitably missing some positive instances. Then the SFM applies a high threshold  $\sigma^{dt}$  for SFM to complement the missing instances adaptively, where  $\sigma^{dt}$  is dynamically adapted based on current scenes. This complementary mining approach ensures an optimal balance between the quality and quantity of pseudo labels. Subsequently, we merge the positive instances identified by both MFM and SFM, and use NAS to select neighbor instances. Similar to the warm-up stage, we use EMA strategy to update the dynamic teacher.

In the refinement stage, the student and dynamic teacher are optimized via mutual learning, where the dynamic teacher generates part of high-quality pseudo labels to train the student, and the student updates the knowledge it learned back to the dynamic teacher. With the above interaction, both models can evolve jointly and continuously to improve detection accuracy, which also means that the dynamic teacher can generate more accurate and stable pseudo labels. Furthermore, due to continuous supervision by pseudo labels, both student and dynamic teacher are able to detect a large number of high-quality objects effectively.

## Experimental Results

### Datasets and Evaluation Metrics

We conduct experiments on four large-scale collaborative perception datasets and evaluate the methods with average precision (AP) at IoU thresholds of 0.3, 0.5 and 0.7.

**V2X-Sim** (Li et al. 2022) is a simulated dataset for V2X collaborative perception. We utilize the V2XSim 2.0 version and set the LiDAR range as  $x \in [-32m, 32m], y \in [-32m, 32m]$ . The training set comprises 698,991 fully supervised objects and 29,300 sparsely supervised objects. **OPV2V** (Xu et al. 2022b) is a simulated dataset for V2V collaborative perception. The LiDAR range is set as  $x \in$

---

### Algorithm 1: Staged Training Strategy

---

**input** : Static teacher  $\theta^{st}$ . Dynamic teacher  $\theta^{dt}$ .  
Student  $\theta^s$ . Refinement starts iteration  
 $I_{refine}$ . Max training iteration  $I_{max}$

- 1 **Preparation:**
- 2 Initialize the weight of student model  $\theta^s$
- 3 Copy the weight of student  $\theta^s$  to dynamic teacher  $\theta^{dt}$
- 4 **Do training iteration:**
- 5 **for**  $iter = 1, 2, \dots, I_{max}$  **do**
- 6     **if**  $iter < I_{refine}$  **then**      $\triangleright$  *The Warm-Up Stage*
- 7          $R^{st} = \text{MFM}(\theta^{st}, \sigma_l^{st}, \tau)$
- 8          $R^{nei} = \text{NAS}(R^{st}, \tau^{nei})$
- 9     **else**      $\triangleright$  *The Refinement Stage*
- 10          $R^{st} = \text{MFM}(\theta^{st}, \sigma_h^{st}, \tau)$
- 11         Get  $\sigma^{dt}$  by dynamic threshold adaptation
- 12          $R^{dt} = \text{SFM}(\theta^{dt}, \sigma^{dt}, \tau)$
- 13          $R^{nei} = \text{NAS}(R^{st} \cup R^{dt}, \tau^{nei})$
- 14         Merge pseudo labels and sparse labels
- 15         Calculate loss and update the weight of  $\theta^s$
- 16         Update the weight of  $\theta^{dt}$  by EMA

**output:** Student  $\theta^s$  and dynamic teacher  $\theta^{dt}$

---

$[-140.8m, 140.8m], y \in [-38.4m, 38.4m]$ . The training set includes 358,142 fully supervised objects and 21,139 sparsely supervised objects. **DAIR-V2X** (Yu et al. 2022) is the first real-world V2I collaborative perception dataset. We set the LiDAR range as  $x \in [-100.8m, 100.8m], y \in [-40m, 40m]$ . The training set has 237,725 fully supervised objects and 9,622 sparsely supervised objects. **V2V4Real** (Xu et al. 2023) is the first real-world V2V collaborative perception dataset. The LiDAR range is set as  $x \in [-140.8m, 140.8m], y \in [-38.4m, 38.4m]$ . The training set comprises 145,278 fully supervised objects and 14,210 sparsely supervised objects.

### Experimental Setup

**Implementation Details** We implement the CoDTS with OpenCOD framework (Xu et al. 2022b). Specifically, the collaborative detector is trained using the Adam optimizer with a learning rate of 0.002. The parameters for MFM, SFM, and NAS are set as follows:  $\sigma_l^{st} = 0.15, \sigma_h^{st} = 0.2, \tau = 0.15, \tau^{nei} = 0.6$ . For the EMA strategy, we set  $\alpha = 0.999$ . The refinement starts iteration  $I_{refine}$  is set to  $0.5 \cdot I_{max}$ . All models are trained on NVIDIA RTX A4000.

**Benchmarks** We compare our CoDTS with sparsely supervised methods SSC3OD (Han et al. 2023b) and HINTED (Xia et al. 2024). Our static teacher (ST) is the same as that of SSC3OD, which first applies the self-supervised learning method to pre-train the encoder. We then load the pre-trained encoder into ST and train it with sparse labels. To enable quantitative comparison, we utilize three classical intermediate collaborative detectors: F-Cooper (Chen et al. 2019), AttFusion (Xu et al. 2022b), and DiscoNet (Li et al. 2021).

Datasets			V2X-Sim			OPV2V			DAIR-V2X			V2V4Real		
Detectors	Label	Methods	AP@0.3	AP@0.5	AP@0.7	AP@0.3	AP@0.5	AP@0.7	AP@0.3	AP@0.5	AP@0.7	AP@0.3	AP@0.5	AP@0.7
F-Cooper	Full	-	79.69	71.59	56.75	94.23	88.44	69.16	80.11	73.37	56.51	81.35	72.10	40.73
	Sparse	-	62.62	52.14	38.72	80.83	74.75	51.96	54.41	49.49	32.58	46.61	43.91	23.09
	Sparse	ST	67.49	60.31	45.60	82.32	76.89	54.26	59.68	54.26	35.39	53.62	50.70	30.51
	Sparse	SSC3OD	71.20	62.49	50.09	84.40	79.08	57.54	60.89	56.68	40.53	55.74	52.46	33.22
	Sparse	HINTED	71.08	61.15	46.48	87.23	80.41	55.10	70.68	62.86	38.68	71.90	61.72	35.37
	Sparse	<b>CoDTS</b>	<b>73.77</b>	<b>65.52</b>	<b>53.02</b>	<b>90.83</b>	<b>84.05</b>	<b>62.04</b>	<b>78.17</b>	<b>70.35</b>	<b>50.53</b>	<b>71.91</b>	<b>65.60</b>	<b>38.03</b>
AttFusion	Full	-	79.07	76.47	66.43	95.55	93.49	83.03	79.82	72.89	57.43	79.81	70.37	45.20
	Sparse	-	67.02	65.53	56.53	78.74	77.63	61.98	56.31	51.21	32.93	48.66	46.02	26.09
	Sparse	ST	69.69	68.41	59.06	80.12	79.19	68.86	57.32	52.94	38.05	55.06	50.66	33.56
	Sparse	SSC3OD	74.10	72.68	62.66	82.35	81.64	72.55	62.02	57.17	40.57	56.75	53.88	35.52
	Sparse	HINTED	74.42	72.07	60.37	86.24	81.50	60.76	71.59	65.42	42.37	69.74	61.67	34.77
	Sparse	<b>CoDTS</b>	<b>77.84</b>	<b>74.86</b>	<b>65.66</b>	<b>91.23</b>	<b>89.42</b>	<b>75.77</b>	<b>76.95</b>	<b>69.23</b>	<b>51.24</b>	<b>78.09</b>	<b>69.30</b>	<b>41.30</b>
DiscoNet	Full	-	81.02	77.63	69.29	96.22	94.31	84.42	80.21	73.62	58.00	79.25	68.61	42.61
	Sparse	-	68.56	65.39	55.84	81.82	80.79	69.23	55.90	51.39	32.97	54.15	50.22	33.23
	Sparse	ST	70.97	68.21	59.06	82.08	81.19	72.33	58.10	53.94	38.93	56.41	51.87	31.44
	Sparse	SSC3OD	74.07	70.54	62.29	83.12	82.22	73.56	60.93	56.78	42.63	53.66	49.44	30.44
	Sparse	HINTED	77.12	73.18	60.60	87.20	83.24	69.61	70.19	63.78	41.00	74.87	65.34	37.95
	Sparse	<b>CoDTS</b>	<b>81.23</b>	<b>76.41</b>	<b>66.18</b>	<b>93.41</b>	<b>90.56</b>	<b>76.47</b>	<b>75.82</b>	<b>68.73</b>	<b>51.05</b>	<b>80.45</b>	<b>70.56</b>	<b>38.55</b>

Table 1: Performance comparison among different methods on sparsely supervised collaborative 3D object detection.

Detectors	Methods	Semi-10%		Semi-20%	
		AP@0.5	AP@0.7	AP@0.5	AP@0.7
F-Cooper	3DIoUMatch	56.81	28.83	66.86	45.06
	HSSDA	58.22	30.41	66.99	46.54
	HINTED	58.93	30.42	67.15	46.22
	<b>CoDTS</b>	<b>59.54</b>	<b>32.82</b>	<b>69.31</b>	<b>48.46</b>
AttFusion	3DIoUMatch	54.56	31.64	63.86	45.91
	HSSDA	57.09	34.38	63.95	47.56
	HINTED	57.53	34.88	63.76	47.19
	<b>CoDTS</b>	<b>58.89</b>	<b>37.26</b>	<b>66.22</b>	<b>50.23</b>
DiscoNet	3DIoUMatch	58.78	34.88	65.96	46.81
	HSSDA	59.51	36.56	66.45	48.28
	HINTED	60.21	<b>38.18</b>	66.39	48.56
	<b>CoDTS</b>	<b>60.61</b>	37.94	<b>67.31</b>	<b>49.37</b>

Table 2: Comparison with state-of-the-art semi-supervised methods on DAIR-V2X. All methods are trained on randomly selected 10% or 20% fully annotated frames.

We keep the same supervised loss as these detectors used.

## Quantitative Evaluation

**Comparison with Sparsely Supervised Methods** Table 1 presents the collaboration detection results in both fully and sparsely supervised paradigms. The evaluation is conducted on four large-scale collaborative perception datasets. Due to insufficient supervision, collaborative detectors trained with sparse labels have lower accuracy than those trained with full labels. Specifically, SSC3OD and HINTED partially address this issue, but the performance of these detectors still significantly lags behind those trained with full labels. This is because these methods mainly rely on high-quality pseudo labels filtered by high confidence thresholds.

Our CoDTS mines pseudo labels with both high-quality and high-quantity, enabling us to achieve the highest per-

Detectors	Modules			Student $\theta^s$		Dynamic Teacher $\theta^{dt}$	
	MFM	SFM	NAS	AP@0.5	AP@0.7	AP@0.5	AP@0.7
F-Cooper	-	-	-	49.49	32.58	49.49	32.58
	✓	-	-	66.90	42.92	67.90	47.33
	✓	-	✓	66.86	47.11	68.55	48.72
	✓	✓	✓	68.73	48.05	<b>70.35</b>	<b>50.53</b>
AttFusion	-	-	-	51.21	32.93	51.21	32.93
	✓	-	-	65.58	48.77	67.25	50.35
	✓	-	✓	67.87	50.66	68.51	51.22
	✓	✓	✓	65.46	41.05	<b>69.23</b>	<b>51.24</b>
DiscoNet	-	-	-	51.39	32.97	51.39	32.97
	✓	-	-	65.40	48.61	67.06	49.10
	✓	-	✓	66.16	49.05	67.60	49.77
	✓	✓	✓	67.27	45.73	<b>68.73</b>	<b>51.05</b>

Table 3: Ablation study of main components on DAIR-V2X.

formance in sparsely supervised scenarios. Specifically, in the V2X-Sim, the detector trained with the CoDTS performs similarly to a fully supervised learning approach at IoU=0.3,0.5 and 0.7. In the OPV2V, our CoDTS improves the AP of collaborative detectors by over 10%. In the DAIR-V2X and V2VReal, our CoDTS improves the AP of collaborative detectors by over 20% at IoU=0.3 and 0.5, and by approximately 10% at IoU=0.7. These results demonstrate the effectiveness of the CoDTS in generating pseudo labels.

**Comparison with Semi-supervised Methods** Given the limited research on sparsely supervised scenes and the fact that our proposed CoDTS employs a teacher-student framework similar to those used in semi-supervised methods, we also explore its performance in semi-supervised settings. Table 2 presents a comparison of our results with various semi-supervised methods: 3DIoUMatch (Wang et al. 2021), HSSDA (Liu et al. 2023), and HINTED (Xia et al. 2024). Since it is difficult to pre-train collaborative detectors with

MFM		F-Cooper		AttFusion		DiscoNet	
$\sigma_i^{st}$	$\sigma_h^{st}$	AP@0.5	AP@0.7	AP@0.5	AP@0.7	AP@0.5	AP@0.7
0.1	0.2	69.61	49.80	68.83	51.02	68.52	<b>51.23</b>
0.1	0.25	69.03	49.22	68.72	51.31	67.84	49.33
0.1	0.3	67.93	48.92	67.37	50.11	67.86	49.26
0.15	0.2	70.35	<b>50.53</b>	69.23	<b>51.24</b>	<b>68.73</b>	51.05
0.15	0.25	<b>70.43</b>	50.02	<b>69.56</b>	51.10	67.77	50.49
0.15	0.3	70.13	49.92	68.92	50.74	67.51	50.25

Table 4: Ablation study of  $\sigma_i^{st}$  and  $\sigma_h^{st}$  on DAIR-V2X.

a very low semi-supervised ratio, we randomly choose 10% and 20% of fully annotated frames from the training split. Our method achieves the best performance in both 10% and 20% labeling rates, which demonstrates the superiority of our CoDTS in producing pseudo labels.

## Ablation Studies

**Contribution of Main Components** We conduct ablation studies to assess the effectiveness of the proposed MFM, SFM and NAS modules. We take collaborative detectors trained from scratch as a baseline and gradually incorporate each component. As shown in Table 3, when we use MFM with a high confidence threshold and utilize the identified positive instances as pseudo labels, the performance of the student and dynamic teacher is improved.

The performance is further enhanced when we combine MFM and NAS to generate pseudo labels, emphasizing the importance of sampling neighbor instances. By integrating MFM, SFM, and NAS, and implementing the staged training strategy, we achieve the best performance. This demonstrates the effectiveness of adaptive complementary learning in mining high-quality and high-quantity pseudo labels. In the training process, since the dynamic teacher can be viewed as a temporal ensemble of students across different time steps, the dynamic teacher consistently outperforms the student. Furthermore, as employing varying confidence thresholds at different stages may impede the performance of the student, we choose the dynamic teacher for inference.

**Confidence Threshold for Static Teacher** This section evaluates the influence of static teacher thresholds  $\sigma_i^{st}$  and  $\sigma_h^{st}$  on performance. Given that the OpenCOOD benchmark only considers predictions with scores above 0.2 during inference, and the maximum prediction score of the static teacher is below 0.35, we establish the low confidence threshold interval as (0, 0.2) and the high threshold interval as [0.2, 0.3]. We randomly select thresholds from the above intervals. Specifically, we set low confidence thresholds  $\sigma_i^{st}$  at 0.1 and 0.15, and high confidence thresholds  $\sigma_h^{st}$  at 0.2, 0.25, and 0.3. The results of different combinations of these confidence thresholds are in Table 4.

The detection results across these combinations show no significant difference, with all achieving higher performance than the baseline methods in Table 1. This indicates that our framework does not rely on manually set low or high thresholds. As long as the MFM threshold falls within an appropriate range, our framework performs effectively. Additionally, we tried to apply dynamic threshold adaptation to MFM but

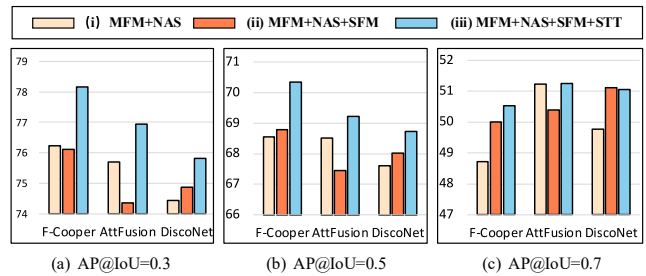


Figure 5: Ablation study of staged training strategy (STT) on DAIR-V2X.

found the training to be unstable. This instability is due to the lower performance of the static teacher and its unstable predictions on sparse labels.

**The Importance of Staged Training Strategy** We evaluate the importance of the staged training strategy (STT) and present the results in Figure 5. The settings for the three methods are as follows: i) involves high confidence threshold-based MFM and NAS; ii) involves high confidence threshold-based MFM, NAS, and SFM. Specifically, we first train student and teacher models using a low-confidence threshold-based MFM and NAS. In the next stage, we take this teacher model as an auxiliary static teacher and apply SFM to identify missing instances. iii) involves MFM, NAS, and SFM, and utilizes the staged training strategy to achieve end-to-end training.

The results from (i) to (ii) indicate that directly employing SFM could improve collaborative detectors when IoU=0.7, since it provides high-quality missing positive instances for the MFM. However, its performance remains inferior to that of using STT, especially at IoU=0.3 and 0.5. This is because, two stages in (ii) are trained separately. In the second stage, the student is optimized from scratch, while the dynamic teacher remains static without further updates. In contrast, STT in (iii) enables end-to-end training, allowing the student to undergo continuous optimization. Additionally, the dynamic teacher is updated using the EMA of the student and generates partial labels to supervise the student, enabling their mutual learning. Therefore, the dynamic teacher can generate more complemented pseudo labels adaptively.

## Conclusion

In this paper, we propose a collaborative perception dual teacher-student framework, CoDTS, which leverages adaptive complementary learning to generate pseudo labels in both high quality and high quantity. Specifically, the MFM generates high-quality pseudo labels, while the SFM adaptively identifies missing instances using a dynamic threshold. Furthermore, the NAS enhances the dense representation of pseudo labels during training. The framework employs a staged training strategy, where the warm-up stage pre-trains the student and dynamic teacher, and the refinement stage produces pseudo labels via adaptive complementary learning. Extensive experiments on four large-scale datasets demonstrate that CoDTS effectively addresses the imbalance of pseudo labels between quality and quantity.

## Acknowledgments

This work was supported by the National Natural Science Foundation of China under Grant U2268203, 62203040 and 62306027.

## References

- Chen, Q.; Ma, X.; Tang, S.; Guo, J.; Yang, Q.; and Fu, S. 2019. F-Cooper: Feature based cooperative perception for autonomous vehicle edge computing system using 3D point clouds. In *SEC*, 88–100.
- Geiger, A.; Lenz, P.; and Urtasun, R. 2012. Are we ready for autonomous driving? the kitti vision benchmark suite. In *CVPR*, 3354–3361.
- Han, Y.; Zhang, H.; Li, H.; Jin, Y.; Lang, C.; and Li, Y. 2023a. Collaborative Perception in Autonomous Driving: Methods, Datasets, and Challenges. *ITSM*, 15(6): 131–151.
- Han, Y.; Zhang, H.; Zhang, H.; and Li, Y. 2023b. SSC3OD: Sparsely supervised collaborative 3d object detection from lidar point clouds. In *SMC*, 1360–1367.
- Hong, S.; Liu, Y.; Li, Z.; Li, S.; and He, Y. 2024. Multi-agent Collaborative Perception via Motion-aware Robust Communication Network. In *CVPR*, 15301–15310.
- Hu, Y.; Fang, S.; Lei, Z.; Zhong, Y.; and Chen, S. 2022. Where2comm: Communication-Efficient Collaborative Perception via Spatial Confidence Maps. In *NeurIPS*, volume 35, 4874–4886.
- Li, Y.; Ma, D.; An, Z.; Wang, Z.; Zhong, Y.; Chen, S.; and Feng, C. 2022. V2X-Sim: Multi-Agent Collaborative Perception Dataset and Benchmark for Autonomous Driving. *RAL*, 7(4): 10914–10921.
- Li, Y.; Ren, S.; Wu, P.; Chen, S.; Feng, C.; and Zhang, W. 2021. Learning distilled collaboration graph for multi-agent perception. In *NeurIPS*, volume 34, 29541–29552.
- Liu, C.; Gao, C.; Liu, F.; Li, P.; Meng, D.; and Gao, X. 2023. Hierarchical Supervision and Shuffle Data Augmentation for 3D Semi-Supervised Object Detection. In *CVPR*, 23819–23828.
- Liu, C.; Gao, C.; Liu, F.; Liu, J.; Meng, D.; and Gao, X. 2022. SS3D: Sparsely-Supervised 3D Object Detection from Point Cloud. In *CVPR*, 8428–8437.
- Luo, G.; Zhang, H.; Yuan, Q.; and Li, J. 2022a. Complementarity-Enhanced and Redundancy-Minimized Collaboration Network for Multi-agent Perception. In *ACM MM*, 3578–3586.
- Luo, G.; Zhang, H.; Yuan, Q.; Li, J.; Wang, W.; and Wang, F.-Y. 2023. One size fits all: A unified traffic predictor for capturing the essential spatial-temporal dependency. *TNNLS*.
- Luo, Y.; Chen, Z.; Wang, Z.; Yu, X.; Huang, Z.; and Bakhtashmotlagh, M. 2022b. Exploring active 3d object detection from a generalization perspective. In *ICLR*.
- MacQueen, J.; et al. 1967. Some methods for classification and analysis of multivariate observations. In *BSMSP*, volume 1, 281–297.
- Meng, Q.; Wang, W.; Zhou, T.; Shen, J.; Van Gool, L.; and Dai, D. 2020. Weakly supervised 3d object detection from lidar point cloud. In *ECCV*, 515–531.
- Neubeck, A.; and Van Gool, L. 2006. Efficient non-maximum suppression. In *ICPR*, volume 3, 850–855.
- Song, S.; Lichtenberg, S. P.; and Xiao, J. 2015. Sun rgb-d: A rgb-d scene understanding benchmark suite. In *CVPR*, 567–576.
- Su, W.; Chen, L.; Bai, Y.; Lin, X.; Li, G.; Qu, Z.; and Zhou, P. 2024. What Makes Good Collaborative Views? Contrastive Mutual Information Maximization for Multi-Agent Perception. In *AAAI*, volume 38, 17550–17558.
- Sun, P.; Kretschmar, H.; Dotiwalla, X.; Chouard, A.; Patnaik, V.; Tsui, P.; Guo, J.; Zhou, Y.; Chai, Y.; Caine, B.; et al. 2020. Scalability in perception for autonomous driving: Waymo open dataset. In *CVPR*, 2446–2454.
- Wang, H.; Cong, Y.; Litany, O.; Gao, Y.; and Guibas, L. J. 2021. 3DIoUMatch: Leveraging IoU Prediction for Semi-Supervised 3D Object Detection. In *CVPR*, 14615–14624.
- Xia, Q.; Ye, W.; Wu, H.; Zhao, S.; Xing, L.; Huang, X.; Deng, J.; Li, X.; Wen, C.; and Wang, C. 2024. HINTED: Hard Instance Enhanced Detector with Mixed-Density Feature Fusion for Sparsely-Supervised 3D Object Detection. In *CVPR*, 15321–15330.
- Xu, R.; Xia, X.; Li, J.; Li, H.; Zhang, S.; Tu, Z.; Meng, Z.; Xiang, H.; Dong, X.; Song, R.; et al. 2023. V2V4Real: A real-world large-scale dataset for vehicle-to-vehicle cooperative perception. In *CVPR*, 13712–13722.
- Xu, R.; Xiang, H.; Tu, Z.; Xia, X.; Yang, M.-H.; and Ma, J. 2022a. V2X-ViT: Vehicle-to-everything cooperative perception with vision transformer. In *ECCV*, 107–124.
- Xu, R.; Xiang, H.; Xia, X.; Han, X.; Li, J.; and Ma, J. 2022b. OPV2V: An open benchmark dataset and fusion pipeline for perception with vehicle-to-vehicle communication. In *ICRA*, 2583–2589.
- Yang, D.; Yang, K.; Wang, Y.; Liu, J.; Xu, Z.; Yin, R.; Zhai, P.; and Zhang, L. 2024. How2comm: Communication-efficient and collaboration-pragmatic multi-agent perception. In *NeurIPS*, volume 36.
- Yu, H.; Luo, Y.; Shu, M.; Huo, Y.; Yang, Z.; Shi, Y.; Guo, Z.; Li, H.; Hu, X.; Yuan, J.; et al. 2022. DAIR-V2X: A large-scale dataset for vehicle-infrastructure cooperative 3d object detection. In *CVPR*, 21361–21370.
- Yu, H.; Tang, Y.; Xie, E.; Mao, J.; Luo, P.; and Nie, Z. 2024. Flow-Based Feature Fusion for Vehicle-Infrastructure Cooperative 3D Object Detection. In *NeurIPS*, volume 36.
- Zhang, H.; Luo, G.; Li, Y.; and Wang, F.-Y. 2022. Parallel vision for intelligent transportation systems in metaverse: Challenges, solutions, and potential applications. *TSMCS*, 53(6): 3400–3413.
- Zhang, J.; Yang, K.; Wang, Y.; Wang, H.; Sun, P.; and Song, L. 2024. ERMVP: Communication-Efficient and Collaboration-Robust Multi-Vehicle Perception in Challenging Environments. In *CVPR*, 12575–12584.