

Learning from Spatio-temporal Correlation for Semi-Supervised LiDAR Semantic Segmentation

Seungho Lee
 Yonsei University
 seungholee@yonsei.ac.kr

Hwijeong Lee
 KAIST
 hjlee0612@kaist.ac.kr

Hyunjung Shim
 KAIST
 kateshim@kaist.ac.kr

Abstract—We address the challenges of the semi-supervised LiDAR segmentation (SSLS) problem, particularly in low-budget scenarios. The two main issues in low-budget SSLS are the poor-quality pseudo-labels for unlabeled data, and the performance drops due to the significant imbalance between ground-truth and pseudo-labels. This imbalance leads to a vicious training cycle. To overcome these challenges, we leverage the spatio-temporal prior by recognizing the substantial overlap between temporally adjacent LiDAR scans. We propose a proximity-based label estimation, which generates highly accurate pseudo-labels for unlabeled data by utilizing semantic consistency with adjacent labeled data. Additionally, we enhance this method by progressively expanding the pseudo-labels from the nearest unlabeled scans, which helps significantly reduce errors linked to dynamic classes. Additionally, we employ a dual-branch structure to mitigate performance degradation caused by data imbalance. Experimental results demonstrate remarkable performance in low-budget settings (i.e., $\leq 5\%$) and meaningful improvements in normal budget settings (i.e., 5 – 50%). Finally, our method has achieved new state-of-the-art results on SemanticKITTI and nuScenes in semi-supervised LiDAR segmentation. With only 5% labeled data, it offers competitive results against fully-supervised counterparts. Moreover, it surpasses the performance of the previous state-of-the-art at 100% labeled data (75.2%) using only 20% of labeled data (76.0%) on nuScenes. The code is available on <https://github.com/halbielee/PLE>.

I. INTRODUCTION

LiDAR segmentation is crucial for autonomous driving, offering robust 3D environment perception and entity identification under various lighting and weather conditions [1], [2]. Current methods employ powerful deep neural networks, requiring extensive labeled data for training [3], [4]. However, annotating LiDAR point clouds is costly and complex, posing a significant challenge. In contrast, unlabeled LiDAR point clouds are abundantly available at no additional cost. Therefore, we focus on semi-supervised LiDAR segmentation (SSLS), leveraging large unlabeled and small labeled LiDAR point clouds.

While semi-supervised semantic segmentation has been extensively studied in 2D images [7], [8], [9], [10], it has been less explored in LiDAR data. Recently, LaserMix [5] introduced a new data augmentation method for SSLS, leveraging the spatial prior in LiDAR. By blending laser beams from different LiDAR scans within the same spatial region, LaserMix generates new training samples, boosting performance across various labeled data ratios (e.g., 1%, 10%, 20%, 50%) and establishing a new state-of-the-art




Fig. 1: Segmentation performances (mIoU) across various labeled ratios. We outperform the state-of-the-art method, LaserMix [5], across all labeled ratios in nuScenes [6], with particularly large margins in low-budget settings. Notably, our method with 5× fewer labels already achieves the performance of full supervision.

(SoTA) in SSLS.

Despite the impressive achievements, existing studies suffer from a substantial performance drop in low-budget settings (e.g., $\leq 5\%$). Fig. 1 shows more than 18%p drop between 0.5% and 5%. Reducing labeled data without compromising performance has a significant practical impact, considering the complexity of labeling LiDAR data compared to 2D images. Therefore, our work focuses on enhancing the practicality of SSLS in low-budget settings.

We identify two main problems of SSLS in low-budget scenarios. (P1) The pseudo-labels for unlabeled data suffer from poor quality. (P2) A significant imbalance between the amount of labeled and unlabeled data (e.g., 1% vs 99%) causes extra performance degradation. To address these problems, (A1) we propose proximity-based label estimation (PLE) that utilizes the unique spatio-temporal characteristics inherent to LiDAR data to generate high-quality pseudo-labels. Unlike 2D images, LiDAR scans possess temporal continuity, sharing rich semantic information across closely-timed scans [11]. By aligning unlabeled and labeled scans into the same coordinate, we can accurately assign pseudo-labels to unlabeled LiDAR points using the nearest labeled scan. This strategy can efficiently capture dynamic objects (e.g., vehicles) by progressively extending labels from closer to more distant unlabeled data. Overall, our PLE effectively

addresses the challenge of labeling in LiDAR datasets by exploiting their inherent spatio-temporal characteristics.

Additionally, (A2) we employ a dual-branch architecture [9], [12], [13], [14] to mitigate performance degradation caused by data imbalance. Unlike previous methods [5], which rely on a single branch model and struggle with the extreme imbalance between labeled and unlabeled data, our approach offers a robust solution. Specifically, we recognize the popular baseline for SSL, MeanTeacher [15], falls into a vicious cycle; a poor teacher network produces poor pseudo-labels, which in turn hampers the network training. To break this cycle, we revise the network’s final layer into two parallel branches: a labeled branch and an unlabeled branch, each responsible for their respective data. This separation protects the labeled branch from noisy pseudo-labels, leading to stable performance improvement.

In our study, extensive experiments on SSSL benchmarks achieved new SoTA performances with significant improvements on both SemanticKITTI (up to 5.6%p) and nuScenes (up to 6.6%p) datasets, particularly in low-budget scenarios. We achieve competitive performance to fully supervised counterparts (100% labeled) by only utilizing 5% of full labels and the previous SoTA LaserMix’s 100% labeled data performance using only 20% of labeled data. Moreover, our experiments demonstrate the effectiveness of PLE for generating accurate pseudo-labels and a dual-branch architecture for improving the training process of the teacher network. Our method proves to be robust and scalable, demonstrating compatibility and substantial performance boosts across different LiDAR representations and backbone networks, underlining the adaptability and potential of our proposed solutions in a variety of settings.

Our key contributions are summarized as follows: (1) We propose a novel approach for generating highly accurate pseudo-labels by utilizing a spatio-temporal prior unique to LiDAR data. (2) We address the performance degradation caused by the data imbalance and break the vicious training cycle in low-budget SSSL with a dual-branch structure. (3) Through extensive experiments, our approach achieves new state-of-the-art results on SemanticKITTI and nuScenes datasets, excelling in low-budget settings, highlighting the effectiveness and practical applicability of our methods.

II. RELATED WORK

A. LiDAR Segmentation

Several studies have explored various representations in LiDAR segmentation. Projection-based methods [5], [16], [17], [18], [19], [20], [21] convert LiDAR point clouds into 2D images for processing with 2D convolution. Voxel-based methods [22], [23], [24] rasterize the LiDAR point clouds into voxels and apply 3D convolution. Multi-view methods [25], [26], [27] integrate data from several perspectives. Despite their promising results, these fully-supervised models heavily rely on large-scale, high-quality datasets. Since their labeling cost is prohibitively expensive, recent studies have focused on developing weakly supervised [4] or semi-supervised methods [5], [28], [29] to significantly

reduce the annotation cost. In this work, we primarily focus on a semi-supervised learning approach that improves segmentation performance by effectively utilizing readily available unlabeled data.

B. Semi-supervised Segmentation

Semi-supervised learning has been extensively explored in 2D image segmentation. One prevalent strategy is the consistency regularization approach, enforcing consistency between predictions for perturbed data [8], [30] or networks [29], [31]. Another approach [32], [33] is to leverage the generative adversarial network (GAN) [34] to provide auxiliary signals for unlabeled data. In contrast to 2D segmentation, there is comparatively less research in 3D. Notably, the studies on outdoor scenes have only recently gained attention. GPC [35] proposes a contrastive learning framework for point clouds but does not differentiate indoor and outdoor scenes. LaserMix [5] exploits the spatial prior in outdoor LiDAR scenes and utilizes a mixture of labeled and unlabeled scans along the inclination axis. Recently, LiM3d [28] introduces the frame downsampling strategy, which determines the most informative keyframes for choosing labeled data. Both LaserMix and LiM3d adopt a pseudo-labeling strategy [8], particularly MeanTeacher [15], repeatedly generating pseudo-labels for unlabeled data. In contrast, we leverage the distinct spatio-temporal prior inherent to LiDAR point clouds and generate highly accurate pseudo-labels for unlabeled data.

C. Temporal information in LiDAR perception

Temporal information of LiDAR data is commonly used in various fields of LiDAR perception, such as motion prediction [36], [37], LiDAR detection [38], [39], [40], [41], sensor calibration [42], and scene flow estimation [43], [44], [45]. Recently, temporal information has been utilized in self-supervised learning [46], [47] and cross-modal 3D understanding [48]. In semantic segmentation, 4D semantic segmentation is actively explored, where multiple scans are taken as input to distinguish between moving and static states of objects [49]. These studies extract 4D spatio-temporal features by stacking frames [50], [51], using k-NN based approaches [52], [53], memory-based technique [54] or voxel-adjacent neighborhood [55]. Recently, LiM3d utilizes temporal information to extract the most informative LiDAR scans. In contrast, our approach focuses on generating accurate pseudo-labels by applying the spatio-temporal prior.

III. METHOD

A. SSSL Baseline

Following the previous work [5], we use the MeanTeacher [15], which involves teacher and student networks. The teacher network is updated from the student network through Exponential Moving Average (EMA) and generates pseudo-labels for unlabeled data. The student network trains on labeled data with ground-truth and unlabeled data with pseudo-labels. We apply cross-entropy and Lovász-softmax losses. We also use the MeanTeacher loss [15] for comparing the two networks’ predictions.

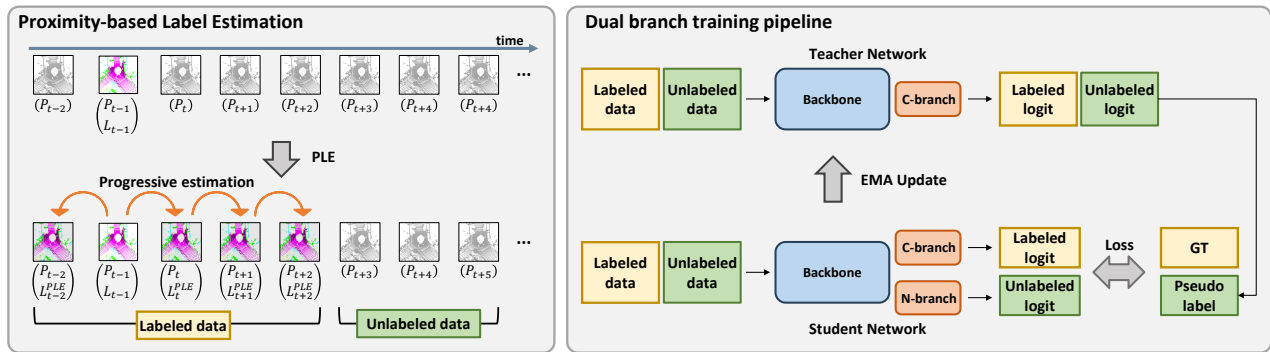


Fig. 2: **Overall framework.** We generate pseudo-labels for unlabeled scans by leveraging spatio-temporal prior in LiDAR. Unlabeled scans with PLE labels are treated as a labeled set during training. We adopt the MeanTeacher model, where the teacher network generates pseudo-labels for remaining unlabeled data. To mitigate the performance degradation caused by the noisy pseudo-labels, we introduce a dual-branch structure where the labeled and unlabeled data are processed separately.

B. Proximity-based Label Estimation

Popular benchmark datasets used in 2D segmentation, such as PASCAL VOC [56], COCO [57], and ADE-20k [58], do not exhibit spatial or temporal correlations between samples. In contrast, outdoor LiDAR segmentation datasets, such as SemanticKITTI [59] and nuScenes [6], present unique characteristics with highly correlated samples in terms of spatial position and sampling time. This is due to the sequential frame-based recording of point clouds as a LiDAR-equipped vehicle moves. Additionally, the high scanning frequency of LiDAR sensors results in substantial overlap between successive scans [11]. For example, a vehicle moving at 100 kilometers per hour with a LiDAR sensor that has a ± 70 meter range would see an overlap of about 80% in the 140-meter range across successive frames, significantly due to the rapid data acquisition rate of these sensors.

By leveraging the inherent spatio-temporal prior in LiDAR data, we propose a Proximity-based Label Estimation (PLE) to produce accurate pseudo-labels for unlabeled scans. The key idea is to reference nearby labeled scans based on the spatio-temporal consistency. The process of pseudo-label generation consists of three steps. (1) *Coordinate transformation*: We transform the coordinate of the referenced labeled scan into the coordinate system of the target unlabeled scan. We utilize pose information, readily available through GPS/IMU [46] or ICP [60]. (2) *Proximate point identification*: We identify the proximate labeled points coinciding with (or closest to) the target unlabeled points in the coordinate system. While the coordinate transformation ensures that the two scans share the same coordinate system, points from the two scans do not exhibit an exact spatial match (*i.e.*, coinciding at the same position). Therefore, we identify the closest labeled counterparts to the unlabeled points. To reduce the searching cost, we build a KD tree [61] for all transformed labeled points to efficiently compute the distances between target unlabeled points. (3) *Label assignment*: We transfer the label of the selected labeled point to the unlabeled point as a pseudo-label. We refer to this pseudo-label as a PLE label to distinguish it from those generated by the teacher network.

The original PLE is effective in handling static scenes since only global transformation is considered in coordinate transformation. Therefore, it has limitations in allocating labels for dynamic objects, especially when the time interval between ground truth and unlabeled data is sufficiently far. To address this performance limit of dynamic objects, we devise a *progressive version of PLE*. In the original PLE, we only reference ground-truth of the labeled data, regardless of the temporal distance between scans. The progressive PLE addresses this by sequentially generating labels, starting from unlabeled scans closest to labeled ones and progressively moving to more distant scans, as illustrated in Fig. 2. This incremental approach significantly reduces the effective time interval between the labeled and unlabeled data, thereby significantly improving accuracies for dynamic objects.

Despite the simplicity, our PLE method efficiently assigns accurate pseudo-labels to unlabeled data. Since these pseudo-labels from PLE are highly accurate compared to those generated by the teacher network during the training process (Fig. 4), it significantly improves the training of semi-supervised models. Also, PLE can seamlessly integrate into existing frameworks as it generates pseudo-labels independently of the training process.

C. Dual-branch for low-budget scenario

In the MeanTeacher training process, the student network learns from both ground-truth of labeled data and the pseudo-labels generated by the teacher network for unlabeled data. However, in low-budget scenarios where a small amount of labeled data is available, there exists a significant imbalance between the amount of labeled and unlabeled data. Consequently, the training of student network is heavily influenced by inaccurate pseudo-labels, leading to degraded performance. This, in turn, results in the generation of inaccurate pseudo-labels, establishing a vicious cycle between the generation of pseudo-labels and the training of the network.

To address this issue, we employ a popular dual-branch architecture [9], [12], [13], [14], [63], which is widely utilized to capture rich context from various sources, including different object sizes [64], adversarial examples [12], [13], different domains [14] and different tasks [63]. Based on the

Method	SemanticKITTI [59]								nuScenes [6]							
	0.5%	1%	2%	5%	10%	20%	50%	100%	0.5%	1%	2%	5%	10%	20%	50%	100%
<i>Sup.-only</i> [†] [23]	40.0	47.1	50.3	55.1	57.6	57.6	58.5	58.5	45.6	52.0	59.2	65.3	68.4	72.0	74.0	75.2
MeanTeacher [†] [15]	42.2	51.0	54.6	55.6	59.1	60.1	60.4	-	43.7	52.2	59.3	63.1	67.2	67.8	68.7	-
CBST [62]	-	48.8	-	-	58.3	59.4	59.7	-	-	53.0	-	-	66.5	69.6	71.6	-
GPC [35]	-	-	-	41.8	49.9	58.8	-	65.8	-	-	-	-	-	-	-	-
CPS [31]	-	46.7	-	-	58.7	59.6	60.5	-	-	52.9	-	-	66.3	70.0	72.5	-
Unal et al. [4]	-	-	-	49.9	58.7	59.1	-	68.2	-	-	-	-	-	-	-	-
LiM3d [28]	-	58.4	-	59.5	62.2	63.1	63.6	69.5	-	-	-	-	-	-	-	-
LaserMix [†] [5]	47.3	55.5	59.2	61.7	62.4	62.4	62.1	63.8	51.4	58.4	63.9	69.7	71.6	73.7	73.7	75.2
Our method	52.2	61.1	62.9	62.8	63.1	64.1	64.3	-	58.0	62.9	67.2	72.8	74.3	76.0	76.1	-
Δ ↑	+4.9	+5.6	+3.7	+1.1	+0.7	+1.7	+2.2	-	+6.6	+4.5	+3.3	+3.1	+2.7	+2.3	+2.4	-

TABLE I: **Segmentation results (mIoU) on benchmark datasets.** Our method achieves state-of-the-art performance across all labeled ratios. The **best** and second best scores for each ratio are highlighted in **bold** and underline, respectively. † denotes the reproduced result. Δ indicates an improvement over LaserMix.

observation that the latter layer of the network is susceptible to noisy labels [65], [66], [67], as shown in Fig. 2, we employ a shared backbone followed by two branches with identical structures to segregate the noisy pseudo-labels. The first branch, called the clean branch (C-branch), learns from accurate labels. The second branch, called the noisy branch (N-branch), learns from noisy pseudo-labels. Since the dual-branch effectively separates the training of clean labels (*i.e.*, ground truth labels and PLE labels) from noisy labels (*i.e.*, pseudo-labels from the teacher network), it prevents the quality degradation of pseudo-labels from the teacher. As a result, the vicious cycle between the training of the network and noisy pseudo-labels is resolved. This approach is particularly effective in low-budget scenarios. Moreover, it makes the network robust to the confidence threshold that determines the quality of pseudo-labels.

D. Overall Pipeline

Our pipeline consists of PLE and the dual-branch structure. Prior to training, we generate PLE labels for some of the unlabeled data. Specifically, we generate pseudo-labels only for frames within a 1-second timeframe of a labeled sample. This is because we have observed that the precision of PLE labels decreases as the temporal gap from the labeled sample increases, as shown in Fig. 4 (b).

As detailed in Sec. III-A and III-C, our approach involves two networks: the student network and the teacher network. During the student network’s training, the C-branch is trained using the labeled data and some of the unlabeled data with PLE labels, while the N-branch focuses on the rest of the unlabeled data without PLE labels. The teacher network is updated through EMA and generates the pseudo-labels using its C-branch, with the N-branch remaining unused.

IV. EXPERIMENT

A. Settings

Datasets. We conduct an empirical study on two popular benchmark datasets: SemanticKITTI [59] and nuScenes [6]. SemanticKITTI contains 19 classes with 19,130 training scans and 4,071 validation scans. nuScenes consists of 16

classes with 29,130 scans for training and 6,019 scans for validation. For the low-budget semi-supervised LiDAR semantic segmentation, we adopt a uniform sampling strategy, selecting 0.5%, 1%, 2%, and 5% of training scans for labeled data, designating the remainder as unlabeled data. Also, we select 10%, 20%, and 50% of training scans for the normal budget setting. This selection strategy aligns with the conventional semi-supervised settings [5], [8], [31].

Implementation details. We select the voxel-based network, Cylinder3D [23], as the segmentation backbone, configuring the voxel resolution as [240, 180, 20] to accommodate LaserMix [5]. Additionally, the hyperparameters for MeanTeacher [15] and LaserMix are adopted based on the settings of LaserMix. To validate the effectiveness of our proposed method, we measure the intersection-over-union (IoU) for each class and reported the mean IoU (mIoU). All experiments were conducted on four NVIDIA A6000 GPUs with 48GB memory using PyTorch.

B. Comparison with SoTA

Tab. I compares our method with current SoTA methods [5], [15], [28], [31], [35], [62] on the SemanticKITTI and nuScenes datasets. Our method demonstrates clear improvements across all labeled ratios on both datasets, especially under low ratio settings. We see over over 3%p increases at 0.5%–2% labeled ratios. In the normal budget settings (10%–50%), our method exhibits moderate improvements of approximately 0.7–2.7%p. Notably, our method achieves over 64% and 75% mIoU on the 20% ratio in both datasets, outperforming LaserMix’s full-label benchmark with only 20% of labels. This highlights the effectiveness of our PLE and dual-branch architecture in generating accurate pseudo-labels and reducing performance degradation, respectively. Our method establishes new SoTA benchmarks in semi-supervised LiDAR segmentation, surpassing other leading techniques across all labeling ratios.

Fig 3 presents qualitative examples in a bird’s-eye view from SemanticKITTI. Correct predictions are marked in green, while errors are marked in red. Notably, our results exhibit greater accuracy than LaserMix. We specifically emphasize *person* categories with dashed circles, underscoring

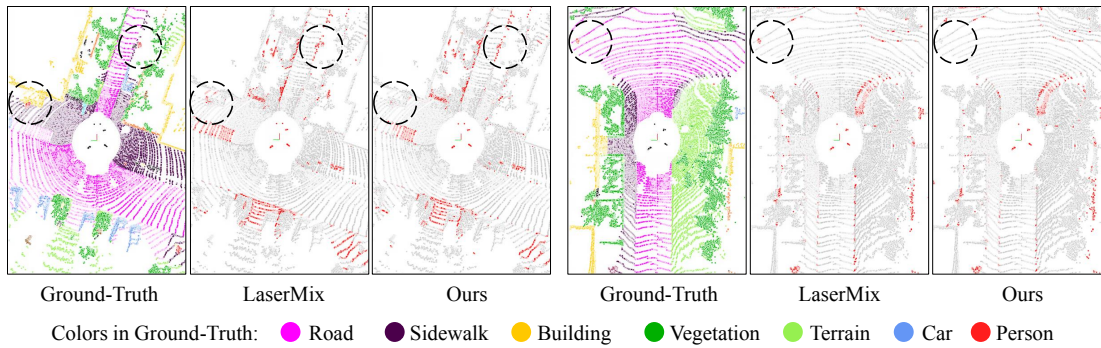


Fig. 3: **Qualitative comparisons between our method and LaserMix.** All samples are visualized from LiDAR bird’s-eye view on *val* set of SemanticKITTI [59]. The **correct** and **incorrect** are painted in **gray** and **red** to highlight the difference. Dashed circles highlight the misprediction of LaserMix. Best viewed in color.

	0.5%	1%	2%	5%
# GT labels	95	191	382	956
# PLE labels	1,805	3,404	6,247	11,490
# Unlabeled data	17,230	15,535	12,501	6,684
<i>Naive PLE</i>				
mIoU (%)	79.5	79.9	79.8	81.1
mPrecision (%)	86.5	87.6	87.4	88.4
<i>Progressive PLE</i>				
mIoU (%)	81.1	81.5	81.7	83.0
mPrecision (%)	88.7	88.9	89.3	89.9

TABLE II: **Statistics changes of training data and accuracy of PLE labels on SemanticKITTI [59].** IoU and Precision are evaluated on each ratio of PLE labels.

	mean	car	bicy	moto	ped	b.cyc	m.cyc
Naive	79.9	87.1	61.2	78.2	82.1	34.7	13.8
Prg.	81.5	87.7	66.6	83.4	83.5	54.2	20.6

TABLE III: **Accuracy (IoU) of dynamic classes in PLE labels.** Prg. refers to progressive version of PLE.

the superior accuracy of our method in identifying *person* classes compared to LaserMix. This is also supported by the class-wise performance, where *person* and *motorcycle* show 8.4%p and 44.2% improvements. These classes have a significant impact on real-world scenarios, in terms of the safety and reliability of autonomous driving.

C. Ablation Study

Accuracy of PLE labels. In Tab. II, we detail the impact of applying PLE on the statistics of training data and the accuracy of PLE labels. Using PLE, we search labeled data within a ± 1 -second interval around each unlabeled scan. With the 10Hz acquisition rate of SemanticKITTI, our method assigns pseudo-labels up to 20 unlabeled frames around each labeled frame, increasing the size of the labeled set by about 19 times. At a 5% labeled ratio, our PLE generates over half of the total LiDAR scans. Furthermore, the PLE labels are highly accurate, achieving a mIoU surpassing 79% and outperforming state-of-the-art fully supervised models with approximately 70% mIoU [23], [68]. Notably, the progressive version of PLE yields more accurate labels compared to the naive version, particularly with significant improvements in dynamic classes. As shown in Tab. III,

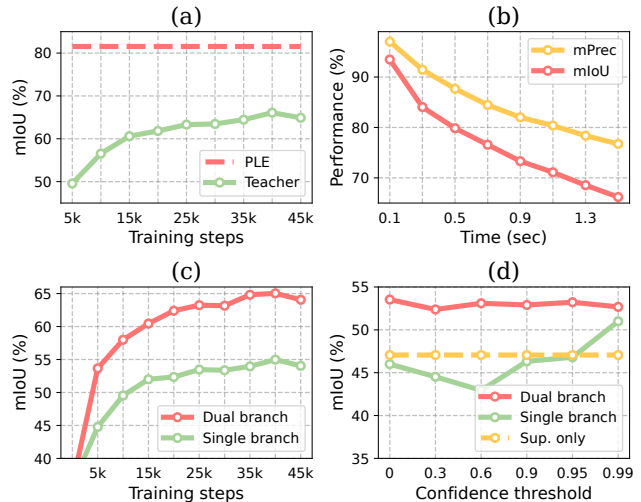


Fig. 4: **Ablation studies.** (a) Accuracy of pseudo-labels from the Teacher network and PLE. (b) Accuracy of PLE-labels over time intervals. (c) Accuracy of pseudo-labels during training. (d) Training results according to confidence threshold. All results are from a 1% ratio of the SemanticKITTI.

PLE	PRG	DB	0.5%	1.0%	2.0%	5.0%
			42.2	51.0	54.6	55.6
✓			45.2	55.7	56.9	57.4
✓	✓		47.1	56.0	57.5	59.3
		✓	47.9	52.9	55.7	57.9
✓	✓	✓	48.2	56.6	58.9	59.9

TABLE IV: **Ablation study on each component.** The accuracy (mIoU) is evaluated on the *val* set in SemanticKITTI. MeanTeacher is used as a baseline. DB and PRG refer to dual-branch and the progressive PLE. The **best score** for each ratio is highlighted in **bold**.

there is a remarkable increase of 5.2%–19.5% in in dynamic classes, highlighting the effectiveness of using intermediate PLE labels during the PLE process instead of solely relying on direct references from labeled samples. This strategy minimizes errors in capturing the movement of dynamic objects, proving its effectiveness in handling the complexities of dynamic classes.

Pseudo-labels from the teacher network. We compare the performance between the PLE labels and the pseudo-

labels generated by the Teacher network in Fig. 4(a). During training, the quality of the pseudo-labels from the Teacher network initially improves but eventually saturates and even slightly declines. In contrast, PLE labels consistently present remarkably higher accuracy (over 10%p) than the pseudo-labels from the Teacher network. This result indicates that PLE labels lead to significant improvement in our semi-supervised training.

PLE labels according to time intervals. Fig. 4(b) depicts the accuracy of PLE labels on SemanticKITTI relative to varying time intervals (frame difference), achieved by selecting adjacent frames at fixed time differences. The quality of PLE labels deteriorates with increasing time gaps between scans. This is due to reduced scan overlap over time, as well as dynamic changes in object movement and occlusion within the scene. Throughout the experiments, we select adjacent scans within a 1-second interval, ensuring reasonable temporal proximity and performance.

Effectiveness of dual-branch. Fig. 4(c) compares the quality of pseudo-labels in single-branch and dual-branch settings during training. The dual-branch consistently enhances the quality of pseudo-labels, particularly presenting strong performance from the early training stages. Also, we compare the segmentation performance of the dual-branch structure with the single branch according to confidence threshold values in Fig. 4(d). The single branch network occasionally underperforms compared to using only labeled data. In contrast, the dual-branch is robust to confidence thresholds and consistently outperforms the single-branch at all thresholds. This improvement is because the dual-branch breaks the vicious training cycle where noisy pseudo-labels affect the training process of the Teacher network and vice versa.

Effect of each component. In Tab. IV, we evaluate the effectiveness of individual components of our approach at various labeled ratios (0.5%, 1%, 2%, and 5%). The results demonstrate a notable performance gain by employing PLE and the dual-branch, respectively. Specifically, PLE enhances accuracy by 1.8–4.7%p over the MeanTeacher, with the progressive reference of the PLE process further boosting performance by 0.3–1.9%p. This improvement is attributed to the superior accuracy of PLE labels compared to those from the teacher network, as detailed in Fig. 4(a). Moreover, the dual-branch consistently improves the performance of the MeanTeacher. Unlike the original MeanTeacher, which uses a single branch, our modification utilizing a dual-branch eliminates noisy training signals and instead focuses on clean labels for training the teacher’s C-branch. This has greater efficacy at lower ratios (e.g., 0.5% ratio translates to 99.5% of noisy labels for training the MeanTeacher), highlighting its suitability for low-budget scenarios. Finally, when we apply PLE and the dual-branch simultaneously, we achieve further performance gain. This is because each component is mutually exclusive when it is applied to the unlabeled data. PLE generates accurate pseudo-labels for the nearby unlabeled scans of the labeled data, while the dual-branch separates the remaining unlabeled data from the C-branch. This highlights the complementary effects of PLE and the

	0.5%	1%	2%	5%
<i>Range View</i>				
CENet	42.3	45.2	52.2	55.1
+ Ours	46.2 +3.9	51.5 +6.3	54.3 +2.1	58.1 +3.0
<i>Voxel View</i>				
Minku18	51.9	57.3	57.6	59.0
+ Ours	56.1 +4.2	60.1 +2.8	60.9 +3.3	63.4 +4.4
SPVCNN	51.5	56.9	57.2	59.3
+ Ours	54.9 +3.4	60.7 +3.8	60.4 +3.2	60.8 +1.5

TABLE V: **Ablation study on other representation and backbones.** The accuracy (mIoU) is evaluated on the *val* set in SemanticKITTI. MT and DB refer to MeanTeacher [15] and dual-branch, respectively.

dual-branch, thereby establishing the effectiveness of each module in our method.

Representation and backbones. To demonstrate the scalability of our method, we evaluate it across other representation and backbones. Specifically, we incorporate the widely used CENet [69] in the range view, and introduce MinkUNet [52] and SPVCNN [22]. As depicted in Tab. V, our approach consistently significantly enhances performance consistently across different types of representations and backbones. Notably, in the range view, we observed substantial improvements ranging from 2.1–6.3%p. Similarly, with MinkUNet18 and SPVCNN, performance improvements ranged from 1.5–4.4%p. The substantial enhancements achieved through the use of PLE and a dual-branch architecture underscore their versatility and applicability across different representations and backbone architectures.

V. CONCLUSION

Our work tackles key challenges in semi-supervised LiDAR segmentation, focusing on low-budget scenarios. We introduce a pseudo-labeling approach, Proximity-based Label Estimation (PLE), leveraging LiDAR’s spatio-temporal prior to enhance the quality of pseudo-labels. Our dual-branch structure overcomes the labeled-unlabeled data imbalance, ensuring strong performance. Through extensive experiments, we demonstrate our method’s efficacy, setting new state-of-the-art results on SemanticKITTI and nuScenes and surpassing the fully supervised counterparts. Based on its effectiveness and simplicity, our approach significantly boosts scalable LiDAR perception, expediting its application in real-world scenarios.

Acknowledgments. This research was supported by the Basic Science Research Program through the National Research Foundation of Korea (NRF) funded by the MSIP (NRF-2022R1A2C3011154, RS-2023-00219019, RS-2023-00240135) and MOE (NRF-2022R1A6A3A13073319), the IITP grant funded by the Korea government(MSIT) (RS-2019-II190075, Artificial Intelligence Graduate School Program(KAIST)) and KEIT grant funded by the Korea government(MOTIE) (No. 2022-0-00680, 2022-0-01045, 2021-0-02068, Artificial Intelligence Innovation Hub) and Samsung Electronics Co., Ltd (IO230508-06190-01).

REFERENCES

- [1] R. Roriz, J. Cabral, and T. Gomes, "Automotive lidar technology: A survey," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 7, pp. 6282–6297, 2021. **1**
- [2] Y. Li and J. Ibanez-Guzman, "Lidar for autonomous driving: The principles, challenges, and trends for automotive lidar and perception systems," *IEEE Signal Processing Magazine*, vol. 37, no. 4, pp. 50–61, 2020. **1**
- [3] Q. Hu, B. Yang, G. Fang, Y. Guo, A. Leonardis, N. Trigoni, and A. Markham, "Sqn: Weakly-supervised semantic segmentation of large-scale 3d point clouds," in *European Conference on Computer Vision*. Springer, 2022, pp. 600–619. **1**
- [4] O. Unal, D. Dai, and L. Van Gool, "Scribble-supervised lidar semantic segmentation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 2697–2707. **1, 2, 4**
- [5] L. Kong, J. Ren, L. Pan, and Z. Liu, "Lasermix for semi-supervised lidar semantic segmentation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 21 705–21 715. **1, 2, 4**
- [6] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "nusenes: A multimodal dataset for autonomous driving," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 11 621–11 631. **1, 3, 4**
- [7] Y. Wang, H. Wang, Y. Shen, J. Fei, W. Li, G. Jin, L. Wu, R. Zhao, and X. Le, "Semi-supervised semantic segmentation using unreliable pseudo-labels," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 4248–4257. **1**
- [8] Y. Zou, Z. Zhang, H. Zhang, C.-L. Li, X. Bian, J.-B. Huang, and T. Pfister, "Pseudoseg: Designing pseudo labels for semantic segmentation," *arXiv preprint arXiv:2010.09713*, 2020. **1, 2, 4**
- [9] W. Luo and M. Yang, "Semi-supervised semantic segmentation via strong-weak dual-branch network," in *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part V 16*. Springer, 2020, pp. 784–800. **1, 2, 3**
- [10] M. Lee, S. Lee, J. Lee, and H. Shim, "Saliency as pseudo-pixel supervision for weakly and semi-supervised semantic segmentation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023. **1**
- [11] N. P. A. Duong, A. Almin, L. Lemarié, and B. R. Kiran, "Reducing redundancy in semantic-kitti: Study on data augmentations within active learning," *Adv. Neural Inform. Process. Syst. Worksh.*, p. 8, 2017. **1, 3**
- [12] C. Xie and A. Yuille, "Intriguing properties of adversarial training at scale," *arXiv preprint arXiv:1906.03787*, 2019. **2, 3**
- [13] C. Xie, M. Tan, B. Gong, J. Wang, A. L. Yuille, and Q. V. Le, "Adversarial examples improve image recognition," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 819–828. **2, 3**
- [14] R. Bermúdez-Chacón, M. Salzmann, and P. Fua, "Domain adaptive multibranch networks," in *International Conference on Learning Representations*, 2019. **2, 3**
- [15] A. Tarvainen and H. Valpola, "Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results," *Advances in neural information processing systems*, vol. 30, 2017. **2, 4, 6**
- [16] T. Cortinhal, G. Tzelepis, and E. Erdal Aksoy, "Salsanext: Fast, uncertainty-aware semantic segmentation of lidar point clouds," in *Advances in Visual Computing: 15th International Symposium, ISVC 2020, San Diego, CA, USA, October 5–7, 2020, Proceedings, Part II 15*. Springer, 2020, pp. 207–222. **2**
- [17] A. Milioto, I. Vizzo, J. Behley, and C. Stachniss, "Rangenet++: Fast and accurate lidar semantic segmentation," in *2019 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE, 2019, pp. 4213–4220. **2**
- [18] C. Xu, B. Wu, Z. Wang, W. Zhan, P. Vajda, K. Keutzer, and M. Tomizuka, "Squeezesegv3: Spatially-adaptive convolution for efficient point-cloud segmentation," in *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXVIII 16*. Springer, 2020, pp. 1–19. **2**
- [19] Y. Zhang, Z. Zhou, P. David, X. Yue, Z. Xi, B. Gong, and H. Foroosh, "Polarnet: An improved grid representation for online lidar point clouds semantic segmentation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 9601–9610. **2**
- [20] A. Ando, S. Gidaris, A. Bursuc, G. Puy, A. Boulch, and R. Marlet, "Rangevit: Towards vision transformers for 3d semantic segmentation in autonomous driving," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 5240–5250. **2**
- [21] L. Kong, Y. Liu, R. Chen, Y. Ma, X. Zhu, Y. Li, Y. Hou, Y. Qiao, and Z. Liu, "Rethinking range view representation for lidar segmentation," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 228–240. **2**
- [22] H. Tang, Z. Liu, S. Zhao, Y. Lin, J. Lin, H. Wang, and S. Han, "Searching efficient 3d architectures with sparse point-voxel convolution," in *European conference on computer vision*. Springer, 2020, pp. 685–702. **2, 6**
- [23] H. Zhou, X. Zhu, X. Song, Y. Ma, Z. Wang, H. Li, and D. Lin, "Cylinder3d: An effective 3d framework for driving-scene lidar semantic segmentation," *arXiv preprint arXiv:2008.01550*, 2020. **2, 4, 5**
- [24] H. Thomas, C. R. Qi, J.-E. Deschaud, B. Marcotegui, F. Goulette, and L. J. Guibas, "Kpconv: Flexible and deformable convolution for point clouds," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 6411–6420. **2**
- [25] V. E. Liong, T. N. T. Nguyen, S. Widjaja, D. Sharma, and Z. J. Chong, "Amvnet: Assertion-based multi-view fusion network for lidar semantic segmentation," *arXiv preprint arXiv:2012.04934*, 2020. **2**
- [26] J. Xu, R. Zhang, J. Dou, Y. Zhu, J. Sun, and S. Pu, "Rpvnet: A deep and efficient range-point-voxel fusion network for lidar point cloud segmentation," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 16 024–16 033. **2**
- [27] L. Zhao, H. Zhou, X. Zhu, X. Song, H. Li, and W. Tao, "Lif-seg: Lidar and camera image fusion for 3d lidar semantic segmentation," *IEEE Transactions on Multimedia*, 2023. **2**
- [28] L. Li, H. P. Shum, and T. P. Breckon, "Less is more: Reducing task and model complexity for 3d point cloud semantic segmentation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 9361–9371. **2, 4**
- [29] Z. Ke, D. Qiu, K. Li, Q. Yan, and R. W. Lau, "Guided collaborative training for pixel-wise semi-supervised learning," in *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIII 16*. Springer, 2020, pp. 429–445. **2**
- [30] G. French, T. Aila, S. Laine, M. Mackiewicz, and G. Finlayson, "Semi-supervised semantic segmentation needs strong, high-dimensional perturbations," 2019. **2**
- [31] X. Chen, Y. Yuan, G. Zeng, and J. Wang, "Semi-supervised semantic segmentation with cross pseudo supervision," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 2613–2622. **2, 4**
- [32] N. Souly, C. Spampinato, and M. Shah, "Semi supervised semantic segmentation using generative adversarial network," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 5688–5696. **2**
- [33] W.-C. Hung, Y.-H. Tsai, Y.-T. Liou, Y.-Y. Lin, and M.-H. Yang, "Adversarial learning for semi-supervised semantic segmentation," *arXiv preprint arXiv:1802.07934*, 2018. **2**
- [34] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial networks," *Communications of the ACM*, vol. 63, no. 11, pp. 139–144, 2020. **2**
- [35] L. Jiang, S. Shi, Z. Tian, X. Lai, S. Liu, C.-W. Fu, and J. Jia, "Guided point contrastive learning for semi-supervised point cloud semantic segmentation," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 6423–6432. **2, 4**
- [36] P. Wu, S. Chen, and D. N. Metaxas, "Motionnet: Joint perception and motion prediction for autonomous driving based on bird's eye view maps," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 11 385–11 395. **2**
- [37] M. Ye, T. Cao, and Q. Chen, "Tpcn: Temporal point cloud networks for motion forecasting," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 11 318–11 327. **2**
- [38] R. Huang, W. Zhang, A. Kundu, C. Pantofaru, D. A. Ross, T. Funkhouser, and A. Fathi, "An lstm approach to temporal 3d object detection in lidar point clouds," in *Computer Vision—ECCV*

- 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, *Proceedings, Part XVIII 16*. Springer, 2020, pp. 266–282. 2
- [39] S. McCrae and A. Zakhor, “3d object detection for autonomous driving using temporal lidar data,” in *2020 IEEE International Conference on Image Processing (ICIP)*. IEEE, 2020, pp. 2661–2665. 2
- [40] Z. Yuan, X. Song, L. Bai, Z. Wang, and W. Ouyang, “Temporal-channel transformer for 3d lidar-based video object detection for autonomous driving,” *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 32, no. 4, pp. 2068–2078, 2021. 2
- [41] K. C. Kumar and S. Al-Stouhi, “Real-time spatial-temporal context approach for 3d object detection using lidar,” in *VEHITS*, 2020, pp. 432–439. 2
- [42] C. Park, P. Moghadam, S. Kim, S. Sridharan, and C. Fookes, “Spatiotemporal camera-lidar calibration: A targetless and structureless approach,” *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 1556–1563, 2020. 2
- [43] X. Liu, C. R. Qi, and L. J. Guibas, “Flownet3d: Learning scene flow in 3d point clouds,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 529–537. 2
- [44] G. Puy, A. Boulch, and R. Marlet, “Flot: Scene flow on point clouds guided by optimal transport,” in *European conference on computer vision*. Springer, 2020, pp. 527–544. 2
- [45] R. Rishav, R. Battraw, R. Schuster, O. Wasenmüller, and D. Stricker, “DeepLidarFlow: A deep learning architecture for scene flow estimation using monocular camera and sparse lidar,” in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 10460–10467. 2
- [46] L. Nunes, L. Wiesmann, R. Marcuzzi, X. Chen, J. Behley, and C. Stachniss, “Temporal consistent 3d lidar representation learning for semantic perception in autonomous driving,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 5217–5228. 2, 3
- [47] Y. Wu, T. Zhang, W. Ke, S. Süssstrunk, and M. Salzmann, “Spatiotemporal self-supervised learning for point clouds in the wild,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 5251–5260. 2
- [48] R. Chen, Y. Liu, L. Kong, X. Zhu, Y. Ma, Y. Li, Y. Hou, Y. Qiao, and W. Wang, “Clip2scene: Towards label-efficient 3d scene understanding by clip,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 7020–7030. 2
- [49] J. Sun, Y. Dai, X. Zhang, J. Xu, R. Ai, W. Gu, and X. Chen, “Efficient spatial-temporal information fusion for lidar-based 3d moving object segmentation,” in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 11456–11463. 2
- [50] M. Aygun, A. Osep, M. Weber, M. Maximov, C. Stachniss, J. Behley, and L. Leal-Taixé, “4d panoptic lidar segmentation,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 5527–5537. 2
- [51] L. Kreuzberg, I. E. Zulfikar, S. Mahadevan, F. Engelmann, and B. Leibe, “4d-stop: Panoptic segmentation of 4d lidar using spatio-temporal object proposal generation and aggregation,” in *European Conference on Computer Vision*. Springer, 2022, pp. 537–553. 2
- [52] C. Choy, J. Gwak, and S. Savarese, “4d spatio-temporal convnets: Minkowski convolutional neural networks,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 3075–3084. 2, 6
- [53] H. Shi, G. Lin, H. Wang, T.-Y. Hung, and Z. Wang, “Spsequencenet: Semantic segmentation network on 4d point clouds,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 4574–4583. 2
- [54] P. Schutt, R. A. Rosu, and S. Behnke, “Abstract flow for temporal semantic segmentation on the permutohedral lattice,” in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 5139–5145. 2
- [55] X. Chen, S. Xu, X. Zou, T. Cao, D.-Y. Yeung, and L. Fang, “Svqnet: Sparse voxel-adjacent query network for 4d spatio-temporal lidar semantic segmentation,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 8569–8578. 2
- [56] M. Everingham, S. A. Eslami, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman, “The pascal visual object classes challenge: A retrospective,” *International Journal of Computer Vision*, vol. 111, no. 1, pp. 98–136, 2015. 3
- [57] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, “Microsoft coco: Common objects in context,” in *Proceedings of the European Conference on Computer Vision*. Springer, 2014, pp. 740–755. 3
- [58] B. Zhou, H. Zhao, X. Puig, S. Fidler, A. Barriuso, and A. Torralba, “Scene parsing through ade20k dataset,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 633–641. 3
- [59] J. Behley, M. Garbade, A. Milioto, J. Quenzel, S. Behnke, C. Stachniss, and J. Gall, “Semantickitti: A dataset for semantic scene understanding of lidar sequences,” in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 9297–9307. 3, 4, 5
- [60] I. et al., “KISS-ICP: In Defense of Point-to-Point ICP – Simple, Accurate, and Robust Registration If Done the Right Way,” *RA-L*, vol. 8, no. 2, pp. 1029–1036, 2023. 3
- [61] J. L. Bentley, “Multidimensional binary search trees used for associative searching,” *Communications of the ACM*, vol. 18, no. 9, pp. 509–517, 1975. 3
- [62] Y. Zou, Z. Yu, B. Kumar, and J. Wang, “Unsupervised domain adaptation for semantic segmentation via class-balanced self-training,” in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 289–305. 4
- [63] S. Ren, K. He, R. Girshick, and J. Sun, “Faster r-cnn: Towards real-time object detection with region proposal networks,” *Advances in neural information processing systems*, vol. 28, 2015. 3
- [64] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, “DeepLab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 40, no. 4, pp. 834–848, 2017. 3
- [65] Y. Bai, E. Yang, B. Han, Y. Yang, J. Li, Y. Mao, G. Niu, and T. Liu, “Understanding and improving early stopping for learning with noisy labels,” *Advances in Neural Information Processing Systems*, vol. 34, pp. 24392–24403, 2021. 4
- [66] B. Kang, S. Xie, M. Rohrbach, Z. Yan, A. Gordo, J. Feng, and Y. Kalantidis, “Decoupling representation and classifier for long-tailed recognition,” *arXiv preprint arXiv:1910.09217*, 2019. 4
- [67] K. Baek, S. Lee, and H. Shim, “Learning from better supervision: Self-distillation for learning with noisy labels,” in *2022 26th International Conference on Pattern Recognition (ICPR)*. IEEE, 2022, pp. 1829–1835. 4
- [68] Y. Hou, X. Zhu, Y. Ma, C. C. Loy, and Y. Li, “Point-to-voxel knowledge distillation for lidar semantic segmentation,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 8479–8488. 5
- [69] H.-X. Cheng, X.-F. Han, and G.-Q. Xiao, “Cenet: Toward concise and efficient lidar semantic segmentation for autonomous driving,” in *2022 IEEE International Conference on Multimedia and Expo (ICME)*. IEEE, 2022, pp. 01–06. 6