

LiDAR Distillation: Bridging the Beam-Induced Domain Gap for 3D Object Detection

Yi Wei^{1,2}, Zibu Wei¹, Yongming Rao^{1,2}, Jiaxin Li³, Jie Zhou^{1,2}, Jiwen Lu^{1,2*}

¹Department of Automation, Tsinghua University, China

²Beijing National Research Center for Information Science and Technology, China

³Gaussian Robotics, China

y-wei19@mails.tsinghua.edu.cn,

weizb18@mails.tsinghua.edu.cn, raoyongming95@gmail.com

lijx1992@gmail.com, jzhou@tsinghua.edu.cn, lujiwen@tsinghua.edu.cn

Abstract. In this paper, we propose the LiDAR Distillation to bridge the domain gap induced by different LiDAR beams for 3D object detection. In many real-world applications, the LiDAR points used by mass-produced robots and vehicles usually have fewer beams than that in large-scale public datasets. Moreover, as the LiDARs are upgraded to other product models with different beam amount, it becomes challenging to utilize the labeled data captured by previous versions' high-resolution sensors. Despite the recent progress on domain adaptive 3D detection, most methods struggle to eliminate the beam-induced domain gap. We find that it is essential to align the point cloud density of the source domain with that of the target domain during the training process. Inspired by this discovery, we propose a progressive framework to mitigate the beam-induced domain shift. In each iteration, we first generate low-beam pseudo LiDAR by downsampling the high-beam point clouds. Then the teacher-student framework is employed to distill rich information from the data with more beams. Extensive experiments on Waymo, nuScenes and KITTI datasets with three different LiDAR-based detectors demonstrate the effectiveness of our LiDAR Distillation. Notably, our approach does not increase any additional computation cost for inference. Code is available at <https://github.com/weiyithu/LiDAR-Distillation>.

Keywords: Unsupervised domain adaptation, 3D object detection, Knowledge distillation

1 Introduction

Recently, LiDAR-based 3D object detection [20,34,35,33,47,58] has attracted more and more attention due to its crucial role in autonomous driving. While the great improvement has been made, most of the existing works focus on the single domain, where the training and testing data are captured with the

* Corresponding author

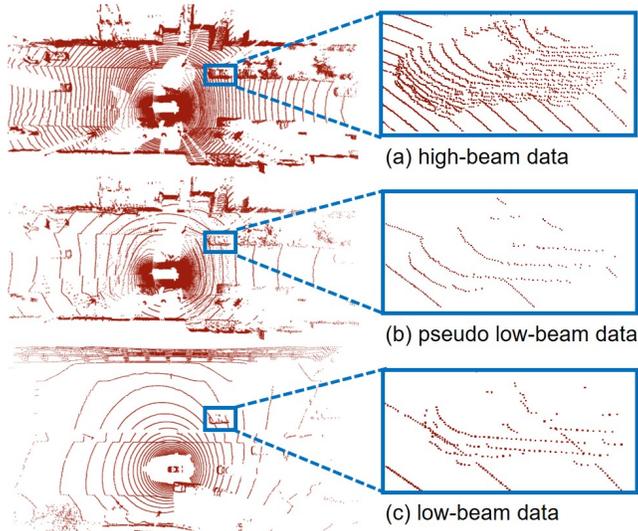


Fig. 1. The beam-induced domain gap. The high-beam data (source domain) and the low-beam data (target domain) are from Waymo [36] and nuScenes[2] datasets respectively. We generate pseudo low-beam data by downsampling high-beam data to align the point cloud density of the source domain to the target domain. The models are trained on this synthetic data with a teacher-student framework.

same LiDAR sensors. In many real-world scenarios, we can only get low-beam (*e.g.* 16 beams) point clouds since high-resolution LiDAR (*e.g.* 64 beams) is prohibitively expensive. For example, the price of a 64-beam Velodyne LiDAR is even higher than a cleaning robot or an unmanned delivery vehicle. However, since 3D detectors [20,47,33] cannot simply transfer the knowledge from the high-beam data to the low-beam one, it is challenging to make use of large-scale datasets [9,36,18] collected by high-resolution sensors. Moreover, with the update of product models, the robots or vehicles will be equipped with the LiDAR sensors with different beams. It is unrealistic to collect and annotate massive data for each kind of product. One solution is to use the highest-beam LiDAR to collect informative training data and adapt them to the low-beam LiDARs. Thus, it is an important direction to bridge the domain gap induced by different LiDAR beams.

Since beam-induced domain gap directly comes from the sensors, it is specific and important among many domain-variant factors. Different with research-guided datasets, the training data in industry tends to have similar environments with that during inference, where the environmental domain gaps such as object sizes, weather conditions are marginal. The main gap comes from the beam difference of LiDARs used during training and inference. Moreover, although LiDAR Distillation is designed for beam-induced domain gap, we can easily

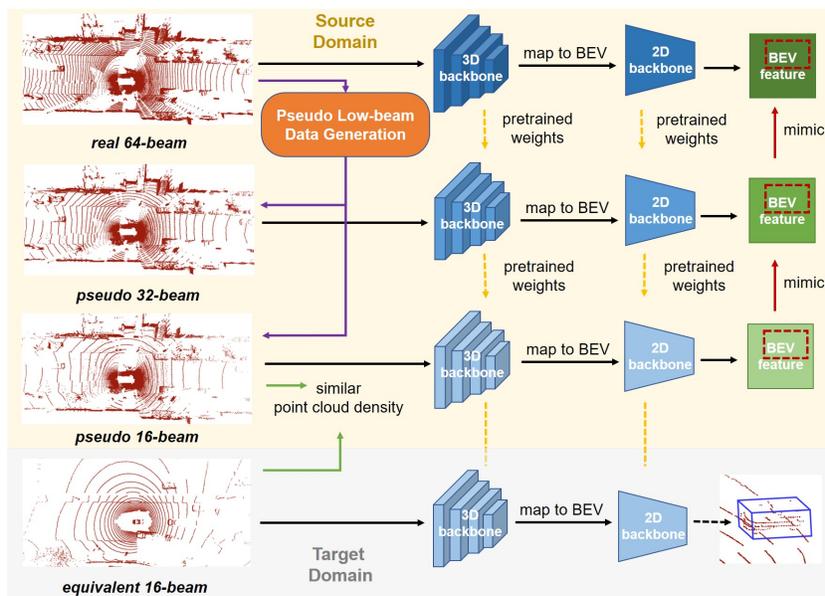


Fig. 2. An overview of LiDAR Distillation. Our method aligns the point cloud density of the source domain with that of the target domain in a progressive way. In each iteration, we first generate pseudo low-beam data by downsampling high-beam point clouds. Then we employ a teacher-student framework to improve the model trained on synthetic low-beam data. Specifically, teacher and student networks are trained on high-beam and low-beam data respectively, which have the same architecture. The student network is initialized with the pretrained weights of the teacher model. As summarized in [37], many 3D detectors employ 3D backbones and 2D backbones, and predict dense BEV features. Based on this fact, we conduct the mimicking operation on the ROI regions of BEV feature maps.

combine our method with other 3D unsupervised domain adaptation methods to solve general domain gaps.

While there are several works [23,32,40,45,49,55] that concentrate on the unsupervised domain adaptation (UDA) task in 3D object detection, most of them aim to address general UDA issues. However, due to the fact that most 3D detectors are sensitive to point cloud densities, the beam-induced domain gap cannot be handled well with these methods. As mentioned in ST3D [49], it is difficult for their UDA method to adapt detectors from the data with more beams to the data with fewer beams. Moreover, since the structure of point clouds is totally different from images, the UDA methods [6,60,29,31,30,56,3,44,21,59] designed for image tasks are not suitable for 3D object detection.

To address the issue, we propose the LiDAR Distillation to bridge the beam-induced domain gap in 3D object detection task. The key insight is to align the point cloud density and distill the rich knowledge from the high-beam data in

a progressive way. First, we downsample high-beam data to pseudo low-beam point clouds, where downsampling operations are both conducted on beams and points in each beam. To achieve this goal, we split the point clouds to each beam via a clustering algorithm. Then, we present a teacher-student pipeline to boost the performance of the model trained on low-beam pseudo data. Unlike knowledge distillation [15,28,28,4,22] used in model compression, we aim to reduce the performance gap caused by data compression. Specifically, the teacher and student networks have the same structure, and the only difference comes from the different beams of training data. Initialized by the weights of the teacher model, the student network mimics the region of interest (ROI) of bird’s eye view (BEV) feature maps predicted by the teacher model.

Experimental results on both cross-dataset [36,2] and single-dataset [9] settings demonstrate the effectiveness of LiDAR Distillation. For cross-dataset adaptation [36,2], although there exists many other kinds of domain gap (*e.g.* object sizes), our method still outperforms state-of-the-art methods [49,40]. In addition, our method is complementary to other general UDA approaches [23,32,40,49,55] and we can get more promising results combining LiDAR Distillation with ST3D [49]. To exclude other domain-variant factors, we further conduct single-dataset experiments on KITTI [9] benchmark. The proposed method significantly improves the direct transfer baseline with a great margin while the general UDA method ST3D [49] surprisingly degrades the performance.

2 Related Work

LiDAR-based 3D Object Detection: LiDAR-based 3D object detection [20,34,35,33,47,58,50,25] focuses on the problem of localizing and classifying objects in 3D space. It has attracted increasing attention in recent works due to its eager demand in computer vision and robotics. As the pioneer works, [5,48,19] directly project point clouds to 2D BEV maps and adopt 2D detection methods to extract features and localize objects. Beyond these works, as a commonly used 3D detector in the industry due to its good trade-off of efficiency and accuracy, PointPillars [20] leverages an encoder to learn the representation of point clouds organized in pillars. Recently, voxel-based methods [47,58,33,8] achieve remarkable performances thanks to the development of 3D sparse convolution [10,7]. As a classical backbone, SECOND [47] investigates an improved sparse convolution method embedded with voxel feature encoding layers. Combining voxel-based networks with point-based set abstraction via a two-step strategy, PV-RCNN [33] becomes one of the state-of-the-art methods. In our work, we adopt PointPillars [20], SECOND [47] and PV-RCNN [33] to demonstrate the effectiveness of our LiDAR Distillation.

Unsupervised Domain Adaptation on Point Clouds: The goal of unsupervised domain adaptation [6,60,29,31,30,56,3,44,21,59,52] is to transfer the model trained on source domain to unlabeled target domains. Recently, some works begin to concentrate on the UDA problem in 3D tasks. [26,57] explore object-level feature alignment between global features and local features. Among the

UDA methods for 3D semantic segmentation [43,17,51], [51] tries to solve the domain shift induced by different LiDAR sensors. However, their method adopts 3D convolutions for scene completion, which will bring huge additional computation cost. Moreover, the complete scenes can be easily used in 3D segmentation but they are not suitable for 3D object detection. Compared with great development in 3D object detection, only a few methods [23,32,40,45,49,55] have been proposed for the UDA of 3D detection. Wang *et al.*[40] conclude that the key factor of the domain gap is the difference in car size and they propose SN to normalize the object sizes of the source and target domains. Yang *et al.*[49] propose ST3D which redesigns the self-training pipeline in 2D UDA methods for 3D object detection task. While most of these works target at general UDA methods, we find that they struggle to bridge the beam-induced domain gap.

Knowledge Distillation: Knowledge distillation (KD) [15,28,4,22] has become one of the most effective techniques for model compression. KD leverages the teacher-student framework and distills the knowledge from teacher models to improve the performance of student models. As one of the pioneer works, Hinton *et al.*[15] transfer the knowledge through soft targets. Compared with the KD in classification task [1,14,38,53,54], compressing object detectors [4,22,41,11,39] is more challenging. Chen *et al.*[4] improve the students by mimicking all components, including intermediate features and soft outputs. Beyond this work, Li *et al.*[22] only mimic the areas sampled from region proposals. Recently, Guo *et al.*[11] point out that features from background regions are also important. Different from above methods, we utilize knowledge distillation for data compression (high-beam data to low-beam data) instead of model compression. In addition, few works study the problem of knowledge distillation in 3D object detection. Due to the differences in backbones, it is not trivial to directly apply 2D methods to 3D tasks. In our method, we adopt feature imitation on bird’s eye view (BEV) feature maps.

3 Approach

3.1 Problem Statement

The goal of unsupervised domain adaptation is to mitigate the domain shift between the source domain and the target domain and maximize the performance on the target domain. Specifically, in this work, we aim to transfer a 3D object detector trained on high-beam source labeled data $\{(X_i^s, Y_i^s)\}_{i=1}^{N_s}$ to low-beam unlabeled target domain $\{X_i^t\}_{i=1}^{N_t}$, where s and t represent source and target domains respectively. The N_s is the samples number of the source domain while X_i^s and Y_i^s mean the i th source point cloud and its corresponding label. The 3D bounding box label Y_i^s is parameterized by its center location of (c_x, c_y, c_z) , the size (d_x, d_y, d_z) , and the orientation θ .

3.2 Overview

We propose the LiDAR Distillation to bridge the domain gap induced by LiDAR beam difference between source and target domains. Figure 2 shows the pipeline

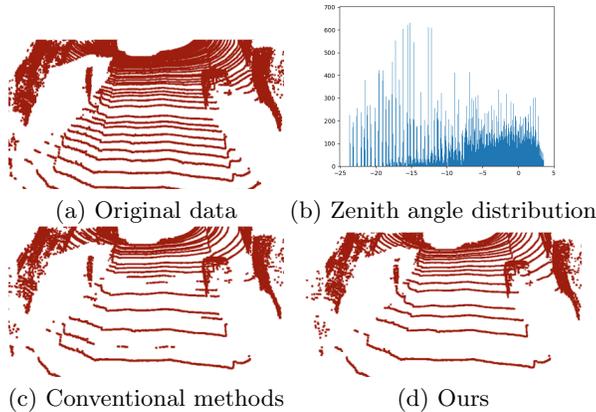


Fig. 3. Pseudo low-beam data generation. **Top row:** (a) original 64-beam LiDAR points. (b) the distribution of zenith angles θ . **Bottom row:** (c) pseudo 32-beam data generated by conventional methods [42,24]. (d) pseudo 32-beam data generated by our clustering method. Although we can see the distinct peaks in the distribution of θ , there also exists some noise. Conventional methods will produce short broken lines with the wrong beam labels while our method can generate more realistic data.

of our method. A key factor of our approach is to progressively generate low-beam pseudo LiDAR $\{X_i^m\}_{i=1}^{N_s}$ in the source domain, which has the similar density with the point clouds of the target domain. The density indicates the density of beams and point number per beam. Section 3.3 introduces this part in details. Then we leverage the teacher-student framework to improve the performance of the model trained on pseudo low-beam data, which is discussed in Section 3.4. Finally, Section 3.5 shows how we progressively distill knowledge from the informative high-beam real point clouds.

We find that the general UDA method [49] will surprisingly degrade the performance when dealing with the domain shift only caused by LiDAR beams. The observation indicates that it is hard to directly transfer the knowledge from the high-beam source domain to the low-beam target domain. Although subsampling high-beam source domain data to low-beam ones will lose rich data knowledge, it guarantees the point cloud density of the source domain is similar to that of the target domain, which is more important compared with data information gain. In this work, we focus on the margin between real high-beam data and generated low-beam data without leveraging the training set in the target domain. Therefore, our approach can be easily integrated with other UDA methods by substituting the source domain with the generated middle domain.

3.3 Pseudo Low-beam Data Generation

The objective of this stage is to close the gap of input point cloud densities between two domains. An alternative solution is to increase the beams of the

point clouds in the target domain. This can be easily implemented by upsampling the range images converted from point clouds. However, our experiments show that the upsampled point clouds have many noisy points that heavily destroy objects’ features. In this work, we propose to downsample the real high-beam data in the source domain to realize point cloud density’s alignment. To this end, we first need to split point clouds into different beams. Although some datasets have beam annotations, many LiDAR points do not have these labels. Here we present a clustering algorithm to get the beam label of each point. We assume the beam number and mean point number in each beam of source and target domains are B_s, P_s and B_t, P_t respectively.

We first transfer cartesian coordinates (x, y, z) to spherical coordinates:

$$\begin{aligned}\theta &= \arctan \frac{z}{\sqrt{x^2 + y^2}} \\ \phi &= \arcsin \frac{y}{\sqrt{x^2 + y^2}}\end{aligned}\tag{1}$$

where ϕ and θ are azimuth and zenith angles. Figure 3 shows one case in KITTI dataset [9] and the subfigure (b) is the distribution of zenith angles θ . Conventional methods [42,24] assume that beam angles are uniformly distributed in a range (*e.g.* $[-23.6^\circ, 3.2^\circ]$ in KITTI dataset) and assign a beam label to each point according to the distance between its zenith angle and beam angles. However, this does not always hold true due to the noise in raw data. To tackle the problem, we apply K-Means algorithm on zenith angles to find B_s cluster centers as beam angles. As shown in Figure 3 (c)(d), compared to the conventional methods, our method is more robust to the noise and can generate more realistic pseudo data. Moreover, from Table 2, we can see that the vertical field of view (the range of beam angles) of datasets are different. Intuitively aligning B_s beams to B_t beams is not correct, which will cause different beam densities. We define the equivalent beams B'_t to the source domain as follows:

$$B'_t = \left[\frac{\alpha_s - \beta_s}{\alpha_t - \beta_t} B_t \right]\tag{2}$$

where $[\alpha_s, \beta_s]$ and $[\alpha_t, \beta_t]$ represent the vertical field of view in source and target domains. With the beam label of each point, we can easily downsample B_s beams to B'_t beams. Moreover, we should also align P_s to P_t by sampling points in each beam. Instead of random sampling, we sort points according to their azimuth angles ϕ and subsample them orderly.

3.4 Distillation from High-beam Data

Knowledge distillation usually employs a teacher-student framework, where the teacher network is a large model while the student network is a small model. The goal is to improve the performance of the student network by learning the prediction of teacher network. There are two kinds of mimic targets: classification logits and intermediate feature maps. In object detection task, since feature maps

are crucial for both object classification and localization, they are the commonly used mimic targets. Formally, regression loss is applied between the feature maps predicted by student and teacher networks:

$$L_m = \frac{1}{N} \sum_i \|f_t^i - r(f_s^i)\|_2 \quad (3)$$

where f_t^i and f_s^i are the teacher and student networks' feature maps of the i th sample. r is a transformation layer to align the dimensions of feature maps. During the mimic training, the weights of teacher are fixed and we only train the student network.

Different from knowledge distillation methods aiming at model compression, our goal is to distill the knowledge from the high-beam data $\{X_i^s\}_{i=1}^{N_s}$ to the pseudo low-beam data $\{X_i^m\}_{i=1}^{N_s}$, which can be regarded as a data compression problem. Unlike model compression, our student and teacher models employ the same network architecture. The only difference between them is that the input to the student and teacher models are X^m and X^s respectively. To get higher accuracy, we use the pretrained teacher model's parameters as the initial weights of the student network.

However, compared with 2D detectors, the architectures of 3D detectors are more diverse. To the best of our knowledge, few works study the mimicking methods for 3D object detection. Moreover, it is not trivial to mimic sparse 3D features since the occupied voxels of student and teacher networks are different. We observe that many 3D object detection methods [20,5,48,19,47,33,8] project 3D features to bird's eye view (BEV). Based on this fact, we conduct mimicking operation on dense BEV features.

As mentioned in [22,41], the dimension of feature maps can be million magnitude and it is difficult to perform regression of two high-dimension features. This phenomenon also exists in BEV maps. In addition, many regions on feature maps have weak responses and not all elements are crucial for final prediction. The student network should pay more attention to the features in important regions. Here, we propose to mimic the BEV features sampled from the region of interest (ROI). The ROIs contain both positive and negative samples with different ratios and sizes, which are informative and have significant guidance effects. These ROIs generated by the teacher model will force the student network to learn more representative features. For the 3D detectors which do not use ROIs (*e.g.* PointPillars [20] and SECOND [47]), we can add a light-weight ROI head assigning negative and positive targets during training and remove it during inference.

The loss function that the student network intends to minimize is defined as follows:

$$L = L_{gt} + \lambda L_m^{3D} \\ L_m^{3D} = \frac{1}{N} \sum_i \frac{1}{M_i} \sum_j \|b_t^{ij} - b_s^{ij}\|_2 \quad (4)$$

where b^{ij} indicates the j th ROI of the i th sample’s BEV feature maps. N and M_i represent sample number and ROI number in the i th sample. L_{gt} is the original objective function for 3D object detection and λ is the mimic loss weight. Note that there is no need for transformation layer r since student and teacher networks have the same architecture.

3.5 Progressive Knowledge Transfer

If the beam difference between real high-beam data and pseudo low-beam data is too large (*e.g.* 128 vs 16), the student network cannot learn from the teacher model effectively. To solve the issue, we further present a progressive pipeline to distill the knowledge step by step, which is described in Algorithm 1. For example, if the source domain is 64-beam and the target domain is 16-beam, we first need to generate pseudo 32-beam data and train a student model on it. This student model will be used as the teacher network in the next iteration. Then we downsample both the beams and points per beam of 32-beam data to 16-beam data. We utilize the student model trained on the generated 16-beam data as the final model for the target domain. We downsample beams for 2 times in each iteration since lots of LiDAR data contains 2^k beams.

4 Experiments

4.1 Experimental Setup

Datasets: We conduct experiments on three popular autonomous driving datasets: KITTI [9], Waymo [36] and nuScenes [2]. Specifically, cross-dataset experiments are conducted from Waymo to nuScenes while single-dataset adaptation uses KITTI benchmark. Table 2 shows an overview of three datasets. Following [49,40], we adopt the KITTI evaluation metric for all datasets on the commonly used car category (the vehicle category in Waymo). We report the average precision (AP) over 40 recall positions. The IoU thresholds are 0.7 for both the BEV IoUs and 3D IoUs evaluation. The models are evaluated on the validation set of each dataset.

Implementation Details: We evaluate our LiDAR Distillation on three commonly used 3D detectors: PointPillars [20], SECOND [47] and PV-RCNN [33]. Following [49], we also add an extra IoU head to SECOND, which is named as SECOND-IoU. To guarantee the same input format in each dataset, we only use (x, y, z) as the raw points’ features. To fairly compare with ST3D [49], we adopt same data augmentation without using GT sampling. The data augmentation is performed simultaneously in a low-beam and high-beam point clouds pair. We utilize 128 ROIs to sample the features on the BEV feature maps and the ratio of positive and negative samples of ROIs was 1:1. The mimic loss weight is set to 1 to balance multiple objective functions. We run full training epochs for each iteration in the progressive knowledge transfer pipeline, *i.e.*, 80 epochs for KITTI and 30 epochs for Waymo. Our work is built upon the 3D object detection codebase OpenPCDet [37].

Algorithm 1 Progressive Knowledge Transfer

Require: Source domain labeled data $\{(X_i^s, Y_i^s)\}_{i=1}^{N_s}$. The (equivalent) beam number and mean point number in each beam of source and target domains are B_s, P_s and B'_t, P_t respectively.

Ensure: The 3D object detection model for target domain.

- 1: Calculate the iteration n for progressive knowledge transfer: $n = \lceil \log_2(B_s/B'_t) \rceil$
- 2: Pretrain 3D detector D_0 on $\{(X_i^s, Y_i^s)\}_{i=1}^{N_s}$.
- 3: **for** $j = 1$ to n : **do**
- 4: Use the method introduced in Section 3.3 to downsample $\{X_i^s\}_{i=1}^{N_s}$ to $\{X_i^{m_j}\}_{i=1}^{N_s}$ which have $\lfloor B_s/2^j \rfloor$ beams. If $j == n$, also downsample the mean point number per beam P_s to P_t .
- 5: Adopt D_{j-1} as the teacher model to train the student model D_j on pseudo low-beam data $\{(X_i^{m_j}, Y_i^s)\}_{i=1}^{N_s}$ with the distillation method described in Section 3.4.
- 6: **end for**
- 7: D_n is the final model for the target domain.

Task	Method	SECOND-IoU	PV-RCNN	PointPillars
		AP _{BEV} / AP _{3D}	AP _{BEV} / AP _{3D}	AP _{BEV} / AP _{3D}
Waymo → nuScenes	Direct Transfer	32.91 / 17.24	34.50 / 21.47	27.8 / 12.1
	SN [40]	33.23 / 18.57	34.22 / 22.29	28.31 / 12.98
	ST3D	35.92 / 20.19	36.42 / 22.99	30.6 / 15.6
	Hegde <i>et al.</i> [12]	- / 20.47	- / -	- / -
	UMT [13]	35.10 / 21.05	- / -	- / -
	3D-CoCo [52]	- / -	- / -	33.1 / 20.7
	Ours	40.66 / 22.86	43.31 / 25.63	40.23 / 19.12
Ours (w/ ST3D)	42.04 / 24.50	44.08 / 26.37	40.83 / 20.97	

Table 1. Results of cross-dataset adaptation experiments. **Direct Transfer** indicates that the model trained on the Waymo dataset is directly tested on nuScenes. We report AP_{BEV} and AP_{3D} over 40 recall positions of the car category at IoU = 0.7.

4.2 Cross-dataset Adaptation

Table 1 shows the results of cross-dataset experiments. We compare our method with two state-of-the-art general UDA methods SN [40] and ST3D [49]. The reported numbers of these two methods are from [49]. Although our LiDAR Distillation only aims to bridge the domain gap caused by LiDAR beams and there exist many other domain-variant factors (*e.g.* car sizes, geography and weather) in two datasets, our method still outperforms other general UDA methods for a large margin. The improvements over Direct Transfer baseline are also significant. Moreover, our LiDAR Distillation does not use the training set of the target domain and is complementary to other general UDA methods. Take the ST3D as an example: we can simply replace the model pretrained on the source domain with the model pretrained on pseudo low-beam data. From Table 1, we can see that our method greatly boosts the performance of ST3D (AP_{BEV} : 35.92 → 42.04), which demonstrates that point cloud density is the key factor among many domain-variant factors.

Dataset	LiDAR Type	VFOV	Points Per Beam
Waymo [36]	64-beam	$[-17.6^\circ, 2.4^\circ]$	2258
KITTI [9]	64-beam	$[-23.6^\circ, 3.2^\circ]$	1863
nuScenes [2]	32-beam	$[-30.0^\circ, 10.0^\circ]$	1084

Table 2. Datasets overview. We use **version 1.0** of Waymo Open Dataset. **VFOV** means vertical field of view (the range of beam angles). Statistical information is computed from whole dataset.

4.3 Single-dataset Adaptation

To decouple the influence of other factors, we further conduct single-dataset adaptation experiments, where the domain gap is only induced by LiDAR beams. Different with research-guided datasets, the training data in industry tends to have similar environments with that during inference, where the environmental domain gaps such as object sizes are marginal. The main gap is the beam difference of LiDARs used during training and inference. Thus, it is valuable to investigate the performance drop only caused by this certain domain gap. Unfortunately, the popular autonomous driving datasets [9,2,36] are captured by only one type of LiDAR sensor. To solve the problem, we build the synthetic low-beam target domain by uniformly downsampling the validation set of KITTI. We adopt widely used 3D detector PointPillars [20] in this experiment.

Table 3 shows the results of single-dataset adaptation on KITTI dataset. 32* and 16* mean that we not only reduce LiDAR beams but also subsample 1/2 points in each beam. We do not compare with SN since car sizes are the same in training and validation sets. We can see that our LiDAR Distillation significantly improves Direct Transfer baseline. As the difference between point cloud densities increases, the performance gap becomes larger. Surprisingly, we find that the general UDA method ST3D [49] will degrade the performance. This is mainly due to different point cloud densities in source and target domains, which are not friendly to the self-training framework.

4.4 Ablation Studies

In this section, we conduct sufficient experiments to verify the effectiveness of each proposed component. To exclude the effects of other domain variables, all experiments are conducted under the single-dataset setting with PointPillars [20] backbones.

Component Analysis of Teacher-student Framework: The experiment uses 32-beam data as the target domain. As shown in Table 4, all components contribute to the final performance. While high-beam data is informative and we can get high-performance model, it is hard to directly transfer the knowledge to the low-beam domain. With the alignment of the point cloud density, 3D detectors can learn more robust domain-invariant features in the source domain training set. The distillation from real high-beam data will introduce rich knowledge, which can be maximumly preserved in teacher pretrained weights.

Target Domain Beams	Method	PointPillars					
		AP _{BEV}			AP _{3D}		
		Easy	Moderate	Hard	Easy	Moderate	Hard
32	Direct Transfer	91.19	79.81	76.53	80.79	65.91	61.09
	ST3D [49]	86.65	71.29	66.77	75.36	57.57	52.88
	Ours	91.59	82.22	79.57	86.00	70.15	66.86
	<i>Improvement</i>	+0.40	+2.41	+3.04	+5.21	+4.24	+5.77
32*	Direct Transfer	88.39	73.56	68.22	74.59	57.77	51.45
	ST3D [49]	83.81	67.08	62.57	71.09	53.30	48.34
	Ours	90.60	79.47	76.58	82.83	66.96	62.51
	<i>Improvement</i>	+2.21	+5.91	+8.36	+8.24	+9.19	+11.06
16	Direct Transfer	83.11	64.91	58.32	67.64	47.48	41.41
	ST3D [49]	75.49	57.58	52.33	60.36	42.40	37.93
	Ours	89.98	74.32	71.54	80.21	59.87	55.32
	<i>Improvement</i>	+6.87	+9.41	+13.22	+12.57	+12.39	+13.91
16*	Direct Transfer	77.22	56.32	49.51	57.36	38.75	32.88
	ST3D [49]	76.04	55.63	49.18	55.17	37.02	31.84
	Ours	87.44	70.43	66.35	75.35	55.24	50.96
	<i>Improvement</i>	+10.22	+14.11	+16.84	+17.99	+16.49	+18.08

Table 3. Results of single-dataset adaptation on KITTI dataset [9]. **Direct Transfer** indicates that the model trained on the original KITTI training set is directly evaluated on the low-beam validation set. For **32*** and **16***, we not only reduce LiDAR beams but also subsample 1/2 points in each beam. We report AP_{BEV} and AP_{3D} over 40 recall positions of the car category at IoU = 0.7. The improvement is calculated according to Direct Transfer baseline. We mark the best result of each target domain in **bold**.

To better understand the effects of point cloud density alignment and knowledge distillation, we further conduct experiments of all synthetic target domains in supplementary material, where we find that knowledge distillation plays an important role when the domain gap is small while density alignment is more crucial when the domain gap is large. The experiments also show that our method significantly outperforms point density alignment baseline, i.e. taking the low-beam point clouds as training data, which verify the effectiveness of the proposed teacher-student pipeline.

Effectiveness of Progressive Knowledge Transfer: We further investigate the effectiveness of progressive knowledge transfer. From Table 5, we can see that if the teacher model is trained on original 64-beam data, it cannot provide effective guidance for the student model since there exist large gaps between 64-beam and 16-beam data. Moreover, compared with the teacher model directly trained on 32-beam data, the teacher network obtained from the proposed progressive pipeline gets better results.

Analysis for Mimicking Regions: The experimental result is shown in Table 6. We find that it is not effective to mimic all elements of BEV feature maps. We also try mimicking the regions inside groundtruth bounding boxes. However, the results are not promising since some background features are also useful. With a

density alignment	knowledge distillation	pretrained	AP _{3D}		
			Easy	Moderate	Hard
			80.79	65.91	61.09
✓			82.80	67.01	63.82
✓		✓	83.05	68.66	64.79
✓	✓		83.71	69.01	64.97
✓	✓	✓	86.00	70.15	66.86

Table 4. The ablation study of the teacher-student framework. **Density alignment** means that we use generated low-beam data to train the network. **Pretrained** indicates that we employ pretrained teacher weights as the initial weights for the student network. 32-beam data is set as the target domain.

teacher model	AP _{3D}		
	Easy	Moderate	Hard
64	78.75	58.80	54.47
32	78.31	58.74	55.62
32 progressive	80.21	59.87	55.32

Table 5. The ablation study of progressive knowledge transfer. **64** and **32** represent that the teacher model is trained without progressive distillation. The teacher model obtained from our method is named as **32 progressive**. The result is evaluated on 16-beam data.

mimicking regions	AP _{3D}		
	Easy	Moderate	Hard
all	83.83	69.36	64.98
groundtruth	84.46	69.99	65.41
ROI	86.00	70.15	66.86

Table 6. The ablation study of mimicking regions. **all** means that we perform mimicking operation on the whole BEV feature maps while **groundtruth** indicates that we only mimic the regions inside groundtruth bounding boxes.

good balance of foreground and background information, mimicking ROI regions performs best.

4.5 Further Analysis

The Necessity of Point Cloud Density Alignment: In this part, we study whether the point cloud density is a key factor to bridge the domain gap among many domain-variant factors. Although the data in nuScenes [2] is 32-beam, according to the VFOV in Table 2 and Equation 2, its equivalent beam in the source domain is 16. In addition, the points per beam of Waymo data is about twice as much as that of nuScenes data. Therefore, 16* beam is the most similar point cloud density compared with target domain data. Experimental results are shown in Table 7, where evaluation results on target and source domain are illustrated in the first and second columns respectively. We observe that we can get better results when point cloud densities of source and target domain become closer. Indeed, both the beams of Waymo and nuScenes obtain non-uniform distribution and our method does not require them to strictly align. While the alignment of the point cloud density will sacrifice the performance on the source domain due to the information loss, it will bring huge improvements on the target domain.

Source Domain Beams	AP _{BEV} /AP _{3D}	
	nuScenes	Waymo
64	32.91/17.24	67.72/54.01
32	37.35/20.66	64.67/51.18
16	40.08/21.51	57.04/43.41
16*	40.66/22.86	53.23/39.97

Table 7. The Necessity of Point Cloud Density Alignment. We downsample Waymo [36] data to different beams and evaluate models on nuScenes [2] (target domain) and Waymo [36] (source domain). Notice that the equivalent beam of nuScenes is 16*. The experiment is conducted with SECOND-IoU backbone.

Method	AP _{3D}		
	Easy	Moderate	Hard
interpolation	74.43	54.56	49.72
super resolution	79.78	62.58	58.04
Ours	86.00	70.15	66.86

Table 8. Comparison with LiDAR upsampling methods. We first convert point clouds of the target domain to range images. Then we upsample range images with bilinear interpolation and a light-weight super-resolution network [16], and convert them back to point clouds. 32-beam KITTI data is used as the target domain.

Comparison with LiDAR Upsampling Methods: An alternative solution for point cloud density alignment is to upsample the low-beam data in the target domain. Different from scene completion methods [46,27] which utilize 3D convolutions, we do not need to complete the whole scenes but need to process point clouds efficiently. To solve the problem, we first convert point clouds to range images. Then we upsample range images by naive non-parameter bilinear interpolation and a super-resolution network [16]. As shown in Table 8, these two upsampling methods fail to boost the accuracy, which demonstrates that it is not trivial to accurately upsample LiDAR points with high efficiency. This is mainly because there exist some noisy points, which will lead to incorrect object shapes.

5 Conclusion

In this work, we propose the LiDAR Distillation to bridge the domain gap caused by different LiDAR beams in 3D object detection task. Inspired by the discovery that point cloud density is an important factor for this problem, we first generate pseudo low-beam data by downsampling real high-beam LiDAR points. Then knowledge distillation technique is used to progressively boost the performance of the model trained on synthetic low-beam point clouds. Experimental results on three popular autonomous driving datasets demonstrate the effectiveness of our method.

Acknowledgments

This work was supported in part by the National Key Research and Development Program of China under Grant 2017YFA0700802, in part by the National Natural Science Foundation of China under Grant 62125603 and Grant U1813218, in part by a grant from the Beijing Academy of Artificial Intelligence (BAAI).

References

1. Anil, R., Pereyra, G., Passos, A., Ormandi, R., Dahl, G.E., Hinton, G.E.: Large scale distributed neural network training through online distillation. arXiv preprint arXiv:1804.03235 (2018) [5](#)
2. Caesar, H., Bankiti, V., Lang, A.H., Vora, S., Liong, V.E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., Beijbom, O.: nuscenes: A multimodal dataset for autonomous driving. In: CVPR. pp. 11621–11631 (2020) [2](#), [4](#), [9](#), [11](#), [13](#), [14](#)
3. Chen, C., Zheng, Z., Ding, X., Huang, Y., Dou, Q.: Harmonizing transferability and discriminability for adapting object detectors. In: CVPR. pp. 8869–8878 (2020) [3](#), [4](#)
4. Chen, G., Choi, W., Yu, X., Han, T., Chandraker, M.: Learning efficient object detection models with knowledge distillation. NeurIPS (2017) [4](#), [5](#)
5. Chen, X., Ma, H., Wan, J., Li, B., Xia, T.: Multi-view 3d object detection network for autonomous driving. In: CVPR. pp. 1907–1915 (2017) [4](#), [8](#)
6. Chen, Y., Li, W., Sakaridis, C., Dai, D., Van Gool, L.: Domain adaptive faster r-cnn for object detection in the wild. In: CVPR. pp. 3339–3348 (2018) [3](#), [4](#)
7. Choy, C., Gwak, J., Savarese, S.: 4d spatio-temporal convnets: Minkowski convolutional neural networks. In: CVPR. pp. 3075–3084 (2019) [4](#)
8. Deng, J., Shi, S., Li, P., Zhou, W., Zhang, Y., Li, H.: Voxel r-cnn: Towards high performance voxel-based 3d object detection. In: AAAI. pp. 1201–1209 (2021) [4](#), [8](#)
9. Geiger, A., Lenz, P., Urtasun, R.: Are we ready for autonomous driving? the kitti vision benchmark suite. In: CVPR (2012) [2](#), [4](#), [7](#), [9](#), [11](#), [12](#)
10. Graham, B., Engelcke, M., Van Der Maaten, L.: 3d semantic segmentation with submanifold sparse convolutional networks. In: CVPR. pp. 9224–9232 (2018) [4](#)
11. Guo, J., Han, K., Wang, Y., Wu, H., Chen, X., Xu, C., Xu, C.: Distilling object detectors via decoupled features. In: CVPR. pp. 2154–2164 (2021) [5](#)
12. Hegde, D., Patel, V.: Attentive prototypes for source-free unsupervised domain adaptive 3d object detection. arXiv preprint arXiv:2111.15656 (2021) [10](#)
13. Hegde, D., Sindagi, V., Kilic, V., Cooper, A.B., Foster, M., Patel, V.: Uncertainty-aware mean teacher for source-free unsupervised domain adaptive 3d object detection. arXiv preprint arXiv:2109.14651 (2021) [10](#)
14. Heo, B., Kim, J., Yun, S., Park, H., Kwak, N., Choi, J.Y.: A comprehensive overhaul of feature distillation. In: ICCV. pp. 1921–1930 (2019) [5](#)
15. Hinton, G., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network. arXiv preprint arXiv:1503.02531 (2015) [4](#), [5](#)
16. Hui, Z., Gao, X., Yang, Y., Wang, X.: Lightweight image super-resolution with information multi-distillation network. In: ACM-MM. pp. 2024–2032 (2019) [14](#)
17. Jaritz, M., Vu, T.H., Charette, R.d., Wirbel, E., Pérez, P.: xmuda: Cross-modal unsupervised domain adaptation for 3d semantic segmentation. In: CVPR. pp. 12605–12614 (2020) [5](#)
18. Kesten, R., Usman, M., Houston, J., Pandya, T., Nadhamuni, K., Ferreira, A., Yuan, M., Low, B., Jain, A., Ondruska, P., Omari, S., Shah, S., Kulkarni, A., Kazakova, A., Tao, C., Platinsky, L., Jiang, W., Shet, V.: Lyft level 5 perception dataset 2020. <https://level5.lyft.com/dataset/> (2019) [2](#)
19. Ku, J., Mozifian, M., Lee, J., Harakeh, A., Waslander, S.L.: Joint 3d proposal generation and object detection from view aggregation. In: IROS. pp. 1–8 (2018) [4](#), [8](#)

20. Lang, A.H., Vora, S., Caesar, H., Zhou, L., Yang, J., Beijbom, O.: Pointpillars: Fast encoders for object detection from point clouds. In: CVPR. pp. 12697–12705 (2019) [1](#), [2](#), [4](#), [8](#), [9](#), [11](#)
21. Li, C., Du, D., Zhang, L., Wen, L., Luo, T., Wu, Y., Zhu, P.: Spatial attention pyramid network for unsupervised domain adaptation. In: ECCV. pp. 481–497 (2020) [3](#), [4](#)
22. Li, Q., Jin, S., Yan, J.: Mimicking very efficient network for object detection. In: CVPR. pp. 6356–6364 (2017) [4](#), [5](#), [8](#)
23. Luo, Z., Cai, Z., Zhou, C., Zhang, G., Zhao, H., Yi, S., Lu, S., Li, H., Zhang, S., Liu, Z.: Unsupervised domain adaptive 3d detection with multi-level consistency. In: ICCV. pp. 8866–8875 (2021) [3](#), [4](#), [5](#)
24. Milioto, A., Vizzo, I., Behley, J., Stachniss, C.: Rangenet++: Fast and accurate lidar semantic segmentation. In: 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 4213–4220. IEEE (2019) [6](#), [7](#)
25. Qi, C.R., Liu, W., Wu, C., Su, H., Guibas, L.J.: Frustum pointnets for 3d object detection from rgb-d data. In: CVPR. pp. 918–927 (2018) [4](#)
26. Qin, C., You, H., Wang, L., Kuo, C.C.J., Fu, Y.: Pointdan: A multi-scale 3d domain adaption network for point cloud representation. In: NeurIPS. pp. 7192–7203 (2019) [4](#)
27. Roldão, L., de Charette, R., Verroust-Blondet, A.: Lmscnet: Lightweight multiscale 3d semantic completion. In: 3DV (2020) [14](#)
28. Romero, A., Ballas, N., Kahou, S.E., Chassang, A., Gatta, C., Bengio, Y.: Fitnets: Hints for thin deep nets. arXiv preprint arXiv:1412.6550 (2014) [4](#), [5](#)
29. Saito, K., Ushiku, Y., Harada, T.: Asymmetric tri-training for unsupervised domain adaptation. In: ICML. pp. 2988–2997 (2017) [3](#), [4](#)
30. Saito, K., Ushiku, Y., Harada, T., Saenko, K.: Strong-weak distribution alignment for adaptive object detection. In: CVPR. pp. 6956–6965 (2019) [3](#), [4](#)
31. Saito, K., Watanabe, K., Ushiku, Y., Harada, T.: Maximum classifier discrepancy for unsupervised domain adaptation. In: CVPR. pp. 3723–3732 (2018) [3](#), [4](#)
32. Saltori, C., Lathuilière, S., Sebe, N., Ricci, E., Galasso, F.: SF-UDA 3D: Source-Free Unsupervised Domain Adaptation for LiDAR-Based 3D Object Detection. In: 3DV. pp. 771–780 (2020) [3](#), [4](#), [5](#)
33. Shi, S., Guo, C., Jiang, L., Wang, Z., Shi, J., Wang, X., Li, H.: Pv-rcnn: Point-voxel feature set abstraction for 3d object detection. In: CVPR. pp. 10529–10538 (2020) [1](#), [2](#), [4](#), [8](#), [9](#)
34. Shi, S., Wang, X., Li, H.: Pointrenn: 3d object proposal generation and detection from point cloud. In: CVPR. pp. 770–779 (2019) [1](#), [4](#)
35. Shi, S., Wang, Z., Shi, J., Wang, X., Li, H.: From points to parts: 3d object detection from point cloud with part-aware and part-aggregation network. arXiv preprint arXiv:1907.03670 (2019) [1](#), [4](#)
36. Sun, P., Kretschmar, H., Dotiwalla, X., Chouard, A., Patnaik, V., Tsui, P., Guo, J., Zhou, Y., Chai, Y., Caine, B., et al.: Scalability in perception for autonomous driving: Waymo open dataset. In: CVPR. pp. 2446–2454 (2020) [2](#), [4](#), [9](#), [11](#), [14](#)
37. Team, O.D.: Openpcdet: An open-source toolbox for 3d object detection from point clouds. <https://github.com/open-mmlab/OpenPCDet> (2020) [3](#), [9](#)
38. Tung, F., Mori, G.: Similarity-preserving knowledge distillation. In: ICCV. pp. 1365–1374 (2019) [5](#)
39. Wang, T., Yuan, L., Zhang, X., Feng, J.: Distilling object detectors with fine-grained feature imitation. In: CVPR. pp. 4933–4942 (2019) [5](#)

40. Wang, Y., Chen, X., You, Y., Li, L.E., Hariharan, B., Campbell, M., Weinberger, K.Q., Chao, W.L.: Train in germany, test in the usa: Making 3d object detectors generalize. In: CVPR. pp. 11713–11723 (2020) [3](#), [4](#), [5](#), [9](#), [10](#)
41. Wei, Y., Pan, X., Qin, H., Ouyang, W., Yan, J.: Quantization mimic: Towards very tiny cnn for object detection. In: ECCV. pp. 267–283 (2018) [5](#), [8](#)
42. Wu, B., Wan, A., Yue, X., Keutzer, K.: Squeezeseg: Convolutional neural nets with recurrent crf for real-time road-object segmentation from 3d lidar point cloud. In: 2018 IEEE International Conference on Robotics and Automation (ICRA). pp. 1887–1893. IEEE (2018) [6](#), [7](#)
43. Wu, B., Zhou, X., Zhao, S., Yue, X., Keutzer, K.: Squeezesegv2: Improved model structure and unsupervised domain adaptation for road-object segmentation from a lidar point cloud. In: ICRA. pp. 4376–4382 (2019) [5](#)
44. Xu, C.D., Zhao, X.R., Jin, X., Wei, X.S.: Exploring categorical regularization for domain adaptive object detection. In: CVPR. pp. 11724–11733 (2020) [3](#), [4](#)
45. Xu, Q., Zhou, Y., Wang, W., Qi, C.R., Anguelov, D.: SPG: Unsupervised Domain Adaptation for 3D Object Detection via Semantic Point Generation. In: ICCV. pp. 15446–15456 (2021) [3](#), [5](#)
46. Yan, X., Gao, J., Li, J., Zhang, R., Li, Z., Huang, R., Cui, S.: Sparse single sweep lidar point cloud segmentation via learning contextual shape priors from scene completion. In: AAAI. pp. 3101–3109 (2021) [14](#)
47. Yan, Y., Mao, Y., Li, B.: Second: Sparsely embedded convolutional detection. *Sensors* **18**(10), 3337 (2018) [1](#), [2](#), [4](#), [8](#), [9](#)
48. Yang, B., Luo, W., Urtasun, R.: Pixor: Real-time 3d object detection from point clouds. In: CVPR. pp. 7652–7660 (2018) [4](#), [8](#)
49. Yang, J., Shi, S., Wang, Z., Li, H., Qi, X.: ST3D: Self-training for Unsupervised Domain Adaptation on 3D Object Detection. In: CVPR. pp. 10368–10378 (2021) [3](#), [4](#), [5](#), [6](#), [9](#), [10](#), [11](#), [12](#)
50. Yang, Z., Sun, Y., Liu, S., Shen, X., Jia, J.: Std: Sparse-to-dense 3d object detector for point cloud. In: ICCV. pp. 1951–1960 (2019) [4](#)
51. Yi, L., Gong, B., Funkhouser, T.: Complete & label: A domain adaptation approach to semantic segmentation of lidar point clouds. In: CVPR. pp. 15363–15373 (2021) [5](#)
52. Yihan, Z., Wang, C., Wang, Y., Xu, H., Ye, C., Yang, Z., Ma, C.: Learning transferable features for point cloud detection via 3d contrastive co-training. *NeurIPS* **34** (2021) [4](#), [10](#)
53. Yim, J., Joo, D., Bae, J., Kim, J.: A gift from knowledge distillation: Fast optimization, network minimization and transfer learning. In: CVPR. pp. 4133–4141 (2017) [5](#)
54. You, S., Xu, C., Xu, C., Tao, D.: Learning from multiple teacher networks. In: SIGKDD. pp. 1285–1294 (2017) [5](#)
55. Zhang, W., Li, W., Xu, D.: Srdan: Scale-aware and range-aware domain adaptation network for cross-dataset 3d object detection. In: CVPR. pp. 6769–6779 (2021) [3](#), [4](#), [5](#)
56. Zheng, Y., Huang, D., Liu, S., Wang, Y.: Cross-domain object detection through coarse-to-fine feature adaptation. In: CVPR. pp. 13766–13775 (2020) [3](#), [4](#)
57. Zhou, X., Karpur, A., Gan, C., Luo, L., Huang, Q.: Unsupervised domain adaptation for 3d keypoint estimation via view consistency. In: ECCV. pp. 137–153 (2018) [4](#)
58. Zhou, Y., Tuzel, O.: Voxelnet: End-to-end learning for point cloud based 3d object detection. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 4490–4499 (2018) [1](#), [4](#)

59. Zhu, X., Pang, J., Yang, C., Shi, J., Lin, D.: Adapting object detectors via selective cross-domain alignment. In: CVPR. pp. 687–696 (2019) [3](#), [4](#)
60. Zou, Y., Yu, Z., Vijaya Kumar, B., Wang, J.: Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In: ECCV. pp. 289–305 (2018) [3](#), [4](#)