

# Commonsense Prototype for Outdoor Unsupervised 3D Object Detection

Hai Wu<sup>1</sup> Shijia Zhao<sup>1</sup> Xun Huang<sup>1</sup> Chenglu Wen<sup>1\*</sup> Xin Li<sup>2</sup> Cheng Wang<sup>1</sup>

<sup>1</sup>Fujian Key Laboratory of Sensing and Computing for Smart Cities, Xiamen University

<sup>2</sup>Section of Visual Computing and Interactive Media, Texas A&M University

## Abstract

The prevalent approaches of unsupervised 3D object detection follow cluster-based pseudo-label generation and iterative self-training processes. However, the challenge arises due to the sparsity of LiDAR scans, which leads to pseudo-labels with erroneous size and position, resulting in subpar detection performance. To tackle this problem, this paper introduces a Commonsense Prototype-based Detector, termed CPD, for unsupervised 3D object detection. CPD first constructs Commonsense Prototype (CProto) characterized by high-quality bounding box and dense points, based on commonsense intuition. Subsequently, CPD refines the low-quality pseudo-labels by leveraging the size prior from CProto. Furthermore, CPD enhances the detection accuracy of sparsely scanned objects by the geometric knowledge from CProto. CPD outperforms state-of-the-art unsupervised 3D detectors on Waymo Open Dataset (WOD), PandaSet, and KITTI datasets by a large margin. Besides, by training CPD on WOD and testing on KITTI, CPD attains **90.85%** and **81.01%** 3D Average Precision on easy and moderate car classes, respectively. These achievements position CPD in **close proximity to fully supervised detectors**, highlighting the significance of our method. The code will be available at <https://github.com/hailanyi/CPD>.

## 1. Introduction

Autonomous driving requires reliable detection of 3D objects (e.g. vehicle and cyclist) in urban scenes for safe path planning and navigation. Thanks to the power of neural networks, numerous studies have developed high-performance 3D detectors through fully supervised approaches [4, 15, 30–33]. However, these models heavily depend on human annotations from diverse scenes to guarantee their effectiveness across various scenarios. This data labeling process is typically laborious and time-consuming, limiting the wide deployment of detectors in practice [40].

\*Corresponding author

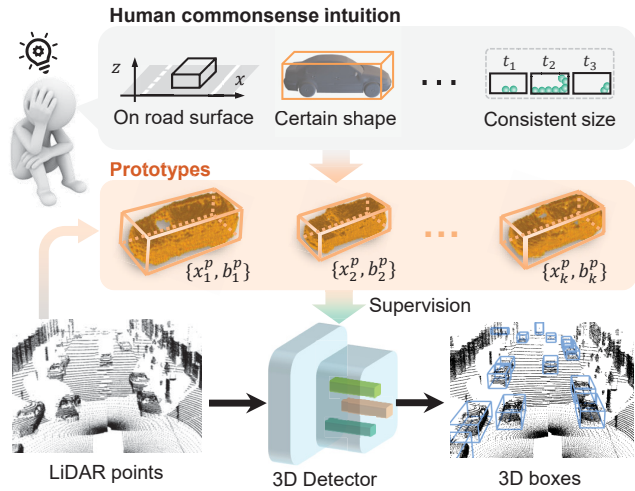


Figure 1. Illustration of commonsense prototypes for unsupervised 3D object detection in autonomous driving scenes.

Several studies have explored approaches to reduce labeling requirements by weakly supervised learning [3, 26, 46], decreasing the label cost by over 80%. Notably, the objects within a 3D scene exhibit distinguishable attributes and can be easily identified through certain commonsense reasoning (see Fig. 1). For example, the objects are usually located on the ground surface with a certain shape; the object sizes are fixed across frames. This insight has prompted us to develop an *unsupervised 3D detector* that operates without using human annotations.

In recent years, traditional methods leveraged ground removal [9] and clustering technique [42] for unsupervised 3D object detection. However, these methods often struggle to achieve satisfactory performance due to the sparsity and occlusion of objects in 3D scenes. Advanced methods create initial pseudo-labels from point cloud sequences by clustering and bootstrap a good detector by iteratively training a deep network [41]. Nevertheless, the sparse and view-limited nature of LiDAR scanning leads to pseudo-labels with inaccurate sizes and positions, misleading the network convergence and resulting in suboptimal detection performance. A subset of objects, denoted as *complete objects*  $\mathcal{T}$ ,

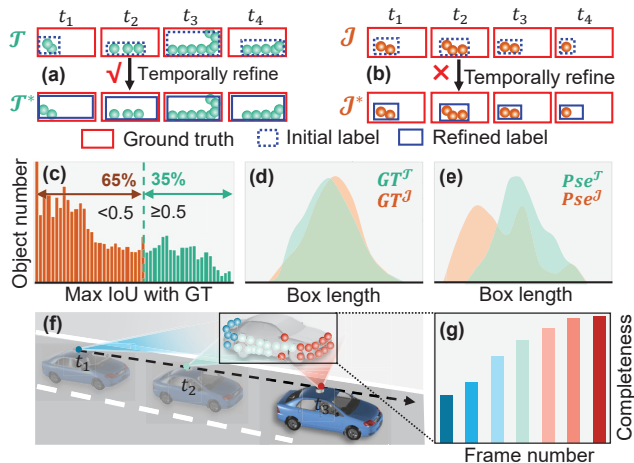


Figure 2. Illustration and statistics of complete and incomplete objects on WOD [25] validation set (large enough to demonstrate the general problem). (a) Pseudo-labels of *complete object*  $\mathcal{T}$  are refined by temporal consistency. (b) Pseudo-labels of *incomplete object*  $\mathcal{J}$  fail to be refined by temporal consistency. (c) 65% objects lack full scan coverage and generate inaccurate pseudo-labels (Max IoU (Intersection over Union)  $< 0.5$  with GT (Ground Truth)). (d) The vehicle GT of *complete object*  $GT^{\mathcal{T}}$  and *incomplete object*  $GT^{\mathcal{J}}$  have similar size distributions. (e) The pseudo-label of *complete object*  $Pse^{\mathcal{T}}$  and *incomplete object*  $Pse^{\mathcal{J}}$  have different size distributions. (f)(g) The nearby stationary objects are with high completeness in consecutive frames.

benefit from having at least one complete scan across the entire point cloud sequence, allowing their pseudo-labels to be refined through temporal consistency [41] (see Fig. 2 (a)). However, the majority of objects (e.g. 65% on WOD [25], as shown in Fig. 2 (c)), termed *incomplete objects*  $\mathcal{J}$ , lack full scan coverage (see Fig. 2 (b)), and cannot be recovered by temporal consistency.

To tackle this issue, this paper proposes a Commonsense Prototype-based Detector, termed **CPD**, for *unsupervised 3D object detection*. CPD is built upon two key insights: (1) The ground truth of intra-class objects keeps a similar size (length, width, and height) distribution between incomplete objects and complete objects (see Fig. 2 (d)). (2) The nearby stationary objects are very complete in consecutive frames and can be recognized accurately by commonsense intuition (see Fig. 2 (f)(g)). Our idea is to construct a Commonsense Prototype (**CProto**) set representing accurate geometry and size from complete objects to refine the pseudo-labels of incomplete objects and improve the detection accuracy. To this end, we first design an unsupervised Multi-Frame Clustering (**MFC**) method that yields high-recall initial pseudo-labels. Subsequently, we introduce an unsupervised Completeness and Size Similarity (**CSS**) score that selects high-quality labels to construct the CProto set. Furthermore, we design a CProto-constrained Box Regulariza-

tion (**CBR**) method to refine the pseudo-labels by incorporating the size prior from CProto. In addition, we develop CProto-constrained Self-Training (**CST**) that improves the detection accuracy of sparsely scanned objects by the geometry knowledge from CProto.

The effectiveness of our design is verified by experiments on widely used WOD [25], PandaSet [35], and KITTI dataset [6]. Besides, the individual components of our design are also verified by extensive experiments on WOD [25]. The main contributions of this work include:

- We propose a Commonsense Prototype-based Detector (**CPD**) for unsupervised 3D object detection. CPD outperforms state-of-the-art unsupervised 3D detectors by a large margin.
- We propose Multi-Frame Clustering (**MFC**) and CProto-constrained Box Regularization (**CBR**) for pseudo-label generation and refinement, greatly improving the recall and precision of pseudo-label.
- We propose CProto-constrained Self-Training (**CST**) for unsupervised 3D detection. It improves the recognition and localization accuracy of sparse objects, boosting the detection performance significantly.

## 2. Related Work

**Fully/weakly supervised 3D object detection.** Recent fully-supervised 3D detectors build single-stage [8, 10, 27, 39, 48, 49], two-stage [4, 20–22, 31–33, 37] or multiple stage [2, 30] deep networks for 3D object detection. However, these methods heavily rely on a large amount of precise annotations. Some weakly supervised methods replace the box annotation with low-cost click annotation [17]. Other methods decrease the supervision by only annotating a part of scenes [3, 26, 45, 46] or a part of instances [34]. Unlike all of the above works, we aim to design a 3D detector that does not require human-level annotations.

**Unsupervised 3D object detection.** Previous unsupervised pre-training methods discern latent patterns within the unlabeled data by masked labels [36] or contrastive loss [14, 38]. But these methods require human labels for fine-tuning. Traditional methods [1, 19, 24] employ ground removal and clustering for 3D object detection without human labels, but suffer from poor detection performance. Some deep learning-based methods generate pseudo-labels by clustering and use the pseudo-labels to train a 3D detector [40] iteratively. Recent OYSTER [41] improves pseudo-label quality with temporal consistency. However, most pseudo-labels of incomplete objects cannot be recovered by temporal consistency. Our CPD addresses this problem by leveraging the geometry prior from CProto to refine the pseudo-label and guide the network convergence.

**Prototype-based methods.** The prototype-based methods are widely used in 2D detection [11, 12, 16, 29, 44] when novel classes are incorporated. Inspired by these

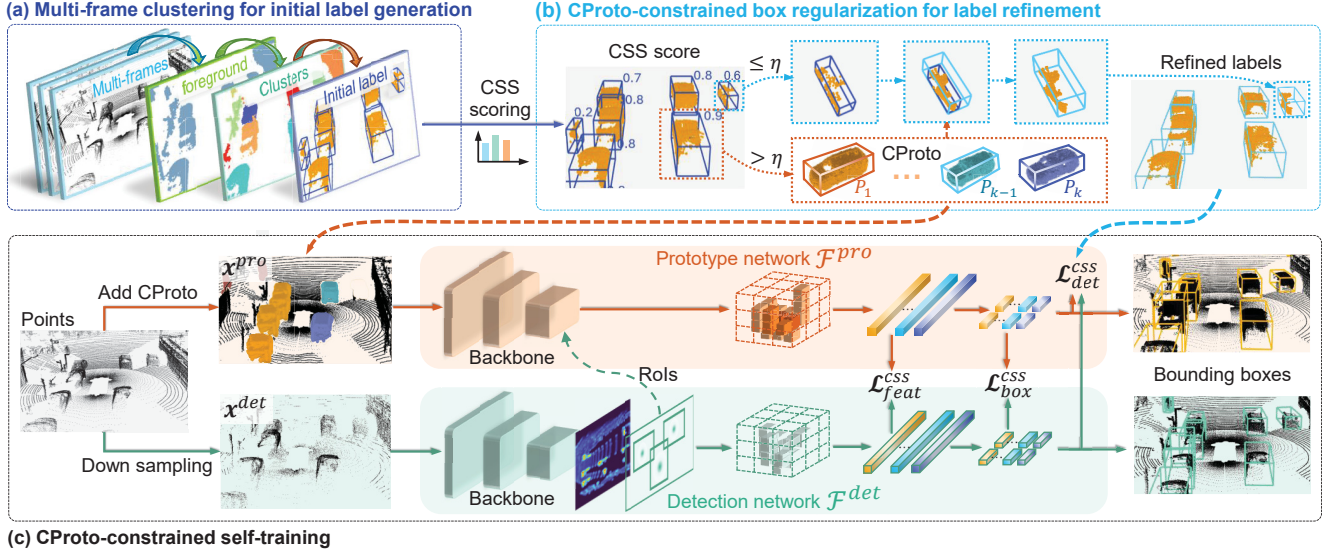


Figure 3. CPD framework. (a) Initial pseudo-labels are generated by multi-frame clustering. (b) The commonsense prototype (CProto) is constructed from high-quality pseudo-labels based on CSS score. The low-quality labels are further refined by the shape prior from CProto. (c) A prototype network fed with dense points from CProto produces high-quality features to guide the detection network convergence.

methods, Prototypical VoteNet [47] constructs geometric prototypes learned from basic classes for few-shot 3D object detection. GPA-3D [13] and CL3D [18] build geometric prototypes from a source-domain model for domain adaptive 3D detection. However, both the learning from basic class and training on the source domain require high-quality annotations. Unlike that, we construct CProto using commonsense knowledge and detect 3D objects in a zero-shot manner without human-level annotations.

### 3. CPD Method

This paper introduces the Commonsense Prototype-based Detector (CPD), a novel approach for unsupervised 3D object detection. As shown in Fig. 3, CPD consists of three main parts: (1) initial label generation; (2) label refinement; (3) self-training. We detail the designs as follows.

#### 3.1. Initial Label Generation

Recent unsupervised methods [40, 41] detect 3D objects in a class-agnostic way. How to classify objects (e.g. vehicle and pedestrian) without annotation is still an unsolved challenge. Our observations indicate that some stationary objects in consecutive frames, appear more complete (see Fig.2 (f)) and can be classified by predefined sizes. This motivates us to design a Multi-Frame Clustering (MFC) method to generate initial labels. MFC involves motion artifact removal, clustering, and post-processing.

**Motion Artifact Removal (MAR).** Directly transforming and concatenating  $2n + 1$  consecutive frames  $\{x_{-n}, \dots, x_n\}$  (i.e., past  $n$ , future  $n$ , and the current frame)

into a single point cloud  $x_0^*$  introduces motion artifacts from moving objects, leading to increased label errors as the  $n$  grows (see Fig. 4(a)). To mitigate this issue, we first transform the consecutive frames to global system and calculate the Persistence Point Score (PPScore)[40] by consecutive frames to identify the points in motion. We keep all the points from  $x_0$  and remove moving points from the other frames  $x_{-n}, \dots, x_{-1}, x_1, \dots, x_n$ . After this removal, we concatenate the frames to obtain dense points  $x_0^*$ .

**Clustering and post-processing.** In line with recent study [41], we apply the ground removal[9], DBSCAN [5] and bounding box fitting [43] on  $x_0^*$  to obtain a set of class-agnostic bounding boxes  $\hat{b}$ . We observe that the objects of the same class typically have similar sizes in 3D space. Therefore, we pre-define class-specific size thresholds (e.g. the length of vehicle is generally larger than 0.5m) based on human commonsense to classify  $\hat{b}$  into different categories. We then apply class-agnostic tracking to associate the small background objects with foreground trajectories, and enhance the consistency of objects' sizes by using temporal coherency [41]. This process results in a set of initial pseudo-labels  $\mathbf{b} = \{b_j\}_j$ , where  $b_j = [x, y, z, l, w, h, \alpha, \beta, \tau]$  represents position, width, length, height, azimuth angle, class identity, and tracking identity, respectively.

#### 3.2. CProto-constrained Box Regularization for Label Refinement

As noted in Section 1, initial labels for incomplete objects often suffer from inaccuracies in sizes and positions. To






Figure 4. (a) Length absolute error with different frames. (b) Multi-level occupancy score. (c) Mean size error of initial labels.

tackle this issue, we introduce the CProto-constrained Box Regularization (CBR) method. The key idea is to construct a high-quality CProto set based on unsupervised scoring from complete objects to refine the pseudo-labels of incomplete objects. Different from OYSTER [41], which can only refine the pseudo-labels of objects having at least one complete scan, our CBR can refine pseudo-labels of all objects, significantly decreasing the overall size and position errors.

**Completeness and Size Similarity (CSS) scoring.** Existing label scoring methods such as IoU scoring [21] are designed for fully supervised detectors. In contrast, we introduce an *unsupervised* Completeness and Size Similarity scoring (CSS) method. It aims to approximate the IoU score using commonsense knowledge alone (see Fig. 5).

*Distance score.* CSS first assesses the object completeness based on distance, assuming labels closer to the ego vehicle are likely to be more accurate. For an initial label  $b_j$ , we normalize the distance to the ego vehicle within the range [0,1] to compute the distance score as

$$\psi^1(b_j) = 1 - \mathcal{N}(\|c_j\|), \quad (1)$$

where  $\mathcal{N}$  is the normalization function and  $c_j$  is the location of  $b_j$ . However, this distance-based approach has its limitations. For example, occluded objects near the ego vehicle, which should receive lower scores, are inadvertently assigned high scores due to their proximity. To mitigate this issue, we introduce a Multi-Level Occupancy (MLO) score, further detailed in Fig. 4 (b).

*MLO score.* Considering the diverse sizes of objects, we divide the bounding box of the initial label into multiple grids with different length and width resolutions. The MLO score is then calculated by determining the proportion of grids occupied by cluster points, via

$$\psi^2(b_j) = \frac{1}{N^o} \sum_k \frac{O^k}{(r^k)^2}, \quad (2)$$

where  $N^o$  denotes resolution number,  $O^k$  is the number of occupied grids under  $k$ -th resolution, and  $r^k$  is the grid number of  $k$ -th resolution.

*Size Similarity (SS) score.* While the distance and MLO scores effectively evaluate the localization and size quality,




Figure 5. Completeness and size similarity scoring.

they fall short in assessing classification quality. To bridge this gap, we introduce the SS score. This score utilizes a class-specific template box  $a$  (average size of typical objects in Wikipedia) and calculates a truncated KL divergence [7]. Note that, this score is decided by ratio difference, rather than their specific values. Simple commonsense of  $l, w, h$  ratios (2:1:1 for Vehicle, 1:1:2 for Pedestrian, 2:1:2 for Cyclist) can also be used here.

$$\psi^3(b_j) = 1 - \min(0.05, \sum_{\sigma} q_{\sigma}^b \log(\frac{q_{\sigma}^b}{q_{\sigma}^a}))/0.05, \quad (3)$$

where  $q_{\sigma}^a \in \{l^a, w^a, h^a\}$ ,  $q_{\sigma}^b \in \{l^b, w^b, h^b\}$  refer to the normalized length, width, and height of the template and label.

We linearly combine the three metrics  $\mathcal{S}(b_j) = \sum_i \omega^i \psi^i(b_j)$  to produce final scoring, where  $\omega^i$  is the weighting factor (in this study we adopt a simple average,  $\omega^i = 1/3$ ). For each  $b_j \in \mathbf{b}$ , we compute its CSS score  $s_j^{css} = \mathcal{S}(b_j)$  and obtain a set of scores  $s = \{s_j^{css}\}_j$ .

**CProto set construction.** Regular learnable prototype-based methods require annotations [13, 47], which are unavailable in the unsupervised problem. We construct a high-quality CProto set  $\mathbf{P} = \{P_k\}_k$ , representing geometry and size centers based on the *unsupervised* CSS score. Here,  $P_k = \{x_k^p, b_k^p\}$ , where  $x_k^p$  indicates the inside points, and  $b_k^p$  refers to the bounding box. Specifically, we first categorize the initial labels  $\mathbf{b}$  into different groups based on their tracking identity  $\tau$ . Within each group, we select the high-quality boxes and inside points that meet a high CSS score threshold  $\eta$  (determined on validation set, using 0.8 in this study). Then, we transform all points and boxes into a local coordinate system, and obtain  $x_k^p$  by averaging the high-quality boxes and  $b_k^p$  by concatenating all the points.

**Box regularization.** We next regularize the initial labels by the size prior from CProto. Based on the statistics on WOD validation set [25], we observe that the height of the initial labels is relatively correct than length and width (see Fig. 4 (c)). Intuitively, the intra-class 3D objects with the same height have similar length and width. Therefore, we associate the initial label  $b_j$  with CProto  $P_k$  by the minimum difference in box height. The initial pseudo-labels with the same  $P_k$  and similar length and width are naturally classified into the same group. We then perform re-size and re-localization for each group to refine the pseudo-




Figure 6. The size of the initial label is replaced by the CProto box, and the position is also corrected.

labels. (1) *Re-size*. We directly replace the size of  $b_j$  using the length, width, and height of  $b_k^p \in P_k$ . (2) *Re-location*. Since points are mostly on the object’s surface and boundary, we divide the object into different bins and align the box boundary and orientation to the boundary point of the densest part (see Fig. 6). Finally, we obtain improved pseudo-labels  $\mathbf{b}^* = \{b_j^*\}_j$ .

### 3.3. CProto-constrained Self-Training (CST)

Recent methods [40, 41] utilize pseudo-labels for training 3D detectors. However, even after refinement, some pseudo-labels remain inaccurate, diminishing the effectiveness of correct supervision and potentially misleading the training process. To tackle these issues, we propose two designs: (1) *CSS-Weighted Detection Loss*, which assigns different training weights based on label quality to suppress false supervision signals. (2) *Geometry Contrast Loss*, which aligns predictions of sparsely scanned points with the dense CProto, thereby improving feature consistency.

**Network architecture.** We adopt a dense-sparse alignment architecture (Fig. 3 (c)), consisting of a prototype network  $\mathcal{F}^{pro}$  and a detection network  $\mathcal{F}^{det}$ , constructed from two-stage CenterPoint [39]. During training, for each  $b_j^*$ , we add its corresponding points  $x_k^p$  from CProto  $P_k$  to the scene to obtain a dense point cloud  $\mathbf{x}^{pro}$ . We feed  $\mathbf{x}^{pro}$  to  $\mathcal{F}^{pro}$  to produce relatively good features and detections. We then feed randomly downsampled points  $\mathbf{x}^{det}$  as a sparse sample to the  $\mathcal{F}^{det}$ . We align the features and detections from two branches by the detection loss and contrast loss. During testing, we feed points without downsampling to the detection network  $\mathcal{F}^{det}$  to perform detection.

**CSS weight.** Considering that the false pseudo-labels may mislead the network convergence, we first calculate a loss weight based on different label qualities. Formally, we convert a CSS score  $s_i^{css}$  of a pseudo-label to

$$\omega_i = \begin{cases} 0 & s_i^{css} < S_L \\ \frac{s_i^{css} - S_L}{S_H - S_L} & S_L < s_i^{css} < S_H \\ 1 & s_i^{css} > S_H \end{cases}, \quad (4)$$

where  $S_H$  and  $S_L$  are high/low-quality thresholds (we empirically set 0.7 and 0.4, respectively).

**CSS-weighted detection loss.** To decrease the influence of false labels, we formulate the CSS-weighted detection loss to refine  $N$  proposals

$$\mathcal{L}_{det}^{css} = \frac{1}{N} \sum_i \omega_i (\mathcal{L}_i^{pro} + \mathcal{L}_i^{det}), \quad (5)$$

where  $\mathcal{L}_i^{pro}$  and  $\mathcal{L}_i^{det}$  are detection losses [4] of  $\mathcal{F}^{pro}$  and  $\mathcal{F}^{det}$ , respectively. The losses are calculated by pseudo-labels  $\mathbf{b}^*$  and network predictions.

**Geometry contrast loss.** We formulate two contrast losses that minimize the feature and predicted box difference between the prototype and detection network. (1) *Feature contrast loss*.

For a foreground RoI  $r_i$  from the detection network, we extract features  $\mathbf{f}_i^p$  from the prototype network by voxel set abstract [4], and extract features  $\mathbf{f}_i^d$  from detection network. We then formulate the contrast loss by cosine distance:

$$\mathcal{L}_{feat}^{css} = \frac{1}{N_f} \sum_i \omega_i \frac{\mathbf{f}_i^d \cdot \mathbf{f}_i^p}{\|\mathbf{f}_i^d\| \|\mathbf{f}_i^p\|}, \quad (6)$$

where  $N_f$  is the foreground proposal number. (2) *Box contrast loss*. For a box prediction  $d_i^p$  from the prototype network and a box prediction  $d_i^d$  from the detection network. We then formulate the box contrast loss by IoU, location difference, and angle difference:

$$\mathcal{L}_{box}^{css} = \frac{1}{N_f} \sum_i \omega_i [1 - I(d_i^d, d_i^p) + \|c_i^d - c_i^p\| + |\sin(\alpha_i^d - \alpha_i^p)|], \quad (7)$$

where  $I$  denote IoU function;  $c_i^d, \alpha_i^d$  refers to position and angle of  $d_i^d$ ;  $c_i^p, \alpha_i^p$  refers to position and angle of  $d_i^p$ . We finally summat all losses to training the detector.

## 4. Experiments

### 4.1. Datasets

**Waymo Open Dataset (WOD).** We conducted extensive experiments on the WOD [25] due to its diverse scenes. The WOD contains 798, 202 and 150 sequences for training, validation and testing, respectively. We adopted similar metrics (3D AP L1 and L2) as fully/weakly supervised methods [31, 34]. *No annotations* were used for training.

**PandaSet dataset.** To compare with recent unsupervised methods [41], we also conducted experiments on the PandaSet [35]. Like [41], we split the dataset into 73 training and 30 validation snippets and use class-agnostic BEV AP and recall metrics with 0.3, 0.5, and 0.7 IoU thresholds.

**KITTI dataset.** Since the KITTI detection dataset [6] did not provide consecutive frames, we only tested our method on the 3769 val split [4]. We used similar metrics (Car 3D AP R40 with 0.5 and 0.7 IoU thresholds) as employed in fully/weakly supervised methods [32, 34].

Method	Vehicle 3D AP				Pedestrian 3D AP				Cyclist 3D AP			
	L1		L2		L1		L2		L1		L2	
	$IoU_{0.5}$	$IoU_{0.7}$	$IoU_{0.5}$	$IoU_{0.7}$	$IoU_{0.3}$	$IoU_{0.5}$	$IoU_{0.3}$	$IoU_{0.5}$	$IoU_{0.3}$	$IoU_{0.5}$	$IoU_{0.3}$	$IoU_{0.5}$
DBSCAN [5]	2.32	0.29	1.94	0.25	0.51	0.00	0.19	0.00	0.28	0.03	0.20	0.00
DBSCAN init-train [40]	17.36	2.65	14.87	2.29	1.65	0.00	1.35	0.00	0.48	0.25	0.43	0.20
MODEST [40]	18.51	6.46	15.83	5.48	11.83	0.17	8.96	0.10	1.47	1.14	1.17	1.01
OYSTER [41]	30.48	14.66	26.21	14.10	4.33	0.18	3.52	0.14	1.27	0.33	1.24	0.32
Proto-vanilla	35.22	20.19	31.58	18.36	17.60	10.34	14.62	8.59	4.21	3.45	3.80	3.31
<b>CPD (Ours)</b>	<b>57.79</b>	<b>37.40</b>	<b>50.18</b>	<b>32.13</b>	<b>21.91</b>	<b>16.31</b>	<b>18.01</b>	<b>13.22</b>	<b>5.83</b>	<b>5.06</b>	<b>5.61</b>	<b>4.87</b>

Table 1. Unsupervised 3D object detection results on WOD validation set. The results of previous methods are reproduced by us.

Method	3D AP L1 ( $IoU_{0.7,0.5,0.5}$ )			3D AP L2 ( $IoU_{0.7,0.5,0.5}$ )		
	Vehicle	Ped.	Cyclist	Vehicle	Ped.	Cyclist
MODEST [40]	7.5	0.0	0.0	6.5	0.0	0.0
OYSTER [41]	21.6	0.6	0.0	18.7	0.5	0.0
<b>CPD (Ours)</b>	<b>37.2</b>	<b>18.6</b>	<b>5.7</b>	<b>32.4</b>	<b>16.5</b>	<b>5.5</b>

Table 2. Unsupervised 3D detection results on WOD test set.

Method	BEV AP			BEV Recall		
	$IoU_{0.3}$	$IoU_{0.5}$	$IoU_{0.7}$	$IoU_{0.3}$	$IoU_{0.5}$	$IoU_{0.7}$
MODEST [40]	22.0	7.5	2.8	49.7	28.9	14.9
OYSTER [41]	43.5	29.5	18.1	62.8	44.8	28.1
<b>CPD (Ours)</b>	<b>50.7</b>	<b>41.0</b>	<b>24.6</b>	<b>63.1</b>	<b>54.8</b>	<b>37.4</b>

Table 3. The class-agnostic comparison results on the PandaSet dataset, evaluated on the 0-80m detection range.

## 4.2. Implementation Details

**Network details.** Both prototype and detection networks adopt the same 3D backbone as CenterPoint [39] and the same RoI refinement network as Voxel-RCNN [4]. For the WOD and KITTI datasets, we use the same detection range and voxel size as CenterPoint [39]. For the Pandaset, we use the same detection range as OYSTER [41].

**Training details.** We adopt the widely used global scaling and rotation data augmentation. We trained our network on 8 Tesla V100 GPUs with the ADAM optimizer. We used a learning rate of 0.003 with a one-cycle learning rate strategy. We trained the CPD for 20 epochs.

## 4.3. Comparison with Unsupervised Detectors

**Results on WOD.** The results on the WOD validation set and test set are presented in Table 1 and Table 2. All methods use identical size thresholds to define the object classes and use single traversal. Our method significantly outperforms existing unsupervised methods. Notably, under the 3D AP L2 with IoU thresholds of 0.7, 0.5, and 0.5, our CPD outperforms OYSTER [41] by 18.03%, 13.08%, and 4.55% on Vehicle, Pedestrian, and Cyclist, respectively. These advancements come from our MFC, CBR, and CST designs, which yield superior pseudo-labels and enhanced detection accuracy. CPD also surpasses the Proto-vanilla method, which uses class-specific prototype [23].

Method	Labels	3D AP @ $IoU_{0.5}$			3D AP @ $IoU_{0.7}$		
		Easy	Mod.	Hard	Easy	Mod.	Hard
<i>CenterPoint</i> [39]	100%	97.07	89.23	81.81	88.55	78.38	71.43
<i>Sparsely-sup.</i> [34]	2%	-	-	-	49.69	31.55	25.91
MODEST [40]	0	47.56	33.43	30.57	12.65	11.14	10.60
OYSTER [41]	0	65.33	54.82	43.59	23.22	20.31	19.97
<b>CPD (Ours)</b>	0	<b>90.85</b>	<b>81.01</b>	<b>79.80</b>	<b>72.98</b>	<b>55.07</b>	<b>53.94</b>

Table 4. Car detection comparison with fully/weakly supervised detectors on KITTI val set. The models are trained on WOD.

Method	Labels	3D AP L1		3D AP L2	
		$IoU_{0.5}$	$IoU_{0.7}$	$IoU_{0.5}$	$IoU_{0.7}$
<i>CenterPoint</i> [39]	100%	89.23	73.72	78.52	65.52
<i>Sparsely-sup.</i> [34]	2%	-	32.15	-	27.97
MODEST [40]	0	18.51	6.46	15.83	5.48
OYSTER [41]	0	30.48	14.66	26.21	14.60
<b>CPD (Ours)</b>	0	<b>57.79</b>	<b>37.40</b>	<b>50.18</b>	<b>32.13</b>

Table 5. Vehicle detection comparison with fully/weakly supervised detectors on WOD validation set.

**Results on PandaSet.** The class-agnostic results on PandaSet are presented in Table 3. Our method outperforms OYSTER by 6.5% AP and 9.3% Recall under 0.7 IoU threshold. This improvement is largely due to our CPD’s enhanced label quality. Unlike OYSTER, which suffers from the misleading effects of false labels during training, our CPD leverages the size prior from CProto to significantly improve these labels.

## 4.4. Comparison with Fully/Weakly Supervised Detectors

**Results on KITTI dataset.** To further validate our method, we pre-trained our CPD, along with OYSTER [41] and MODEST [40], on WOD and tested them on the KITTI dataset using Statistical Normalization (SN) [28]. The car detection results are in Table 4. We first compared our method with a sparsely supervised method (weakly supervised with 2% labels) [34] that annotates a single instance per frame for training. Our unsupervised CPD outperforms this sparsely supervised method by 23.52% 3D AP @  $IoU_{0.7}$  on moderate car class. Additionally, our method attains 90.85% and 81.01% 3D AP for the easy and moderate

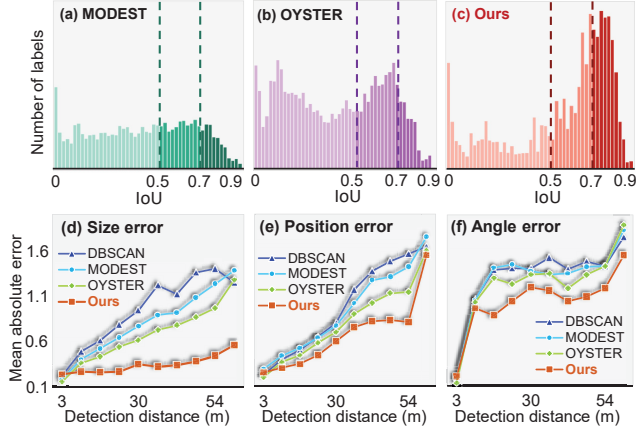


Figure 7. (a-c) IoU distribution between pseudo-labels and ground truth. (d-f) Mean absolute error associated with the size, position, and angle of pseudo-labels generated by different methods.

car classes at a 0.5 IoU threshold. Notably, this performance is comparable to that of the fully supervised method CenterPoint [39], demonstrating the advancement of our method.

Method	3D Recall			3D Precision		
	$IoU_{0.3}$	$IoU_{0.5}$	$IoU_{0.7}$	$IoU_{0.3}$	$IoU_{0.5}$	$IoU_{0.7}$
DBSCAN [5]	22.85	16.44	6.52	29.41	21.16	8.39
MODEST [40]	17.35	12.04	4.89	32.28	22.81	10.05
OYSTER [41]	31.10	21.01	11.12	31.22	21.09	9.45
<b>Ours</b>	<b>45.66</b>	<b>39.33</b>	<b>20.54</b>	<b>34.17</b>	<b>28.22</b>	<b>14.74</b>

Table 6. Pseudo-label comparison results on WOD validation set.

**Results on WOD.** We also compared our method with fully/weakly supervised methods on the WOD validation set [25]. The vehicle detection results are in Table 5. Our unsupervised CPD outperforms the sparsely supervised method (2% annotation) by 5.25% and 4.16% in terms of 3D AP L1 and L2 respectively.

#### 4.5. Pseudo-label Comparison

To validate our pseudo-labels, we analyzed their 3D recall and precision on the WOD validation set. As shown in Table 6, our method surpasses the previous best-performing OYSTER with a 9.42% recall and 5.29% precision improvement (under a 0.7 IoU threshold). To understand the sources of this improvement, we examined the IoU between the pseudo-labels and ground truth, and compared the IoU distributions in Fig. 7 (a)(b)(c). We also present the mean absolute error of size, position, and angle between different pseudo-labels in Fig. 7 (d)(e)(f). The IoU distribution of our method is much closer to 1 than other methods, and it also exhibits lower errors in size, position, and angle. These results verify that our MFC and CBR significantly reduce label errors.




Figure 8. (a)(b)The recall and precision of initial pseudo-labels by using different frames. (c) The recall-precision curve of initial pseudo-labels by using different scores.

Components				3D AP L1		3D AP L2	
SFC	MFC	CBR	CST	$IoU_{0.5}$	$IoU_{0.7}$	$IoU_{0.5}$	$IoU_{0.7}$
✓				17.36	2.65	14.87	2.29
	✓			19.91	5.01	18.31	4.77
		✓		48.26	28.01	41.69	24.04
			✓	<b>57.79</b>	<b>37.40</b>	<b>50.18</b>	<b>32.13</b>

Table 7. CPD component analysis results on WOD validation set.

CSS Components	BEV AP		
	$IoU_{0.3}$	$IoU_{0.5}$	$IoU_{0.7}$
Distance	32.70	24.57	14.80
Distance+MLO	34.40	26.25	15.91
Distance+MLO+SS	<b>38.95</b>	<b>31.06</b>	<b>19.49</b>

Table 8. CSS component analysis results on WOD validation set.

#### 4.6. Ablation Study

**Components analysis of CPD.** To evaluate the individual contributions of our designs, we incrementally added each component and assessed their impact on vehicle detection using the WOD validation set. The results are shown in Table 7. Our MFC method surpasses Single Frame Clustering (SFC) by 2.52% in AP, attributed to the more complete point representation of objects across consecutive frames compared to a single frame. The CBR further enhances performance by 19.27% in AP, as it reduces size and location errors in pseudo-labels. The CST contributes an 8.09% increase in AP, demonstrating the effectiveness of geometric features from CProto in detecting sparse objects.

**Frame number of MFC.** To examine the effect of frame count on initial pseudo-label quality, we experimented with different numbers of past and future point cloud frames on the WOD validation set. The BEV results, shown in Fig. 8 (a)(b), indicate optimal performance with [-5, 5] frames (five past, five future, and the current frame). Additional frames did not significantly improve recall or precision. Consequently, we used 11 frames for initial pseudo-label generation in this study.

**Component analysis of CSS Scoring.** To assess the effectiveness of our scoring system, we calculated the BEV AP of initial pseudo-labels with different scores. These






Figure 9. Visualization comparison of different detection results on WOD validation set.

CBR Components	BEV Recall		BEV Precision	
	$IoU_{0.5}$	$IoU_{0.7}$	$IoU_{0.5}$	$IoU_{0.7}$
MFC	26.88	18.14	26.98	18.16
MFC+Re-size	30.79	21.54	29.43	21.33
MFC+Re-size+Re-localization	<b>43.47</b>	<b>27.97</b>	<b>30.90</b>	<b>21.62</b>

Table 9. CBR component analysis results on WOD validation set.

evaluations, reported in Table 8, show that incorporating all components (distance, MLO, and SS) yields the highest AP. The recall-precision curve, plotted in Fig. 8 (c), also supports this finding. These indicate the significance of each component in accurately measuring pseudo-label quality.

**Components analysis of CBR.** To evaluate the impact of re-sizing and re-localization in CBR, we conducted experiments and analyzed pseudo-label performance. As shown in Table 9, re-sizing results in a 3.91% and 3.4% increase in BEV recall at the 0.5 and 0.7 IoU thresholds, respectively; re-localization further enhances recall by 12.68% and 6.43% at these thresholds, while also increasing precision. These results indicate the importance of both components, which effectively refine pseudo-labels.

**Components analysis of CST.** To assess the effectiveness of each component in CST, we established a baseline using only CBR-generated pseudo-labels for training a two-stage CenterPoint detector, then incrementally added our loss components and evaluated vehicle detection performance on the WOD validation set. As shown in Table 10, all loss components contribute to performance improvement. Specifically, our  $\mathcal{L}_{det}^{css}$  mitigates the influence of false pseudo-label using CSS weight, and improves the 3D AP L2 at  $IoU_{0.7}$  by 4.79%. Our  $\mathcal{L}_{feat}^{css}$  and  $\mathcal{L}_{box}^{css}$  improve the 3D AP L2 at  $IoU_{0.7}$  by 0.75% and 2.55% respectively, through leveraging geometric knowledge from dense CProto for more effective sparse object detection.

#### 4.7. Visualization Comparison

To provide a more intuitive understanding of how our method improves detection performance, we visually compare our results with those of MODEST [40] and OYSTER [41], as shown in Fig. 9. MODEST often misses

CST Components	3D AP L1		3D AP L2	
	$IoU_{0.5}$	$IoU_{0.7}$	$IoU_{0.5}$	$IoU_{0.7}$
CBR-only	48.26	28.01	41.69	24.04
CBR+ $\mathcal{L}_{det}^{css}$	49.31	29.78	42.50	28.83
CBR+ $\mathcal{L}_{det}^{css}+\mathcal{L}_{feat}^{css}$	52.01	32.17	44.12	29.58
CBR+ $\mathcal{L}_{det}^{css}+\mathcal{L}_{feat}^{css}+\mathcal{L}_{box}^{css}$	<b>57.79</b>	<b>37.40</b>	<b>50.18</b>	<b>32.13</b>

Table 10. CST component analysis results on WOD validation set.

distant, sparse objects (Fig. 9(1.1)), while OYSTER detects them but inaccurately reports their sizes and positions (Fig. 9(2.1)). In contrast, CPD, using our CProto-based design, not only recognizes these objects but also accurately predicts their sizes and positions (Fig. 9(3.1)). Furthermore, since our CST reduces the influence of false pseudo-labels, the false positives (Fig. 9(3.2)) are also much fewer than the previous methods (Fig. 9(1.2)(2.2)).

## 5. Conclusion

This paper presents the CPD framework, a novel approach for accurate unsupervised 3D object detection. First, we develop an MFC method to generate initial pseudo-labels. Then, a CProto set is constructed using CSS scoring. Next, we introduce a CBR method to refine these pseudo-labels. Lastly, a CST is designed to enhance detection accuracy for sparse objects. Extensive experiments have verified the effectiveness of our design. Notably, for the first time, our unsupervised CPD method surpasses some weakly supervised methods, demonstrating the advancement of our approach.

**Limitations.** One notable limitation of our work is the significantly lower Average Precision (AP) for minority classes, such as cyclists (Table 1), compared to more prevalent classes like vehicles. This disparity is largely due to the scarce instances of these minority classes within the dataset. Future efforts to collect such objects could be a promising avenue to tackle this issue.

**Acknowledgements.** This work was supported in part by the National Natural Science Foundation of China (No.62171393), the Fundamental Research Funds for the Central Universities (No.20720220064).



## References

- [1] Asma Azim and Olivier Aycard. Detection, classification and tracking of moving objects in a 3d environment. In *2012 IEEE Intelligent Vehicles Symposium*, pages 802–807. IEEE, 2012. [2](#)
- [2] Qi Cai, Yingwei Pan, Ting Yao, and Tao Mei. 3d cascade rcnn: High quality object detection in point clouds. *IEEE Transactions on Image Processing*, 31:5706–5719, 2022. [2](#)
- [3] Zehui Chen, Zhenyu Li, Shuo Wang, Dengpan Fu, and Feng Zhao. Learning from noisy data for semi-supervised 3d object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6929–6939, 2023. [1](#), [2](#)
- [4] Jiajun Deng, Shaoshuai Shi, Peiwei Li, Wen gang Zhou, Yanyong Zhang, and Houqiang Li. Voxel r-cnn: Towards high performance voxel-based 3d object detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2021. [1](#), [2](#), [5](#), [6](#)
- [5] Martin Ester, Hans-Peter Kriegel, Jörg Sander, and Xiaowei Xu. A density-based algorithm for discovering clusters in large spatial databases with noise. In *Knowledge Discovery and Data Mining*, 1996. [3](#), [6](#), [7](#)
- [6] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3354–3361, 2012. [2](#), [5](#)
- [7] Goldberger, Gordon, and Greenspan. An efficient image similarity measure based on approximations of kl-divergence between two gaussian mixtures. In *Proceedings Ninth IEEE International conference on computer vision*, pages 487–493. IEEE, 2003. [4](#)
- [8] Chenhang He, Hui Zeng, Jianqiang Huang, Xian-Sheng Hua, and Lei Zhang. Structure aware single-stage 3d object detection from point cloud. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11873–11882, 2020. [2](#)
- [9] Michael Himmelsbach, Felix V Hundelshausen, and H-J Wuensche. Fast segmentation of 3d point clouds for ground vehicles. In *2010 IEEE Intelligent Vehicles Symposium*, pages 560–565. IEEE, 2010. [1](#), [3](#)
- [10] 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 conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12697–12705, 2019. [2](#)
- [11] Huifang Li, Yidong Li, Yuanzhouhan Cao, Yushan Han, Yi Jin, and Yunchao Wei. Weakly supervised object detection with class prototypical network. *IEEE Transactions on Multimedia*, 2022. [2](#)
- [12] Tianqin Li, Zijie Li, Harold Rockwell, Amir Farimani, and Tai Sing Lee. Prototype memory and attention mechanisms for few shot image generation. In *Proceedings of the Eleventh International Conference on Learning Representations*, 2022. [2](#)
- [13] Ziyu Li, Jingming Guo, Tongtong Cao, Liu Bingbing, and Wankou Yang. Gpa-3d: Geometry-aware prototype alignment for unsupervised domain adaptive 3d object detection from point clouds. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6394–6403, 2023. [3](#), [4](#)
- [14] Hanxue Liang, Chenhan Jiang, Dapeng Feng, Xin Chen, Hang Xu, Xiaodan Liang, Wei Zhang, Zhenguo Li, Luc Van Gool, and Sun Yat-sen. Exploring geometry-aware contrast and clustering harmonization for self-supervised 3d object detection. In *International Conference on Computer Vision (ICCV)*, pages 3273–3282, 2021. [2](#)
- [15] Zhijian Liu, Haotian Tang, Alexander Amini, Xinyu Yang, Huizi Mao, Daniela Rus, and Song Han. Bevfusion: Multi-task multi-sensor fusion with unified bird’s-eye view representation. *ArXiv*, 2022. [1](#)
- [16] Xiaonan Lu, Wenhui Diao, Yongqiang Mao, Junxi Li, Peijin Wang, Xian Sun, and Kun Fu. Breaking immutable: information-coupled prototype elaboration for few-shot object detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1844–1852, 2023. [2](#)
- [17] Qinghao Meng, Wenguan Wang, Tianfei Zhou, Jianbing Shen, Yunde Jia, and Luc Van Gool. Towards a weakly supervised framework for 3d point cloud object detection and annotation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(8):4454–4468, 2022. [2](#)
- [18] Xidong Peng, Xinge Zhu, and Yuexin Ma. Cl3d: Unsupervised domain adaptation for cross-lidar 3d detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 2047–2055, 2023. [3](#)
- [19] Gheorghii Postica, Andrea Romanoni, and Matteo Matteucci. Robust moving objects detection in lidar data exploiting visual cues. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1093–1098. IEEE, 2016. [2](#)
- [20] Shaoshuai Shi, Xiaogang Wang, and Hongsheng Li. Pointcnn: 3d object proposal generation and detection from point cloud. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–779, 2019. [2](#)
- [21] Shaoshuai Shi, Chaoxu Guo, Li Jiang, Zhe Wang, Jianping Shi, Xiaogang Wang, and Hongsheng Li. Pv-rcnn: Point-voxel feature set abstraction for 3d object detection. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10526 – 10535, 2020. [4](#)
- [22] Shaoshuai Shi, Zhe Wang, Jianping Shi, Xiaogang Wang, and Hongsheng Li. From points to parts: 3d object detection from point cloud with part-aware and part-aggregation network. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 43:2647–2664, 2021. [2](#)
- [23] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. *Advances in neural information processing systems*, 30, 2017. [6](#)
- [24] Muhammad Sualeh and Gon-Woo Kim. Dynamic multi-lidar based multiple object detection and tracking. *Sensors*, 19(6): 1474, 2019. [2](#)
- [25] Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, and al. et. Scalability in perception for autonomous driving: Waymo open dataset. In *Proceedings of the IEEE conference on*

- Computer Vision and Pattern Recognition (CVPR)*, pages 2443–2451, 2020. 2, 4, 5, 7
- [26] He Wang, Yezhen Cong, Or Litany, Yue Gao, and Leonidas J. Guibas. 3dioumatch: Leveraging iou prediction for semi-supervised 3d object detection. *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition CVPR*, pages 14610–14619, 2021. 1, 2
- [27] Haiyang Wang, Chen Shi, Shaoshuai Shi, Meng Lei, Sen Wang, Di He, Bernt Schiele, and Liwei Wang. DSVT: Dynamic Sparse Voxel Transformer With Rotated Sets. In *CVPR*, 2023. 2
- [28] Yan Wang, Xiangyu Chen, Yurong You, Li Erran Li, Bharath Hariharan, Mark Campbell, Kilian Q Weinberger, and Wei-Lun Chao. Train in germany, test in the usa: Making 3d object detectors generalize. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11713–11723, 2020. 6
- [29] Aming Wu, Yahong Han, Linchao Zhu, and Yi Yang. Universal-prototype enhancing for few-shot object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9567–9576, 2021. 2
- [30] Hai Wu, Jinhao Deng, Chenglu Wen, Xin Li, and Cheng Wang. Casa: A cascade attention network for 3d object detection from lidar point clouds. *IEEE Transactions on Geoscience and Remote Sensing*, 2022. 1, 2
- [31] Hai Wu, Chenglu Wen, Wei Li, Ruigang Yang, and Cheng Wang. Learning transformation-equivariant features for 3d object detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2023. 2, 5
- [32] Hai Wu, Chenglu Wen, Shaoshuai Shi, Xin Li, and Cheng Wang. Virtual sparse convolution for multimodal 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023. 5
- [33] Xiaopei Wu, Liang Peng, Honghui Yang, Liang Xie, Chenxi Huang, Chengqi Deng, Haifeng Liu, and Deng Cai. Sparse fuse dense: Towards high quality 3d detection with depth completion. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 1, 2
- [34] Qiming Xia, Jinhao Deng, Chenglu Wen, Hai Wu, Shaoshuai Shi, Xin Li, and Cheng Wang. Coin: Contrastive instance feature mining for outdoor 3d object detection with very limited annotations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023. 2, 5, 6
- [35] Pengchuan Xiao, Zhenlei Shao, Steven Hao, Zishuo Zhang, Xiaolin Chai, Judy Jiao, Zesong Li, Jian Wu, Kai Sun, Kun Jiang, Yunlong Wang, and Diange Yang. Pandaset: Advanced sensor suite dataset for autonomous driving. *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*, 2021. 2, 5
- [36] Honghui Yang, Tong He, Jiaheng Liu, Huaguan Chen, Boxi Wu, Binbin Lin, Xiaofei He, and Wanli Ouyang. Gd-mae: Generative decoder for mae pre-training on lidar point clouds. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9403–9414, 2023. 2
- [37] Zetong Yang, Yanan Sun, Shu Liu, Xiaoyong Shen, and Jiaya Jia. Std: Sparse-to-dense 3d object detector for point cloud. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 1951–1960, 2019. 2
- [38] Junbo Yin, Dingfu Zhou, Liangjun Zhang, Jin Fang, Cheng-Zhong Xu, Jianbing Shen, and Wenguan Wang. Proposal-contrast: Unsupervised pre-training for lidar-based 3d object detection. In *European conference on computer vision*, pages 17–33, 2022. 2
- [39] Tianwei Yin, Xingyi Zhou, and Philipp Krähenbühl. Center-based 3d object detection and tracking. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2, 5, 6, 7
- [40] Yurong You, Katie Luo, Cheng Perng Phoo, Wei-Lun Chao, Wen Sun, Bharath Hariharan, Mark E. Campbell, and Kilian Q. Weinberger. Learning to detect mobile objects from lidar scans without labels. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 1, 2, 3, 5, 6, 7, 8
- [41] Lunjun Zhang, Anqi Joyce Yang, Yuwen Xiong, Sergio Casas, Bin Yang, Mengye Ren, and Raquel Urtasun. Towards unsupervised object detection from lidar point clouds. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 1, 2, 3, 4, 5, 6, 7, 8
- [42] Quanshi Zhang, Xuan Song, Xiaowei Shao, Huijing Zhao, and Ryosuke Shibasaki. Unsupervised 3d category discovery and point labeling from a large urban environment. In *2013 IEEE International Conference on Robotics and Automation*, pages 2685–2692. IEEE, 2013. 1
- [43] Xiao Zhang, Wenda Xu, Chiyu Dong, and John M Dolan. Efficient l-shape fitting for vehicle detection using laser scanners. In *2017 IEEE Intelligent Vehicles Symposium (IV)*, pages 54–59. IEEE, 2017. 3
- [44] Yixin Zhang, Zilei Wang, and Yushi Mao. Rpn prototype alignment for domain adaptive object detector. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12425–12434, 2021. 2
- [45] Zehan Zhang, Yang Ji, Wei Cui, Yulong Wang, Hao Li, Xian Zhao, Duo Li, Sanli Tang, Ming Yang, Wenming Tan, et al. Atf-3d: Semi-supervised 3d object detection with adaptive thresholds filtering based on confidence and distance. *IEEE Robotics and Automation Letters*, 7(4):10573–10580, 2022. 2
- [46] Na Zhao, Tat-Seng Chua, and Gim Hee Lee. Sess: Self-ensembling semi-supervised 3d object detection. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 1, 2
- [47] Shizhen Zhao and Xiaojuan Qi. Prototypical votenet for few-shot 3d point cloud object detection. *Advances in Neural Information Processing Systems*, 35:13838–13851, 2022. 3, 4
- [48] Wu Zheng, Weiliang Tang, Li Jiang, and Chi-Wing Fu. Sessd: Self-ensembling single-stage object detector from point cloud. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14494–14503, 2021. 2
- [49] Yin Zhou and Oncel Tuzel. Voxelnet: End-to-end learning for point cloud based 3d object detection. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4490–4499, 2018. 2