Semi-supervised 3D Object Detection with Proficient Teachers

Junbo Yin^{1*}, Jin Fang^{2,3,4*}, Dingfu Zhou^{2,3}, Liangjun Zhang^{2,3}, Cheng-Zhong Xu⁴, ⊠Jianbing Shen⁴, and ⊠Wenguan Wang⁵

¹School of Computer Science, Beijing Institute of Technology ²Baidu Research ³National Engineering Laboratory of Deep Learning Technology and Application, China ⁴SKL-IOTSC, CIS, University of Macau ⁵ReLER, AAII, University of Technology Sydney {yinjunbocn,wenguanwang.ai}@gmail.com https://github.com/yinjunbo/ProficientTeachers

Abstract. Dominated point cloud-based 3D object detectors in autonomous driving scenarios rely heavily on the huge amount of accurately labeled samples, however, 3D annotation in the point cloud is extremely tedious, expensive and time-consuming. To reduce the dependence on large supervision, semi-supervised learning (SSL) based approaches have been proposed. The Pseudo-Labeling methodology is commonly used for SSL frameworks, however, the low-quality predictions from the teacher model have seriously limited its performance. In this work, we propose a new Pseudo-Labeling framework for semi-supervised 3D object detection, by enhancing the teacher model to a proficient one with several necessary designs. First, to improve the recall of pseudo labels, a Spatialtemporal Ensemble (STE) module is proposed to generate sufficient seed boxes. Second, to improve the precision of recalled boxes, a Clusteringbased Box Voting (CBV) module is designed to get aggregated votes from the clustered seed boxes. This also eliminates the necessity of sophisticated thresholds to select pseudo labels. Furthermore, to reduce the negative influence of wrongly pseudo-labeled samples during the training, a soft supervision signal is proposed by considering Box-wise Contrastive Learning (BCL). The effectiveness of our model is verified on both ONCE and Waymo datasets. For example, on ONCE, our approach significantly improves the baseline by 9.51 mAP. Moreover, with half annotations, our model outperforms the oracle model with full annotations on Waymo.

Keywords: 3D Object Detection, Semi-supervised Learning, Point Cloud

1 Introduction

With the rapid development of range sensors (e.g., LiDAR) and their wide application in the field of robotics and autonomous driving, point cloud-based scene

 $[\]boxtimes$: Corresponding author.

^{*:} Equal contribution. Work done when J. Yin was an intern at Baidu Research.



Fig. 1. The main idea of our *ProficientTeachers*. Given an unlabeled point cloud, the spatial-temporal ensemble (STE) module first produces (a) seed boxes by combing predictions from multiple augmented views. Then, the clustering-based box voting (CBV) module adaptively aggregates these boxes to get (b) the final high-quality pseudo labels. In this way, we not only achieve better detection results, but also remove the necessity of sophisticated thresholds for selecting pseudo labels, as seen in (c).

understanding such as 3D object detection has received great attention recently. With the great capabilities of deep neural networks (DNNs) and a huge number of annotated samples, impressive performances have been achieved on different public benchmarks [2, 8, 9, 26]. VoxelNet [41], PointRCNN [22], PointPillars [11], PV-RCNN [21] and CenterPoint [38] are several representative 3D object detection frameworks. Nevertheless, the results highly rely on huge annotations, while the annotation in 3D data is extremely expensive and time-consuming, *e.g.*, a skilled worker may take hundreds of hours to annotate just one hour of driving data [7, 12–14].

Semi-supervised learning (SSL) techniques, which train a model with a small number of labeled samples together with an abundance of unlabeled data, are a promising alternative to the fully-supervised learning frameworks. Compared to labeled data, unlabeled data is obtained very conveniently and cheaply. However, due to the inherent difficulty of point cloud (e.g., orderless, textureless, and sparsity), only a few SSL-based 3D object detection frameworks have been proposed. Up to now, SESS [39] and 3DIoUMatch [29] are two pioneers of this domain. To handle the unlabeled data, SESS leverages asymmetric data augmentation and enforces consistency regularization between the predictions of teacher and student models. Although noticeable improvements have been achieved upon a vanilla VoteNet [16] on indoor datasets, other researchers [18,29] found that the consistency regularization is suboptimal if it is uniformly enforced on all the student and teacher predictions because the quality of these predictions may be quite different. To well handle this limitation, 3DIoUMatch seeks a pseudolabeling approach and applies a confidence-based filtering strategy for pseudolabel selection, where the confidence is defined as a combination of classifica-

3

tion score and IoU estimation. Though achieving better performance, it takes tremendous effort to select a suitable confidence score. Moreover, the pseudo labels produced by its plain teacher model limit the final detection performance.

To address these challenges, we propose a new pseudo-labeling SSL framework, ProficientTeachers, that not only provides high-quality pseudo labels via an enhanced teacher model, but also reduces the necessity of deliberately selected thresholds. To be specific, false negative (FN) and false positive (FP) in the pseudo labels are two main challenges. The LiDAR point cloud is sparse and noisy, and some street objects are of small sizes (e.g., less than 2m for pedestrian) and unevenly distributed across a considerably wide range (e.g., $150m \times 150m$ in Waymo [26]). Thus it is prone to cause FN detections from only a single point cloud view. To handle this, a spatial-temporal ensemble (STE) module is proposed to generate sufficient seed boxes from spatially and temporally augmented views. Aggregating predictions from different views reduces the prediction bias of the teacher model, thus essentially boosting the *recall*. The redundant seed boxes from STE inevitably involve FP detections. To further resolve the FP problem, we propose a clustering-based box voting (CBV) module. Our CBV module groups the seed boxes into different clusters and generates votes (*i.e.*, a refined bounding box) for each box in a cluster. These votes are then aggregated to produce a more accurate box for each cluster. In this way, more clean and precise pseudo labels can be obtained by a simple NMS, removing the need of selecting thresholds. Our CBV module significantly improves the *precision* of pseudo labels produced by the STE module, without losing the *recall*. By equipping the STE and CBV modules, the vanilla teacher model has become to be proficient teachers model. Furthermore, we find that the original pseudolabeling method enforces a hard training target, where the inaccurate pseudo labels will undermine the performance of the student. To alleviate this problem, a soft training target is proposed by box-wise contrastive learning (BCL), which aims to learn the cross-view feature consistency based on the informative boxes.

To summarize, we propose a new 3D SSL framework, *ProficientTeachers*, for LiDAR-based 3D object detection, which is achieved by promoting the plain teacher model to proficient teachers inspired by ensemble learning. Our framework not only performs better results, but also removes the necessity of confidence based thresholds for filtering pseudo labels. In our model, spatial-temporal ensemble (STE) and clustering-based box voting (CBV) modules are developed to improve the recall and precision of pseudo labels. Furthermore, a box-wise contrastive learning (BCL) strategy has been advocated to explore representation learning based on expressive 3D boxes. Comprehensive evaluations have been conducted on ONCE and Waymo. Our *ProficientTeachers* can improve the baseline detector by 9.51 mAP on ONCE, and save half annotations on Waymo.

2 Related Work

3D Object Detection. 3D object detection from LiDAR point cloud has been studied for decades with various approaches being proposed. The mainstream

methodologies can be divided into two categories: voxel-based [4, 35, 36, 38, 41] and point-based [22, 23, 34, 40]. Voxel-based expression is a popular way of processing the point cloud in deep learning, and has been widely applied for 3D object detection with the development of sparse convolution [33]. Voxelnet [41] splits the LiDAR data into voxels and sends the points in each voxel into a voxel feature encoding layer to get voxel-wise features. To accelerate the speed, PointPillars [11] uses a pillar representation to replace the voxel representation. Different from the above approaches, the point-based approaches take the point cloud directly into the DNNs. PointRCNN [22] is representative of this, which employs the Pointnet++ [17] as the backbone for semantic segmentation first, and then regions of interest (RoIs) are generated based on foreground points. Besides this, PV-RCNN [21] proposes to combine both point cloud and multi-scale voxel representation and achieve high performance.

Semi-supervised Learning (SSL). SSL has been studied for a long time and many approaches have been proposed [28, 42]. The recently popular SSL approaches such as Temporal Ensembling [19], Mean Teacher [27] and Noisy Student [31] have achieved impressive performance on the 2D tasks. The temporal model, which explores consistency in the prediction level, tries to minimize the difference between the predictions from the current step and the EMA (an exponential moving average) predictions over multiple previous training epochs. The EMA predictions can largely improve the quality of the predictions. Mean Teacher [27] further improves the temporal model by replacing network prediction average with network parameter average. It contains two network branches, *i.e.*, teacher and student, with the same architecture. The parameters of the teacher are the EMA of the student, while the parameters of the student are updated by stochastic gradient descent. The student network is trained to yield consistent predictions with the teacher network. Noisy Student [31] deliberately injects noise to the student model from both input level and model level to strengthen the student. Then it makes the student a new teacher and performs iterative self-training to further improve performance.

SSL on 3D Object Detection. Most of the SSL approaches are proposed for classification tasks [1,24] and a few SSL approaches have been proposed to leverage the object detection task [10,25], especially 3D object detection. SESS [39] and 3DIoUMatch [29] are two typical recently proposed approaches for 3D object detection from point cloud data. The SESS, which is built based on the Mean Teacher paradigm by updating the parameters of the teacher network with an EMA technique, employs asymmetric data augmentation and enforces three kinds of consistency losses between the teacher and student predictions. Different from SESS, 3DIoUMatch proposes a series of handcrafted strategies to achieve better pseudo labels such as joint class, objectness, localization confidences-based pseudo-label filtering, and IoU-guided lower-half suppression for deduplication, etc. The proposed framework works well on PV-RCNN [21]. By contrast, our method benefits from an enhanced teacher model, without sophisticated thresholds selection, and can be flexibly applied on popular LiDAR-based 3D detectors, *e.g.*, SECOND [33], CenterPoint [38] and PV-RCNN [21].



Fig. 2. The framework of our *ProficientTeachers* model. It consists of a spatialtemporal ensemble (STE) module and a clustering-based box voting (CBV) module. STE produces sufficient seed boxes and CBV adaptively fuses them to obtain better pseudo labels. Besides, a consistency loss based on box-wise contrastive learning (BCL) is applied on the student model to explicitly learn from informative box features.

3 Proposed Approach

3.1 SSL Framework

Semi-supervised 3D object detection is very valuable and practical in real selfdriving scenarios. Here, we explore the pseudo-labeling framework, which is a popular methodology in semi-supervised learning. Typical pseudo-labeling methods [31] exploit a teacher model to produce pseudo labels which are used to supervise a student model. We argue that the low-quality predictions produced by the vanilla teacher model limit the performance. In this work, we propose to promote the vanilla teacher model to a proficient one, as seen in Fig. 2.

Our *ProficientTeachers* model, as described in Sec. 3.2, contains a spatialtemporal ensemble (STE) module and a clustering-based box voting (CBV) module, which are designed to handle false negative (FN) and false positives (FP) in pseudo labels, respectively. In particular, STE produces multi-group boxes based on augmented and assembled views to recall the missed objects. To further remove the redundant FP boxes, the original pseudo-labeling approaches apply a fixed threshold to filter out boxes with lower confidence scores, which is unstable and inefficient. In contrast, our CBV adaptively aggregates the seed boxes to reduce FP and also improves the precision by voting the boxes within a cluster. Furthermore, we propose a soft supervision signal in Sec. 3.3, *i.e.*, *Contrastive Student* model, by involving box-wise contrastive learning (BCL). Next, we elaborate on each module in subsequent sections. J. Yin et al.

3.2 Proficient Teachers Model

Assuming that we have total N training samples, including N_l labeled samples $\mathcal{P}^L = \{\mathbf{p}_i^L, \mathbf{y}_i^L\}_{i=1}^{N_l}$ and N_u unlabeled samples $\mathcal{P}^U = \{\mathbf{p}_i^U\}_{i=1}^{N_u}$, where $\mathbf{p}_i \in \mathcal{R}^{n \times \{3+r\}}$ represents a point cloud sample \mathbf{p}_i that has n points with 3-dimensional coordinates and other r-dimensional attributes such as intensity, timestamps, etc. The core idea of our *ProficientTeachers* model is to produce high-quality pseudo labels $\{\mathbf{y}_i^U\}_{i=1}^{N_u}$ for the unlabeled point clouds $\{\mathbf{p}_i^U\}_{i=1}^{N_u}$.

Spatial-Temporal Ensemble Module. As described above, predictions produced by vanilla teacher inevitably encounter the FN problem, which will cause the student model to treat a foreground object as a negative example, thus degenerating the detection performance. Motivated from this observation, we attempt to improve the recall of the teacher model with the STE module, which generates sufficient candidate pseudo boxes with a spatial ensemble module and a temporal ensemble module.

The spatial ensemble in the STE aims to produce multi-group detections based on differently augmented point cloud views, which is inspired by test time augmentation [20]. Detections on a certain view may exist bias and miss objects. By combining detections from multiple views after reverse transformation, less FN will happen, thus improving the recall. Formally, given an unlabeled data $\mathbf{p}_i^U \in \mathcal{P}^U$, we first apply *fixed* data augmentation to spatially transform the input point cloud into different views. This can be deemed as a form of weak augmentation compared with the *random* train-time augmentation policies, which is:

$$\{\mathbf{p}_i^{U_1}, \mathbf{p}_i^{U_2}, \dots, \mathbf{p}_i^{U_K}\} = \mathcal{T}(\mathbf{p}_i^U), \tag{1}$$

where $\mathcal{T}(\cdot)$ is the augmentation function and K is the number of augmented views. More details about our data augmentation can be found in Sec. 4.2.

Then, we use the teacher network $f_{\text{TEA}}(\cdot)$ to give predictions for each view:

$$\{\mathbf{y}_{i}^{U_{1}}, \mathbf{y}_{i}^{U_{2}}, \dots, \mathbf{y}_{i}^{U_{K}}\} = f_{\text{TEA}}(\{\mathbf{p}_{i}^{U_{1}}, \mathbf{p}_{i}^{U_{2}}, \dots, \mathbf{p}_{i}^{U_{K}}\}),$$
(2)

With the spatial ensemble module, objects not detected in one view may be detected in another view. All detection results are then projected to the original coordinate with reverse transformations.

Moreover, we can also incorporate a temporal ensemble module. In particular, previous pseudo label methods update weight of teacher model by student model from a certain epoch. We argue that models from different epochs exist bias, and temporal aggregation of models from different epochs can alleviate this problem. This is presented as:

$$\{\mathbf{y}_{i}^{U_{1}^{T}}, \mathbf{y}_{i}^{U_{2}^{T}}, \dots, \mathbf{y}_{i}^{U_{K}^{T}}\} = f_{\text{TEA}}^{T}(\{\mathbf{p}_{i}^{U_{1}}, \mathbf{p}_{i}^{U_{2}}, \dots, \mathbf{p}_{i}^{U_{K}}\}),$$
(3)

where $f_{\text{TEA}}^T(\cdot)$ denotes the teacher model updated by student model from epoch T. We consider the models from recent epochs for temporal ensemble. Afterward, we collect all the predictions produced by both the spatial and temporal ensemble modules. As seen in Table 1, our STE improves the recall by 5.6%, compared with the original teacher model. Here, simple post-processing like Non-maximum Suppression (NMS) can be directly applied to aggregate these predictions.



Fig. 3. Illustration of clustering-based box voting (CBV) module, which contains four steps: seed box clustering, RoI feature extraction, box voting and votes aggregation. By adaptively fusing the high-recall candidate boxes output by the STE, CBV significantly improves the precision of pseudo labels.

Clustering-based Box Voting Module. Though the proposed STE module has ensured a higher recall by combing sufficient pseudo boxes, some inaccurate candidate boxes will also lead to a lower precision (*i.e.*, from 47.2% to 27.4%, as seen in Table. 1), which undermines the quality of pseudo labels. Thus, the main purpose of the CBV module is to address these redundant candidate boxes.

An intuition way is to define a score to formulate the quality of the pseudo boxes, and filter the low-quality boxes under a certain threshold. For instance, FixMatch [24] directly uses the classification score to filter the pseudo boxes. Later, 3DIoUMatch [29] leverages a predicted IoU score for filtering objects. However, there are potential limitations in these threshold-based filtering methods. First of all, it requires manual experience and elaborate ex-

Table 1. Recall and precision of "Vehicle" class based on SECOND detector trained with different SSL settings. "c" represents the threshold of the confidence score.

Methods	Recall Precision
SECOND [33] (Baseline)	78.3% 47.2%
SECOND $[33] + STE$	83.9% 27.4%
SECOND $[33] + STE (c = 0.3)$	80.4% 74.6%
SECOND $[33] + STE + CBV$	83.6% $76.5%$

periments to choose a suitable threshold for a class, and different datasets or detectors might not share the same thresholds. Secondly, the predicted box confidence score either formulates class probability, e.g., classification score, or the localization accuracy, e.g., IoU, individually, which requires a sophisticated combination from different aspects. Thirdly, they separately consider each individual box, ignoring the relations of boxes in a cluster covering the same ground truth. To this end, we propose a Clustering-based Box Voting (CBV) module to obtain more clean and accurate pseudo labels via a learnable voting process, as well as eliminate the need of sophisticated threshold selection.

The detailed pipeline of the CBV module can be found in Fig. 3. Specifically, given the pseudo labels produced by the STE module, we first cluster these boxes based on the IoU criterion. For example, the boxes are first arranged by the scores. Then, for the box with the highest score, we select other boxes that have a larger IoU with it as the same cluster. This process is iteratively performed for all the boxes to produce all the clusters. Assuming the k-th cluster has M

8 J. Yin et al.

boxes, *i.e.*, $\{\mathbf{b}_m^k\}_{m=1}^{M'}$, we aim to aggregate these boxes and get a refined box \mathbf{b}^k via box voting and votes aggregation. For box voting, we first obtain features of each box via a pre-trained RoI network [5], *i.e.*, $\tilde{\mathbf{b}}_m^k = f_{\text{RoI}}(\mathbf{b}_m^k)$. Then, we let each box in a cluster predict a vote, based on the context-aware features of the box. The voting network is realized by a shared-weight two-layer MLP, and is trained to regress the offset to the ground-truth box. This can be denoted as $\mathbf{v}_m^k = f_{\text{Vote}}(\tilde{\mathbf{b}}_m^k)$. In this way, we have the vote results $\{\mathbf{v}_m^k\}_{m=1}^{M'}$, where each vote \mathbf{v}_m^k presents a refined bounding box. Furthermore, we also predict an objectness \mathbf{s}_m^k to describe the quality of a box. This is achieved by training another two-layer MLP head for foreground and background classification. Accordingly, we can aggregate these votes to get a more accurate box \mathbf{b}^k for the k-th cluster, by weighting the votes \mathbf{v}_m^k with the objectness \mathbf{s}_m^k , which is denoted as:

$$\mathbf{b}^{k} = \frac{\sum_{m=1}^{M'} \mathbf{s}_{m}^{k} \cdot \mathbf{v}_{m}^{k}}{\sum_{m=1}^{M'} \mathbf{s}_{m}^{k}},\tag{4}$$

In this way, more clean pseudo labels are produced. The objectness \mathbf{s}^k is also obtained by considering the average scores and the detection number in this cluster. Then, NMS can be directly applied to get the final pseudo labels, *i.e.*, $\mathbf{y}^U = {\mathbf{b}^k}_{k=1}^M$, where M is the total pseudo box number.

In Table. 1, we find that the CBV module significantly improves the precision from 27.4% to 76.5%, and even improves the recall by 5.3%, due to the votes aggregation in clusters. We also implement a heuristic method by selecting better confidence threshold, *i.e.*, c = 0.3, to filter boxes as in [24]. As seen in Table. 1, such a method is inferior to ours according to both recall and precision metrics.

3.3 Contrastive Student Model

In this section, we introduce the learning of the student model. Concretely, our student model takes as input both the labeled data $\{\mathbf{p}^L, \mathbf{y}^L\} \in \mathcal{P}^L$ and the unlabeled data $\mathbf{p}^U \in \mathcal{P}^U$. For training the labeled data \mathbf{p}^L , we directly obtain the predictions of the student model, and optimize with the detection loss:

$$\mathcal{L}_{det}^{L} = \mathcal{L}_{cls}^{L} + \mathcal{L}_{smooth-\ell_{1}}^{L}, \qquad (5)$$

For training the unlabeled data \mathbf{p}^U , we seek a soft learning target besides the pseudo label-based detection loss. To be specific, though our *ProficientTeach*ers model has significantly improved the quality of pseudo labels, it inevitably contains inaccurate predictions. Learning towards these noisy pseudo targets will affect the student performance. To alleviate this problem, we propose to further mine the information in unlabeled data \mathbf{p}^U , which is achieved by boxwise contrastive learning (BCL). Though contrastive learning [3, 30, 32, 37] has been explored in 3D point cloud, our BCL is different from those by directly contrasting expressive box-level features based on the pseudo predictions.

In particular, our BCL module enforces the feature consistency of the same box instance from different augmented views. Formally, given an unlabeled point cloud \mathbf{p}^{U} , we first apply two random augmentations to generate different views \mathbf{p}^{U_1} and \mathbf{p}^{U_2} . Then, we get the pseudo labels of the two views:

$$\mathbf{y}^{U_1} = f_{\text{STU}}(\mathbf{p}^{U_1}), \mathbf{y}^{U_2} = f_{\text{STU}}(\mathbf{p}^{U_1}),$$
 (6)

where $\mathbf{y}^{U_1} \in \mathbb{R}^{M_1 \times 7}$, $\mathbf{y}^{U_2} \in \mathbb{R}^{M_2 \times 7}$ and M_1, M_2 are the number of the predicted bounding boxes in each view with 7-d attributes. Next, we transform \mathbf{y}^{U_1} and \mathbf{y}^{U_2} to the same view, and build positive and negative sample pairs by a greedy matching, *i.e.*, box pairs with the smallest distance are treated as positives and other boxes are viewed as negatives. This results in M matched positive box pairs. Afterwards, we extract box features by a point-wise interpolation method, which is inspired by CenterPoint [38]. More precisely, we use the center points features from the six faces of a 3D bounding box to present a box, by applying bilinear interpolation on the bird-eye-view (BEV) feature maps \mathbf{F}^U . Due to operating on BEV, the center points of the six faces of a 3D box are equal to the center points of the four sides of a 2D box plus one box center. This essentially simplifies the box feature extraction process. We formulate this as:

$$\mathbf{h}^{U_1} = I(\mathbf{y}^{U_1}, \mathbf{F}^{U_1}), \mathbf{h}^{U_2} = I(\mathbf{y}^{U_2}, \mathbf{F}^{U_2}),$$
(7)

where $I(\cdot)$ is the bilinear interpolation function, and \mathbf{h}^{U_1} and \mathbf{h}^{U_1} are resultant box features from different views. These features are then collected and arranged as $\{\mathbf{h}_k\}_{k=1}^{2M}$, where the even indexes denote the boxes from \mathbf{h}^{U_1} and the odd ones present boxes from \mathbf{h}^{U_2} . Later, a projection head $\phi(\cdot)$ that contains two 1×1 convolutional layers is used to map the box features to an embedding space, such that $\mathbf{z}_k = \phi(\mathbf{h}_k)$. Then, InfoNCE loss [15] is exploited to build the soft target:

$$\ell(p,q) = -\log \frac{\exp(\mathbf{z}_p \cdot \mathbf{z}_q/\tau)}{\sum_{k=1}^{2M} \mathbb{1}_{[k \neq q]} \exp(\mathbf{z}_p \cdot \mathbf{z}_k/\tau)},\tag{8}$$

$$\mathcal{L}_{\rm con}^U = \frac{1}{2M} \sum_{k=1}^M [\ell(2k-1,2k) + \ell(2k,2k-1)],\tag{9}$$

where $p, q \in [1, ..., 2M]$ and $p \neq q$. τ is a temperature hyper-parameter that is set to 0.1. Furthermore, we leverage the high-quality pseudo labels produced by the *ProficientTeachers* to construct the detection loss, which is denoted as:

$$\mathcal{L}_{det}^{U} = \frac{1}{2} (\mathcal{L}_{cls}^{U_1} + \mathcal{L}_{smooth-\ell_1}^{U_1} + \mathcal{L}_{cls}^{U_2} + \mathcal{L}_{smooth-\ell_1}^{U_2}).$$
(10)

Finally, the overall loss of the student model is defined as:

$$\mathcal{L} = \mathcal{L}_{det}^{L} + \mathcal{L}_{det}^{U} + \alpha \mathcal{L}_{con}^{U}, \qquad (11)$$

where α is a coefficient and we set it to 0.05 empirically.

4 Experimental Results

In this work, we first introduce the datasets and the implementation details of our model in Sec. 4.1 and Sec. 4.2, respectively. Then, we report the main evaluation results by comparing with other SSL approaches in Sec. 4.3. Finally, the ablation studies are presented in Sec. 4.4.

10 J. Yin et al.

4.1 Datasets

ONCE Dataset. ONCE [12] is a large-scale autonomous driving dataset with 1 million LiDAR point cloud samples. Only 15,000 samples are with annotations which have been divided into training, validation, and testing split with 5K, 3K, and 8K samples, respectively. In this dataset, five kinds of foreground objects have been annotated i.e., "Car", "Bus", "Truck", "Pedestrian" and "Cyclist", while "Car", "Bus" and "Truck" are merged into one class "Vehicle" during evaluation. In particular, a specific setting is designed for SSL approaches evaluation, *i.e.*, 5K labeled samples and all the unlabeled samples have been divided into 3 subsets: Small, Medium and Large to explore the effects of different data amounts for SSL-based 3D detection. The small unlabeled set *Small* contains 70 sequences (100k samples), the medium set *Medium* contains 321 sequences (500k samples) and the large set *Large* contains 560 sequences (about 1M samples) in total. Similar to other 3D object detection benchmarks, mean AP (average precision) [6] over all the classes is employed for evaluation, based on the 3D IoU thresholds 0.7, 0.3 and 0.5 for "Vehicle", "Pedestrian" and "Cyclist", respectively. In addition, three different perception ranges, '0-30m", "30-50m", and "50m-inf", are specified to well evaluate the performance of 3D detectors.

Waymo Open Dataset. Waymo [26] provides a large-scale LiDAR point cloud dataset that contains 798 sequences (158,361 frames) for training and 202 sequences (40,077 frames) for validation. Since it does not provide additional unlabeled raw data for semi-supervised training. We thus manually tailor a semi-supervised learning dataset following the setting in ONCE [12]. Specifically, we divide the 798 Waymo training sequences equally into two splits, *i.e.*, labeled split \mathcal{P}^L and unlabeled split \mathcal{P}^U (without using the original labels), with each containing 399 sequences. Then, we randomly sample 5%, 10%, 20% and 50% sequences from \mathcal{P}^L , which lead to the ratio of labeled data and unlabeled data $\mathcal{P}^L : \mathcal{P}^U$ as 1:20, 1:10, 1:5 and 1:2, respectively. mAP and mAPH under LEVEL_2 metric are used to evaluate the 3D object detection performance on the full validation set, where 3D IoU thresholds for "Vehicle", "Pedestrian" and "Cyclist" are 0.7, 0.5 and 0.5, respectively.

4.2 Implementation Details

For the STE module in the teacher model, we have empirically defined the fixed (weak) augmentation types as rotation and double flip. Grid search is conducted to find the more effective rotation parameters, *i.e.*, $\{0^{\circ}, +22.5^{\circ}, -22.5^{\circ}\}$. The final number of augmented views K is computed by the product of flip times and rotation times, *i.e.*, $4 \times 3 = 12$ in our case. The strong augmentation includes *random* rotation and flip and scaling, which is the same as that in [33]. For the CBV module, we offline train the RoI network [5] for 10 epochs based only on the labeled data, with a stop-gradient operation to detach from the backbones. We define positive and negative box samples according to the boxes classification scores and the IoU between ground truths. For example, boxes with classification scores below 0.1 or IoU below 0.3 are treated as negatives, and other boxes are

Table 2. Evaluation results on ONCE validation set with different amounts of unlabeled samples (*e.g.*, "Small", "Medium" and "Large") following the official implementation in ONCE [12]. For better understanding, the best overall result in each class has been highlighted in **bold** and the relative gains of each SSL method compared to the baseline model (*i.e.*, SECOND [33] trained with only labeled samples) have been illustrated in colors where the positive gains in **blue** and negative gains are in **green**.

Mathada	\ \	/ehicle	AP (%)	Pe	destria	an AP	(%)	(Cyclist	AP (%)	$\mathbf{m} \mathbf{A} \mathbf{P} (\mathbf{\emptyset})$
Methous	overall	0-30m	30-50m	50m-inf	overall	0-30m	30-50m	50m-inf	overall	0-30m	30-50m	50m-inf	\mathbf{mar} (70)
Baseline [33]	71.19	84.04	63.02	47.25	26.44	29.33	24.05	18.05	58.04	69.96	52.43	34.61	51.89
Small (100K unlabeled Samples)													
Pseudo Label	72.80	84.46	64.97	51.46	25.50	28.36	22.66	18.51	55.37	65.95	50.34	34.42	51.22 (- 0.67)
Noisy Student [31]	73.69	84.69	67.72	53.41	28.81	33.23	23.42	16.93	54.67	65.58	50.43	32.65	52.39 (+ 0.50)
Mean Teacher [27]	74.46	86.65	68.44	53.59	30.54	34.24	26.31	20.12	61.02	72.51	55.24	39.11	55.34 (+ 3.45)
SESS [39]	73.33	84.52	66.22	52.83	27.31	31.11	23.94	19.01	59.52	71.03	53.93	36.68	53.39 (+ 1.50)
3DIoUMatch [29]	73.81	84.61	68.11	54.48	30.86	35.87	25.55	18.30	56.77	68.02	51.80	35.91	53.81 (+ 1.92)
Our Mehtod	76.07	86.78	70.19	56.17	35.90	39.98	31.67	24.37	61.19	73.97	55.13	36.98	57.72 (+ 5.83)
				Me	dium ((500K u	ınlabele	ed Sampl	es)				
Pseudo Label	73.03	86.06	65.96	51.42	24.56	27.28	20.81	17.00	53.61	65.26	48.44	33.58	50.40 (- 1.49)
Noisy Student [31]	75.53	86.52	69.78	55.05	31.56	35.80	26.24	21.21	58.93	69.61	53.73	36.94	55.34 (+ 3.45)
Mean Teacher [27]	76.01	86.47	70.34	55.92	35.58	40.86	30.44	19.82	63.21	74.89	56.77	40.29	58.27 (+ 6.38)
SESS [39]	72.11	84.06	66.44	53.61	33.44	38.58	28.10	18.67	61.82	73.20	56.60	38.73	55.79 (+ 3.90)
3DIoUMatch [29]	75.69	86.46	70.22	56.06	34.14	38.84	29.19	19.62	58.93	69.08	54.16	38.87	56.25 (+ 4.36)
Our Mehtod	78.07	87.43	72.5	59.51	38.38	42.45	34.62	25.58	63.23	74.70	58.19	40.73	59.89 (+ 8.00)
	Large (1M unlabeled Samples)												
Pseudo Label	72.41	84.06	64.54	50.05	23.62	26.80	20.13	16.66	53.25	64.69	48.52	33.47	49.76 (- 2.13)
Noisy Student [31]	75.99	86.67	70.48	55.60	33.31	37.81	28.19	21.39	59.81	70.01	55.13	38.33	56.37 (+ 4.48)
Mean Teacher [27]	76.38	86.45	70.99	57.48	35.95	41.76	29.05	18.81	65.50	75.72	60.07	43.66	59.28 (+ 7.39)
SESS [39]	75.95	86.83	70.45	55.76	34.43	40.00	27.92	19.20	63.58	74.85	58.88	39.51	57.99 (+ 6.10)
3DIoUMatch [29]	75.81	86.11	71.82	57.84	35.70	40.68	30.34	21.15	59.69	70.69	54.92	39.08	57.07 (+ 5.18)
Our Mehtod	78.12	87.22	72.74	59.58	41.95	48.09	35.13	26.01	64.12	75.85	58.04	41.45	61.40 (+ 9.51)

viewed as positives. The votes objectness is optimized by a sigmoid focal loss and the votes localization is optimized by smooth- ℓ_1 loss. The IoU score for clustering boxes is fixed as 0.5. For the semi-supervised training configuration, we follow the Noisy Student implementation provided by ONCE official benchmark [12], *i.e.*, a model pre-trained on the full training set is used to warm up both the student and teacher model. Then the student is trained for 25 to 75 epochs depending on the amount of unlabeled data with learning rate 0.001, and the teacher is updated every 25 epochs.

4.3 Main Results

ONCE Results. First of all, we aim to compare our *ProficientTeachers* with other state-of-the-art SSL approaches on the ONCE dataset [12]. Here, we borrow all the evaluation results from the official benchmark¹ [12], where five typical SSL approaches have been included, *i.e.*, Pseudo Label, Noisy Student [31], Mean Teacher [27], SESS [39], 3DIoUMatch [29]. For a fair comparison, the SECOND [33] detector trained with only the labeled samples has been adopted as the baseline. All the comparison results are given in Tab. 2. Compared to the baseline model that is trained with only the labeled samples, all these SSL frameworks can obtain positive gains, with the help of large amounts of unlabeled samples except for the Pseudo Label method. This may be because the

 $^{^{1}\} https://once-for-auto-driving.github.io/benchmark.html\#benchmark$

Table 3. Generalizability on differ-ent detectors with our SSL method.

Methods	mAP(%)	Gain	Vehicle	AP (%) Pedestrian	Cvclist
SECOND [33]	51.89	-	71.19	26.44	58.04
Our Method	59.89	+8.00	78.07	38.38	63.23
PV-RCNN [21]	57.24	-	79.35	29.64	62.73
Our Method	63.40	+6.16	81.09	41.55	67.57
CenterPoint [38]	62.99	-	75.26	51.65	65.79
Our Method	68.22	+5.23	77.77	56.34	70.55

Table 4. Performance on the theONCE test set.

Methods	mAP (%)	AP (%)						
		Vehicle	Pedestrian	Cyclist				
SECOND [33]	51.90	69.71	26.09	59.92				
Pseudo Label	49.29	70.29	21.85	55.72				
Noisy Student [31]	56.61	74.50	33.28	62.05				
Mean Teacher [27]	59.99	76.60	36.37	66.99				
SESS [39]	58.78	74.52	36.29	65.52				
3DIoUMatch [29]	57.43	74.48	35.74	62.06				
Our Method	61 44	76.85	41.27	66 19				

official implementation of Pseudo Label includes no augmentation when training the student. Interestingly, the improvements increase gradually with the increase of the number of unlabeled samples. Compared to the other SSL methods, our framework achieves the best performance among all the three splits, which obtains 5.83, 8.00, 9.51 mAP improvements, respectively.

Since our method removes the necessity of threshold selection, we compare it with the confidence-based filtering method, *i.e.*, Noisy Student [31] with different pseudo-label threshold c. In our implementation, pseudo boxes with scores above c are viewed as positive samples, meanwhile, we also ignore the pseudo boxes with scores below c to avoid taking potential true positives as negatives. As shown in Fig. 1 (c), c = 0.3 achieves better performance (55.75 mAP), which is even stronger than Mean Teacher [27], while our method still surpasses it a lot.

We also verify our SSL method on two more baselines, which are PV-RCNN [21] and CenterPoint [38]. PV-RCNN is a hybrid point-voxel approach that inherits the advantages from both point [17] and voxels features [41]. CenterPoint is a voxel-based anchor-free framework that has achieved SOTA results in several benchmarks. All the methods are trained with the "Medium" unlabeled subset. From Tab. 3, we can see that our method works well on these different detectors, *e.g.*, still yielding 5.23 points improvements over the strong detector CenterPoint. This indicates the good generalizability of our SSL method.

The results on the testing split of the ONCE benchmark are presented in Table. 4, where the results of other SSL methods are reported in terms of the official benchmark. All the SSL methods are based on the SECOND detector for a fair comparison. Among all the competitors, our method shows remarkable superiority. It outperforms the strong competitor Mean Teacher by 1.45 points, and also exceeds the recently proposed SESS and 3DIoUMatch by a large margin. Waymo Results. For evaluating on Waymo, we compare our model with the strongest competitor verified on ONCE, *i.e.*, Noisy Student [31] with a carefully selected threshold c = 0.3. Since this filtering strategy derives from Fixmatch [24], we name it Fixmatch for simplicity. The SECOND [33] detector trained with different amounts of labeled data is used as the baseline. As shown in Table 5, impressive results are obtained. Both Fixmatch and our *Proficient*-*Teachers* achieve better results than the full-supervised baseline, which proves the advantage of semi-supervised learning. In particular, Fixmatch surpasses the baseline by 1.90 to 2.95 mAPH. By contrast, our method obtains much better results than Fixmatch, *i.e.*, improving the baseline by 4.34 to 5.35 mAPH. More-

Table 5. Semi-supervised 3D Object Detection on Waymo dataset. We train the SECOND [33] baseline with different fractions of labeled data. Then, we compare our *ProficientTeachers* model with a strong competitor, Fixmatch [24]. It shows that our model can consistently improve the detection performance.

Different	Training Davadigm	Performance	31	D AP/APH	@0.7 (LEV	EL 2)
Label Amounts	Training Faradigin	Gain	Overall	Vehicle	Pedestrian	Cyclist
5% (a. 4k Labola)	Baseline [33]	-/-	45.78/40.40	50.03/49.52	45.77/34.98	41.53/36.69
J/0 (~ 4K Labels)	Fixmatch [24]	+3.02/+2.95	48.80/43.35	51.87/51.27	48.28/36.56	46.26/42.21
$\mathcal{P}^L:\mathcal{P}^0=1:20$	ProficientTeacher	+5.32/+5.35	51.10/45.75	53.04/52.54	50.33/38.67	49.92/46.03
10% (a. Sk Labola)	Baseline [33]	-/-	50.00/45.83	54.90/54.25	48.45/38.44	46.66/44.79
10% (~ ok Labels)	Fixmatch [24]	+2.28/+1.90	52.28/47.73	56.60/55.99	51.60/40.63	48.63/46.56
$\mathcal{P}^L:\mathcal{P}^0=1:10$	ProficientTeacher	+5.01/+4.60	55.01/50.43	57.59/56.92	54.28/43.19	53.15/51.18
20% (16h Labela)	Baseline [33]	-/-	53.09/49.11	57.40/56.81	51.54/41.91	50.33/48.62
20% (~ Tok Labers)	Fixmatch [24]	+2.72/+2.34	55.81/51.45	58.94/58.37	54.37/44.23	54.11/51.75
$\mathcal{P}^L:\mathcal{P}^0=1:5$	ProficientTeacher	+5.50/+5.05	58.59/54.16	59.97/59.36	57.88/46.97	57.93/56.15
50% (40k Labela)	Baseline [33]	-/-	57.07/53.26	60.93/60.37	55.98/46.68	54.31/52.74
50% (~ 40k Labers)	Fixmatch [24]	+2.72/+2.45	59.79/55.71	61.88/61.34	58.64/49.00	58.85/56.78
$\mathcal{P}^L: \mathcal{P}^U = 1:2$	ProficientTeacher	+4.57/+4.34	61.64/57.60	63.06/62.50	61.53/51.33	60.33/58.97
	Baseline [33]	-/-	59.63/55.94	62.78/62.24	59.45/50.44	56.67/55.13
100% (~ 80k Labels)	Fixmatch [24]	+2.43/+2.02	62.06/57.96	63.50/62.98	62.00/52.52	60.69/58.37
\mathcal{D}^L , $\mathcal{D}^U = 1 \cdot 1$	ProficientTeacher	+3.33/+3.20	62.96/59.14	63.56/63.06	62.34/53.19	62.97/61.18
P : P = 1:1	Oracle Model [33]	+2.80/+2.93	62.43/58.87	65.43/64.91	61.82/52.93	60.03/58.76

over, we also run an oracle model that trains the detector with the full 158,361 labels, while our model with only half labels outperforms it. This further demonstrates the generalization of our method over different datasets.

Qualitative Results. We visualize some examples of pseudo labels on ONCE in Fig. 4. The false positives and false negatives are highlighted in circles. We compare our STE and CBV modules with the threshold-based method Fix-Match [24]. It incurs FN when detecting distant objects with sparse points, and it gives FP when detecting a hard distractor as shown in the circles in (b) (plz zoom in for a better view). In contrast, by gathering boxes from different views, our STE successfully addresses the FN object. As for the FP box, its score will be refined and re-scaled in our CBV such that a box with low detection frequency will be removed accordingly. Moreover, for boxes that are close to GTs, their box coordinates will also be refined after voting in a cluster. All these designs lead to high-quality pseudo labels.

4.4 Ablation Studies

A series of ablation studies are set to verify the effectiveness of different modules and all the results are given in Tab. 6. All the experiments are conducted based on the SECOND detector with the "Small" unlabeled subset. Our *ProficientTeachers* model is mainly based on the Noisy Student implementation in ONCE, and the default confidence threshold c for filtering the noisy pseudo labels is 0.1. To explore a more effective threshold, we perform grid search and find that c = 0.3 gives better detection results, as shown in Fig. 1 (c). Thus we use this threshold for the subsequent experiments. STE, CBV and BCL are the three necessary modules proposed in this work, and we ablate the contribution of each module. First, the STE module has been used to generate more pseudo



Fig. 4. Visualization of pseudo labels produced by our STE and CBV modules, or by threshold-based FixMatch [24]. Predictions are in red and GTs are in green.

Table 6. Ablation studies to verify the effect of different modules. The experiments are conducted based on the SECOND with the "Small" unlabeled subset.

Module	Aspect	mAP (%)	Gain	Vehicle	$\begin{array}{c} \mathbf{AP} \ (\%) \\ \mathrm{Pedestrian} \end{array}$	Cyclist
Noisy Student (Baseline)	c = 0.3	55.75	-	74.66	32.25	60.34
BCL Module	c = 0.3, w/o STE or CBV	56.54	+0.79	75.70	33.65	60.28
STE Module	c = 0.3, w/o BCL or CBV	56.21	+0.46	75.62	33.14	59.87
STE+CBV Modules	w/o BCL	57.17	+1.42	75.98	34.86	60.67
Full Model	w/ STE, CBV and BCL	57.72	+1.97	76.07	35.90	61.19

boxes. We evaluate it by using c = 0.3 to filter the boxes. This improves the baseline by 0.46 points. Next, we replace the confidence-based thresholding with the proposed CBV module to adaptively aggregate these pseudo boxes. This exceeds the baseline by 1.42 points. Finally, thanks to the BCL module, our full model achieves 57.72% mAP, improving the baseline by 1.97 points.

5 Discussion and Conclusion

In this work, we proposed a new SSL framework for LiDAR-based 3D object detection. In particular, our work focuses on improving the FP and FN in the pseudo labels produced by the teacher model. First, to address the FN, a spatial-temporal ensemble (STE) module is introduced to produce sufficient seed boxes and ensure a high recall. This is realized by a spatial data augmentation and a temporal model ensemble. Second, to resolve the FP predictions and improve the precision, we developed a clustering-based box voting (CBV) module that performs voting and aggregating based on boxes in a cluster. More importantly, our CBV can yield high-quality pseudo labels without the need of deliberately selecting thresholds. The STE and CBV modules enhance the original teacher to proficient teachers. Finally, we proposed a box-wise contrastive learning (BCL) strategy to optimize the student towards cross-view feature consistency, reducing the effect of inaccurate pseudo labels. Experiments on the large-scale ONCE and Waymo datasets demonstrated the superiority of our method.

Acknowledgements. This work was partially supported by Zhejiang Lab's International Talent Fund for Young Professionals (ZJ2020GZ023), ARC DECRA DE220101390, FDCT under grant 0015/2019/AKP, and the Start-up Research Grant (SRG) of University of Macau.

References

- 1. Berthelot, D., Carlini, N., Goodfellow, I., Papernot, N., Oliver, A., Raffel, C.A.: Mixmatch: A holistic approach to semi-supervised learning. In: NeurIPS (2019)
- Caesar, H., Bankiti, V., Lang, A.H., Vora, S., Liong, V.E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., Beijbom, O.: nuscenes: A multimodal dataset for autonomous driving. In: CVPR (2020)
- 3. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: ICML (2020)
- 4. Chen, X., Ma, H., Wan, J., Li, B., Xia, T.: Multi-view 3d object detection network for autonomous driving. In: CVPR (2017)
- 5. Deng, J., Shi, S., Li, P., gang Zhou, W., Zhang, Y., Li, H.: Voxel r-cnn: Towards high performance voxel-based 3d object detection. In: AAAI (2021)
- Everingham, M., Van Gool, L., Williams, C.K., Winn, J., Zisserman, A.: The pascal visual object classes (voc) challenge. IJCV 88(2), 303–338 (2010)
- Fang, J., Zhou, D., Yan, F., Zhao, T., Zhang, F., Ma, Y., Wang, L., Yang, R.: Augmented lidar simulator for autonomous driving. IEEE Robotics and Automation Letters 5(2), 1931–1938 (2020)
- 8. Geiger, A., Lenz, P., Urtasun, R.: Are we ready for autonomous driving? the kitti vision benchmark suite. In: CVPR (2012)
- Huang, X., Wang, P., Cheng, X., Zhou, D., Geng, Q., Yang, R.: The apolloscape open dataset for autonomous driving and its application. PAMI 42(10), 2702–2719 (2019)
- Jeong, J., Lee, S., Kim, J., Kwak, N.: Consistency-based semi-supervised learning for object detection. In: NeurIPS (2019)
- 11. Lang, A.H., Vora, S., Caesar, H., Zhou, L., Yang, J., Beijbom, O.: Pointpillars: Fast encoders for object detection from point clouds. In: CVPR (2019)
- Mao, J., Niu, M., Jiang, C., Liang, H., Chen, J., Liang, X., Li, Y., Ye, C., Zhang, W., Li, Z., et al.: One million scenes for autonomous driving: Once dataset. In: NeurIPS Datasets and Benchmarks (2021)
- Meng, Q., Wang, W., Zhou, T., Shen, J., Jia, Y., Van Gool, L.: Towards a weakly supervised framework for 3d point cloud object detection and annotation. TPAMI (2021)
- Meng, Q., Wang, W., Zhou, T., Shen, J., Van Gool, L., Dai, D.: Weakly supervised 3d object detection from lidar point cloud. In: ECCV (2020)
- Oord, A.v.d., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. arXiv preprint (2018)
- Qi, C.R., Litany, O., He, K., Guibas, L.J.: Deep hough voting for 3d object detection in point clouds. In: CVPR (2019)
- 17. Qi, C.R., Yi, L., Su, H., Guibas, L.J.: Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In: NeurIPS (2017)
- Rizve, M.N., Duarte, K., Rawat, Y.S., Shah, M.: In defense of pseudo-labeling: An uncertainty-aware pseudo-label selection framework for semi-supervised learning. In: ICLR (2021)
- Samuli, L., Timo, A.: Temporal ensembling for semi-supervised learning. In: ICLR (2017)
- 20. Shanmugam, D., Blalock, D., Balakrishnan, G., Guttag, J.: Better aggregation in test-time augmentation. In: ICCV (2021)
- Shi, S., Guo, C., Jiang, L., Wang, Z., Shi, J., Wang, X., Li, H.: Pv-rcnn: Point-voxel feature set abstraction for 3d object detection. In: CVPR (2020)

- 16 J. Yin et al.
- 22. Shi, S., Wang, X., Li, H.: Pointrcnn: 3d object proposal generation and detection from point cloud. In: CVPR (2019)
- 23. Shi, W., Rajkumar, R.: Point-gnn: Graph neural network for 3d object detection in a point cloud. In: CVPR (2020)
- Sohn, K., Berthelot, D., Carlini, N., Zhang, Z., Zhang, H., Raffel, C.A., Cubuk, E.D., Kurakin, A., Li, C.L.: Fixmatch: Simplifying semi-supervised learning with consistency and confidence. In: NeurIPS (2020)
- 25. Sohn, K., Zhang, Z., Li, C.L., Zhang, H., Lee, C.Y., Pfister, T.: A simple semisupervised learning framework for object detection. arXiv preprint (2020)
- Sun, P., Kretzschmar, H., Dotiwalla, X., Chouard, A., Patnaik, V., Tsui, P., Guo, J., Zhou, Y., Chai, Y., Caine, B., et al.: Scalability in perception for autonomous driving: Waymo open dataset. In: CVPR (2020)
- Tarvainen, A., Valpola, H.: Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In: NeurIPS (2017)
- Van Engelen, J.E., Hoos, H.H.: A survey on semi-supervised learning. Machine Learning 109(2), 373–440 (2020)
- Wang, H., Cong, Y., Litany, O., Gao, Y., Guibas, L.J.: 3dioumatch: Leveraging iou prediction for semi-supervised 3d object detection. In: CVPR (2021)
- Wang, W., Zhou, T., Yu, F., Dai, J., Konukoglu, E., Van Gool, L.: Exploring cross-image pixel contrast for semantic segmentation. In: ICCV (2021)
- Xie, Q., Luong, M.T., Hovy, E., Le, Q.V.: Self-training with noisy student improves imagenet classification. In: CVPR (2020)
- 32. Xie, S., Gu, J., Guo, D., Qi, C.R., Guibas, L., Litany, O.: Pointcontrast: Unsupervised pre-training for 3d point cloud understanding. In: ECCV (2020)
- Yan, Y., Mao, Y., Li, B.: Second: Sparsely embedded convolutional detection. Sensors 18(10), 3337 (2018)
- Yang, Z., Sun, Y., Liu, S., Jia, J.: 3dssd: Point-based 3d single stage object detector. In: CVPR (2020)
- Yin, J., Shen, J., Gao, X., Crandall, D., Yang, R.: Graph neural network and spatiotemporal transformer attention for 3d video object detection from point clouds. TPAMI (2021)
- Yin, J., Shen, J., Guan, C., Zhou, D., Yang, R.: Lidar-based online 3d video object detection with graph-based message passing and spatiotemporal transformer attention. In: CVPR (2020)
- Yin, J., Zhou, D., Zhang, L., Fang, J., Xu, C.Z., Shen, J., Wang, W.: Proposalcontrast: Unsupervised pre-training for lidar-based 3d object detection. In: ECCV (2022)
- Yin, T., Zhou, X., Krahenbuhl, P.: Center-based 3d object detection and tracking. In: CVPR (2021)
- Zhao, N., Chua, T.S., Lee, G.H.: Sess: Self-ensembling semi-supervised 3d object detection. In: CVPR (2020)
- 40. Zhou, D., Fang, J., Song, X., Liu, L., Yin, J., Dai, Y., Li, H., Yang, R.: Joint 3d instance segmentation and object detection for autonomous driving. In: CVPR (2020)
- Zhou, Y., Tuzel, O.: Voxelnet: End-to-end learning for point cloud based 3d object detection. In: CVPR (2018)
- Zhu, X., Goldberg, A.B.: Introduction to semi-supervised learning. Synthesis lectures on artificial intelligence and machine learning 3(1), 1–130 (2009)