

GPA-3D: Geometry-aware Prototype Alignment for Unsupervised Domain Adaptive 3D Object Detection from Point Clouds

Ziyu Li^{1*} Jingming Guo² Tongtong Cao² Liu Bingbing² Wankou Yang^{1†}

¹ School of Automation, Southeast University ² Huawei Noah's Ark Lab

{liziyu, wkyang}@seu.edu.cn

{guojingming, caotongtong, liu.bingbing}@huawei.com

Abstract

LiDAR-based 3D detection has made great progress in recent years. However, the performance of 3D detectors is considerably limited when deployed in unseen environments, owing to the severe domain gap problem. Existing domain adaptive 3D detection methods do not adequately consider the problem of the distributional discrepancy in feature space, thereby hindering generalization of detectors across domains. In this work, we propose a novel unsupervised domain adaptive 3D detection framework, namely **Geometry-aware Prototype Alignment (GPA-3D)**, which explicitly leverages the intrinsic geometric relationship from point cloud objects to reduce the feature discrepancy, thus facilitating cross-domain transferring. Specifically, GPA-3D assigns a series of tailored and learnable prototypes to point cloud objects with distinct geometric structures. Each prototype aligns BEV (bird's-eye-view) features derived from corresponding point cloud objects on source and target domains, reducing the distributional discrepancy and achieving better adaptation. The evaluation results obtained on various benchmarks, including Waymo, nuScenes and KITTI, demonstrate the superiority of our GPA-3D over the state-of-the-art approaches for different adaptation scenarios. The MindSpore version code will be publicly available at <https://github.com/Liz66666/GPA3D>.

1. Introduction

As a fundamental research in 3D scene understanding, 3D detection from point clouds has attracted increasing attention due to its essential role in intelligent robotics, augmented reality and autonomous driving [7, 21, 18, 8, 1]. Despite significant progress, state-of-the-art 3D detectors still




Figure 1. The performance comparison with previous works [33, 36, 37]. The detection architecture is SECOND-IoU [34, 36].

suffer from dramatic performance degradation when training data and test data are from different environments, *i.e.*, domain shift problem [33]. Various factors, such as diverse weather conditions, object sizes, laser beams, and scanning patterns, lead to substantial discrepancies across different domains, hindering the transferability of existing LiDAR-based 3D detectors. Intuitively, fine-tuning the detectors with adequate data from the target domain could alleviate this issue. However, manually annotating a large amount of point cloud scenes is a prohibitively expensive task. Therefore, the research on unsupervised domain adaptation (UDA) for LiDAR-based 3D detection is essential.

Although many works have been proposed to deal with the UDA for image-based detection [42, 11, 24, 17, 10, 14, 13, 25, 3], directly applying these methods to 3D point cloud detection is insufficient for tackling the domain shifts. These approaches mainly concentrate on the gaps of lighting and texture variations, which could not be obtained from point clouds. While there is only a limited number of literature [33, 36, 26, 9, 39, 37, 20] dealing with the UDA on LiDAR-based 3D detection. Prior work [33] utilizes the statistics from target annotation to perform data-level normalization. MLC-Net [20] designs a mean-teacher framework to provide reliable pseudo-labels to facilitate transferring. ST3D [36] and ST3D++ [37] propose a self-training pipeline with a memory bank to collect and refine pseudo-labels. Despite their great success, these methods do not adequately consider the problem of distributional discrepancy in feature space, hampering the adaptation performance.

*Work was jointly done when Ziyu was an intern at Huawei Noah's Ark Lab

†Corresponding Author

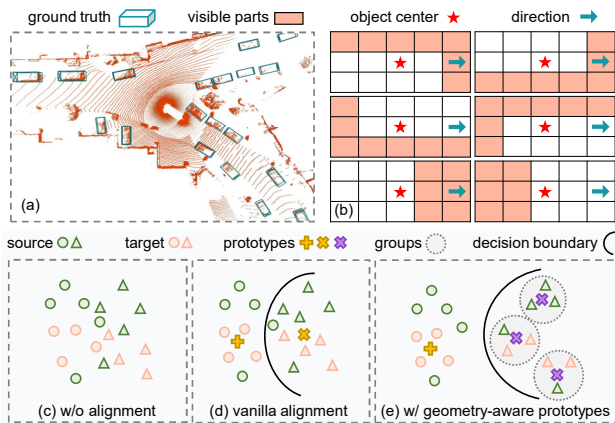


Figure 2. (a) Point cloud scene on BEV (bird’s-eye-view). (b) Distinct geometric structures of point cloud objects. (c-e) Illustration of the distributional discrepancy. With explicit geometric constraints, the features from different domains are better aligned.

To reduce this discrepancy in 2D UDA task, some approaches utilize the class-wise prototypes align features from different domains [12, 31, 40, 19]. In these works, a universal prototype is employed to enforce high representational similarity among features belonging to the same category. However, in the case of 3D scenes, such as vehicles on the road, diverse locations and directions can result in distinct geometric structures, *i.e.*, distributional patterns of point clouds, as presented in Fig. 2 (a) and (b). If a uniform prototype is applied to objects with completely different geometric structures, the efficacy of feature alignment might be hindered, as illustrated in Fig. 2 (c-d). We argue that adopting different prototypes to point cloud objects with distinct geometric structures could deal with the problem of distributional discrepancy, but more attention should also be paid to model these geometric structures during adaptation.

Based on the considerations, we propose a novel UDA framework for LiDAR-based 3D detectors, namely Geometry-aware Prototype Alignment (GPA-3D). Concretely, we first explore the potential relationships between the geometric structures of point cloud objects. During training, we randomly extract the BEV features of point clouds from both the source and target domains, and subsequently divide them into distinct groups based on their geometric structures. In this process, BEV features derived from point clouds with the similar geometric structures will be classified into the same group. Each group is then assigned a unique prototype, which enforces high representational similarity among the BEV features within that group, as illustrated in Fig. 2 (e). To this end, the soft contrast loss is devised to pull the intra-group feature-prototype pairs closer in the representational space and push the inter-group pairs farther away. Additionally, we develop the framework with two components, namely noise sample suppression (NSS) and instance replacement augmentation (IRA). NSS utilizes the similarities between foreground areas and

the background prototype, to produce a mask for decreasing the impact of noise. IRA displaces pseudo-labels with high-quality samples that have similar geometric structures, enriching the diversity on the target domain.

The main contributions of this paper include:

- We propose a novel UDA framework for LiDAR-based 3D detectors, namely Geometry-aware Prototype Alignment (GPA-3D). It explicitly integrates geometric associations into feature alignment, effectively decreasing the distributional discrepancy and facilitating the adaptation of existing point cloud detectors.
- Noise sample suppression and instance replacement augmentation are designed to enhance pseudo-labels in terms of reliability and versatility, respectively.
- We conduct comprehensive experiments on Waymo, nuScenes, and KITTI. The encouraging results demonstrate the GPA-3D outperforms state-of-the-art methods in various adaptation scenarios. More importantly, thanks to the architecture-agnostic design, GPA-3D is flexible to be applied to point cloud detectors.

2. Related Work

LiDAR-based 3D Detection. Mainstream point cloud detectors can be broadly divided into two categories: point-based and grid-based. Point-based methods mainly adopt the architectures of PointNet [22] and PointNet++ [23] to extract features from raw point clouds. PointRCNN [28] designs an encoder-decoder backbone to learn the point-wise representation. 3DSSD [38] improves the point sampling operator from the aspect of feature distance. IA-SSD [41] utilizes the instance-aware downsampling to preserve more foregrounds. On the other hand, grid-based methods first divide point clouds into fixed-size voxels, which are then processed via 2D/3D CNN. SECOND [34] adopts sparse 3D convolution for efficient feature learning. PointPillars [16] proposes a pillar encoding method and achieves a good trade-off between speed and performance. PV-RCNN [27] incorporates the voxel backbone with the keypoint branch to learn the representative scene features. Other approaches [35, 4, 30] project the point clouds into certain kinds of 2D views, and employ 2D CNN to extract the features. In this work, we conduct focused discussions with SECOND [34] and PointPillars [16] as base detectors. To demonstrate the generalization ability of our method, we also provide the comparisons with PV-RCNN [27] detector.

Domain Adaptive Object Detection. A large amount of literature has been presented in UDA for 2D-image detection, which can be roughly classified into two groups: distribution alignment and self-training. Alignment-based

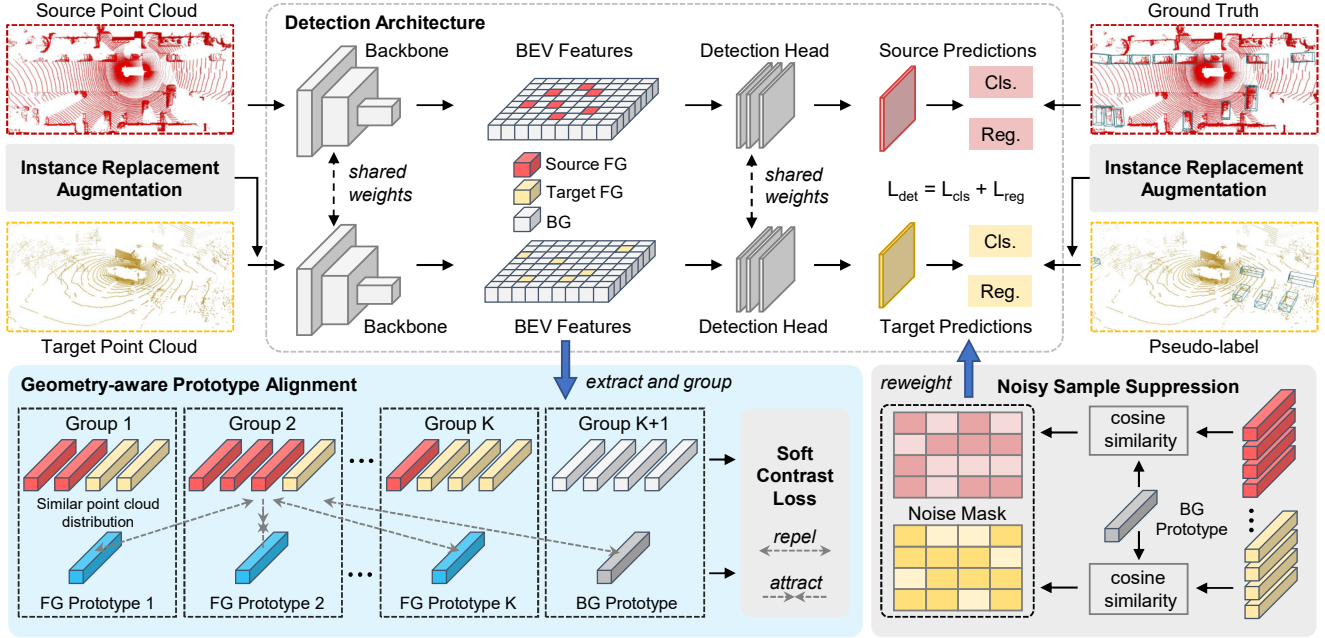


Figure 3. Overview of our proposed GPA-3D framework. It adopts a basic co-training manner to adapt the 3D detector with the point clouds from source and target domains. The BEV features are processed via geometry-aware prototype alignment, which reduces the distributional discrepancy and enables the learning of general representation across domains. To this end, the soft contrast loss is devised for jointly optimizing the prototypes and network parameters. Besides, the noisy sample suppression is proposed to alleviate the impact of noisy samples during training, and instance replacement augmentation is designed to enhance the diversity on the target domain.

methods [3, 25] leverage the adversarial training [5] to learn aligned features across domains. Self-training approaches [13, 14] utilize a multi-phase strategy to generate pseudo-labels on unlabeled data. Besides, some works [10, 17, 24, 11] adopt the CycleGAN [42] to generate training samples with styles of source and target domains.

Similarly, several recent works also aim to address the domain bias for 3D point cloud detectors. Wang *et al.* [33] investigate the domain bias of popular autonomous driving 3D datasets, and propose to alleviate the gaps via three techniques, *i.e.*, output transformation, statistical normalization and few shot. SF-UDA^{3D} [26] adopts a mature 3D tracker to find the best scaling parameter, which is further used to re-scale the target point clouds for producing high-quality pseudo-labels. MLC-Net [20] designs a mean-teacher paradigm to provide pseudo-labels for facilitating smooth learning of the student model. ST3D [36] and ST3D++ [37] build a self-training pipeline to produce pseudo-labels for fine-tuning model and update pseudo-labels via memory bank. 3D-CoCo [39] devises the domain-specific encoders with a hard sample mining strategy to learn transferable representations. Compared with previous works, our method explicitly embraces the geometric relationship to reduce the distributional discrepancy during adaptation.

3. The Proposed Method

In the following, we present GPA-3D to mitigate the domain gap for LiDAR-based detectors. Fig. 3 illustrates the

whole pipeline. Sec. 3.1 formulates the UDA task for point cloud detectors. Sec. 3.2 introduces the detection architecture in our method. In Sec. 3.3, we explain the details of the geometry-aware prototype alignment, followed by the soft contrast loss, which is discussed in Sec. 3.4. Finally, we present the noise sample suppression and instance replacement augmentation in Sec. 3.5 and Sec. 3.6, respectively.

3.1. Problem Statement

In this work, we focus on the problem of unsupervised domain adaptation on 3D detection. Concretely, given the labeled source domain point clouds $\mathbb{D}^s = \{(P_i^s, L_i^s)\}_{i=1}^{N^s}$, as well as unlabeled target domain point clouds $\mathbb{D}^t = \{(P_i^t)\}_{i=1}^{N^t}$, our goal is to train a 3D detector based on \mathbb{D}^s and \mathbb{D}^t and maximize its performance on \mathbb{D}^t . Here, N is the total number of scenes, and P_i indicates the i -th point cloud scene, where each point has the 3-dim spatial coordinates and an extra intensity. The corresponding label L_i represents a series of 3D bounding boxes, each of them can be parameterized by the center location (c_x, c_y, c_z) , spatial dimension (h, w, l) and rotation r . Note that the superscripts s and t stand for source and target domain respectively.

3.2. Detection Architecture

The input point cloud P_i is first sent to a backbone network with 3D sparse convolutions or 2D convolutions to extract the point cloud representation as following:

$$F_i = h_1(P_i; \theta_1), \quad (1)$$

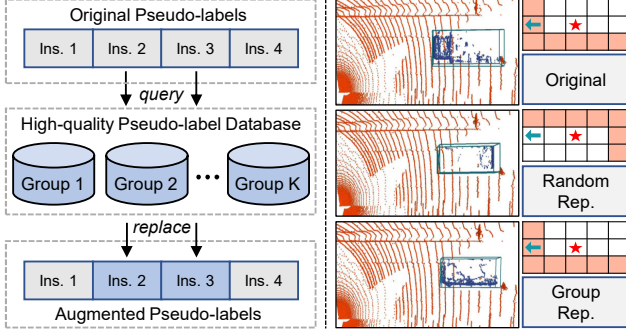


Figure 5. Illustration of instance replacement augmentation (IRA). Left: IRA leverages a group mechanism to displace the original instances with high-quality candidates. Right: Compared with random replacing, our group mechanism does not interfere the spatial context of the point cloud scene.

i.e., $A_{i,j} = Q_{i,j} \pm 1$. The overall soft contrast loss \mathcal{L}_{contra} can be formulated as:

$$\mathcal{L}_{contra} = \mathcal{L}_{att}^+ + \mathcal{L}_{att}^- + \beta_1 \mathcal{L}_{rep}^{+adj} + \beta_2 \mathcal{L}_{rep}^{+other} + \beta_3 \mathcal{L}_{rep}^-, \quad (9)$$

where β_1 , β_2 and β_3 are the balance coefficients.

3.5. Noise Sample Suppression

The pseudo-labels used on the target domain are noisy and can lead to the accumulation of errors. To mitigate the impact of noise, we propose the noise sample suppression (NSS) approach, which generates a noise mask to suppress the magnitude of the gradient descent for the foreground areas that might be underlying noise. The noise mask can be represented as $S \in \{\alpha, 1.0\}^{H \times W}$, where α ($\alpha < 1.0$) is the suppression factor to decrease the contribution of low-quality samples. In S , the foreground areas that have high similarities with the background prototype, *i.e.*, $\text{sim}(\mathbf{F}_{i,j}^+, \mathbf{g}_{K+1}) > 0.3$, are assigned to α , while rest foreground and background areas are assigned to 1.0.

During training, the noise mask S is multiplied to the co-training loss $\mathcal{L}_{co-train}$, elaborated in Sec. 3.7. With the progress of training, prototypes will be optimized with better representative capability, which enables NSS to suppress the noise more reliably and facilitate the training procedure.

3.6. Instance Replacement Augmentation

Those uncertain pseudo-labels (with scores of 0.2 ~ 0.5) are usually ignored in training. Despite inaccurate, they might provide partial localization information. To this end, we devise the instance replacement augmentation (IRA) module. As shown in Fig. 5 (left), we first pick the pseudo-labels with scores over 0.5 to construct a high-quality database, which utilizes the group mechanism as Eq. 4 to divide the picked instances into groups belonging to different geometric structures. During training, we calculate the group indexes for the uncertain pseudo-labels, and replace them with instances having same group indexes from the

Algorithm 1: The learning procedure of GPA-3D

input : labeled source domain \mathbb{D}^s , unlabeled target domain \mathbb{D}^t , 3D detector θ , total epochs T , steps per epoch N , and list of update epochs U

output: adapted 3D detector θ^t

- 1 Pre-train the network $\theta^s \leftarrow \mathbb{D}^s$ according to Eq. (10);
- 2 Initialize $\theta \leftarrow \theta^s$;
- 3 Generate pseudo-labels and database $(\hat{L}^t, \hat{D}^t) \leftarrow (\theta, \mathbb{D}^t)$;
- 4 **for** $epoch \leftarrow 1$ to T **do**
- 5 **for** $step \leftarrow 1$ to N **do**
- 6 Sample mini-batches $(\beta^s, \beta^t) \leftarrow (\mathbb{D}^s, \mathbb{D}^t)$;
- 7 Instance replacement $\beta_{aug}^t \leftarrow \text{IRA}(\beta^t, \hat{D}^t)$;
- 8 update $\theta \leftarrow (\beta^s, \beta_{aug}^t)$ according to Eq. (12);
- 9 **if** $epoch \in U$ **then**
- 10 Update pseudo-labels $\hat{L}^t \leftarrow (\theta, \mathbb{D}^t)$;
- 11 $\theta^t \leftarrow \theta$;

database. In this procedure, a parameter p_{IRA} is adopted to regulate the probability of the replacement operation.

There are two main merits of IRA. First, the quantity of target data is maintained and the diversity is also enhanced. Second, benefiting from the group mechanism, the spatial contexts around the replaced instances are unchanged and no ambiguous or unreasonable case is introduced, as devised in Fig. 5 (right).

3.7. Overall Training Procedure

The overall training procedure of GPA-3D is illustrated in Alg. 1. Following previous works [36, 37], the 3D detector is first trained on the labeled source domain \mathbb{D}^s via minimizing the detection loss \mathcal{L}_{det}^s as:

$$\mathcal{L}_{det}^s = \mathcal{L}_{reg}^s + \mathcal{L}_{cls}^s, \quad (10)$$

where the \mathcal{L}_{reg}^s and \mathcal{L}_{cls}^s indicate the regression and classification errors respectively. Next, we use the pre-trained detector to generate pseudo-labels \hat{L}_i^t and the database of IRA on the unlabeled target domain \mathbb{D}^t . Finally, the co-training paradigm is employed to further fine-tune the model as:

$$\mathcal{L}_{co-train} = \mathcal{L}_{det}^s + \mathcal{L}_{det}^t, \quad (11)$$

where \mathcal{L}_{det}^t is the detection loss on target data, same as in Eq. 10. The overall adaptation loss \mathcal{L}_{adapt} is calculated via:

$$\mathcal{L}_{adapt} = \beta \cdot \mathcal{L}_{contra} + S \cdot \mathcal{L}_{co-train}, \quad (12)$$

where β is the total weight of the soft contrast loss, and S is the noise mask of NSS. For more details of the training procedure, please refer to the supplements.

4. Experiments

4.1. Experimental Setup

Datasets. We evaluate the GPA-3D on widely used autonomous driving benchmarks including Waymo [29], nuScenes [2], and KITTI [6]. These datasets exhibit significant diversities in foreground patterns and LiDAR beams,

Table 1. Comparison with the state-of-the-art methods on the Waymo \rightarrow KITTI adaptation scenario, with BEV and 3D average precisions of 40 recall positions. In addition, we also report the Closed Gap from ST3D[36], which is defined as $\frac{AP_{\text{model}} - AP_{\text{source}}}{AP_{\text{oracle}} - AP_{\text{source}}} \times 100\%$. For fair comparison, the results with the detector of SECOND-IoU are obtained from the original paper of ST3D++ [37], while the performances with PointPillars are cited from 3D-CoCo [39]. The best result is indicated by **bold**.

Methods	SECOND-IoU				PointPillars			
	AP _{BEV}	Closed Gap	AP _{3D}	Closed Gap	AP _{BEV}	Closed Gap	AP _{3D}	Closed Gap
Source Only	67.64	-	27.48	-	47.8	-	11.5	-
SN [33]	78.96	+72.33%	59.20	+69.00%	27.4	-55.14%	6.4	-8.49%
UMT [9]	77.79	+64.86%	64.56	+80.66%	-	-	-	-
3D-CoCo [39]	-	-	-	-	76.1	+76.49%	42.9	+52.25%
ST3D [36]	82.19	+92.97%	61.83	+74.72%	58.1	+27.84%	23.2	+19.47%
ST3D++ [37]	80.78	+83.96%	65.64	+83.01%	-	-	-	-
GPA-3D (ours)	83.79	+103.19%	70.88	+94.41%	77.29	+79.70%	50.84	+65.46%
<i>Improvement</i>	+1.6	+10.22%	+5.24	+11.4%	+1.19	+3.21%	+7.94	+13.21%
Oracle	83.29	-	73.45	-	84.8	-	71.6	-

Table 2. Adaptation performance on the Waymo \rightarrow nuScenes in comparison with different base detectors and state-of-the-art approaches.

Methods	SECOND-IoU				PointPillars			
	AP _{BEV}	Closed Gap	AP _{3D}	Closed Gap	AP _{BEV}	Closed Gap	AP _{3D}	Closed Gap
Source Only	32.91	-	17.24	-	27.8	-	12.1	-
SN [33]	33.23	+1.69%	18.57	+7.54%	28.31	+2.41%	12.98	+4.58%
UMT [9]	35.10	+11.54%	21.05	+21.61%	-	-	-	-
3D-CoCo [39]	-	-	-	-	33.1	+25.00%	20.7	+44.79%
ST3D [36]	35.92	+15.87%	20.19	+16.73%	30.6	+13.21%	15.6	+18.23%
ST3D++ [37]	35.73	+14.87%	20.90	+20.76%	-	-	-	-
GPA-3D (ours)	37.25	+22.88%	22.54	+30.06%	35.47	+36.18%	21.01	+46.41%
<i>Improvement</i>	+1.33	+7.01%	+1.49	+8.45%	+2.37	+11.18%	+0.31	+1.62%
Oracle	51.88	-	34.87	-	49.0	-	31.3	-

which can lead to severe domain bias when transferring 3D detectors from one dataset to another. Detailed information about datasets is available in the supplementary material.

Implementation Details. We verify the GPA-3D with two popular LiDAR-based detectors, namely SECOND-IoU [36] and PointPillars [16]. All the parameter settings for network architecture are set the same with OpenPCDet [32] and ST3D [36]. We perform all experiments using 8 NVIDIA V100 GPU cards. For the pre-training step, the model is trained for 30 epochs using the ADAM [15] optimizer and the total batch size of 32 on the source domain. Next, we utilize the pre-trained model to generate pseudo-labels on the target domain with a score threshold of 0.2. Note that instances with scores over 0.5 are retained and subsequently utilized to establish the high-quality pseudo-label database for IRA. Finally, we further fine-tune the model with our proposed approach for 30 epochs. To avoid local minima, we employ the cosine annealing strategy to adjust the learning rate, which was set to 0.003 for pre-training and 0.0015 for fine-tuning. Please refer to the sup-

plements for more implementation details.

Compared Methods. As shown in Tab. 2, GPA-3D is first compared with the Source Only method, which trains the model on the source domain and evaluates it on the target domain without any adaptation. Next, 5 existing works are included in the comparison, namely, SN [33], UMT [9], 3D-CoCo [39], ST3D [36], and ST3D++ [37]. SN utilizes the statistics from target annotations to normalize the foreground objects on the source domain. UMT employs a mean-teacher framework to filter inaccurate pseudo-labels. 3D-CoCo learns the instance-level transferable features for better generalization. ST3D and ST3D++ adopt a memory bank to produce high-quality pseudo-labels. Additionally, we also compare GPA-3D with the Oracle method, which trains the model on the labeled target data, serving as an upper bound for performance.

4.2. Comparison with State-of-the-art Methods

Waymo \rightarrow KITTI Adaptation. To validate the effectiveness to the domain shift about object size, we con-




Figure 6. Qualitative results of GPA-3D on Waymo \rightarrow KITTI. For each box, we use the X to specify the orientation. The predicted results and ground truths are painted in blue and green, respectively.

Table 3. Component ablation studies in GPA-3D. **Proto** indicates the geometry-aware prototype alignment. **Soft** is the soft contrast loss. **NSS** means the noise sample filtering. **IRA** represents the instance replacement augmentation.

Setting	Proto	Soft	NSS	IRA	AP _{BEV}	AP _{3D}
(a)					77.87	60.36
(b)	✓				80.49	66.28
(c)	✓	✓			80.51	67.34
(d)	✓	✓	✓		83.07	69.45
(e)	✓	✓		✓	81.94	67.79
(f)	✓	✓	✓	✓	83.79	70.88

duct a comprehensive comparison on Waymo \rightarrow KITTI. As demonstrated in Tab. 1, with the 3D detector SECOND-IoU, our proposed GPA-3D outperforms ST3D++ [37] with a large margin, and significant performance gains are obtained compared with previous best results, *i.e.*, 5.24% of AP_{3D} and 1.6% of AP_{BEV}. Note that the AP_{BEV} of GPA-3D is also higher than Oracle method, indicating the effectiveness of incorporating the geometric structure information into UDA on 3D detection task. Even switching the base detector to PointPillars, our method still exceeds previous SOTA 3D-CoCo [39] by 7.94% and 1.19% in terms of AP_{3D} and AP_{BEV}, respectively.

Waymo \rightarrow nuScenes Adaptation. For the domain gap of LiDAR beams, we select Waymo \rightarrow nuScenes as representatives due to their different LiDAR sensors, *i.e.*, 64-beam vs 32-beam. As shown in Tab. 2, GPA-3D improves the adaptation performances to 37.25% AP_{BEV} and 22.54% AP_{3D} with the SECOND-IoU detector, surpassing previous SOTA methods. Compared with ST3D++ [37], 1.52% and 1.64% gains separately in terms of AP_{BEV} and AP_{3D} are achieved. Based on PointPillars, our approach exceeds the best method 3D-CoCo [39] by 2.37% in AP_{BEV}, and outperforms ST3D [36] with 4.87% and 5.41% respectively in terms of AP_{BEV} and AP_{3D}. These improvements demonstrate the advancement of our GPA-3D to mitigate the more challenging domain shift of cross-beam scenarios.

4.3. Ablation Studies

All ablation studies are conducted on Waymo \rightarrow KITTI with SECOND-IoU as the base detector.




Figure 7. Ablations on the geometry-aware prototype alignment. **Baseline** is the co-training method without any feature alignment. **V-Proto** refers to the vanilla alignment with a pair of fore/background prototypes. **G-Proto[n]** indicates that n prototypes are employed in GPA-3D.

Component Analysis in GPA-3D. We assess the effectiveness of each component in GPA-3D, as presented in Tab. 3. Baseline (a) represents self-training via pseudo-labels on the target domain. The application of geometry-aware prototype alignment provides 5.92% and 2.62% gains separately in terms of AP_{3D} and AP_{BEV}, and the soft contrast loss brings an improvement of 1.06% on AP_{3D}. The improvements demonstrate that incorporating the geometric relationship into domain adaptation is feasible and effective. In addition, NSS and IRA boost the performance by around 2.5% and 1.5% respectively, which indicates the efficacy of enhancing the quality of supervision on target data.

Effectiveness of Geometry-aware Prototype Alignment.

We further investigate the effects of the geometry-aware prototype alignment. As illustrated in Fig. 7, the vanilla alignment with one pair of fore/background prototypes performs better than the co-training baseline, implying that the misalignment of features distribution affects the performance. Applying two prototypes yields 3.57% and 5.05% gains of AP_{BEV} and AP_{3D} respectively, compared to the co-training baseline. The performance reaches to the peak of 84.44% AP_{BEV} when 4 foreground prototypes are employed, indicating the advancement of combining geometric information with feature alignment. However, we observe minor performance degradation when too many prototypes are used, which we attribute to redundant prototypes leading

Table 4. Ablations on noise sample suppression. The symbols -T/-S denote that NSS is applied solely on the target/source domain, while -TS performs NSS on both target and source domains. -TSH additionally adopts a hard truncated factor, *i.e.*, $\alpha = 0$.

Methods	Filter Domain	α	AP _{BEV}	AP _{3D}
GPA-3D (w/o NSS)	-	-	81.94	67.79
GPA-3D (w/ NSS-T)	Target	0.5	83.37	68.24
GPA-3D (w/ NSS-S)	Source	0.5	82.33	67.93
GPA-3D (w/ NSS-TS)	Target + Source	0.5	83.45	69.77
GPA-3D (w/ NSS-TSH)	Target + Source	0.0	83.79	70.88

Table 5. Effects of the instance replacement augmentation. **RandRep** discards the group mechanism in IRA.

Method	w/o IRA	RandRep	w/ IRA
AP _{BEV} / AP _{3D}	83.07 / 69.45	82.99 / 69.59	83.79 / 70.88

Table 6. Comparison with different adaptation frameworks. **Source** refers to the Source Only method. **Self-T.** is the self-training framework. **Co-T.** symbolizes the co-training pipeline. **Mean T.** represents the mean teacher paradigm.

Framework	Source	Self-T.	Co-T.	Mean T.	GPA-3D
AP _{BEV} / AP _{3D}	67.64 / 27.48	77.87 / 60.36	80.06 / 61.67	80.01 / 64.62	83.79 / 70.88

to indistinguishable features in the representational space.

Effectiveness of Noise Sample Suppression. We conduct ablations on the noise sample filter (NSS) with various settings. As shown in Tab. 4, the detection performance drops to 67.79% AP_{3D}, when we remove the NSS from GPA-3D. Only applying the NSS on target domain achieves the gains of 1.43% and 0.45% on AP_{BEV} and AP_{3D}, respectively. We could see that using NSS on the source domain could also bring improvements. We think this is due to the fact that NSS suppresses those source samples with only a few points, which are very similar to the background noise. When the hard truncated α is adopted, AP_{3D} is further improved to 70.88%, indicating the effectiveness of NSS.

Effectiveness of Instance Replacement Augmentation. Also, we compare different policies in instance replacement augmentation (IRA). We can see from Tab. 5 that our proposed IRA attains 0.72% and 1.43% gains in terms of AP_{BEV} and AP_{3D}, respectively. Without the group mechanism in IRA, *i.e.*, randomly replacing pseudo-labels with instances the database, only marginal gains are obtained in AP_{3D}, and even degradation in AP_{BEV}. This highlights the significance of maintaining the consistency between instances and their contextual environments.

Domain Adaptation Frameworks. We compare our proposed GPA-3D with several adaptation frameworks, as presented in Tab. 6. The results confirm the effectiveness of GPA-3D, which leverages the geometric association to




Figure 8. Comparisons of self-training baseline and our GPA-3D.

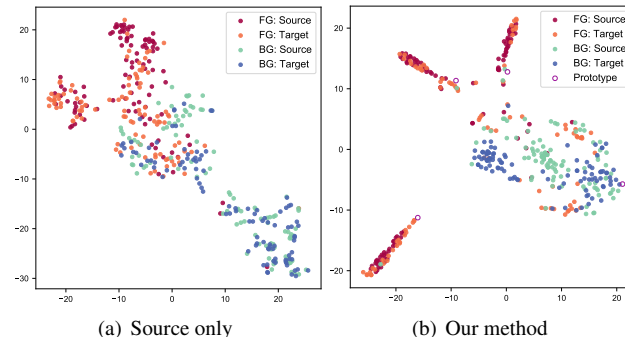


Figure 9. The t-SNE visualization of BEV features. The feature points are obtained by SECOND-IoU on Waymo \rightarrow nuScenes.

transfer 3D detectors across different domains. Fig. 8 further illustrates that, despite all models fluctuate at early epochs, our GPA-3D steadily and consistently enhances the detection performance in later training stages.

Visualization. We exhibit some qualitative results of cross-domain adaptation in Fig. 6. Additionally, in Fig. 9, we visualize the distribution of BEV features. It is obvious that GPA-3D aggregates foreground samples into different prototypes, and separates them from the backgrounds. Further visualizations can be found in the supplements.

5. Conclusion

This paper presents a novel framework for unsupervised domain adaptive 3D detection. Our proposed GPA-3D leverages the underlying geometric relationship to reduce the distributional discrepancy in the feature space, thus mitigating the domain shift problems. Comprehensive experiments demonstrate that our method is effective and can be easily incorporated into mainstream LiDAR-based 3D detectors. For future work, we plan to extend GPA-3D to support multi-modal 3D detectors. This requires a more efficient alignment mechanism to process feature streams from both point clouds and images.

Acknowledgement

This work was supported by the National Natural Science Foundation of China under No.62276061 and 62006041.

References

- [1] Eduardo Arnold, Omar Y Al-Jarrah, Mehrdad Dianati, Saber Fallah, David Oxtoby, and Alex Mouzakitis. A survey on 3D object detection methods for autonomous driving applications. *TITS*, 20(10):3782–3795, 2019. 1
- [2] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. In *CVPR*, pages 11621–11631, 2020. 5
- [3] Yuhua Chen, Wen Li, Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Domain adaptive faster r-cnn for object detection in the wild. In *CVPR*, pages 3339–3348, 2018. 1, 3
- [4] Lue Fan, Xuan Xiong, Feng Wang, Naiyan Wang, and Zhaoxiang Zhang. Rangedet: In defense of range view for lidar-based 3D object detection. In *ICCV*, pages 2918–2927, 2021. 2
- [5] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *JMLR*, 17(1):2096–2030, 2016. 3
- [6] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *CVPR*, pages 3354–3361, 2012. 5
- [7] Yulan Guo, Hanyun Wang, Qingyong Hu, Hao Liu, Li Liu, and Mohammed Bannamoun. Deep learning for 3D point clouds: A survey. *TPAMI*, 43(12):4338–4364, 2020. 1
- [8] Yulan Guo, Hanyun Wang, Qingyong Hu, Hao Liu, Li Liu, and Mohammed Bannamoun. Deep learning for 3D point clouds: A survey. *TPAMI*, 2020. 1
- [9] Deepti Hegde, Vishwanath Sindagi, Velat Kilic, A Brinton Cooper, Mark Foster, and Vishal Patel. Uncertainty-aware mean teacher for source-free unsupervised domain adaptive 3D object detection. *arXiv preprint arXiv:2109.14651*, 2021. 1, 6
- [10] Han-Kai Hsu, Chun-Han Yao, Yi-Hsuan Tsai, Wei-Chih Hung, Hung-Yu Tseng, Maneesh Singh, and Ming-Hsuan Yang. Progressive domain adaptation for object detection. In *WCCV*, pages 749–757, 2020. 1, 3
- [11] Sheng-Wei Huang, Che-Tsung Lin, Shu-Ping Chen, Yen-Yi Wu, Po-Hao Hsu, and Shang-Hong Lai. Auggan: Cross domain adaptation with gan-based data augmentation. In *ECCV*, pages 718–731, 2018. 1, 3
- [12] Zhengkai Jiang, Yuxi Li, Ceyuan Yang, Peng Gao, Yabiao Wang, Ying Tai, and Chengjie Wang. Prototypical contrast adaptation for domain adaptive semantic segmentation. In *ECCV*, pages 36–54. Springer, 2022. 2
- [13] Mehran Khodabandeh, Arash Vahdat, Mani Ranjbar, and William G Macready. A robust learning approach to domain adaptive object detection. In *ICCV*, pages 480–490, 2019. 1, 3
- [14] Seunghyeon Kim, Jaehoon Choi, Taekyung Kim, and Chang-ick Kim. Self-training and adversarial background regularization for unsupervised domain adaptive one-stage object detection. In *ICCV*, pages 6092–6101, 2019. 1, 3
- [15] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6
- [16] Alex H Lang, Sourabh Vora, Holger Caesar, Lubing Zhou, Jiong Yang, and Oscar Beijbom. Pointpillars: Fast encoders for object detection from point clouds. In *CVPR*, pages 12697–12705, 2019. 2, 6
- [17] Guofa Li, Zefeng Ji, and Xingda Qu. Stepwise domain adaptation (sda) for object detection in autonomous vehicles using an adaptive centernet. *TITS*, 2022. 1, 3
- [18] Ying Li, Lingfei Ma, Zilong Zhong, Fei Liu, Michael A Chapman, Dongpu Cao, and Jonathan Li. Deep learning for lidar point clouds in autonomous driving: a review. *TNLLS*, 2020. 1
- [19] Hongbin Lin, Yifan Zhang, Zhen Qiu, Shuaicheng Niu, Chuang Gan, Yanxia Liu, and Mingkui Tan. Prototype-guided continual adaptation for class-incremental unsupervised domain adaptation. In *ECCV*, pages 351–368. Springer, 2022. 2
- [20] Zhipeng Luo, Zhongang Cai, Changqing Zhou, Gongjie Zhang, Haiyu Zhao, Shuai Yi, Shijian Lu, Hongsheng Li, Shanghang Zhang, and Ziwei Liu. Unsupervised domain adaptive 3D detection with multi-level consistency. In *ICCV*, pages 8866–8875, 2021. 1, 3
- [21] Jiageng Mao, Shaoshuai Shi, Xiaogang Wang, and Hongsheng Li. 3D object detection for autonomous driving: a review and new outlooks. *arXiv preprint arXiv:2206.09474*, 2022. 1
- [22] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3D classification and segmentation. In *CVPR*, pages 652–660, 2017. 2
- [23] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *NIPS*, pages 5099–5108, 2017. 2
- [24] Adrian Lopez Rodriguez and Krystian Mikolajczyk. Domain adaptation for object detection via style consistency. *arXiv preprint arXiv:1911.10033*, 2019. 1, 3
- [25] Kuniaki Saito, Yoshitaka Ushiku, Tatsuya Harada, and Kate Saenko. Strong-weak distribution alignment for adaptive object detection. In *CVPR*, pages 6956–6965, 2019. 1, 3
- [26] Cristiano Saltori, Stéphane Lathuilière, Nicu Sebe, Elisa Ricci, and Fabio Galasso. Sf-uda 3D: Source-free unsupervised domain adaptation for lidar-based 3D object detection. In *3DV*, pages 771–780, 2020. 1, 3
- [27] Shaoshuai Shi, Chaoxu Guo, Li Jiang, Zhe Wang, Jianping Shi, Xiaogang Wang, and Hongsheng Li. Pv-rcnn: Point-voxel feature set abstraction for 3D object detection. In *CVPR*, pages 10529–10538, 2020. 2
- [28] Shaoshuai Shi, Xiaogang Wang, and Hongsheng Li. Point-rcnn: 3D object proposal generation and detection from point cloud. In *CVPR*, pages 770–779, 2019. 2
- [29] Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, et al. Scalability in perception for autonomous driving: Waymo open dataset. In *CVPR*, pages 2446–2454, 2020. 5

- [30] Pei Sun, Weiyue Wang, Yuning Chai, Gamaleldin Elsayed, Alex Bewley, Xiao Zhang, Cristian Sminchisescu, and Dragomir Anguelov. Rsn: Range sparse net for efficient, accurate lidar 3D object detection. In *CVPR*, pages 5725–5734, 2021. [2](#)
- [31] Korawat Tanwisuth, Xinjie Fan, Huangjie Zheng, Shujian Zhang, Hao Zhang, Bo Chen, and Mingyuan Zhou. A prototype-oriented framework for unsupervised domain adaptation. *Advances in Neural Information Processing Systems*, 34:17194–17208, 2021. [2](#)
- [32] OpenPCDet Development Team. Openpcdet: An open-source toolbox for 3D object detection from point clouds. <https://github.com/open-mmlab/OpenPCDet>, 2020. [6](#)
- [33] Yan Wang, Xiangyu Chen, Yurong You, Li Erran Li, Bharath Hariharan, Mark Campbell, Kilian Q Weinberger, and Wei-Lun Chao. Train in germany, test in the usa: Making 3D object detectors generalize. In *CVPR*, pages 11713–11723, 2020. [1](#), [3](#), [6](#)
- [34] Yan Yan, Yuxing Mao, and Bo Li. Second: Sparsely embedded convolutional detection. *Sensors*, 18(10):3337, 2018. [1](#), [2](#)
- [35] Bin Yang, Wenjie Luo, and Raquel Urtasun. Pixor: Real-time 3D object detection from point clouds. In *CVPR*, pages 7652–7660, 2018. [2](#)
- [36] Jihan Yang, Shaoshuai Shi, Zhe Wang, Hongsheng Li, and Xiaojuan Qi. St3d: Self-training for unsupervised domain adaptation on 3D object detection. In *CVPR*, pages 10368–10378, 2021. [1](#), [3](#), [5](#), [6](#), [7](#)
- [37] Jihan Yang, Shaoshuai Shi, Zhe Wang, Hongsheng Li, and Xiaojuan Qi. St3d++: Denoised self-training for unsupervised domain adaptation on 3D object detection. *TPAMI*, 2022. [1](#), [3](#), [5](#), [6](#), [7](#)
- [38] Zetong Yang, Yanan Sun, Shu Liu, and Jiaya Jia. 3dssd: Point-based 3D single stage object detector. In *CVPR*, pages 11040–11048, 2020. [2](#)
- [39] Zeng Yihan, Chunwei Wang, Yunbo Wang, Hang Xu, Chaoqiang Ye, Zhen Yang, and Chao Ma. Learning transferable features for point cloud detection via 3D contrastive co-training. *NIPS*, 34:21493–21504, 2021. [1](#), [3](#), [6](#), [7](#)
- [40] Jinze Yu, Jiaming Liu, Xiaobao Wei, Haoyi Zhou, Yohei Nakata, Denis Gudovskiy, Tomoyuki Okuno, Jianxin Li, Kurt Keutzer, and Shanghang Zhang. Cross-domain object detection with mean-teacher transformer. In *ECCV*, 2022. [2](#)
- [41] Yifan Zhang, Qingyong Hu, Guoquan Xu, Yanxin Ma, Jianwei Wan, and Yulan Guo. Not all points are equal: Learning highly efficient point-based detectors for 3D lidar point clouds. In *CVPR*, pages 18953–18962, 2022. [2](#)
- [42] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, pages 2223–2232, 2017. [1](#), [3](#)