

# Density-guided Translator Boosts Synthetic-to-Real Unsupervised Domain Adaptive Segmentation of 3D Point Clouds

Zhimin Yuan<sup>1</sup> Wankang Zeng<sup>1</sup> Yanfei Su<sup>1</sup> Weiquan Liu<sup>1</sup> Ming Cheng<sup>1\*</sup> Yulan Guo<sup>2</sup> Cheng Wang<sup>1</sup>  
<sup>1</sup> Fujian Key Laboratory of Sensing and Computing for Smart Cities, Xiamen University  
<sup>2</sup> National University of Defense Technology

## Abstract

3D synthetic-to-real unsupervised domain adaptive segmentation is crucial to annotating new domains. Self-training is a competitive approach for this task, but its performance is limited by different sensor sampling patterns (i.e., variations in point density) and incomplete training strategies. In this work, we propose a density-guided translator (DGT), which translates point density between domains, and integrates it into a two-stage self-training pipeline named DGT-ST. First, in contrast to existing works that simultaneously conduct data generation and feature/output alignment within unstable adversarial training, we employ the non-learnable DGT to bridge the domain gap at the input level. Second, to provide a well-initialized model for self-training, we propose a category-level adversarial network in stage one that utilizes the prototype to prevent negative transfer. Finally, by leveraging the designs above, a domain-mixed self-training method with source-aware consistency loss is proposed in stage two to narrow the domain gap further. Experiments on two synthetic-to-real segmentation tasks (SynLiDAR  $\rightarrow$  semanticKITTI and SynLiDAR  $\rightarrow$  semanticPOSS) demonstrate that DGT-ST outperforms state-of-the-art methods, achieving 9.4% and 4.3% mIoU improvements, respectively. Code is available at <https://github.com/yuan-zm/DGT-ST>.

## 1. Introduction

3D point cloud segmentation is crucial owing to its diverse applications, e.g., autonomous driving and robotics. Although supervised methods [1, 5, 11, 37, 47] have made substantial strides, they need costly human-annotated data. In contrast, we can obtain massive synthetic labeled data from simulation platforms. However, the domain gap between the synthetic (source) and real-world (target) data makes training directly using synthetic data infeasible. One alternative approach is unsupervised domain adaptation (UDA),

\*Corresponding author.

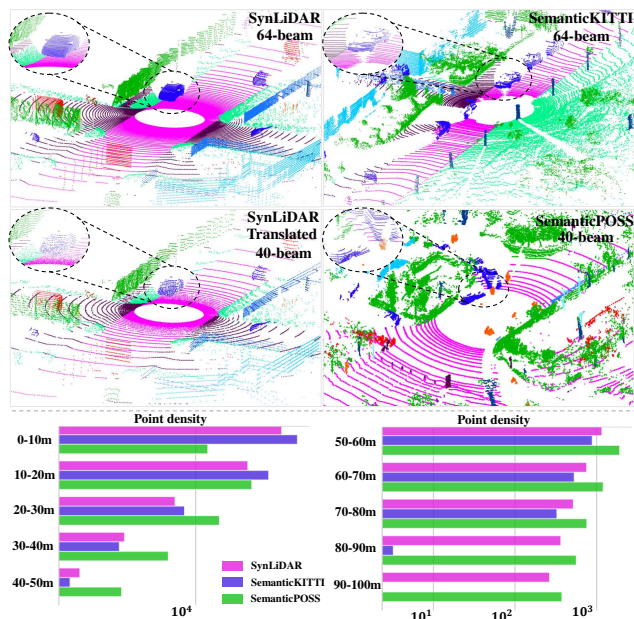


Figure 1. Distinct sampling patterns between synthetic and real-world scans. The synthetic scan (upper left) is integral and clean, whereas the real-world data (upper and middle right) contains unexpected and irregular noise. DGT enhances the realism of synthetic scan (middle left). Point densities of three datasets at various distances from the LiDAR center are shown at the bottom.

which transfers the learned source domain knowledge to make models perform better on the unlabeled target domain.

UDA seeks to acquire invariant knowledge across domains. As shown in Fig. 1, the sampling pattern mismatch between different sensors is the primary cause of the 3D domain gap, directly resulting in distinct point densities (i.e., the number of beams and the point number per beam) between domains. Besides, the synthetic data is integral and clean, whereas real-world data contains varying degrees of noise. Contemporary 3D synthetic-to-real UDA segmentation methods can generally be categorized into two groups: (1) Adversarial training [6, 13, 38–40, 44], which adopts a discriminator to ensure features/predictions of the segmentor domain-invariant. The limitation of this line of approaches

lies in the tendency of aligning the distributions of the two domains as a whole (*i.e.*, global-level alignment), resulting in suboptimal performance; (2) Self-training [9, 22, 23, 33], which usually performs better than the former. It adopts the data mixing techniques to construct an intermediate domain and uses the high confidence pseudo-labels to gradually learn the target domain knowledge. However, it heavily relies on a well-initialized model to provide confident pseudo-labels.

After dissecting the existing works, we find two key problems not adequately addressed. (1) *Input data*: although the existing approaches [13, 34] notice the problem of different sampling patterns, they strive to generate target-like source data and fulfill feature alignment simultaneously in the unstable adversarial training, which are difficult to reconcile. (2) *Pretrained model*: the existing 3D UDA self-training methods [9, 22, 33] completely overlook the importance of the well-initialized pretrained model. They directly employ the model trained on the source domain as the pretrained model, and the unsatisfactory pseudo-labels significantly limit their performance. Drawing inspiration from the 2D UDA counterparts [41, 42], we conjecture that adversarial training warm-up is essential for 3D self-training.

In this work, we propose a LiDAR scan translation strategy and integrate it into a two-stage self-training pipeline named DGT-ST to address the two issues above respectively. For (1), to solve the distinct point density across domains, we propose a non-learnable *density-guided translator* (DGT). It is a statistical-based module that narrows the domain gap at the input level by generating the other domain-like scans for each domain. Specifically, we divide a scan into discrete areas and use point density of each area to determine the location and number of points to be discarded, matching the density of the corresponding area in the other domain.

For (2), to provide a well-initialized model for self-training, we propose a *prototype-guided category-level adversarial network* (PCAN) in DGT-ST stage one. We use prototypes to dynamically measure the aligning confidence of points, and propose a self-adaptive reweighting strategy to reduce the impact of adversarial loss on those well-aligned points. This strategy effectively prevents negative transfer. We leverage this well-initialized model to perform self-training in stage two. During the self-training process, we propose a *source-aware consistency LaserMix* (SAC-LM) to learn the source knowledge from the target data, which enforces the segmentor to give consistent predictions on the target scans and the corresponding scans translated by DGT. Moreover, we employ the teacher-student training strategy to provide robust pseudo-labels and extend LaserMix [10] into UDA segmentation to fully utilize the spatial prior of both domains and bridge the domain gap.

Our main contributions can be summarized as follows:

- We propose a statistical-based density-guided translator (DGT) that directly bridges the domain gap at the in-

put level. Based on DGT, we propose DGT-ST training pipeline for 3D synthetic-to-real UDA segmentation.

- We design PCAN and SAC-LM to constitute the two stages of DGT-ST, respectively. For the former, we use the prototype to perform category-level adversarial alignment to prevent negative transfer and provide a well-initialized model for subsequent self-training. For the latter, we propose a self-training method that extracts source knowledge from target data, further improving the domain-invariant feature extraction power of the segmentor.
- Extensive experiments on two synthetic-to-real tasks verify the effectiveness of DGT-ST, which outperforms the state-of-the-art UDA approaches by a large margin.

## 2. Related work

**Point clouds semantic segmentation.** 3D semantic segmentation aims to give each point a semantic label. PointNet [19] is the pioneering work in point-based [5, 11, 17, 20, 21, 36, 43] methods, employing multilayer perceptrons to extract point features. Then, numerous point-based methods have been proposed and achieved outstanding results. However, these methods generally demand extensive computational resources, which makes them difficult to deploy in practical applications. In contrast, some approaches [1, 4, 8, 30, 31] project the 3D point clouds into 2D grids and leverage the 2D network to perform segmentation tasks. These approaches are efficient as they eliminate the demand for sampling and neighbor search operations. However, the loss of 3D geometric/topological information limits their performance. Currently, the prevailing approach is the voxel-based method [3, 12, 24, 35, 47]. These methods convert points into voxels and use sparse convolutions to extract geometric relationships between voxels. Due to its efficiency and promising result, we select MinkUnet [3] in this work.

**Point clouds UDA semantic segmentation.** UDA segmentation aims to use the labeled source data and unlabeled target data to train a model to perform well on the target. Following the spirit of the 2D counterparts, the dominant 3D UDA segmentation methods can be roughly divided into two groups: adversarial training and self-training. ePointDA [44] employs CycleGAN [46] to render dropout noise for domain alignment explicitly. Complete & Label [38] attempts to complete the two different domain data to a canonical domain, which bridges the domain gap at the input level. LiDARNet [6] simultaneously extracts domain-shared and domain-private features while employing two discriminators that jointly adapt for semantic and boundary predictions. PCT [34] employs two generators and discriminators to translate the point cloud appearance and sparsity, respectively. ASM [13] designs a novel learnable masking module to mimic the pattern of irregular noise and mitigate the domain gap. However, the global-level adversarial alignment can easily cause negative transfer. Two category-level ad-

versarial alignment methods [39, 40] are proposed and show encouraging performance to solve this problem.

Self-training [9, 22, 23, 33] is another line for this task, which leverages pseudo-labels to learn target knowledge gradually. ConDA [9] proposes an image concatenation-based framework for interchanging signals from both domains. PolarMix [33] proposes cutting, editing, and blending of two domain scans to enrich the data distribution for alignment. CosMix [22] proposes a domain-mixing strategy that harnesses semantic and structural information to reduce the domain gap. They adopt a mean-teacher [25] paradigm to get robust pseudo-labels. However, these methods employ a model trained only on the labeled source domain as the pretrained model, and the unreliable pseudo-labels significantly limit their performance. LiDAR-UDA [23] proposes a two-stage method that exploits random discarding source domain beams to obtain a pretrained model, and utilizes the temporal consistency of consecutive frames to generate reliable pseudo-labels. However, it does not take into account the number of points per beam.

### 3. Methodology

In the following, we first provide the necessary preliminaries for 3D UDA segmentation (Sec. 3.1). Then, we introduce the scan translation strategy DGT in Sec. 3.2. After that, we elaborate on the two-stage training pipeline DGT-ST in Secs. 3.3 and 3.4, which is also shown in Fig. 2.

#### 3.1. Preliminaries

In 3D synthetic-to-real UDA segmentation, we have the source  $\mathcal{X}^S = \{x_i^s\}_{i=1}^{N^s}$  and target  $\mathcal{X}^T = \{x_i^t\}_{i=1}^{N^t}$  dataset of  $N^s$  and  $N^t$  scans, respectively.  $\mathcal{X}^S$  has point-wise semantic labels  $\mathcal{Y}^S = \{y_i^s\}_{i=1}^{N^s}$ , while we lack labels for  $\mathcal{X}^T$ . We aim to leverage these datasets to train a segmentation network  $G$  that can provide accurate results on  $\mathcal{X}^T$ . Due to the domain gap,  $G$  trained only on the source data cannot generalize well to the target. Thus,  $G$  needs to have both discriminability and transferability. Generally, to ensure discriminability,  $G$  is optimized on source by cross-entropy (CE) loss:

$$\mathcal{L}_{ce}^s = -\frac{1}{N} \sum_{i=1}^N \sum_{k=0}^K y_i^s \log P_{i,k}^s, \quad (1)$$

where  $N$  and  $K$  are the number of points and classes in the current training batch, respectively.  $y_i^s$  and  $P_{i,k}^s$  are the ground-truth and probability of the  $i$ -th point, respectively.

For transferability, the prevalent strategies are adversarial training and self-training. To ensure the features/predictions of  $G$  domain-invariant, the adversarial methods adopt a domain discriminator  $D$  plus an auxiliary adversarial loss. Here, we show the LS-GAN [15] and self-information

$S_i^t = -P_i^t \log P_i^t$  adopted by ADVENT [27]:

$$\mathcal{L}_{adv}^t = -\frac{1}{N^t} \sum_{i=1}^{N^t} \|D(S_i^t) - 0\|_2, \quad (2)$$

$$\mathcal{L}_D = -\frac{1}{N^s} \sum_{i=1}^{N^s} \|D(S_i^s) - 0\|_2 - \frac{1}{N^t} \sum_{i=1}^{N^t} \|D(S_i^t) - 1\|_2, \quad (3)$$

where  $0/1$  denotes the source/target domain label.

On the other hand, self-training uses the pseudo-label  $\hat{y}^t$  to optimize  $G$  and gain knowledge from the target domain:

$$\mathcal{L}_{ce}^t = -\frac{1}{N} \sum_{i=1}^N \sum_{k=0}^K \hat{y}_i^t \log P_{i,k}^t. \quad (4)$$

Typically, the mean-teacher framework [25] is adopted to provide accurate and robust  $\hat{y}^t$ . The student model  $G$ , parameterized by  $\theta^{stu}$ , is trained by gradient descent. And  $\theta^{tea}$ , the weights of the teacher model  $G^{tea}$ , is updated every  $t$  iterations with an exponential moving average (EMA) manner:

$$\theta_i^{tea} = \alpha \theta_{i-t}^{tea} + (1 - \alpha) \theta_i^{stu}, \quad (5)$$

where  $i$  denotes the current training iteration and  $\alpha$  is a smoothing coefficient that determines the update speed. Finally,  $\hat{y}^t$  is obtained by a confidence strategy:

$$\hat{y}_i^t = \begin{cases} \arg \max_k p_i^{t,tea}, & \max(p_i^{t,tea}) > Th_p, \\ 0, & \text{otherwise,} \end{cases} \quad (6)$$

where  $Th_p$  is the threshold to obtain reliable pseudo-labels.

#### 3.2. Density-guided translator

The sensor sampling pattern mismatch between the source and target domains is the primary cause of the domain gap. By taking a closer look at the synthetic and real-world scans, as shown in Fig. 1, we observe two significant disparities: (1) The density varies across domains, *i.e.*, the number of beams  $D_b$  and the point number per beam  $D_p$ ; (2) The synthetic scan is integral and clean, whereas the real-world data contains a varying amount of noise. To mitigate the domain discrepancy at the input level, *i.e.*, obtain  $D_b^s \approx D_b^t$  and  $D_p^s \approx D_p^t$ , we can accomplish this in two ways: complete surface [38] or discard points [13, 32, 34]. However, surface completion inevitably increases the computational overhead and may bring in points with inaccurate labels. Thus, we choose to discard points and propose a density-guided translator (DGT), a non-learnable and statistical-based translator.

Inspired by LiDAR-Distillation [29], we use K-means to label beams in each scan and discard beams to make the two domains have a similar number of beams. However, it does



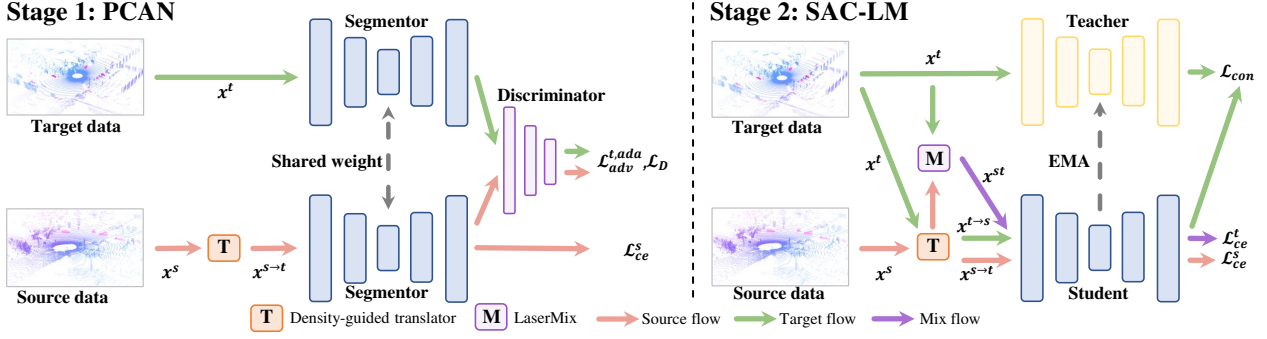


Figure 2. Overview of our two-stage DGT-ST. We propose DGT to bridge the domain gap at the input level and be integrated into both stages. In stage one, we propose PCAN with a segmentor  $G$  and a discriminator  $D$ . We take the target-like source  $x^{s \rightarrow t}$  and raw target data  $x^t$  as input to perform the category-level adversarial alignment. In stage two, SAC-LM, a teacher-student learning architecture is employed and loads the stage one trained model. We use LaserMix [10] to mix two domain scans  $x^{s \rightarrow t}$  and  $x^t$  and obtain the mixed scan  $x^{st}$ . Finally, the student model is trained by  $x^{s \rightarrow t}$  and  $x^{st}$ . Moreover, we enforce the student model to give consistent predictions on  $x^t$  and  $x^{t \rightarrow s}$ .

not work on two domains with the same number of beams while  $D_p$  is still different between domains.  $D_p$  exhibits substantial variation according to the distance to the LiDAR center, where  $D_p$  in the nearby area is much greater than the farther-away area. Two critical problems arise: *where* and *how many* points to discard? To tackle these issues, we propose a statistical-based random discarding strategy to balance the point number for both domains. As shown in Fig. 3, it mainly consists of the following steps:

(1) **Partition.** For both domains, we evenly partition all points within each scan into  $m$  non-overlapping areas  $A = \{a_1, a_2, \dots, a_m\}$  by their distance to the LiDAR center.

(2) **Calculation.** For each domain, we count the number of points in each area on the entire dataset to obtain  $NA^s = \{na_1^s, na_2^s, \dots, na_m^s\}$  and  $NA^t = \{na_1^t, na_2^t, \dots, na_m^t\}$ . The location and number of points that are chosen to be discarded are determined by calculating  $R = [r_1, r_2, \dots, r_m]$ , which can be formulated as:

$$\begin{aligned} R^{s \rightarrow t} &= [na_1^t/na_1^s, na_2^t/na_2^s, \dots, na_m^t/na_m^s], \\ R^{t \rightarrow s} &= [na_1^s/na_1^t, na_2^s/na_2^t, \dots, na_m^s/na_m^t], \end{aligned} \quad (7)$$

where  $s \rightarrow t$  denotes generating a target-like source scan, vice versa.  $R$  directly reveals the point number differences within each area for both domains.  $R$  will be further clipped to  $[0, 1]$ , *i.e.*,  $R = \text{np.clip}(R, a\_min=0, a\_max=1)$ .

(3) **Translation.** For simplicity, we give an example of translating a target-like source scan  $x^{s \rightarrow t}$ . We first count the number of points within each area and obtain  $A^{x^s} = \{a_1^{x^s}, a_2^{x^s}, \dots, a_m^{x^s}\}$ . Then, the discarding point number  $Del_{a_i}$  of each area is:

$$Del_{a_i} = \begin{cases} a_i^{x^s} * (1 - R_i^{s \rightarrow t}), & \text{if } R_i^{s \rightarrow t} \leq 1, \\ 0, & \text{otherwise.} \end{cases} \quad (8)$$

Finally, to obtain  $x^{s \rightarrow t}$ , we randomly select  $Del_{a_i}$  points within the area where  $R_i^{s \rightarrow t} \leq 1$  and random noise is addi-

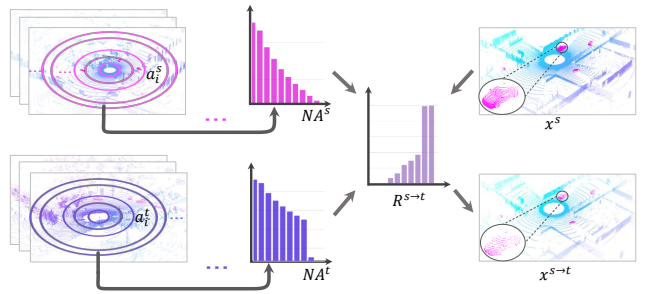


Figure 3. Illustration of the density-guided translator (DGT).

tionally added to the  $X$  and  $Y$  axes to enhance its realism. This strategy can also be employed to acquire  $x^{t \rightarrow s}$ .

**Discussion: how about random discarding?** Random discarding does not align with our aim for the following reasons: (1) Points located in farther-away areas already exhibit high sparsity, and discarding these points would impede model discriminability; (2) In certain regions,  $D_p^s/D_p^t$  may be smaller than  $D_p^t/D_p^s$ , and discarding these points may enlarge the domain gap. We verify this in Sec. 4.3.

### 3.3. Category-level adversarial with prototype

Traditional adversarial methods [26, 27] adopt a discriminator to bridge the domain gap. However, striving to make the source and target marginal distributions match globally is prone to negative transfer, *i.e.*, well-aligned points mapped to incorrect semantic categories. We resort to prototype and propose PCAN to tackle this issue, as shown in Fig. 2 (left).

Inspired by previous works [41, 42], points belonging to the same category tend to cluster together, and the prototype (class centroid) can represent each class in the feature space. Since the source data is fully labeled without noise, we only use source data to calculate the prototype. We define  $\lambda_k^s$  as

the prototype of class  $k$ , which can be calculated by:

$$\lambda_k^s = \frac{1}{|\mathcal{X}_k^S|} \sum_{x_i^s \in \mathcal{X}_k^S} G(x_i^s), \quad (9)$$

where  $\mathcal{X}_k^S$  denotes all points whose labels are  $k$  in  $\mathcal{X}^S$ .

Instead of globally matching the source and target marginal distributions, we propose to perform category-level alignment. Specifically,  $G$  should make high-confidence predictions for those well-aligned points, and we can easily obtain their pseudo-labels through Eq. (6). Thus, we treat those points with high confidence as well-aligned points  $x_{wa}^t$ , and evaluate how well  $x_{wa}^t$  are semantically matched based on the similarity distance  $\mathcal{M}$  (e.g., cosine similarity) between their feature and the corresponding prototype:

$$\mathcal{M}(x_{wa,k}^t, \lambda_k^s) = \begin{cases} 1 - \frac{\langle G(x_{wa,k}^t), \lambda_k^s \rangle}{\|G(x_{wa,k}^t)\| \cdot \|\lambda_k^s\|}, & \text{if } x_i^t \in x_{wa}^t, \\ 1, & \text{otherwise.} \end{cases} \quad (10)$$

$\mathcal{M}$  reveals the alignment degree between the well-aligned points and their corresponding source domain class in the feature space. The more similar  $x_{wa,i}^t$  and  $\lambda^s$  are, the smaller the weight is. For the remaining points (i.e.,  $x^t \notin x_{wa}^t$  in the current batch), we treat them as the unlabel class and calculate their adversarial loss in the traditional manner. Then, we leverage  $\mathcal{M}$  to reweight the adversarial loss, adaptively reduce their impact on  $x_{wa}^t$ , and prevent negative transfer.

Besides, we adopt a class-wise aggregation strategy to individually calculate each appearing class in  $x_{wa}^t$  to reduce interference with the other classes. With the help of the above design, we explicitly incorporate the class information into the adversarial loss. The new category-level adaptive reweight adversarial loss  $\mathcal{L}_{adv}^{t,ada}$  can be written as:

$$\mathcal{L}_{adv}^{t,ada} = - \sum_{k=0}^K \frac{1}{N_k} \mathcal{M} \|D(S_i^t) - 0\|_2. \quad (11)$$

In this adversarial training stage, we use the translated source scan  $x^{s \rightarrow t}$  and target scan  $x^t$  to perform the category-level adversarial alignment. The final loss is:

$$\mathcal{L}_{adv}^{total} = \mathcal{L}_{ce}^s(x^{s \rightarrow t}) + \gamma_1 \mathcal{L}_{adv}^{t,ada}(x^t), \quad (12)$$

where  $\gamma_1$  is a balance parameter. Following previous adversarial training works [14, 26, 27, 40], we fixed it as 0.001.

### 3.4. Source-aware consistency LaserMix

Although the source and target domains have distinct sampling patterns, they are collected in driving scenarios with similar scene layout. Inspired by LaserMix [10], initially proposed for semi-supervised learning (SSL), we conjecture that the spatial positions of objects/background in both domain scans correlate with their respective distributions.

Thus, as shown in Fig. 2 (right), we extend it into UDA and propose SAC-LM to mitigate the domain discrepancy.

**LaserMix (LM).** Given two scans  $x^s$  and  $x^t$  from two domains, we first partition all points from each scan based on their inclination angles and form  $n$  non-overlapping areas  $A\_M = \{a\_m1, a\_m2, \dots, a\_mn\}$ . Then, we mix  $A\_M^s$  and  $A\_M^t$  in an intertwining manner and result in two mixed scans  $x_1^{st}$  and  $x_2^{st}$ , which can be formulated as:

$$\begin{aligned} x_1^{st} &= a\_m1^s \cup a\_m2^t \cup \dots \cup a\_m_{n-1}^s \cup a\_m_n^t, \\ x_2^{st} &= a\_m1^t \cup a\_m2^s \cup \dots \cup a\_m_{n-1}^t \cup a\_m_n^s. \end{aligned} \quad (13)$$

Since the target domain is unlabeled, we adopt the teacher model and Eq. (6) to obtain  $\hat{y}^t$  for  $x^t$ . The semantic labels (i.e.,  $y^s$  and  $\hat{y}^t$ ) of the two scans are mixed similarly.

**Source-aware consistency (SAC) regularization.** In Fig. 1 bottom, in certain areas, the target domain point number is larger than the source. We can easily generate the source-like target scan by DGT. When  $G$  is domain-invariant,  $G$  should give similar predictions on raw target input  $x^t$  and its source-like input  $x^{t \rightarrow s}$ . We minimize the Kullback–Leibler divergence between  $P^{tea}(x^t)$  and  $P(x^{t \rightarrow s})$ :

$$\mathcal{L}_{sac} = -\frac{1}{N} \sum_{i=1}^N P(x^{t \rightarrow s}) \log \frac{P(x^{t \rightarrow s})}{P^{tea}(x^t)}. \quad (14)$$

$\mathcal{L}_{sac}$  enforces  $G$  to give consistent predictions for a target input under two views in the mean-teacher framework. It enables the segmentor to acquire source knowledge from the target data, thereby enhancing the capability of domain-invariant feature extraction and reducing the domain gap.

We use PCAN to provide a well-initialized model for generating high-quality pseudo-labels in this stage. We use the translated source scan  $x^{s \rightarrow t}$  and the mixed scans  $x_1^{st}$  to perform self-training. Since DGT is used in both stages, we name this overall pipeline as DGT-ST, and the final loss is:

$$\mathcal{L}_{st} = \mathcal{L}_{ce}^s(x^{s \rightarrow t}) + \mathcal{L}_{ce}^t(x_1^{st}) + \gamma_2 \mathcal{L}_{sac}, \quad (15)$$

where  $\gamma_2$  is a balance parameter, and we fixed it as 0.001.

## 4. Experiments

### 4.1. Setup

**Datasets.** We perform two synthetic-to-real UDA tasks.

*SynLiDAR* [34] is a recently published synthetic dataset generated by a LiDAR simulator identical to the Velodyne HDL-64E with a 100-meter working range and contains a variety of realistic virtual scenarios. Following the official recommendation, we use all subdataset, which contains 13 sequences of about 19840 scans.

*SemanticKITTI* [2] is the most prevalent real-world dataset for evaluating the large-scale 3D segmentation method, collected in Germany by a Velodyne HDL-64E

Methods	Mech.	car	bi.cle	mt.cle	truck	oth-v.	pers.	bi.clst	mt.clst	road	parki.	sidew.	other-g.	build.	fence	veget.	trunk	terr.	pole	traf.	mIoU	gain	
Source only	-	35.9	7.5	10.7	0.6	2.9	13.3	44.7	3.4	21.8	6.9	29.6	0.0	34.1	7.4	62.9	26.0	35.5	30.3	14.1	20.4	+0.0	
AdaptSegNet [26]	A	52.1	10.8	11.2	2.6	9.6	15.1	35.9	2.6	62.2	10.4	41.3	0.1	58.1	17.1	68.0	38.4	38.7	35.9	20.4	27.9	+7.5	
CLAN [14]	A	51.0	<u>15.8</u>	16.8	2.2	7.8	18.7	46.8	3.0	68.9	11.1	44.9	0.1	59.6	17.5	71.7	41.1	44.0	<u>37.7</u>	19.8	30.5	+10.1	
ADVENT [27]	A	59.9	13.8	14.6	<u>3.0</u>	8.0	17.7	45.8	3.0	67.6	<u>11.3</u>	45.6	0.1	<u>61.7</u>	15.8	<u>72.4</u>	<b>41.5</b>	<u>47.0</u>	34.5	15.3	30.5	+10.1	
FADA [28]	A	49.9	6.7	5.1	2.5	<u>10.0</u>	5.7	26.6	2.3	65.8	10.8	37.8	0.1	60.3	<b>21.5</b>	60.4	37.2	31.9	35.4	17.4	25.6	+5.2	
MRNet [45]	A	49.5	11.0	12.2	2.2	8.6	16.0	46.4	2.7	60.0	10.5	41.9	0.1	55.1	16.5	68.1	38.0	40.7	36.5	<u>20.8</u>	28.3	+7.9	
PMAN [40]	A	<u>71.0</u>	14.9	<u>24.8</u>	1.6	6.6	<u>23.6</u>	<u>61.1</u>	<b>5.5</b>	<u>75.3</u>	10.5	<b>54.1</b>	0.1	47.9	17.4	69.6	38.6	<b>61.5</b>	37.0	18.6	<u>33.7</u>	+13.3	
PCAN (Ours)	A	<b>85.0</b>	<b>17.5</b>	<b>27.4</b>	<b>10.4</b>	<b>11.9</b>	<b>27.5</b>	<b>63.7</b>	2.6	<b>78.1</b>	<b>13.5</b>	<u>50.1</u>	0.1	<b>68.5</b>	<u>20.0</u>	<b>76.2</b>	<u>41.3</u>	45.7	<b>41.0</b>	<b>21.8</b>	<b>37.0</b>	+16.6	
CoSMix [22]	S	56.4	<u>10.2</u>	20.8	2.1	<b>13.0</b>	25.6	41.3	2.2	67.4	8.2	43.4	0.0	<u>57.9</u>	<b>12.2</b>	68.4	44.8	35.0	<u>42.1</u>	<u>17.0</u>	29.9	+9.5	
PolarMix [33]	S	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	31.0	+10.6
LaserMix [10]	S	<u>90.3</u>	7.8	<u>37.2</u>	<u>2.3</u>	2.4	<u>40.6</u>	<u>49.1</u>	<b>5.1</b>	<u>80.5</u>	<u>9.9</u>	<u>57.4</u>	0.0	57.6	3.4	<u>77.6</u>	<u>46.6</u>	<b>60.1</b>	42.0	13.6	<u>36.0</u>	+15.6	
DGT-ST (Ours)	S	<b>92.9</b>	<b>17.3</b>	<b>43.4</b>	<b>15.0</b>	<u>6.1</u>	<b>49.2</b>	<b>54.2</b>	<u>4.2</u>	<b>86.4</b>	<b>19.1</b>	<b>62.3</b>	0.0	<b>78.2</b>	<u>9.2</u>	<b>83.3</b>	<b>56.0</b>	<u>59.1</u>	<b>51.2</b>	<b>32.3</b>	<b>43.1</b>	+22.7	

Table 1. Comparison results of SynLiDAR  $\rightarrow$  semanticKITTI adaptation in terms of mIoU. A/S denotes adversarial training/self-training.

LiDAR sensor. Following [22, 33, 34, 40], we choose sequences 00-10 for training (19130 scans) except sequence 08 (4071 scans) for validation.

*SemanticPOSS* [16] is a real-world dataset, collected in Peking University by a Pandora 40-line LiDAR sensor. Following [33, 34, 40], sequence 03 (500 scans) is used for validation and the rest (2488 scans) for training.

**Evaluation protocol.** The SynLiDAR provides the mapping details to pair with SynLiDAR  $\rightarrow$  SemanticKITTI (Syn  $\rightarrow$  Sk) and SynLiDAR  $\rightarrow$  SemanticPOSS (Syn  $\rightarrow$  Sp). We can fairly compare DGT-ST with other methods. Following [22, 33, 40], the segmentation performance is reported using the mean Intersection over Union (mIoU, as %) metric.

**Implementation.** All experiments are implemented in PyTorch [18] and MinkowskiEngine [3] on a single NVIDIA RTX 3090 GPU. For a fair comparison with other methods, we adopt MinkUNet34 as the segmentation network and set the voxel size as  $0.05m$ . For PCAN, the discriminator consists of 5 sparse convolution layers with kernel size 4 and stride 2, where the channel number is  $\{32, 64, 64, 128, 1\}$ . A Leaky-ReLU activation layer follows each convolution layer except the last. To obtain the same size as the input, we upsample the discriminator output by interpolation, *i.e.*, `ME.MinkowskiInterpolation()`. We adopt the Adam [7] optimizer with the initial learning rate of  $2.5e-4$  and  $1e-4$  respectively for the segmentation network and discriminator and is decayed by a poly learning rate policy with power of 0.9. The batch size is set to 2 and the max training iteration of all experiments is set as 100K. The input feature of all methods is XYZ coordinates. Following CoSMix [22], we set  $Th_p$ ,  $\alpha$  and  $t$  as 0.9, 0.99 and 100, respectively.

## 4.2. Comparisons with previous methods

We comprehensively compare our proposed method with the recent state-of-the-art approaches. These methods could be divided into two groups: (1) adversarial training methods, including AdaptSeg [26], CLAN [14], ADVENT [27], FADA [28], MRNet [45], and PMAN [40]; (2) self-training

methods, including CoSMix [22], PolarMix [33] and LaserMix (LM) [10]. Source only denotes the model trained on the source without adaptation. The mechanism ‘‘A’’ and ‘‘S’’ denotes adversarial training and self-training, respectively.

**SynLiDAR  $\rightarrow$  SemanticKITTI.** In Tab. 1, we show the comparison results. PCAN and DGT-ST significantly outperform the other methods, yielding an accuracy of 37.0% and 43.1% in mIoU. Compared with the non-adapted source only, PCAN and DGT-ST offer mIoU gains of 16.6% and 22.7%. PCAN outperforms the second-best adversarial method PMAN by 3.3%. Among these adversarial training methods, PCAN and PMAN are two specifically designed category-level adversarial networks for 3D UDA segmentation. Their results show that class-level adversarial alignment is effective for this task. For self-training methods, LM significantly outperforms other methods, which means that the spatial prior helps bridge the domain gap. Unlike CoSMix, which requires hyperparameters to select the number of categories to mix and then select and paste the reliable points to boost the performance of rare classes (*e.g.*, oth-v.), LM is more convenient. Among all the 19 classes, DGT-ST obtains the best results in 14 classes with significant improvements.

**SynLiDAR  $\rightarrow$  SemanticPOSS.** We present the results in Tab. 2. PCAN achieves competitive performance in adversarial-based methods, yielding an accuracy of 44.4% in mIoU and outperforming the source-only model by 6.1%. Compared with PMAN, which utilizes a multi-task network and is trained with specially designed losses, PCAN is much simpler. DGT-ST achieves an mIoU score 50.8 % and outperforms the source only model by 12.5%. DGT-ST outperforms all compared methods by a large margin and achieves the top performance in 6 out of the 13 categories.

## 4.3. Ablation Studies

We report DGT-ST with certain components ablated here. Since the SynLiDAR and semanticKITTI are collected by 64-beam LiDAR, discarding beams does not work on Syn  $\rightarrow$  Sk. Thus, we validate the effectiveness of discarding

Methods	Mech.	bi.clst	car	trunk	veget.	traf.	pole	garb.	build.	cone.	fence	bi.cle	ground	pers.	mIoU	gain
Source only	-	47.2	43.6	37.8	70.3	11.1	33.8	19.5	67.9	11.2	19.9	9.6	<u>77.9</u>	47.8	38.3	+0.0
AdaptSegNet [26]	A	43.9	48.2	39.0	69.6	15.5	33.6	21.3	64.3	12.7	25.0	11.6	76.0	49.9	39.3	+1.0
CLAN [14]	A	43.9	46.6	<u>41.3</u>	71.0	15.1	34.3	20.4	69.6	9.5	23.2	12.0	75.1	51.3	39.5	+1.2
ADVENT [27]	A	44.6	47.6	40.3	71.2	15.6	35.6	<u>22.0</u>	68.4	10.6	<u>25.9</u>	10.4	76.7	52.3	40.1	+1.8
FADA [28]	A	39.6	41.2	38.8	69.2	16.3	32.1	18.1	67.9	11.5	22.0	<u>13.0</u>	71.4	47.9	37.6	-0.7
MRNet [45]	A	43.5	47.2	39.1	70.4	15.5	32.8	<u>22.0</u>	66.1	<u>13.2</u>	24.2	11.2	76.8	50.0	39.4	+1.1
PMAN [40]	A	<b>52.6</b>	<u>61.5</u>	<b>46.8</b>	<b>75.1</b>	<u>18.8</u>	<b>36.5</b>	21.4	<u>74.7</u>	<b>18.3</b>	25.8	<b>37.5</b>	73.7	<b>61.9</b>	<b>46.5</b>	+8.2
PCAN (Ours)	A	<u>48.6</u>	<b>62.1</b>	37.5	<u>74.0</u>	<b>23.9</b>	<u>31.4</u>	<b>22.2</b>	<b>76.9</b>	6.5	<b>41.9</b>	11.9	<b>79.1</b>	<u>61.2</u>	<u>44.4</u>	+6.1
CoSMix [22]	S	<u>53.6</u>	47.6	44.8	<b>75.1</b>	16.8	<u>37.9</u>	25.3	<u>72.7</u>	<b>19.9</b>	<u>39.7</u>	<b>10.8</b>	<b>80.0</b>	56.5	44.6	+6.3
PolarMix [33]	S	-	-	-	-	-	-	-	-	-	-	-	-	-	30.4	-8.3
LaserMix [10]	S	<b>58.4</b>	<u>61.3</u>	<b>47.7</b>	69.0	<u>21.9</u>	<b>39.5</b>	<u>30.9</u>	61.0	<u>16.1</u>	36.5	7.1	<u>79.5</u>	<u>62.6</u>	<u>45.5</u>	+7.2
DGT-ST (Ours)	S	55.1	<b>70.7</b>	<u>46.1</u>	<u>74.2</u>	<b>30.1</b>	36.3	<b>44.1</b>	<b>81.0</b>	4.3	<b>62.8</b>	<u>10.3</u>	78.5	<b>67.2</b>	<b>50.8</b>	+12.5

Table 2. Comparison results of SynLiDAR  $\rightarrow$  semanticPOSS adaptation in terms of mIoU. A/S denotes adversarial training/self-training.

	Baseline	PCAN	XY-noise	Random $D_p$	Density $D_p$	mIoU	gain
0						30.5	+0.0
1	ADVENT	✓				35.4	+4.9
2		✓	✓			35.8	+5.3
3		✓	✓	✓		34.9	+4.4
4		✓	✓		✓	<b>37.0</b>	<b>+6.5</b>
		LM	DGT	SAC	PCAN model		
5	Source only	✓				20.4	+0.0
6		✓				36.0	+15.6
7		✓	✓			37.5	+17.1
8		✓	✓	✓		38.7	+18.3
9		✓	✓		✓	39.8	+19.4
10		✓	✓	✓	✓	<b>43.1</b>	<b>+22.7</b>

Table 3. Ablation of each component in DGT-ST on Syn  $\rightarrow$  Sk. The whole training consists of adversarial training (upper) and self-training (bottom). The PCAN model represents the initialization model of the self-training provided by PCAN.

beams and points of DGT on Syn  $\rightarrow$  Sp. For others, we only present the results on Syn  $\rightarrow$  Sk due to its much larger validation set than semanticPOSS and limited space.

**Density-guided translator.** To validate the effectiveness of DGT, we conduct experiments in adversarial training and self-training in Tabs. 3 and 4. As shown in Tab. 3, with DGT, we respectively obtain +1.6% (rows 1 and 4) mIoU gain in PCAN, and +1.5% (rows 6 and 7) mIoU gain in LM, on Syn  $\rightarrow$  Sk task. Specifically, DGT consists of injecting noise on the X and Y axes (XY-noise), discarding beams  $D_b$  and discarding points  $D_p$ . Comparing rows 1 and 2, we get +0.4% mIoU increase by adding XY-noise. Discarding points based on the point density (Density  $D_p$ ) brings +1.2% (rows 2 and 4) mIoU improvement. However, randomly discarding points reduces the performance, which drops the mIoU by 0.9% (rows 2 and 3). These results confirm the discussion part in Sec. 3.2. In Tab. 4, DGT also brings +2.1% (rows 1 and 3) and +2.2% (rows 5 and 7) mIoU gains on Syn  $\rightarrow$  Sp task. Besides, when the source and target domain are collected by sensors with different beams, DGT will discard both beams and points. We validate each of them

	Baseline	PCAN	Discard beams	Discard points	mIoU	gain
0					40.1	+0.0
1	ADVENT	✓			42.3	+2.2
2		✓	✓		43.6	+3.5
3		✓	✓	✓	<b>44.4</b>	<b>+4.3</b>
		LM	Discard beams	Discard points		
4	source only				38.3	+0.0
5		✓			45.5	+7.2
6		✓	✓		46.7	+8.4
7		✓	✓	✓	<b>47.7</b>	<b>+9.4</b>

Table 4. Ablations of PCAN and LaserMix on Syn  $\rightarrow$  Sp. In addition to XY-noise, DGT needs to discard beams and points, we ablate them in adversarial training and self-training, respectively.

on Syn  $\rightarrow$  Sp. In Tab. 4, compared to PCAN and LM, we respectively obtain +1.3% (rows 1 and 2) and +1.2% (rows 5 and 6) mIoU gain by discarding beams. Discarding points further improves the performance, which brings +0.8% and +1.0% mIoU gain in PCAN and LM, respectively.

We further verify the effectiveness of DGT in CoSMix. In Tab. 5, DGT brings +0.6% (rows 1 and 2) mIoU gain. Notably, CoSMix is already equipped with some data augmentation strategies, including random subsample and scan rotation. DGT still improves its performance. All these experimental results prove that DGT can boost the performance of the prevalent 3D UDA segmentation methods.

**PCAN.** We investigate the effectiveness of using the prototype to perform category-level adversarial alignment. In Tab. 3, the baseline method ADVENT only gives 30.5% mIoU score on the target domain (row 0). PCAN brings +4.9% mIoU gain on Syn  $\rightarrow$  Sk task (rows 0 and 1). Its effectiveness is also proved in Tab. 4, where PCAN brings +2.2% mIoU gain on Syn  $\rightarrow$  Sp task (rows 0 and 1).

**LaserMix.** Although LaserMix (LM) is proposed for 3D SSL, we extend it to 3D UDA segmentation and shows encouraging performance. LM brings +15.6% (rows 5 and 6 in Tab. 3) and +7.2% (rows 4 and 5 in Tab. 4) mIoU gain on Syn  $\rightarrow$  Sk and Syn  $\rightarrow$  Sp, respectively. Compared



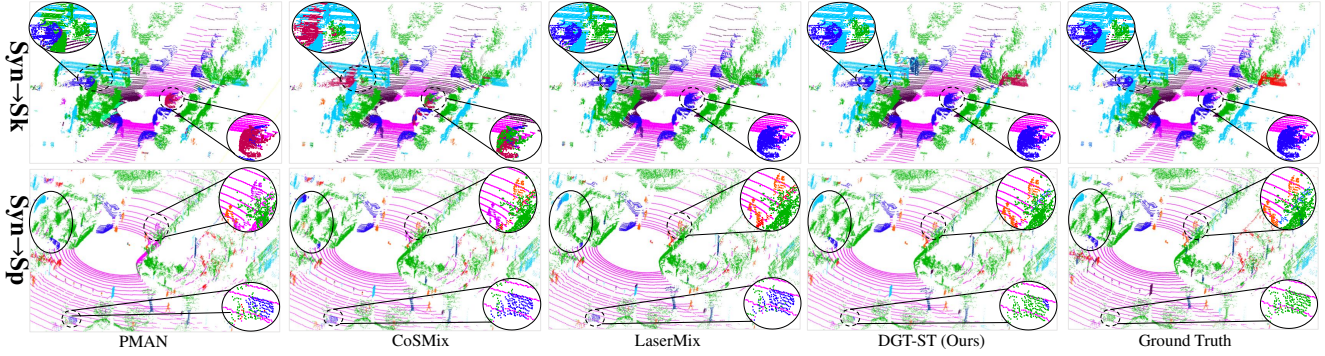


Figure 4. Visual results of UDA segmentation for SynLiDAR  $\rightarrow$  SemanticKITTI and SynLiDAR  $\rightarrow$  SemanticPOSS tasks. Black circles highlight some regions of interest. Best viewed in color.

	Method	mIoU	gain
0	Source only	20.4	+0.0
1	CoSMix [22]	29.9	+9.5
2	CoSMix [22] + DGT	30.5	+10.1
3	CoSMix [22] + DGT + PCAN model	39.1	+18.7
4	CoSMix [22] + DGT + PCAN model+ SAC	<b>40.0</b>	<b>+19.6</b>

Table 5. Validation of the effectiveness of each proposed component with CoSMix [22] on Syn  $\rightarrow$  Sk.

$\gamma_2$	0	0.1	0.01	0.001	0.0001
mIoU	37.5	37.4	38.3	<b>38.7</b>	38.6

Table 6. The effect of the hyper-parameter  $\gamma_2$  in the second stage of DGT-ST on Syn  $\rightarrow$  Sk.

with CoSMix (results in Tabs. 1 and 2), it respectively outperforms CoSMix by 6.1% and 0.9% on two UDA tasks.

**Source-aware consistency (SAC) regularization.** We aim to use DGT to generate source-like target scans to bridge the domain gap. SAC is employed to give consistent predictions about a target scan with and without DGT. In Tabs. 3 and 5, its effectiveness is proved by the fact that adding this term offers +1.2% (rows 7 and 8 in Tab. 3) and +0.9% (rows 3 and 4 in Tab. 5) mIoU gains for LM and CoSMix on Syn  $\rightarrow$  Sk task, respectively. These results also validate the effectiveness of our method in the single-stage setting.

**PCAN pretrained model.** In Tabs. 3 and 5, we conduct experiments with and without PCAN pretrained model for LM and CoSMix. The pretrained model brings +3.8% (rows 6 and 9 in Tab. 3) and +8.6% (rows 2 and 3 in Tab. 5) mIoU gains for them, respectively. It also brings +4.4% (rows 8 and 10 in Tab. 3) mIoU gain for our final model. These results are consistent with 2D counterparts, *i.e.*, the well-initialized model helps the self-training methods.

**Parameter sensitivity of  $\gamma_2$ .**  $\gamma_2$  is a hyper-parameter in Eq. (15) to balance the SAC loss and the other two CE losses. A larger  $\gamma_2$  would force the student model to pay more attention to make consistent predictions with the teacher model,

thereby influencing the student to learn the target domain knowledge. Thus, we conduct experiments to observe the impact of changing this trade-off value, and the results are shown in Tab. 6. The final result is stable when  $\gamma_2$  is smaller than 0.1. A proper choice of  $\gamma_2$  is between 0.001 and 0.0001.

#### 4.4. Qualitative Results

Fig. 4 presents a visual comparison of DGT-ST against the previous two lines of methods, including adversarial training (PMAN) and self-training (CoSMix and LM). From the first row, we can see that PMAN and CoSMix incorrectly classify the car as tree or road while DGT-ST identifies them precisely. DGT-ST gives more accurate results (top black circle) than LM. In the bottom row, DGT-ST can not only accurately classify the person walking on the road, but also identify the plants that are easily misclassified as cars (bottom black circle). Thus, we can conclude that DGT-ST can extract more discriminative features than the previous works.

## 5. Conclusions

In this paper, we present a LiDAR scan translation strategy DGT and a two-stage training pipeline DGT-ST for 3D synthetic-to-real UDA segmentation. DGT leverages the point density to discard points of each scan, thereby alleviating the domain gap at the input level. It can be integrated into the prevalent UDA methods and boost their final performance. In the first stage of DGT-ST, we propose PCAN to provide a well-initialized pretrained model for self-training. It is a category-level adversarial network and uses prototypes to prevent negative transfer. By leveraging DGT and PCAN, in the second stage of DGT-ST, we use LaserMix to construct an intermediate domain and propose SAC-LM to perform self-training. A source-aware consistent loss is proposed to empower the segmentor to learn the source knowledge from the target data, thereby mitigating the domain discrepancy further. Extensive experiments on two prevalent synthetic-to-real tasks demonstrate the superiority of DGT-ST, which outperforms the previous approaches by a large margin.



## References

- [1] Angelika Ando, Spyros Gidaris, Andrei Bursuc, Gilles Puy, Alexandre Boulch, and Renaud Marlet. RangeViT: Towards vision transformers for 3d semantic segmentation in autonomous driving. In *CVPR*, pages 5240–5250, 2023. 1, 2
- [2] Jens Behley, Martin Garbade, Andres Milioto, Jan Quenzel, Sven Behnke, Cyrill Stachniss, and Jurgen Gall. SemanticKITTI: A dataset for semantic scene understanding of LiDAR sequences. In *ICCV*, pages 9297–9307, 2019. 5
- [3] Christopher Choy, JunYoung Gwak, and Silvio Savarese. 4D Spatio-Temporal ConvNeTs: Minkowski convolutional neural networks. In *CVPR*, pages 3075–3084, 2019. 2, 6
- [4] Tiago Cortinhal, George Tzelepis, and Eren Erdal Aksoy. SalsaNext: Fast, uncertainty-aware semantic segmentation of lidar point clouds for autonomous driving. *arXiv preprint arXiv:2003.03653*, 2020. 2
- [5] Qingyong Hu, Bo Yang, Linhai Xie, Stefano Rosa, Yulan Guo, Zhihua Wang, Niki Trigoni, and Andrew Markham. RandLA-Net: Efficient semantic segmentation of large-scale point clouds. In *CVPR*, pages 11108–11117, 2020. 1, 2
- [6] Peng Jiang and Srikanth Saripalli. LiDARNet: A boundary-aware domain adaptation model for point cloud semantic segmentation. In *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, pages 2457–2464, 2021. 1, 2
- [7] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2014. 6
- [8] Lingdong Kong, Youquan Liu, Runnan Chen, Yuexin Ma, Xinge Zhu, Yikang Li, Yuenan Hou, Yu Qiao, and Ziwei Liu. Rethinking range view representation for lidar segmentation. In *ICCV*, pages 228–240, 2023. 2
- [9] Lingdong Kong, Niamul Quader, and Venice Erin Liong. ConDA: Unsupervised domain adaptation for lidar segmentation via regularized domain concatenation. In *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, pages 9338–9345. IEEE, 2023. 2, 3
- [10] Lingdong Kong, Jiawei Ren, Liang Pan, and Ziwei Liu. LaserMix for semi-supervised lidar semantic segmentation. In *CVPR*, pages 21705–21715, 2023. 2, 4, 5, 6, 7
- [11] Xin Lai, Jianhui Liu, Li Jiang, Liwei Wang, Hengshuang Zhao, Shu Liu, Xiaojuan Qi, and Jiaya Jia. Stratified transformer for 3d point cloud segmentation. In *CVPR*, pages 8500–8509, 2022. 1, 2
- [12] Xin Lai, Yukang Chen, Fanbin Lu, Jianhui Liu, and Jiaya Jia. Spherical transformer for lidar-based 3d recognition. In *CVPR*, pages 17545–17555, 2023. 2
- [13] Guangrui Li, Guoliang Kang, Xiaohan Wang, Yunchao Wei, and Yi Yang. Adversarially Masking Synthetic to Mimic Real: Adaptive noise injection for point cloud segmentation adaptation. In *CVPR*, pages 20464–20474, 2023. 1, 2, 3
- [14] Yawei Luo, Liang Zheng, Tao Guan, Junqing Yu, and Yi Yang. Taking a closer look at domain shift: Category-level adversaries for semantics consistent domain adaptation. In *CVPR*, pages 2507–2516, 2019. 5, 6, 7
- [15] Xudong Mao, Qing Li, Haoran Xie, Raymond YK Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks. In *ICCV*, pages 2794–2802, 2017. 3
- [16] Yancheng Pan, Biao Gao, Jilin Mei, Sibao Geng, Chengkun Li, and Huijing Zhao. SemanticPOSS: A point cloud dataset with large quantity of dynamic instances. In *IEEE Intelligent Vehicles Symposium (IV)*, pages 687–693. IEEE, 2020. 6
- [17] Jinyoung Park, Sanghyeok Lee, Sihyeon Kim, Yunyang Xiong, and Hyunwoo J Kim. Self-positioning point-based transformer for point cloud understanding. In *CVPR*, pages 21814–21823, 2023. 2
- [18] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. PyTorch: An imperative style, high-performance deep learning library. *NeurIPS*, 32, 2019. 6
- [19] 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
- [20] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. PointNet++: Deep hierarchical feature learning on point sets in a metric space. *NeurIPS*, 30, 2017. 2
- [21] Guocheng Qian, Yuchen Li, Houwen Peng, Jinjie Mai, Hasan Hammoud, Mohamed Elhoseiny, and Bernard Ghanem. PointNeXt: Revisiting PointNet++ with improved training and scaling strategies. *NeurIPS*, 35:23192–23204, 2022. 2
- [22] Cristiano Saltori, Fabio Galasso, Giuseppe Fiameni, Nicu Sebe, Elisa Ricci, and Fabio Poiesi. CoSMix: Compositional semantic mix for domain adaptation in 3d lidar segmentation. In *ECCV*, pages 586–602. Springer, 2022. 2, 3, 6, 7, 8
- [23] Amirreza Shaban, JoonHo Lee, Sanghun Jung, Xiangyun Meng, and Byron Boots. LiDAR-UDA: Self-ensembling through time for unsupervised lidar domain adaptation. In *ICCV*, pages 19784–19794, 2023. 2, 3
- [24] Haotian Tang, Zhijian Liu, Shengyu Zhao, Yujun Lin, Ji Lin, Hanrui Wang, and Song Han. Searching efficient 3d architectures with sparse point-voxel convolution. In *ECCV*, pages 685–702. Springer, 2020. 2
- [25] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *NeurIPS*, 30, 2017. 3
- [26] Yi-Hsuan Tsai, Wei-Chih Hung, Samuel Schuster, Kihyuk Sohn, Ming-Hsuan Yang, and Manmohan Chandraker. Learning to adapt structured output space for semantic segmentation. In *CVPR*, pages 7472–7481, 2018. 4, 5, 6, 7
- [27] Tuan-Hung Vu, Himalaya Jain, Maxime Bucher, Matthieu Cord, and Patrick Pérez. ADVENT: Adversarial entropy minimization for domain adaptation in semantic segmentation. In *CVPR*, pages 2517–2526, 2019. 3, 4, 5, 6, 7
- [28] Haoran Wang, Tong Shen, Wei Zhang, Ling-Yu Duan, and Tao Mei. Classes Matter: A fine-grained adversarial approach to cross-domain semantic segmentation. In *ECCV*, pages 642–659, 2020. 6, 7
- [29] Yi Wei, Zibu Wei, Yongming Rao, Jiabin Li, Jie Zhou, and Jiwen Lu. LiDAR Distillation: Bridging the beam-induced domain gap for 3d object detection. In *ECCV*, pages 179–195. Springer, 2022. 3
- [30] Bichen Wu, Alvin Wan, Xiangyu Yue, and Kurt Keutzer. SqueezeSeg: Convolutional neural nets with recurrent crf for real-time road-object segmentation from 3d lidar point

- cloud. In *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, pages 1887–1893. IEEE, 2018. [2](#)
- [31] Bichen Wu, Xuanyu Zhou, Sicheng Zhao, Xiangyu Yue, and Kurt Keutzer. SqueezeSegV2: Improved model structure and unsupervised domain adaptation for road-object segmentation from a lidar point cloud. In *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, pages 4376–4382. IEEE, 2019. [2](#)
- [32] Hai Wu, Chenglu Wen, Shaoshuai Shi, Xin Li, and Cheng Wang. Virtual sparse convolution for multimodal 3d object detection. In *CVPR*, pages 21653–21662, 2023. [3](#)
- [33] Aoran Xiao, Jiaying Huang, Dayan Guan, Kaiwen Cui, Shijian Lu, and Ling Shao. PolarMix: A general data augmentation technique for lidar point clouds. In *NeurIPS*, 2022. [2](#), [3](#), [6](#), [7](#)
- [34] Aoran Xiao, Jiaying Huang, Dayan Guan, Fangneng Zhan, and Shijian Lu. Transfer learning from synthetic to real lidar point cloud for semantic segmentation. In *AAAI*, pages 2795–2803, 2022. [2](#), [3](#), [5](#), [6](#)
- [35] Jianyun Xu, Ruixiang Zhang, Jian Dou, Yushi Zhu, Jie Sun, and Shiliang Pu. RPNNet: A deep and efficient range-point-voxel fusion network for lidar point cloud segmentation. In *CVPR*, pages 16024–16033, 2021. [2](#)
- [36] Xu Yan, Chaoda Zheng, Zhen Li, Sheng Wang, and Shuguang Cui. PointASNL: Robust point clouds processing using nonlocal neural networks with adaptive sampling. In *CVPR*, pages 5589–5598, 2020. [2](#)
- [37] Xu Yan, Jiantao Gao, Chaoda Zheng, Chao Zheng, Ruimao Zhang, Shuguang Cui, and Zhen Li. 2DPASS: 2d priors assisted semantic segmentation on lidar point clouds. In *ECCV*, pages 677–695. Springer, 2022. [1](#)
- [38] Li Yi, Boqing Gong, and Thomas Funkhouser. Complete & Label: A domain adaptation approach to semantic segmentation of LiDAR point clouds. In *CVPR*, pages 15363–15373, 2021. [1](#), [2](#), [3](#)
- [39] Zhimin Yuan, Chenglu Wen, Ming Cheng, Yanfei Su, Weiquan Liu, Shangshu Yu, and Cheng Wang. Category-level adversaries for outdoor lidar point clouds cross-domain semantic segmentation. *IEEE Trans. Intell. Transp. Syst.*, 24(2): 1982–1993, 2022. [3](#)
- [40] Zhimin Yuan, Ming Cheng, Wankang Zeng, Yanfei Su, Weiquan Liu, Shangshu Yu, and Cheng Wang. Prototype-guided multitask adversarial network for cross-domain lidar point clouds semantic segmentation. *IEEE Trans. Geosci. Remote Sens.*, 61:1–13, 2023. [1](#), [3](#), [5](#), [6](#), [7](#)
- [41] Pan Zhang, Bo Zhang, Ting Zhang, Dong Chen, Yong Wang, and Fang Wen. Prototypical pseudo label denoising and target structure learning for domain adaptive semantic segmentation. In *CVPR*, pages 12414–12424, 2021. [2](#), [4](#)
- [42] Qiming Zhang, Jing Zhang, Wei Liu, and Dacheng Tao. Category anchor-guided unsupervised domain adaptation for semantic segmentation. In *NeurIPS*, 2019. [2](#), [4](#)
- [43] Hengshuang Zhao, Li Jiang, Jiaya Jia, Philip HS Torr, and Vladlen Koltun. Point transformer. In *ICCV*, pages 16259–16268, 2021. [2](#)
- [44] Sicheng Zhao, Yezhen Wang, Bo Li, Bichen Wu, Yang Gao, Pengfei Xu, Trevor Darrell, and Kurt Keutzer. ePointDA: An end-to-end simulation-to-real domain adaptation framework for lidar point cloud segmentation. In *AAAI*, pages 3500–3509, 2021. [1](#), [2](#)
- [45] Zhedong Zheng and Yi Yang. Unsupervised scene adaptation with memory regularization in vivo. In *IJCAI*, pages 1076–1082, 2020. [6](#), [7](#)
- [46] 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. [2](#)
- [47] Xinge Zhu, Hui Zhou, Tai Wang, Fangzhou Hong, Yuexin Ma, Wei Li, Hongsheng Li, and Dahua Lin. Cylindrical and asymmetrical 3d convolution networks for lidar segmentation. In *CVPR*, pages 9939–9948, 2021. [1](#), [2](#)