

GeoTeacher: Geometry-Guided Semi-Supervised 3D Object Detection

Jingyu Li^{1,2*}, Xiaolong Zhao^{3*}, Zhe Liu⁴, Wenxiao Wu⁵, Li Zhang^{1,2†}

Abstract—Semi-supervised 3D object detection (SS3D), aiming to explore unlabeled data for boosting 3D object detectors, has emerged as an active research area in recent years. Some previous methods have shown substantial improvements by either employing heterogeneous teacher models to provide high-quality pseudo labels or enforcing feature-perspective consistency between the teacher and student networks. However, these methods overlook the fact that the model usually tends to exhibit low sensitivity to object geometries with limited labeled data, making it difficult to capture geometric information, which is crucial for enhancing the student model’s ability in object perception and localization. In this paper, we propose GeoTeacher to enhance the student model’s ability to capture geometric relations of objects with limited training data, especially unlabeled data. We design a keypoint-based geometric relation supervision module that transfers the teacher model’s knowledge of object geometry to the student, thereby improving the student’s capability in understanding geometric relations. Furthermore, we introduce a voxel-wise data augmentation strategy that increases the diversity of object geometries, thereby further improving the student model’s ability to comprehend geometric structures. To preserve the integrity of distant objects during augmentation, we incorporate a distance-decay mechanism into this strategy. Moreover, GeoTeacher can be combined with different SS3D methods to further improve their performance. Extensive experiments on the ONCE and Waymo datasets indicate the effectiveness and generalization of our method and we achieve the new state-of-the-art results.

I. INTRODUCTION

3D object detection, which detects objects of interest within 3D scenes, is essential for autonomous driving, robotics, and various other applications. However, achieving high-performance detection typically depends on large and fully labeled datasets, which are quite costly and time-intensive. In order to solve this problem, semi-supervised 3D object detection (SS3D) is born on demand, which intends to train object detectors based on labeled data as well as easy-to-obtain unlabeled data.

Current SS3D methods generally follow the 2D approach, utilizing a teacher-student framework where the teacher model generates pseudo-labels to train the student model. Some prior methods improve the student model’s performance by incorporating additional information or designing sophisticated networks for the teacher model to generate high-quality pseudo labels. For instance, A-Teacher [1] and X-ray Teacher [2] utilize multi-frame fusion to obtain

pseudo labels, while ProficientTeacher [3] introduces test-time augmentation and a refinement network. Recently, some methods have realized the importance of enforcing feature-level supervision to enhance student model performance. For example, NoiseDet [4] introduces a dense Gaussian feature constraint to improve the transferability of learned features, and DPKE [5] enforces feature consistency on downsampled points for indoor scenarios. However, the aforementioned methods neglect the exploration and integration of higher-order information within the object’s internal geometry, which is particularly crucial for enhancing the student model’s perceptual ability when only limited labeled data is available. This limitation hinders their potential for achieving higher performance.

To this end, we propose a novel semi-supervised 3D object detection method named GeoTeacher, which can be easily combined with existing semi-supervised approaches. GeoTeacher aims to guide the student model learning the inherent information contained within object geometries, particularly geometric information, derived from unlabeled data. Specifically, we design a keypoint-based geometric relation supervision module. We select keypoints (e.g., center points and corner points) that reflect the geometric characteristics of objects and treat the correlations among them as geometric properties. Unlike pseudo labels or features, the geometric relations could effectively represent the unique structure of each object. To mitigate the influence of unreliable knowledge, we employ pseudo-label scores to weigh the geometric knowledge of different objects accordingly.

In addition, we introduce a voxel-wise augmentation data strategy to increase the geometric diversity of objects, applicable to both labeled and unlabeled data. Different from prior approaches [6], [7] that partition the entire point cloud scene, our method focuses on decomposing individual objects into smaller constituent parts. By doing this, we are able to modify their geometry flexibly while ensuring they remain detectable. The student model can achieve more accurate object detection when incorporated with our geometric relation supervision. Since distant objects are often represented by sparse point clouds and are more challenging to detect, we introduce a distance-decay strategy to preserve their geometry during augmentation.

The main contributions of this paper can be summarized as follows: (i) We present a novel SS3D method, named *GeoTeacher*. GeoTeacher guides the student model to learn geometric information from both data and supervision perspectives. (ii) We design a geometric relation supervision model and a distant-decay voxel-wise data augmentation strategy to facilitate the student’s ability to learn geomet-

† Corresponding author, * Equal contribution

¹ The authors are with School of Data Science, Fudan University.

² The authors are with Shanghai Innovation Institute.

³ The author is with Tongji University.

⁴ The author is with Hong Kong University.

⁵ The author is with Huazhong University of Science and technology.

ric knowledge. **(iii)** Extensive experiments demonstrate that our method can effectively combine with existing SS3D approaches and achieve performance on par with or surpassing the current state-of-the-art semi-supervised 3D object detection methods.

II. RELATED WORK

A. LiDAR-based 3D object detection

3D object detection [8], [9], [10], [11], [12], [13], [14], [15] has made rapid progress in the past few years and is widely used in the field of autonomous driving and navigation robot. In terms of point cloud representation: point-based [16], [17], [18] and voxel-based [19], [20], [21], [22], [23]. Point-based methods directly use irregular point cloud to extract features. Voxel-based methods adopt voxelization to make point clouds regular, making it possible to extract features using asymmetric functions, e.g., convolution. Besides this, PV-RCNN [24] achieves higher performance by combining multi-scale voxel and point cloud features. DetZero [25] proposes a combination of an offline tracker and a multi-frame detector, achieving significant performance improvements. However, 3D object detection depends on large-scale labeled data, which is costly to obtain. Our method aims to leverage unlabeled data to enhance detection performance.

B. Semi-supervised learning

Compared to supervised learning, SSL is designed to train models with a small amount of labeled data and abundant unlabeled data, increasing the difficulty of the task. Previous works [26], [27], [28], [29] are mainly divided into two categories: consistency regularization and pseudo-labeling. Temporal Ensembling [30] first proposes consistency regularization, which aims to minimize the difference between predictions from different steps. Mean Teacher [31] contains two branches, i.e., teacher and student, with the same architecture. The teacher model is an EMA of the student model. Pseudo-labeling methods [32], [33] usually first train the model with labeled data and then iteratively generate pseudo labels to unlabeled data through self-training to improve performance. Noisy Student [34] injects noise into the student model, which is then supervised by the teacher model. In order to obtain more accurate pseudo labels.

C. Semi-supervised 3D object detection

Most traditional SSL techniques are proposed for classification tasks, but also benefit semi-supervised object detection [35], [36], [37], [38], [10], especially SS3D [39]. SESS [40] and 3DIoUMatch [41] are two pioneers in exploring SS3D task. SESS uses Mean Teacher paradigm in point-based 3D object detection, introducing asymmetric data augmentation and consistency loss. 3DIoUMatch benefits from a series of handcrafted threshold and IoU-guided deduplicate module. ProficientTeacher [3] introduces the Spatial-temporal Ensemble module and clustering-based box voting module to improve the recall and precision of

pseudo labels. NoiseDet [4] views pseudo-labeling as a noisy learning problem and proposes two constraints supervision. A-Teacher [1] leverages the capability of an offline teacher while jointly updating the entire framework, and achieves significant improvements. X-Ray Teacher [2] uses multi-frame point cloud to reconstruct the complete shapes of objects to enhance the performance of the teacher model. PatchTeacher [7] leverages partial scene detection to build a super high-resolution teacher for high-quality pseudo labels. Unlike previous methods that primarily focus on feature-level consistency or pseudo label quality, our work explicitly explores and leverages the inherent geometric relations within object geometries to guide student learning, offering a complementary perspective for improving SS3D.

III. METHOD

A. Problem definition

The objective of semi-supervised 3D object detection is to leverage both labeled and unlabeled data for training detection models. Specifically, the labeled data is represented as $L = \{(x_i^l, y_i^l)\}_{i=1}^{N_l}$, where x_i^l and y_i^l denote the i -th point clouds and its corresponding labels, and N_l is the number of labeled samples. In addition to the labeled data, we also have a set of unlabeled data, denoted as $U = \{x_i^u\}_{i=1}^{N_u}$, where N_u is the number of unlabeled samples. Typically, $N_u \gg N_l$.

B. Framework Overview

Our semi-supervised framework consists of two phases. In phase 1, we train a high-performance GeoTeacher. In phase 2, as illustrated in Figure 1, GeoTeacher supervises the student with two well-designed modules, Geometric relation supervision and distant-decay voxel-wise data augmentation.

C. Geometric Relation Supervision

Geometric information not only reflects the surface features of objects but also includes internal structures. Geometric relations help the model understand spatial relationships between different parts of an object, such as relative positions and sizes. This information is crucial for accurate detection and localization, especially with limited labeled data. In SS3D, geometric information can help the model capture implicit structural characteristics of objects from unlabeled data, thereby improving detection accuracy. Motivated by this, we propose a geometric relation supervision module (GRS).

To implement geometric relation supervision, we select a set of representative keypoints $\{(x_i, y_i)\}_{i=1}^K$ from each object. Considering the common challenges in point cloud data, such as occlusion and sparsity, we select three types of keypoints: center points, edge midpoints, and corner points based on the bird’s-eye-view (BEV)-projected 2D bounding boxes. The center point offer a stable reference for the object’s localization and is less affected by sparse or occlusion. The edge midpoints capture the object’s orientation and spatial extent along both axes, offering some geometric information. The corner points encode fine-grained boundary information, which is essential for delineating the precise spatial extent

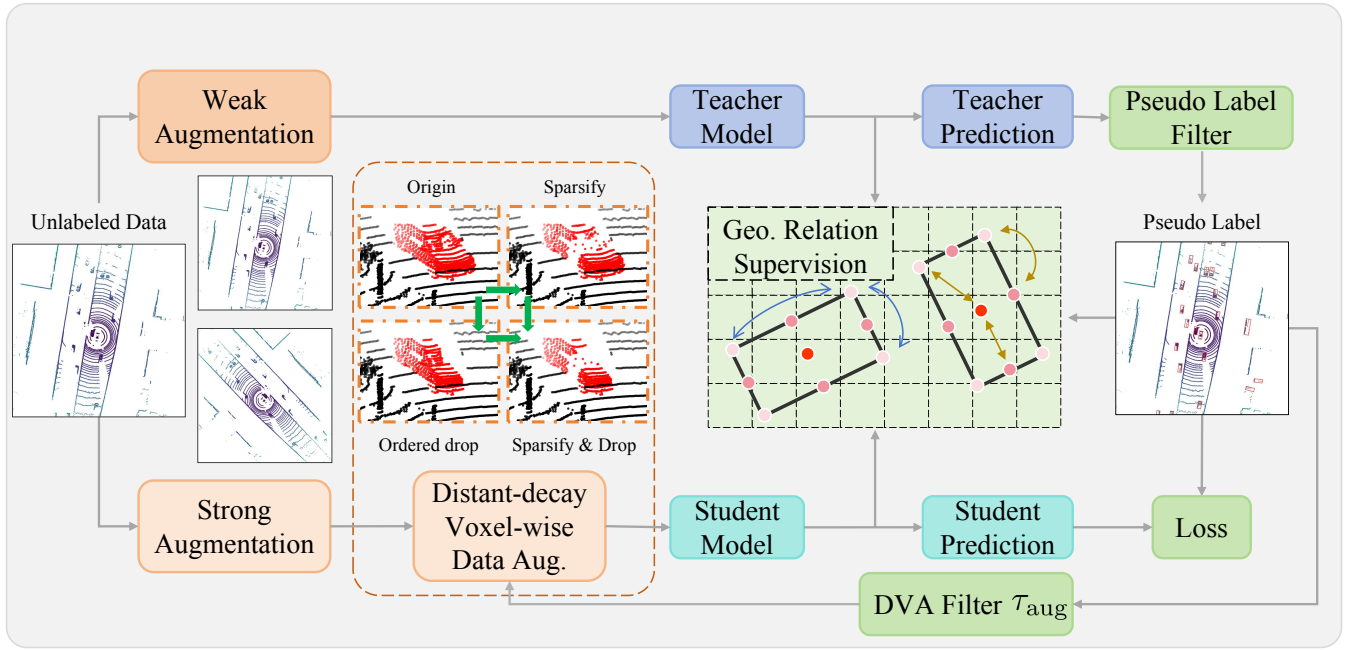


Fig. 1. **The framework of our GeoTeacher.** After applying strong augmentation, we further introduce our proposed data augmentation strategy to enhance the geometric diversity of objects. The student model is supervised with both the standard semi-supervised loss and our geometric relation supervision loss, which effectively improves the student model’s ability to comprehend the intrinsic geometric information of objects.

of the object. While each individual keypoints encodes local geometric attributes, the relationships among these points capture high-order geometric information that reflects the intrinsic geometry of the object.

To enable effective supervision, we first align the BEV feature maps of the teacher and student models by transforming them into the same coordinate space. After alignment, we extract geometry-aware features F^t and F^s from both teacher and student models based on the selected keypoints. These features serve as the basis for computing geometric relations, enabling the student to learn geometric knowledge from the teacher in a spatially consistent manner. We then compute the pairwise geometric relations M_{rel}^s among keypoints by measuring the cosine similarity between their feature representations. This results in a geometric relation matrix that captures the relative structural dependencies among different parts of the object.

$$M_{rel}^s(i, j) = \Phi(F^s(x_i, y_i), F^t(x_j, y_j)), \quad 1 \leq i, j \leq K, \quad (1)$$

where Φ is the cosine similarity function. The relation matrix M_{rel}^t of the teacher can also be calculated based on F^t , and the geometric relation supervision loss is computed as a standard L1 loss:

$$\mathcal{L}_\delta = \frac{1}{K^2} \sum_{1 \leq i, j \leq K} |M_{rel}^s(i, j) - M_{rel}^t(i, j)|. \quad (2)$$

In the semi-supervised training phase, the bounding boxes used for keypoint selection are derived from pseudo labels, which are predicted by the teacher model. However, the quality of these pseudo labels is often suboptimal, especially in challenging scenes with occlusions, sparsity, or

overlapping objects. On the one hand, the inaccurate boxes may lead to the misplacement of the keypoints, resulting in noisy geometric relationships. On the other hand, the feature representations at these keypoints may also be contaminated by background noise, especially when the box includes irrelevant regions. Hence, to mitigate the aforementioned issues, we propose a simple but effective strategy based on confidence-aware weighting. Without bells and whistles, we leverage the classification scores $\{s_k\}_{k=1}^{n_u}$ predicted by the teacher model to assess the reliability of the pseudo labels, where s_k denotes the score of the k -th pseudo-label. Pseudo labels with higher scores are assigned a greater weight in the supervised loss, whereas those with lower scores contribute less. The final loss of geometric relation supervision is computed as follows:

$$\mathcal{L}_{GRS} = \sum_{k=1}^{n_u} s_k \cdot \mathcal{L}_\delta. \quad (3)$$

Furthermore, we introduce an adaptive layer in the student model to ensure it is more complex than the teacher model during training, following [42].

D. Distant-decay Voxel-wise Data Augmentation

Although large-scale unlabeled data is available in semi-supervised settings, object geometric diversity is not fully leveraged. This is mainly because pseudo labels tend to focus on objects with regular and complete shapes, while ignoring those with occlusions or rare structural patterns. As a result, the student model lacks exposure to diverse geometric variations, which may hinder its ability to generalize to unseen object structures. To alleviate this limitation, we propose a

voxel-wise data augmentation strategy (DVA) that explicitly increases geometric diversity at the object level.

Our data augmentation strategy consists of two operations that construct diverse object representations and effectively enhance the student model’s geometric understanding of unlabeled data, which can be effectively combined with our geometric relation supervision module. Meanwhile, considering that point clouds of distant objects are inherently sparse, we introduce a distance-decay strategy to reduce the probability of applying data augmentation to distant objects. This design contributes to the overall stability of the semi-supervised learning framework by improving detection performance in near-range regions while preserving reliability in distant areas.

Specifically, we first partition the objects based on their bounding boxes into $n_l \times n_w \times n_h$ voxels. In the voxel-wise sparsify process, we randomly sample the point clouds within the selected voxels, effectively increasing the geometric diversity of the objects. As for voxel-wise order dropout, we randomly remove points from selected voxels, following a spatial sequence in either a clockwise or counterclockwise order. This approach simulates the point cloud distribution of occluded surfaces. We define that voxels within the same bounding box share the same probability p to participate in the augmentation operations. To ensure that the point clouds of distant objects remain unaffected, we introduce a distance-decay function to adjust the probability p :

$$p = c \cdot \exp\left(-\frac{\sqrt{i^2 + j^2}}{d_{\text{range}}}\right), \quad (4)$$

where p represents the probability of participating in the data augmentation, c is a hyperparameter. i and j are the coordinates of the bounding box center, and d_{range} denotes the range of the scene. Through our distance-decay voxel-wise augmentation strategies, we construct a richer and more diverse training set that reflects the complexities of a real-world environment, thereby enhancing the detection performance of our models. Notably, our method is suitable for both supervised and semi-supervised training settings. For labeled data, we generate the voxels with the ground truth boxes. For unlabeled data, we apply a classification threshold τ_{aug} , set higher than the threshold used for pseudo-label selection, to avoid introducing incorrect information. Furthermore, our data augmentation scheme is plug-and-play, allowing for straightforward integration.

E. Loss Function

Finally, we get the total loss of the student model:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{base}} + \lambda_1 \cdot \mathcal{L}_{\text{GRS}}, \quad (5)$$

where $\mathcal{L}_{\text{base}}$ is the basic semi-supervised loss, consisting of a regression loss and a classification loss. λ_1 is a hyperparameter used to balance the scale of different losses.

A. Datasets

ONCE Dataset. ONCE [43] is a large-scale autonomous driving dataset with 1 million LiDAR point cloud samples. Only 15,000 samples are annotated, which have been divided into training, validation, and testing splits with 5k, 3k, and 8k samples, respectively. According to the official SSL setting, 5k labeled samples and all unlabeled samples have been divided into 3 subsets: Small, Medium, and Large. The small set contains 100k samples (70 sequences), the medium set contains 500k samples (321 sequences), and the large set contains about 1M samples (560 sequences). In addition, three different perception ranges, “0-30m”, “30-50m”, and “50m-inf”, are specified to evaluate the performance of 3D detectors.

Waymo Open Dataset. Waymo [44] is a large-scale LiDAR point cloud dataset, which contains 798 sequences for training and 202 sequences for validation. Following ProficientTeacher [3], we divide the 798 training sequences equally into labeled split $\mathcal{P}^{\mathcal{L}}$ and unlabeled split $\mathcal{P}^{\mathcal{U}}$, with each containing 399 sequences. 5% and 20% sequences are randomly sampled from $\mathcal{P}^{\mathcal{L}}$, leading to the ratio of labeled data and unlabeled data $\mathcal{P}^{\mathcal{L}} : \mathcal{P}^{\mathcal{U}}$ as 1:20 and 1:5, respectively.

B. Implementation details

Since our approach is generic and can be seamlessly combined with other SS3D methods, we select two representative SS3D methods, ProficientTeacher [3] and PTPM [7], to validate its effectiveness. In phase 1, we follow the same training strategy as theirs to train the teacher model. In phase 2, for the training of the student, we sample 1 labeled and 4 unlabeled samples for each mini-batch to avoid training collapse. The student is trained for 50, 75, and 100 epochs for small, medium, and large settings on the ONCE dataset [43]. For DVA, we empirically divide each box into $4 \times 2 \times 1$ voxels, set c in DVA and τ_{aug} to 0.05 and 0.7. While λ_u and λ_1 are set to 1.0, 2.0, respectively. All models are trained on a machine with 4 NVIDIA H100 GPUs.

C. Main results

ONCE dataset. We combine our GeoTeacher with ProficientTeacher [3] and PTPM [7] and evaluate them on the ONCE dataset [43]. As shown in Table I, our method consistently brings significant improvements under different proportions of unlabeled data. For example, under the small protocol, it improves ProficientTeacher by +2.01 mAP and PTPM by +1.92 mAP. Our GeoTeacher also achieves 63.16 mAP and 65.70 mAP under the large protocol, outperforming ProficientTeacher and PTPM by +1.76 and +3.02 mAP, respectively. It is worth noting that when combined with PTPM, our method achieves 62.67 mAP under the small protocol (100K unlabeled data), which is comparable to the performance of PTPM under the large protocol (1M unlabeled data). This indicates that our geometry-guided method can effectively enhance the utilization of unlabeled data.

TABLE I
COMPARISON WITH OTHER SS3D APPROACHES ON ONCE DATASET. WE COMBINE OUR GEOTEACHER WITH PROFICIENTTEACHER [7] AND PROFICIENTTEACHER [3].

Methods	Vehicle AP (%)				Pedestrian AP (%)				Cyclist AP (%)				mAP (%)
	overall	0-30m	30-50m	50m-inf	overall	0-30m	30-50m	50m-inf	overall	0-30m	30-50m	50m-inf	
Baseline	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)													
Noisy Student	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
SESS	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
3DIoUMatch	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
NoiseDet	75.26	86.36	67.52	55.29	37.96	42.36	32.78	23.28	60.77	72.31	55.03	38.87	58.00
ProficientTeacher	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
+ GeoTeacher (Ours)	78.13	86.53	73.00	60.18	38.39	43.49	34.39	21.26	62.67	74.15	57.74	40.06	59.73
PTPM	76.27	86.55	69.61	56.02	44.29	51.95	35.86	20.91	61.70	75.19	54.92	34.57	60.75
+ GeoTeacher (Ours)	79.32	89.61	73.98	58.93	44.60	50.56	38.88	27.03	64.08	76.06	57.54	40.36	62.67
Medium (500K unlabeled Samples)													
Noisy Student	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
SESS	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
3DIoUMatch	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
NoiseDet	77.14	87.21	69.94	58.44	40.45	44.91	35.71	24.11	62.59	73.04	57.98	42.67	60.06
ProficientTeacher	78.07	87.43	72.50	59.51	38.38	42.45	34.62	25.58	63.23	74.70	58.19	40.73	59.89
+ GeoTeacher (Ours)	78.94	87.34	73.41	60.70	42.00	48.11	35.55	21.77	63.94	73.93	58.71	43.38	61.62
PTPM	76.66	86.75	71.30	56.87	45.87	54.98	37.35	20.89	61.88	74.08	56.52	33.30	61.47
+ GeoTeacher (Ours)	79.51	89.77	74.11	59.12	48.33	54.91	41.19	25.83	64.76	76.57	58.57	40.64	64.20
Large (1M unlabeled Samples)													
Noisy Student	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
SESS	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
3DIoUMatch	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
NoiseDet	78.02	87.00	72.55	59.49	42.89	46.52	38.21	26.60	62.74	73.19	58.03	42.88	61.16
ProficientTeacher	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
+ GeoTeacher (Ours)	80.86	89.15	75.77	63.24	43.03	49.93	35.66	22.70	65.59	76.40	61.18	44.30	63.16
PTPM	76.46	86.35	71.31	57.08	45.72	55.00	36.81	20.25	65.87	77.41	59.85	42.39	62.68
+ GeoTeacher (Ours)	79.72	89.86	74.53	59.22	51.36	61.17	42.69	25.06	66.03	77.49	60.04	40.53	65.70

TABLE II
GENERALIZABILITY ON DIFFERENT DETECTORS. WE COMBINE OUR GEOTEACHER WITH PROFICIENTTEACHER [7] AND PROFICIENTTEACHER [3].

Detector	Method	mAP %
PV-RCNN	ProficientTeacher	61.27
	+ GeoTeacher (Ours)	62.83
	PTPM	62.22
	+ GeoTeacher (Ours)	63.48
CenterPoint	ProficientTeacher	66.84
	+ GeoTeacher (Ours)	68.23
	PTPM	68.40
	+ GeoTeacher (Ours)	69.55

Generalization on other detectors. To better assess the generalization capability of our method, we combined GRS and DVA with ProficientTeacher [3] and PTPM [7], then conducted experiments on two additional detectors: PV-RCNN [24], a classic two-stage detector, and CenterPoint [46], a widely adopted single-stage anchor-free detector. Experiments are conducted on the ONCE dataset using the small protocol. As shown in Table II, our method yields significant improvements across different detectors. Using PV-RCNN, our method achieves improvements of 1.56 mAP and 1.26 mAP, respectively. Despite CenterPoint being a stronger detector, our method still delivers a remarkable improvement of +1.39 mAP and +1.15 mAP, respectively. These results demonstrate the robust generalization ability of our approach.

Waymo Open dataset. We further evaluate our method with existing methods on the Waymo Open dataset [44] with Pro-

ficientTeacher [3] and PTPM [7]. The experimental results in Table III show that our GeoTeacher consistently improve two famous methods under different labeling ratios. In the widely used 5% label regime, our approach outperforms PTPM by +0.92 AP and +0.81 APH and ProficientTeacher by +1.28 AP and +1.32 APH, and this advantage remains evident under the 20% labeled setting, further validating the robustness and scalability of our framework. Moreover, we also run an oracle model that trains the detector with the full 798 sequences, while our model with only half labels outperforms it. These demonstrate that by learning geometric information, the student model can achieve superior detection performance even on a complex dataset like Waymo. To further evaluate our method, we use PV-RCNN [24] as baseline detector and compare with HSSDA [6] and PTPM. As shown in Table IV, our approach also delivers a notable improvement over previous methods.

D. Ablation Study

To better understand how the proposed approach works, we conduct a series of ablation studies on the ONCE small protocol training set, with evaluations on the validation set. **The effect of each component.** We conduct ablation studies to evaluate the individual and combined effects of our proposed Distant-decay voxel-wise data augmentation and Geometric Relation Supervision. We combine PTPM with our method, and the result obtained by removing GRS and DVA is used as the baseline. As shown in Table V, each module independently contributes to performance improvement, and their combination yields the highest mAP.

TABLE III

COMPARISON WITH OTHER SS3D APPROACHES ON WAYMO DATASET. WE USE THE SAME DATA SPLIT AS PTPM [7].

Label Amounts	Method	3D AP/APH @0.7 (LEVEL 2)			
		Overall	Vehicle	Pedestrian	Cyclist
5% (~ 4k Labels)	SECOND [21]	45.78/40.40	50.03/49.52	45.77/34.98	41.53/36.69
	HASS [45]	52.64/48.27	53.36/52.73	50.15/39.63	54.42/52.46
	ProficientTeacher [3]	51.10/45.75	53.04/52.54	50.33/38.67	49.92/46.03
	+ GeoTeacher(Ours)	52.38/47.07	55.28/54.69	51.70/41.44	50.18/45.09
$\mathcal{P}^L : \mathcal{P}^U = 1 : 20$	PTPM [7]	54.53/51.07	56.98/56.48	52.64/44.19	53.96/52.55
	+ GeoTeacher(Ours)	55.45/51.88	57.84/57.53	53.65/45.01	54.88/53.10
	SECOND [21]	53.09/49.11	57.40/56.81	51.54/41.91	50.33/48.62
	HASS [45]	59.09/54.22	59.19/58.38	57.75/46.87	60.32/57.41
20% (~ 16k Labels)	ProficientTeacher [3]	58.59/54.16	59.97/59.36	57.88/46.97	57.93/56.15
	+ GeoTeacher (Ours)	59.38/55.60	61.54/61.04	58.51/49.39	58.11/56.37
	PTPM [7]	60.49/57.02	61.62/61.16	59.78/51.17	60.07/58.74
	+ GeoTeacher (Ours)	61.05/57.40	62.30/61.61	60.29/51.49	60.56/59.10
100% (~ 80k Labels)	SECOND [21]	48.80/43.35	51.87/51.27	48.28/36.56	46.26/42.21
	ProficientTeacher [3]	62.96/59.14	63.56/63.06	62.34/53.19	62.97/61.18
	+ GeoTeacher (Ours)	65.21/61.55	66.66/66.19	65.22/56.45	63.77/62.03
	PTPM [7]	65.73/62.13	67.12/66.67	65.85/57.11	64.23/62.60
$\mathcal{P}^L : \mathcal{P}^U = 1 : 5$	+ GeoTeacher (Ours)	66.40/62.52	67.80/66.97	66.59/57.46	64.80/63.14
	Oracle Model	62.43/58.87	65.43/64.91	61.82/52.93	60.03/58.76

TABLE IV

PERFORMANCE COMPARISON ON THE WAYMO OPEN DATASET WITH 202 VALIDATION SEQUENCES FOR THE 3D OBJECT DETECTION WITH PV-RCNN [12] AS THE DETECTOR.

1% Data (~ 1.4k scenes)	3D AP/APH @0.7 (LEVEL 2)			
	Vehicle	Pedestrian	Cyclist	
PV-RCNN	44.7/ 42.2	30.7 / 15.3	2.8 / 1.6	
DetMatch	48.1 / 47.2	35.8 / 17.1	- / -	
HSSDA	49.7 / 47.3	33.5 / 17.5	27.9/ 20.0	
PTPM	53.7 / 52.2	36.3 / 18.8	35.7 / 34.3	
+ GeoTeacher (Ours)	55.3 / 53.6	38.1 / 19.7	37.2 / 35.8	

TABLE V

ABLATION STUDY ON EACH COMPONENT OF GEOTEACHER.

DVA	GRS	mAP (%)			
		Overall	Vehicle	Pedestrian	Cyclist
		60.06	76.45	41.95	61.78
✓		61.45	77.35	43.11	63.91
	✓	61.68	78.14	43.62	63.29
✓	✓	62.67	79.32	44.60	64.08

Specifically, applying DVA alone boosts the overall mAP by +1.39, indicating its effectiveness in enriching the geometric diversity of objects from unlabeled data. GRS alone also brings a comparable gain of +1.62 mAP, demonstrating the benefit of explicitly modeling object-internal geometric relations. When both components are combined, our model achieves the best result, showing that the two components are complementary. Together, they enhance the student model’s ability to perceive and understand geometric cues in the semi-supervised setting. These results verify that both data-level and supervision-level geometry modeling are crucial to the success of GeoTeacher.

Comparison with DVA and other approaches. To validate the superiority of our DVA, we compared it with other related methods and presented the results at different distances. We adopt three well-designed data augmentation strategies for comparisons: 1) SE-SSD [18] divides object’s point cloud into six pyramidal-shape subsets and then undergo random sparsify, random dropout, and random swapping;

TABLE VI

COMPARISON OF OUR DVA AND OTHER APPROACHES.

Method	Overall	mAP (%)		
		0-30m	30-50m	50m-inf
Baseline	53.34	62.44	47.52	33.38
SE-SSD	54.79	64.31	50.18	34.97
TED	55.22	64.44	50.27	35.12
HINTED	54.84	64.52	49.44	35.51
DVA (Ours)	55.45	64.79	50.55	36.19

2) TED [47] augments objects by translating the point clouds along the depth axis and down-sampling the whole object point clouds; 3) HITED [48] down-samples nearby point clouds to simulate distant scenes. We use SECOND [21], trained solely on labeled data, as our baseline and incorporate the aforementioned augmentation modules into it. The results in Table VI show that each approach improves the baseline, while our DVA achieves superior performance. Our DVA brings +2.11 mAP overall and gains significant improvement, which indicates that increasing geometric diversity through our augmentation strategy remains beneficial even in the presence of labeled data. Moreover, our method achieves superior performance in the “50–inf” range, which can be attributed to the proposed distant-decay mechanism.

TABLE VII

COMPARISON OF OUR GRS AND OTHER APPROACHES.

Method	Overall	mAP (%)		
		Vehicle	Pedestrian	Cyclist
Baseline	60.06	76.45	41.95	61.78
SOOD	60.31	76.97	41.94	62.01
NoiseDet	61.10	78.12	42.63	62.55
BoxMask	60.84	78.90	42.39	61.23
GRS (Ours)	61.68	78.14	43.62	63.29

Comparison with GRS and other approaches. Unlike common semi-supervised approaches that rely on feature distillation, our proposed GeoTeacher guides the student model through geometric relation modeling. To demonstrate the superiority of our approach, we compare it with several representative feature-level distillation methods: SOOD [49],

NoisDet [4] and the commonly used box-level masking method. Although our GRS does not directly enforce consistency at the feature level, it transfers high-order geometric knowledge between the teacher and student models through relational modeling among keypoints. Unlike prior methods that treat features in isolation, our approach captures the internal geometric relation within objects. As shown in Table VII, the aforementioned methods have achieved significant improvements compared to the baseline, our GRS attains the highest performance, surpassing the other methods by at least 0.58 mAP overall. These results demonstrate that geometric relations can be more effective than enforcing low-level feature similarity, as it provides richer inductive biases that are crucial for understanding the spatial organization of objects, particularly in sparse or occluded point cloud scenarios.

TABLE VIII
ABLATIONS OF GRID-SET OF DVA.

Grid-Set ($n_l \times n_w \times n_h$)	Overall	Vehicle	Pedestrian	Cyclist
$3 \times 3 \times 2$	55.08	73.17	32.33	59.72
$3 \times 2 \times 1$	55.27	74.67	32.40	58.73
$4 \times 2 \times 1$	55.45	74.71	32.33	59.31
$4 \times 2 \times 2$	54.79	73.68	31.68	59.01

TABLE IX
ABLATIONS OF KEY POINTS SELECTION OF GRS.

Center points	Corner points	Edge midpoints	Overall mAP (%)
✓	✓		62.09
	✓	✓	61.84
✓		✓	62.14
✓	✓	✓	62.67

The effect of hyperparameters. We now investigate the effect of different hyperparameters in our framework. We first ablate the grid-set on our data augmentation strategy. As shown in Table VIII, all grid-set configurations lead to performance improvements compared to the baseline presented in Table VI. Notably, the best performance is achieved with the $4 \times 2 \times 1$ grid-set configuration. Subsequently, we validate the impact of keypoints selection in GRS. The results are shown in Table IX. All combinations of keypoints lead to performance gains, and the best results are achieved when center points, edge midpoints, and corner points are used together. This demonstrates that the selected keypoints provide complementary geometric information, and jointly modeling their relations yields more effective geometric relation supervision. We also conducted hyperparameter experiments for the loss weight λ_1 and the pseudo-label confidence threshold c . As shown in Table X and Table XI determining that the model’s performance is optimal when they are set to 2.0 and 0.3, respectively.

E. Qualitative visualization results

In this section, we present qualitative results to demonstrate the effectiveness of GeoTeacher. As shown in Figure 2, we compare our method with the baseline model. The visualization reveals that our approach significantly reduces false

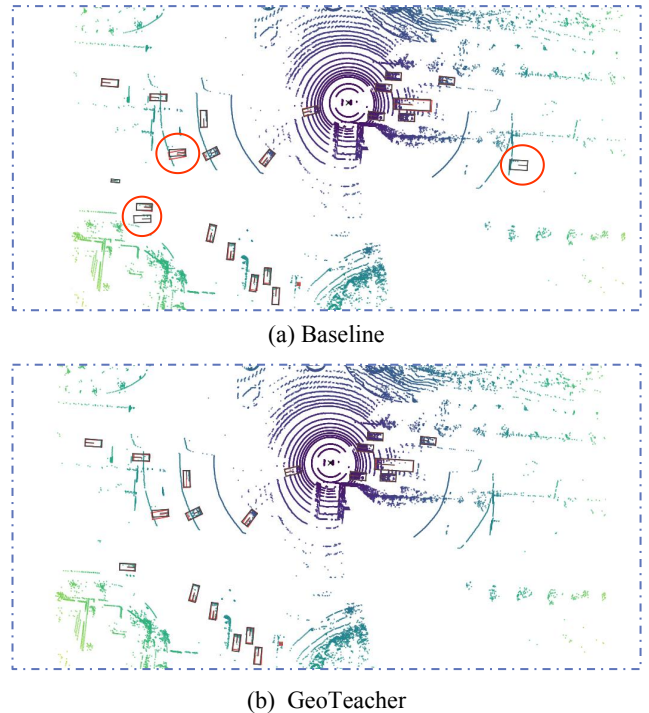


Fig. 2. Some visualization examples from the ONCE dataset [43] are presented. From top to down, each row displays the baseline results and the predictions of our GeoTeacher. Red rectangles denote ground truth, black rectangles represent predictions, and red dashed circles highlight the improvements achieved by our method.

TABLE X
THE EFFECT OF LOSS WEIGHT λ_1 .

λ_1	0.5	1.0	2.0	3.0
Overall mAP (%)	60.97	61.22	61.68	60.54

TABLE XI
THE EFFECT OF CLASSIFICATION THRESHOLD c .

c	0.2	0.3	0.4	0.5
Overall mAP (%)	54.75	55.53	55.21	54.90

positives and better adjusts object orientations, highlighting the benefits of the proposed method.

V. CONCLUSION

In this paper, we present *GeoTeacher*, a novel method for semi-supervised 3D object detection. We analyze the importance of object’s geometric information for SS3D, particularly in the context of limited labeled data and large-scale unlabeled data. To this end, we propose a Geometric Relation Supervision that guides the student model to learn the intrinsic geometric relation of objects, leading to significant performance improvements. In addition, to further enhance geometric diversity, we introduce a distant-decay voxel-wise augmentation strategy that perturbs the geometric relationships of nearby objects more aggressively while preserving the detectability of distant ones. Extensive experiments on the ONCE and Waymo datasets demonstrate the effectiveness and generalization of our approach.

VI. ACKNOWLEDGEMENTS

This work was supported in part by New Generation Artificial Intelligence-National Science and Technology Major Project (2025ZD0123004), Ningbo grant (2025Z038) and National Natural Science Foundation of China (Grant No. 62376060).

REFERENCES

- [1] H. Wang, Z. Zhang, J. Gao, and W. Hu, "A-teacher: Asymmetric network for 3d semi-supervised object detection," in *CVPR*, 2024.
- [2] A. Gambashidze, A. Dadukin, M. Golyadkin, M. Razzhivina, and I. Makarov, "Weak-to-strong 3d object detection with x-ray distillation," in *CVPR*, 2024.
- [3] J. Yin, J. Fang, D. Zhou, L. Zhang, C.-Z. Xu, J. Shen, and W. Wang, "Semi-supervised 3d object detection with proficient teachers," in *ECCV*, 2022.
- [4] Z. Chen, Z. Li, S. Wang, D. Fu, and F. Zhao, "Learning from noisy data for semi-supervised 3d object detection," in *ICCV*, 2023.
- [5] Y. Han, N. Zhao, W. Chen, K. T. Ma, and H. Zhang, "Dual-perspective knowledge enrichment for semi-supervised 3d object detection," in *AAAI*, 2024.
- [6] C. Liu, C. Gao, F. Liu, P. Li, D. Meng, and X. Gao, "Hierarchical supervision and shuffle data augmentation for 3d semi-supervised object detection," in *CVPR*, 2023.
- [7] X. Wu, L. Peng, L. Xie, Y. Hou, B. Lin, X. Huang, H. Liu, D. Cai, and W. Ouyang, "Semi-supervised 3d object detection with patchteacher and pillarmix," *AAAI*, vol. 38, 2024. [Online]. Available: <https://ojs.aaai.org/index.php/AAAI/article/view/28432>
- [8] T. Huang, Z. Liu, X. Chen, and X. Bai, "Epnnet: Enhancing point features with image semantics for 3d object detection," in *ECCV*, 2020.
- [9] Z. Liu, T. Huang, B. Li, X. Chen, X. Wang, and X. Bai, "Epnnet++: Cascade bi-directional fusion for multi-modal 3d object detection," *IEEE TPAMI*, vol. 45, no. 7, pp. 8324–8341, 2022.
- [10] J. Li, Z. Liu, J. Hou, and D. Liang, "Dds3d: Dense pseudo-labels with dynamic threshold for semi-supervised 3d object detection," in *ICRA*, 2023.
- [11] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, "Pointnet: Deep learning on point sets for 3d classification and segmentation," in *CVPR*, 2017.
- [12] S. Shi, L. Jiang, J. Deng, Z. Wang, C. Guo, J. Shi, X. Wang, and H. Li, "Pv-rnn++: Point-voxel feature set abstraction with local vector representation for 3d object detection," *IJCV*, 2023.
- [13] J. Park, C. Xu, Y. Zhou, M. Tomizuka, and W. Zhan, "Detmatch: Two teachers are better than one for joint 2d and 3d semi-supervised object detection," in *ECCV*, 2022.
- [14] H. Wang, C. Shi, S. Shi, M. Lei, S. Wang, D. He, B. Schiele, and L. Wang, "Dsvt: Dynamic sparse voxel transformer with rotated sets," in *CVPR*, 2023.
- [15] Y. Qin, C. Wang, Z. Kang, N. Ma, Z. Li, and R. Zhang, "Supfusion: Supervised lidar-camera fusion for 3d object detection," in *ICCV*, 2023.
- [16] Y. Zhang, Q. Hu, G. Xu, Y. Ma, J. Wan, and Y. Guo, "Not all points are equal: Learning highly efficient point-based detectors for 3d lidar point clouds," in *CVPR*, 2022.
- [17] S. Shi, X. Wang, and H. Li, "Pointcnn: 3d object proposal generation and detection from point cloud," in *CVPR*, 2019.
- [18] W. Zheng, W. Tang, L. Jiang, and C.-W. Fu, "Se-ssd: Self-ensembling single-stage object detector from point cloud," in *CVPR*, 2021.
- [19] Z. Liu, X. Zhao, T. Huang, R. Hu, Y. Zhou, and X. Bai, "Tanet: Robust 3d object detection from point clouds with triple attention," in *AAAI*, vol. 34, no. 07, 2020, pp. 11 677–11 684.
- [20] Y. Zhou and O. Tuzel, "Voxelnet: End-to-end learning for point cloud based 3d object detection," in *CVPR*, 2018.
- [21] Y. Yan, Y. Mao, and B. Li, "Second: Sparsely embedded convolutional detection," *Sensors*, 2018.
- [22] Z. Liu, J. Hou, X. Wang, X. Ye, J. Wang, H. Zhao, and X. Bai, "Lion: Linear group rnn for 3d object detection in point clouds," *NeurIPS*, vol. 37, pp. 13 601–13 626, 2024.
- [23] Z. Liu, J. Hou, X. Ye, J. Wang, H. Zhao, and X. Bai, "Unilion: Towards unified autonomous driving model with linear group rnns," *arXiv*, 2025.
- [24] S. Shi, C. Guo, L. Jiang, Z. Wang, J. Shi, X. Wang, and H. Li, "Pv-rnn: Point-voxel feature set abstraction for 3d object detection," in *CVPR*, 2020.
- [25] T. Ma, X. Yang, H. Zhou, X. Li, B. Shi, J. Liu, Y. Yang, Z. Liu, L. He, Y. Qiao, *et al.*, "Detzero: Rethinking offboard 3d object detection with long-term sequential point clouds," in *CVPR*, 2023.
- [26] B. Zhang, Y. Wang, W. Hou, H. Wu, J. Wang, M. Okumura, and T. Shinozaki, "Flexmatch: Boosting semi-supervised learning with curriculum pseudo labeling," *NeurIPS*, 2021.
- [27] Y. Wang, H. Chen, Q. Heng, W. Hou, Y. Fan, Z. Wu, J. Wang, M. Savvides, T. Shinozaki, B. Raj, *et al.*, "Freematch: Self-adaptive thresholding for semi-supervised learning," *arXiv preprint*, 2022.
- [28] D. Berthelot, N. Carlini, I. Goodfellow, N. Papernot, A. Oliver, and C. A. Raffel, "Mixmatch: A holistic approach to semi-supervised learning," in *NeurIPS*, 2019.
- [29] K. Sohn, D. Berthelot, N. Carlini, Z. Zhang, H. Zhang, C. A. Raffel, E. D. Cubuk, A. Kurakin, and C.-L. Li, "Fixmatch: Simplifying semi-supervised learning with consistency and confidence," in *NeurIPS*, 2020.
- [30] S. Laine and T. Aila, "Temporal ensembling for semi-supervised learning," *arXiv preprint*, 2016.
- [31] A. Tarvainen and H. Valpola, "Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results," *NeurIPS*, 2017.
- [32] J. Li, R. Socher, and S. C. Hoi, "Dividemix: Learning with noisy labels as semi-supervised learning," *arXiv preprint*, 2020.
- [33] E. Arazo, D. Ortego, P. Albert, N. E. O'Connor, and K. McGuinness, "Pseudo-labeling and confirmation bias in deep semi-supervised learning," in *IJCNN*, 2020.
- [34] Q. Xie, M.-T. Luong, E. Hovy, and Q. V. Le, "Self-training with noisy student improves imagenet classification," in *CVPR*, 2020.
- [35] G. Li, X. Li, Y. Wang, Y. Wu, D. Liang, and S. Zhang, "Psec: Pseudo labeling and consistency training for semi-supervised object detection," in *ECCV*, 2022.
- [36] B. Zhu, M. Ding, P. Jacobson, M. Wu, W. Zhan, M. Jordan, and J. Jiao, "Doubly-robust self-training," *NeurIPS*, 2024.
- [37] H. Zhou, Z. Ge, S. Liu, W. Mao, Z. Li, H. Yu, and J. Sun, "Dense teacher: Dense pseudo-labels for semi-supervised object detection," in *ECCV*, 2022.
- [38] M. Xu, Z. Zhang, H. Hu, J. Wang, L. Wang, F. Wei, X. Bai, and Z. Liu, "End-to-end semi-supervised object detection with soft teacher," in *ICCV*, 2021.
- [39] J. Zhan, T. Liu, R. Li, Z. Zhang, and Y. Chen, "Csot: Cross-scan object transfer for semi-supervised lidar object detection," in *ECCV*, 2024.
- [40] N. Zhao, T.-S. Chua, and G. H. Lee, "Sess: Self-ensembling semi-supervised 3d object detection," in *CVPR*, 2020.
- [41] H. Wang, Y. Cong, O. Litany, Y. Gao, and L. J. Guibas, "3dioumatch: Leveraging iou prediction for semi-supervised 3d object detection," in *CVPR*, 2021.
- [42] S. Zhou, W. Liu, C. Hu, S. Zhou, and C. Ma, "Unidistill: A universal cross-modality knowledge distillation framework for 3d object detection in bird's-eye view," in *CVPR*, 2023.
- [43] J. Mao, N. Minzhe, C. Jiang, h. liang, J. Chen, X. Liang, Y. Li, C. Ye, W. Zhang, Z. Li, J. Yu, C. XU, and H. Xu, "One million scenes for autonomous driving: Once dataset," in *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, 2021.
- [44] P. Sun, H. Kretzschmar, X. Dotiwalla, A. Chouard, V. Patnaik, P. Tsui, J. Guo, Y. Zhou, Y. Chai, B. Caine, V. Vasudevan, W. Han, J. Ngiam, H. Zhao, A. Timofeev, S. Ettinger, M. Krivokon, A. Gao, A. Joshi, Y. Zhang, J. Shlens, Z. Chen, and D. Anguelov, "Scalability in perception for autonomous driving: Waymo open dataset," in *CVPR*, 2020.
- [45] S. Zeng, W. Zheng, J. Lu, and H. Yan, "Hardness-aware scene synthesis for semi-supervised 3d object detection," *IEEE Transactions on Multimedia*, 2024.
- [46] T. Yin, X. Zhou, and P. Krahenbuhl, "Center-based 3d object detection and tracking," in *CVPR*, 2021.
- [47] H. Wu, C. Wen, W. Li, X. Li, R. Yang, and C. Wang, "Transformation-equivariant 3d object detection for autonomous driving," in *AAAI*, 2023.
- [48] Q. Xia, W. Ye, H. Wu, S. Zhao, L. Xing, X. Huang, J. Deng, X. Li, C. Wen, and C. Wang, "Hinted: Hard instance enhanced detector with mixed-density feature fusion for sparsely-supervised 3d object detection," in *CVPR*, 2024.
- [49] W. Hua, D. Liang, J. Li, X. Liu, Z. Zou, X. Ye, and X. Bai, "Sood: Towards semi-supervised oriented object detection," in *CVPR*, 2023.