

HiLoTs: High-Low Temporal Sensitive Representation Learning for Semi-Supervised LiDAR Segmentation in Autonomous Driving

R.D. Lin¹, Pengcheng Weng¹, Yinqiao Wang¹, Han Ding², Jinsong Han³, Fei Wang^{1†}

¹ School of Software Engineering, Xi'an Jiaotong University, China

² School of Computer Science and Technology, Xi'an Jiaotong University, China

³ College of Computer Science and Technology, Zhejiang University, China

rdlin@stu.xjtu.edu.cn, hanjinsong@zju.edu.cn, {dinghan, feynmanw}@xjtu.edu.cn

[†]corresponding author and project lead

<https://github.com/rdlin118/HiLoTs>

Abstract

LiDAR point cloud semantic segmentation plays a crucial role in autonomous driving. In recent years, semi-supervised methods have gained popularity due to their significant reduction in annotation labor and time costs. Current semi-supervised methods typically focus on point cloud spatial distribution or consider short-term temporal representations, e.g., only two adjacent frames, often overlooking the rich long-term temporal properties inherent in autonomous driving scenarios. In driving experience, we observe that nearby objects, such as roads and vehicles, remain stable while driving, whereas distant objects exhibit greater variability in category and shape. This natural phenomenon is also captured by LiDAR, which reflects lower temporal sensitivity for nearby objects and higher sensitivity for distant ones. To leverage these characteristics, we propose HiLoTs, which learns high-temporal sensitivity and low-temporal sensitivity representations from continuous LiDAR frames. These representations are further enhanced and fused using a cross-attention mechanism. Additionally, we employ a teacher-student framework to align the representations learned by the labeled and unlabeled branches, effectively utilizing the large amounts of unlabeled data. Experimental results on the SemanticKITTI and nuScenes datasets demonstrate that our proposed HiLoTs outperforms state-of-the-art semi-supervised methods, and achieves performance close to LiDAR+Camera multimodal approaches.

1. Introduction

LiDAR point cloud semantic segmentation is crucial in autonomous driving for tasks such as obstacle avoidance [45], lane detection [13], localization and mapping [5, 14].

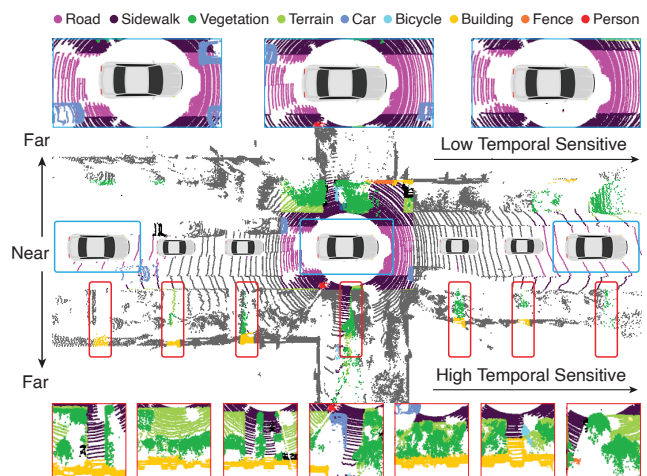


Figure 1. Different semantic classes exhibit varying degrees of sensitivity to temporal changes. Objects farther from the vehicle (e.g., vegetation, building, person, etc.) change more frequently over time, as indicated by the red box. In contrast, objects closer to the vehicle (e.g., road, sidewalk, etc.) are less sensitive to temporal changes, as shown by the blue box.

Most existing segmentation works are fully-supervised approaches [1, 11, 24, 46], which have several drawbacks. They require extensive point-wise annotations, leading to significant labor costs and time consumption. Additionally, the need for large labeled datasets limits their scalability and adaptability to new environments, highlighting the need for more efficient and scalable approaches, such as semi-supervised learning (SSL) methods that can leverage unlabeled data and reduce reliance on costly annotations.

To achieve SSL LiDAR segmentation, many works leverage the spatial distribution information of point clouds [8, 23, 27]. For example, SSPC-Net [8] introduces

an area partition method, constructing a super-point graph structure with both labeled and unlabeled point clouds for semi-supervised learning. DDSemi [23] addresses outlier issues in point-to-point SSL methods by employing a point-density-guided contrastive learning technique. In addition to utilizing the spatial distribution of point clouds, some works exploit temporal information [2, 9, 29, 33]. For instance, Aygün et al. [2] directly fuses multi-frame point clouds and inputs them into a standard encoder-decoder network to produce segmentation results. Shi et al. [33] proposes a temporal matching method, performing one-to-one matching based on differences and similarities between point clouds of two consecutive frames. Although these methods demonstrate strong performance, they tend to focus on either the spatial or the temporal characteristics of point clouds, without fully integrating both aspects.

In driving experience, we observe a phenomenon: objects closer to the vehicle, such as roads and cars, tend to have stable categories and shapes as the vehicle moves, while distant objects, such as pedestrians, guardrails, plants, and buildings, exhibit significant variations in category and shape. Surprisingly, this nature is also reflected in LiDAR point cloud data, as shown in Fig. 1, where the relevant areas are highlighted. To leverage this phenomenon, we propose HiLoTs, which consists of a High Temporal Sensitivity Flow (HTSF) and a Low Temporal Sensitivity Flow (LTSF). The HTSF focuses on regions where distant objects experience significant changes in category and shape, while the LTSF focuses on nearby regions where object categories and shapes remain relatively stable. Furthermore, the features from HTSF and LTSF are fused and interact through a cross-attention mechanism. To better represent near and far objects, we convert the point cloud into cylindrical voxels using a cylindrical voxelization network [46]. To further optimize computational efficiency, we aggregate multiple spatiotemporally neighboring cylindrical voxels, enabling a more efficient computation of HTSF and LTSF.

For semi-supervised point cloud segmentation, we adopt the mainstream Mean Teacher architecture [35], which effectively leverages a small amount of labeled LiDAR frames and a large amount of unlabeled data. In HiLoTs, the labeled LiDAR frames are fed into the student network, while the unlabeled LiDAR frames are processed through the teacher network. In each iteration, a consistency loss is computed to align the predictions made by the teacher network with those from the student network. The student network gradually updates the teacher network’s parameters, a process that can be likened to the student slowly growing into the teacher. After training, the teacher network is used for the LiDAR segmentation during inference.

We evaluate HiLoTs on two widely-used autonomous driving benchmarks, SemanticKITTI and nuScenes. Extensive results show that HiLoTs outperforms the latest

LiDAR-only semi-supervised methods and achieves performance comparable to multimodal approaches such as [6, 19], which combine LiDAR and camera data. Additionally, ablation studies confirm the effectiveness of the proposed HTSF and LTSF components, aligning with our observations from driving experience. In summary, our work claims the following main contributions:

- We observe a natural but often overlooked phenomenon in driving as shown in Fig. 1. We propose HiLoTs, designed to focus on this characteristic. We believe this design could provide valuable insight for future advancements in LiDAR segmentation tasks.
- HiLoTs includes several novel techniques, such as multi-voxel aggregation and temporal sensitivity embedding units, which are efficient and effective in LiDAR spatiotemporal representation learning.
- Experimental results show that HiLoTs surpasses the latest semi-supervised methods and achieves performance comparable to LiDAR+Camera multimodal methods.

2. Related Work

2.1. Semi-supervised LiDAR Segmentation

Various works utilize LiDAR to capture objects’ 3D representation since point clouds can accurately reflect the structural characteristics of objects [20, 22, 24, 25, 36]. However, most of the existing works are based on fully-supervised learning, while it is difficult to annotate point clouds due to the intensive labor and time costs. Consequently, many works have attempted to leverage SSL methods [8, 15, 18, 23, 26, 37, 41]. GPC [15] uses large amounts of unlabeled data as pseudo-label guidance to reduce the negative impact of intra-class negative pairs, achieving good results in both indoor and outdoor scenarios. LaserMix [18] deeply analyzes point cloud distribution priors of various objects and performed fusion between labeled and unlabeled point clouds. Nevertheless, continuous point clouds in autonomous driving usually contain rich temporal features. Current SSL methods mainly focus on spatial feature extraction, neglecting the inherent temporal representations of outdoor point clouds.

2.2. Point Cloud Representation

Mainstream outdoor point cloud representation learning can be analyzed from both spatial and temporal perspectives. For spatial representation, since 3D point clouds are unordered sets [31], it is necessary to first convert the point clouds into volumetric grids, where it can be input into neural networks as tensors. Current mainstream approaches include range mapping [1, 16], cubic-voxelization [44], pillar-based voxelization [22, 25], spherical [20] and cylindrical representation [46]. For temporal representation, commonly used methods can be divided into data-level [2, 9, 29]

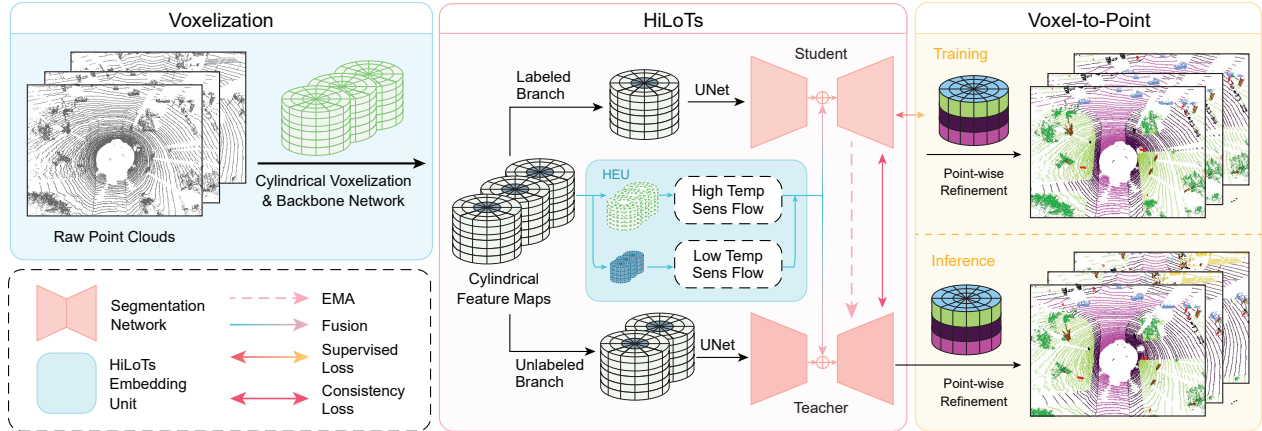


Figure 2. Our segmentation model involves three stages. During voxelization, cylindrical voxelization is applied to transform unordered points into volumetric grids, followed by a spatial feature extraction backbone. Then, HiLoTs processes the labeled and unlabeled cylindrical features through a student-teacher framework. It also integrates the attention map from HiLoTs embedding unit (HEU) to produce voxel-level segmentation maps. Finally, a point-wise refinement network is utilized to obtain point-level segmentation results.

and feature-level [33] multi-frame fusion. However, data-level fusion [2, 9] directly increases computational overhead, making hardware resources a primary bottleneck. Also, current feature-level fusion in semi-supervised learning [33] is based on two adjacent frames, which is insufficient to encode temporal changes compared to multiple frames. We observe that during driving, distant objects show frequent changes in both object categories and shapes over time, while closer objects exhibit more stable distribution. To fully capture this property, we propose HiLoTs.

3. Methods

Fig. 2 illustrates the overall pipeline of the proposed method. We first introduce the preliminary of the task. Next, we provide a detailed description of each component in the proposed HiLoTs. Finally, we present the loss functions for training the overall network.

Preliminary. Suppose $D = \{D_i \mid i = 1, \dots, N\}$ denotes the point cloud dataset, where N is the total number of sequences. Each sequence D_i includes t_i point cloud frames, $\{f_j \mid j = 1, \dots, t_i\}$. $f_j = (x_p, y_p, z_p, r_p) \in \mathbb{R}^{P \times 4}$ represents P point clouds in one frame, where (x_p, y_p, z_p) and r_p denote the Cartesian coordinates and LiDAR intensity, respectively. In the semi-supervised semantic segmentation task, given a supervised ratio of $s\%$, $D_i = [f_1, \dots, f_{t_i}]$ contains uniformly sampled $s\%$ labeled frames, and $(1 - s)\%$ unlabeled frames.

3.1. Voxelization with Cylindrical Network

Common visual modalities such as images and videos are in volumetric grids, which are efficient for neural network processing. However, point clouds are unordered sets [31]. To transform point clouds to volumetric grids for network

processing, we employ a mainstream and efficient feature extraction method, namely cylindrical voxelization [46], which can adequately represent point clouds of different densities at different ranges.

The first step of cylindrical voxelization is to convert the Cartesian coordinates of each point, represented as (x, y, z) , into the corresponding cylindrical coordinates (ρ, θ, z) , where $\rho = \sqrt{x^2 + y^2}$ and $\theta = \arctan(\frac{y}{x})$ represent the radial distance and the azimuth, respectively.

Next, since LiDAR captures dense point clouds in near areas and sparse point clouds in far areas, during voxelization, we set the size of cylindrical cells to increase with the distance, maintaining a balanced number of points in different cylindrical cells. Each cell contains the cylindrical coordinate (ρ, θ, z) of the original point cloud, with the remission r_ϵ . For multiple points projected onto the same cell, we retain the point with the closest range. The resulting volumetric grid of each frame is $x_f \in \mathbb{R}^{R \times \Theta \times H \times C}$, where $C = 4$ containing $(\rho, \theta, z, r_\epsilon)$, and R, Θ and H represent the maximum radius, azimuth, and height, respectively.

Further, we employ 3D ResNet50 [12] to extract the initial cylindrical features for the following LiDAR segmentation, $x_f \rightarrow x_f \in \mathbb{R}^{M \times d}$, where $M = R \times \Theta \times H$, $d = 256$.

3.2. HiLoTs Embedding Unit

We observe that nearby objects, mainly roads and vehicles, remain relatively stable in their spatial and categorical characteristics during driving, while distant objects display greater variability in both category and shape. Leveraging this observation, we propose HiLoTs Embedding Unit (HEU) that distinguishes between high temporally sensitive features for distant scenes and low temporally sensitive features for nearby scenes. By capturing these distinc-




Figure 3. HiLoTs Embedding Unit (HEU). The distant voxel features are passed into high temporal sensitivity flow, while voxel features in closer areas undergo low temporal sensitivity flow. The output map of HEU is fused with the bottleneck feature map from the segmentation model, further passed into the decoder.

tions, HiLoTs can more accurately interpret dynamic scenes and better adapt to the varying spatial-temporal characteristics across different distances.

As shown in Fig. 3, HEU processes two types of initial cylindrical features, which first undergo Multi-Voxel Aggregation. Voxels located at distant ranges are directed through the High Temporal Sensitivity Flow, while those within closer ranges are processed via the Low Temporal Sensitivity Flow. The resulting feature maps from each flow are then integrated, which facilitates interaction between the high- and low-sensitivity features. These enriched features subsequently serve as inputs to the segmentation network.

3.2.1 Multi-Voxel Aggregation

Recent Transformer-based models [10, 21, 34, 38, 42] have demonstrated remarkable performance in various tasks such as LiDAR semantic segmentation [1, 20]. Given that Transformers serve as a unifying foundation for multimodal modeling, we also employ a Transformer-based architecture to implement our HEU model. However, the computational complexity of the attention mechanism is $O(n^2)$, where n represents the number of tokens. With the extensive number of voxels in outdoor scenes, applying the attention mechanism to all voxels becomes computationally prohibitive. To address this, we designed a multi-voxel aggregation method (MVA), which groups neighboring voxels into super-voxels. This approach not only significantly reduces the token count, thus alleviating computational demands, but also aggregates voxel features, enhancing the model’s performance by capturing more coherent semantic information within each super-voxel. As shown in Fig. 3, the proposed MVA includes two steps. **(1) Spatial aggregation:** we aggregate M cylindrical voxels into m super voxels in spatial as:

$$F' = \text{MLP}_{\theta_1}(\text{NNGroup}(F, m)) \quad (1)$$

where $F \in \mathbb{R}^{M \times d \times t}$ represents cylindrical features across t frames; $\text{NNGroup}(\cdot)$ represents nearest neighbor grouping and MLP_{θ_1} is a multi-layer perceptron (MLP) with parameters of θ_1 . $F' \in \mathbb{R}^{m \times d \times t}$ ($m < M$) is the spatially aggregated cylindrical features.

(2) Temporal fusion: we further fuse super-voxel features across temporal dimension as:

$$V = \text{MLP}_{\theta_2}(\text{AvgPool}(F')) \quad (2)$$

where $\text{AvgPool}(\cdot)$ represents temporal average pooling and MLP_{θ_2} is an MLP with parameters of θ_2 . $V \in \mathbb{R}^{m \times d}$ is the final output of the Multi-Voxel Aggregation process.

3.2.2 High Temporal Sensitivity Flow

The High Temporal Sensitivity Flow (HTSF) is designed to process aggregated voxels at greater distances, e.g., those within the farthest 70% of the range. This flow enables point cloud representation learning by focusing on temporally dynamic and spatially variable distant regions, where object categories and shapes are more likely to fluctuate. The process can be represented as follows:

$$V_i = \text{Softmax} \left(\frac{\text{dot}(V_i W_Q, V_i W_K)}{\sqrt{d_k}} \right) V_i W_V \quad (3)$$

$$V_{i+1} = \text{MLP}_i(V_i)$$

where $V_i \in \mathbb{R}^{m \times d}$ represents the far-range voxel features of the i -th encoder layer, $W_Q \in \mathbb{R}^{d \times d_k}$, $W_K \in \mathbb{R}^{d \times d_k}$ and $W_V \in \mathbb{R}^{d \times d_v}$ denotes the weight parameters for Query, Key and Value matrices. $\text{dot}(\cdot, \cdot)$ denotes matrix multiplication. $V_{i+1} \in \mathbb{R}^{m \times d}$ represents the output of the current encoder layer, which is also the input to the next layer.

3.2.3 Low Temporal Sensitivity Flow

As discussed, the point cloud distribution of closer objects, such as roads, remains relatively stable over time in LiDAR point cloud. To capture these low temporally sensitive features accurately, we design the Low Temporal Sensitivity Flow (LTSF), depicted in the lower part of Fig. 3. In this process, all distant voxels are first excluded from the voxel set, leaving only those that represent low temporal sensitivity voxel features. These selected voxels then undergo Multi-Voxel Aggregation and self-attention, similar to the HTSF process. Further, a cross-attention is applied with the attention map generated by HTSF, enabling an interaction between the high and low temporal sensitivity features.

By stacking multiple HTSF and LTSF layers, we obtain the final output of the HiLoTs Embedding Unit. To fuse the HEU output with the bottleneck encoder feature map of the segmentation network, we introduce two learnable parameters, α_1 and α_2 , which control the feature fusion process.

$$S = \alpha_1 \cdot \text{BottleNeck}_{\text{En}}(x_f) + \alpha_2 \cdot \text{HEU}(F) \quad (4)$$

where x_f and F are described in Sec. 3.1 and Equation. 1. Finally, the decoder of the bottleneck network takes S as input and outputs the voxel-level features, denoted as $\text{BottleNeck}_{\text{De}}(S) \rightarrow S \in \mathbb{R}^{M \times d}$. We use Minkowski-UNet [9] as the bottleneck network for segmentation.

3.3. Voxel to Point Cloud Results

Since HiLoTs employs a voxelized approach to convert point clouds into volumetric grids, this inevitably leads to information loss when points from different objects are mapped to the same voxel. We first reverse voxel-level features S to point-level features, $S \rightarrow \mathbb{R}^{P \times d}$, where P represents the number of points in the frame, based on the point-to-voxel mapping table from voxelization process described in Sec. 3.1. Then, we apply a point-wise refinement network [46] to output point-level segmentation results, where S are fused with the point coordinates and LiDAR intensity.

$$\hat{y} = \text{RefineNet}(S, (x_p, y_p, z_p, r_p)) \quad (5)$$

where $\hat{y} \in \mathbb{R}^{P \times K}$ represents K object classification confidences for P points, i.e., semantic segmentation results.

3.4. Semi-supervised Learning and Loss Functions

Our HiLoTs model adopts Mean Teacher architecture [35] to achieve semi-supervised segmentation of point clouds. It includes two segmentation networks, namely the student and teacher network. The student network receives the labeled point clouds and uses the corresponding ground-truth for supervised training. Specifically, we use focal loss [28] as the supervised loss function, which addresses the class imbalance problem in point cloud semantic segmentation.

$$L_{sup} = \text{FocalLoss}(\hat{y}_s, y) \quad (6)$$

where \hat{y}_s represents the student network’s segmentation prediction and y represents the ground-truth labels.

In contrast, the teacher network receives unlabeled data, which are also fed into the student network. The teacher network’s loss function is the consistency loss between the two networks, represented as:

$$L_{con} = \|\hat{y}_s - \hat{y}_t\|_2 \quad (7)$$

where \hat{y}_t denotes the prediction from the teacher network; $\|\cdot\|_2$ represents the \mathcal{L}_2 norm. After obtaining the supervised loss and consistency loss, the final loss of the model is the weighted sum of the two:

$$L = \alpha L_{sup} + \beta L_{con} \quad (8)$$

where α and β are to balance two losses. In our experiments, we set both to 1.

The weights of the teacher network are initialized by exponential moving average (EMA) [35], transferring the student network’s trained parameters to the teacher model, which can be represented as follows:

$$W'_t = \gamma W'_{t-1} + (1 - \gamma)W_t \quad (9)$$

where W'_t and W are for teacher network and student network at the time of t , respectively. The process of mean-teacher semi-supervised learning is akin to a student gradually growing into a teacher, slowly absorbing knowledge and refining their understanding over time. Ultimately, we obtain the teacher network, which is then used for inference.

3.5. Implementation Details

HiLoTs is trained for 50,000 iterations with an early stopping strategy. AdamW optimizer [30] is used with the initial learning rate set to 1e-3. All experiments are conducted on a server with four RTX 3090 GPUs, using a batch size of 4 per GPU, resulting in a total batch size of 16. In cylindrical voxelization, we set the maximum point cloud range, azimuth, and height to (0, 50) meters, $(-\pi, \pi)$, and $(-4, 2)$ meters, respectively. The resolution of voxelized grids is set to (240, 180, 20). Considering both the model performance and the GPU memory cost, we set $t = 5$ as the temporal length in all experiments. The layers of both encoder and decoder of the Transformer are set to $N = 6$.

During training, HiLoTs takes 1 frame (central frame) and its neighboring $t - 1$ frames as input. If the central frame is labeled, it is processed by the labeled branch for supervised segmentation. Otherwise, it is routed to the unlabeled branch. Regardless of label presence, all t frames pass through the HEU module. During inference, HiLoTs takes t -frame LiDAR point clouds as input and leverages the teacher network to estimate semantic segmentation for the central frame.

4. Experiments and Analysis

4.1. Datasets and Evaluation Metric

We evaluate our approach on two widely-used autonomous driving datasets: SemanticKITTI [3] and nuScenes [4]. LiDAR point clouds from SemanticKITTI were collected with a 64-beam Velodyne HDL-64E device at 10Hz, including 19 semantic classes. Sequences 00-07 and 09-10 serve as the training set, with 19120 point cloud scenes, while sequence 08 is used as the validation set, with 4070 scenes. The nuScenes dataset, also widely used, was collected with a 32-beam Velodyne HDL-32E device at 20Hz. It contains 16 semantic classes, with the training and validation sets of 27287 and 5850 point cloud scenes, respectively.

Following prior semi-supervised segmentation works [18, 23, 26], we report the model’s performance on the validation sets of both datasets. We set the labeled ratio

Methods	Modality	SemanticKITTI				nuScenes			
		1%	10%	20%	50%	1%	10%	20%	50%
Cylinder3D [46] (2021)	Lidar	45.4	56.1	57.8	58.7	53.4*	63.4*	67.0*	71.9*
RangeViT [1] (2023)	Lidar	43.8*	53.4*	56.6*	58.8*	53.8*	64.6*	67.8*	73.1*
SphereFormer [20] (2023)	Lidar	41.2*	59.8*	60.6*	62.4*	49.5*	65.3*	69.2*	73.7*
MarS3D [29] (2023)	Lidar	44.5*	58.6*	60.2*	61.7*	51.8*	65.5*	68.4*	72.8*
GPC [15] (2021)	Lidar	41.8	49.9	58.8	59.9	-	-	-	-
PolarMix [39] (2022)	Lidar	50.1	60.9	62.0	63.8	55.6	69.6	71.0	73.8
LaserMix [18] (2023)	Lidar	50.6	60.0	61.9	62.3	55.3	69.9	71.8	73.2
LiM3D [26] (2023)	Lidar	58.4	62.2	63.1	63.6	-	-	-	-
ImageTo360 [32] (2023)	Lidar	54.5	58.6	61.4	64.2	-	-	-	-
IGNet [37] (2024)	Lidar	49.0	61.3	63.1	64.8	-	-	-	-
DDSEmi [23] (2024)	Lidar	59.3	<u>65.1</u>	<u>66.3</u>	<u>67.0</u>	58.1	70.2	74.0	<u>76.5</u>
FRNet [40] (2025)	Lidar	55.8	64.8	65.2	65.4	61.2	72.2	<u>74.6</u>	<u>75.4</u>
CyMix+IPSL [6] (2024)	Lidar+Camera	52.8	64.8	64.9	65.9	59.1	76.0	78.7	80.5
LaserMix++ [19] (2024)	Lidar+Camera	63.2	67.5	67.7	68.6	65.3	75.3	75.2	76.3
HiLoTs (Ours)	Lidar	<u>58.6</u>	65.7	66.5	67.6	<u>58.7</u>	72.2	75.2	76.9

Table 1. Performance comparisons with current state-of-the-art. Methods highlighted in yellow represents fully-supervised methods, while blue denotes semi-supervised methods. Best scores are **bolded** and the second best scores are underlined. Results in * are reproduced.

in $\{1\%, 10\%, 20\%, 50\%\}$ and use mean Intersection-over-Union (mIoU) scores across all semantic classes as the evaluation metric.

4.2. Performance Comparisons

Table 1 presents a comparison of HiLoTs with supervised methods with limited labeled data [1, 20, 29, 46], LiDAR-only semi-supervised methods [15, 23, 26, 32, 37, 39, 40], and LiDAR+Camera semi-supervised methods [6, 19].

Comparisons with Fully-supervised Methods. Cylinder3D [46], RangeViT [1], SphereFormer [20], and MarS3D [29] are fully supervised methods that do not incorporate semi-supervised techniques. We re-train both methods using the same amount of labeled data as our semi-supervised approach, with re-trained results marked with the * in Table 1. HiLoTs consistently outperforms these fully supervised approaches by a notable margin. For instance, using only 1% SemanticKITTI labeled dataset, our method achieves a mIoU of 58.6, while Cylinder3D and MarS3D reach 45.4 and 44.5, respectively. This is because the mean-teacher semi-supervised strategy can effectively leverage unlabeled data to improve segmentation performance. Although involving more labeled data can gradually enhance the performance of fully supervised methods, the high labeling cost in customized point cloud segmentation tasks makes semi-supervised methods more feasible.

Comparisons with Semi-supervised Methods. In mainstream semi-supervised point cloud segmentation for autonomous driving, existing methods can be categorized into single-modal methods that rely solely on LiDAR, and multi-modal methods that incorporate both LiDAR and camera.

As shown in Table 1, HiLoTs outperforms most current state-of-the-art LiDAR-only semi-supervised approaches, including recent methods such as FRNet [40] and DDSEmi [23], across various labeling ratios. This demonstrates the effectiveness of the HiLoTs Embedding Unit module, described in Sec. 3.2, which leverages the distinct temporal and spatial variations of point clouds at different distances, making it highly effective for point cloud segmentation in autonomous driving.

We also compare our method with two recent LiDAR + Camera multimodal semi-supervised approaches, i.e., CyMix+IPSL [6] and LaserMix++ [19]. HiLoTs achieves performance comparable to LaserMix++ on both datasets, performing slightly better than CyMix+IPSL on SemanticKITTI dataset and slightly worse on nuScenes dataset. Overall, HiLoTs demonstrates considerable effectiveness even when compared to multimodal methods. Its advantage lies in not requiring a camera or labeled RGB data from the camera, making the system more cost-effective and reducing the labeling cost.

Class-wise Performance. In terms of class-wise performance, as shown in Table 2, our method performs well on the class “parking” but shows weaker performance on the class “terrain”. This highlights HiLoTs’s capability to identify distant objects effectively. Additionally, all methods demonstrate relatively low performance on the “other-ground”, “bicycle”, and “traffic-sign”, which remains a significant challenge in semi-supervised LiDAR point cloud segmentation. Future strategies could focus on improving recognition for these classes.

Segmentation Visualization. We further present typ-

Methods	car	bicycle	motorcycle	truck	bus	person	bicyclist	road	parking	sidewalk	other-ground	building	fence	vegetation	trunk	terrain	pole	traffic-sign	mIoU
Cylinder3D	94.1	38.8	43.9	42.1	38.2	57.3	72.0	91.4	38.1	76.4	0.4	89.7	56.0	86.4	59.7	70.6	60.5	45.5	55.9
RangeViT	92.0	29.4	37.0	52.6	33.6	43.5	60.4	93.9	41.5	79.9	1.4	83.8	52.1	84.3	57.7	72.4	57.2	39.7	53.4
GPC	94.1	1.4	53.1	47.2	28.1	43.2	0.0	93.1	28.6	77.4	0.1	87.6	39.4	87.8	56.4	77.8	60.8	50.6	48.8
LaserMix	95.8	13.5	52.1	71.4	50.2	66.4	55.9	93.6	48.9	81.8	0.4	91.6	66.1	88.6	66.6	75.4	65.0	49.5	59.7
HiLoTs	96.1	33.1	68.1	63.0	63.4	71.8	81.5	94.7	60.2	83.9	13.1	92.2	70.1	89.1	70.4	76.5	67.0	52.3	65.7

Table 2. Class-wise IoU score on the validation set of SemanticKITTI under 10% supervised ratio. Note that the class ‘motorcycle’ is omitted due to its low distribution in the validation set.

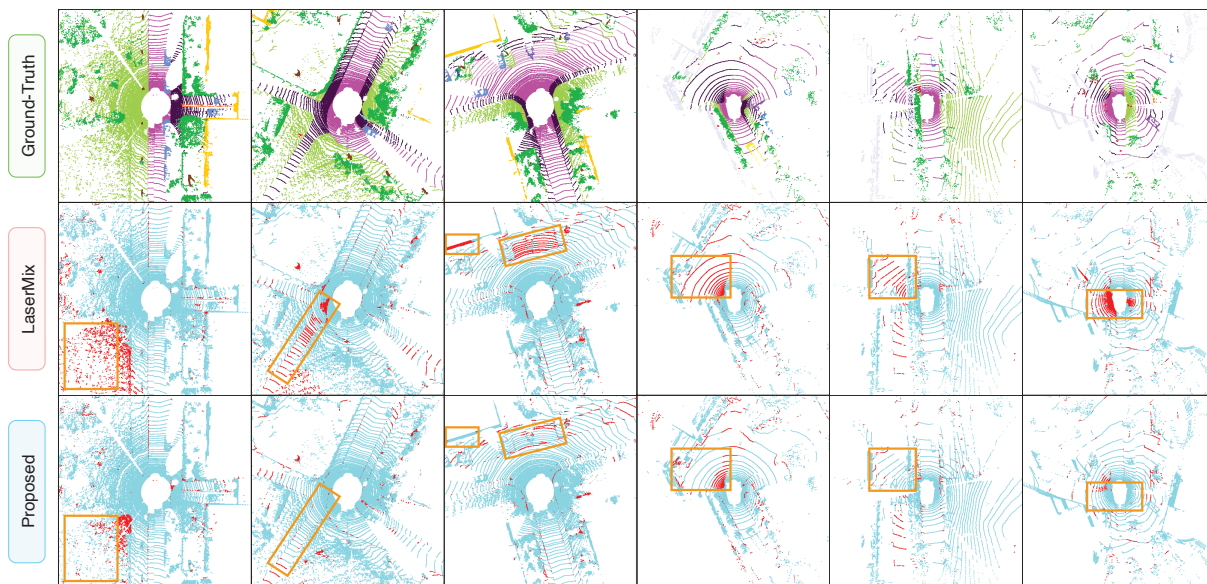


Figure 4. Error maps visualization (blue and red points are for correct predictions and incorrect predictions, respectively.). The left three columns are segmentation results from SemanticKITTI dataset, while the right three columns are from nuScenes. Our HiLoTs method shows a significant improvement in the area of distant objects.

ical examples of LiDAR point cloud semantic segmentation results at a 50% annotation ratio in Fig. 4, with orange boxes highlighting areas of significant differences between HiLoTs and LaserMix [18]. Compared to LaserMix, HiLoTs shows a distinct advantage in accurately segmenting distant areas with objects rapidly changing classes or shapes. This supports the design rationale behind our high and low temporal sensitivity flows, confirming the effectiveness of targeting regions with varying temporal dynamics.

4.3. Ablation Study

(a) **HiLoTs Embedding Unit (HEU).** HEU is the primary innovation of this work, focusing on both nearby regions with minimal class and shape changes and distant regions with more substantial variations. It mainly consists of the high temporal sensitivity flow (HTSF) and low temporal sensitivity flow (LTSF). Table 3(a) presents a performance

comparison for segmentation by showing results with HEU entirely disabled, HTSF-only, LTSF-only, and with the complete HEU module. It shows that HTSF and LTSF significantly improve LiDAR segmentation performance. Furthermore, HTSF, which targets temporally sensitive regions with variations in both class and shape, contributes more prominently to performance enhancement than LTSF. The best results are achieved when both are combined in the HEU module.

(b) **Cross-Attention and Fusion.** As shown in Fig.3, when performing cross-attention between HTSF and LTSF, HTSF is used as the key (K) and value (V), while LTSF is used as the query (Q). We also test by using HTSF as Q and LTSF as Q, as well as simple addition or concatenation of HTSF and LTSF. The experimental results, reported in Table 3(b), show that performing cross-attention between high-sensitive flow and low-sensitive flow improves perfor-

Table 3. Ablation study on the core components of HiLoTs.

Structure	SemanticKITTI			nuScenes		
	10%	20%	50%	10%	20%	50%
None	59.2	60.3	60.9	66.8	68.4	69.2
HTSF	<u>63.4</u>	<u>63.9</u>	<u>64.5</u>	<u>69.2</u>	<u>71.8</u>	<u>74.3</u>
LTSF	62.8	63.5	64.3	68.5	71.2	73.9
HEU	65.7	66.5	67.6	72.2	75.2	76.9

(a) **HEU components.** “None” denotes no HEU is applied. “HTSF” and “LTSF” denote only applying high and low temporal sensitivity flow, respectively. “HEU” represents the original method.

Fusion	SemanticKITTI			nuScenes		
	10%	20%	50%	10%	20%	50%
Add	63.6	64.7	65.3	69.8	72.8	74.6
Concat	64.3	65.1	65.9	70.3	73.5	75.2
High as Q	<u>65.4</u>	<u>66.1</u>	<u>66.9</u>	<u>71.5</u>	75.4	<u>76.5</u>
Low as Q	65.7	66.5	67.6	72.2	<u>75.2</u>	76.9

(b) **Fusion in HEU.** “Add” and “Concat” denote element-wise addition and concatenation, respectively. “High as Q” represents the self-attention map from HTSF serves as Query, and “Low as Q” denotes the original method.

Sampling	SemanticKITTI			nuScenes		
	10%	20%	50%	10%	20%	50%
Random	63.5	64.2	64.9	69.3	72.4	73.5
Density	63.9	64.8	65.4	69.5	73.2	74.1
Aggregate	65.7	66.5	67.6	72.2	75.2	76.9

(c) **Voxel down-sampling strategies in MVA.** “Random” denotes random selection, and “Density” denotes voxels with the most point cloud density. “Aggregate” denotes our approach.

Backbone	SemanticKITTI			nuScenes		
	10%	20%	50%	10%	20%	50%
Cubic	64.3	65.7	66.8	71.4	73.3	75.8
Pillar	62.4	64.1	65.2	69.4	71.5	73.2
Sphere	<u>64.8</u>	<u>65.9</u>	<u>67.2</u>	<u>71.6</u>	<u>73.8</u>	<u>76.2</u>
Cylin.	65.7	66.5	67.6	72.2	75.2	76.9

(d) **Backbone networks.** We test cubic, pillar, spherical, and cylindrical voxelization. All methods perform well.

EMA	SemanticKITTI			nuScenes		
	10%	20%	50%	10%	20%	50%
0.5	65.2	65.9	67.1	71.5	74.5	76.1
0.9	65.5	66.2	67.3	71.7	74.8	76.3
0.99	65.7	66.5	67.6	72.2	75.2	76.9
0.999	64.7	65.3	66.7	71.3	74.1	75.8

(e) **EMA ratio.** As the EMA ratio increases, segmentation performance shows an upward trend, reaching its peak at 0.99.

#	Comp.	mIoU	Fog	Snow	Beam	Echo
S _C	CENet	62.6	42.7	53.6	55.8	53.4
	FRNet	68.7	47.6	57.1	62.5	58.1
	Ours	67.8	56.2	58.0	58.5	57.9
N _{S-C}	CENet	73.3	67.0	61.6	50.0	53.3
	FRNet	79.0	69.1	69.5	68.3	58.7
	Ours	77.3	68.3	70.2	65.7	61.9

(f) **Robustness.** HiLoTs has comparable robustness with SoTA fully-supervised models in “Fog”, “Snow”, “Wet”, and “Echo” conditions.

mance by approximately 2%. This indicates that the cross-attention between high and low temporal sensitivity features allows the model to better capture both dynamic and stable object information, leading to more accurate performance.

(c) **Multi-Voxel Aggregation.** As described in Sec. 3.2.1, we apply multi-voxel aggregation among nearby voxels to reduce computation complexity in HiLoTs Embedding Unit. We also conduct two additional experiments to reduce the number of input voxels. The first method randomly selects m voxels, while the second method chooses the top- m voxels with the highest point cloud density. Since both methods inevitably discard information from the unselected voxels, our proposed method outperforms these two approaches by 2-3%, as reported in Table 3(c).

(d) **Backbone Networks.** In the voxelization step shown in Fig. 2, we can arrange LiDAR point clouds in cubic [45], pillar [22], and spherical [20] formats, and leverage corresponding backbone networks to generate feature maps. As shown in Table 3(d), all these voxelization methods perform well, indicating that HiLoTs can effectively capture LiDAR spatial characteristics at both near and far ranges if the arrangement of point clouds inherently encodes distance-related properties. We choose the cylindrical voxelization in HiLoTs since it performs the best.

(e) **EMA Ratio.** The EMA update ratio is a relatively sensitive factor that impacts segmentation performance in previous semi-supervised mean-teacher architecture [18, 26]. In contrast, as shown in Table 3(e), changes in the update ratio have less impact on the performance outcomes, indicating HiLoTs exhibits a high level of stability and robustness with respect to EMA ratio.

(f) **Performance on Out-of-distribution Datasets.** In addition to conventional segmentation evaluation, we further investigate HiLoTs’s robustness under various data perturbations using SemanticKITTI-C and nuScenes-C datasets [17, 43], and compare it with CENet [7] and FRNet [40]. As shown in Table 3(f), HiLoTs exhibits comparable robustness under three types of perturbations: severe weather conditions (Fog, Snow), external disturbances (Beam-missing), and internal sensor failures (Echo). These results demonstrate the effectiveness of high-low temporal sensitive representation learning in various conditions.

5. Conclusion and Limitation

In this paper, we propose a novel semi-supervised LiDAR point cloud segmentation method, HiLoTs, which effectively leverages temporal dynamics through the High Temporal Sensitivity Flow and Low Temporal Sensitivity Flow. By focusing on different regions with varying temporal characteristics, HiLoTs significantly improves LiDAR semantic segmentation performance, especially for distant objects with rapidly changing shapes and categories. Our experimental results, evaluated on the widely used SemanticKITTI and nuScenes datasets, demonstrate that HiLoTs outperforms both fully-supervised and state-of-the-art semi-supervised methods under limited annotations.

Since HiLoTs is specifically designed for the autonomous driving domain, where it leverages point clouds from consecutive frames, it is not suitable for general object point cloud segmentation task, such as posed in PointNet [31] and KPConv [36].

Acknowledgements: This work was supported by the National Natural Science Foundation of China under grant 62102307, U21A20462, 62372400, and 62372365, and Fundamental Research Funds for the Central Universities.

References

- [1] Angelika Ando, Spyros Gidaris, Andrei Bursuc, Gilles Puy, Alexandre Boulch, and Renaud Marlet. Rangevit: Towards vision transformers for 3d semantic segmentation in autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5240–5250, 2023. 1, 2, 4, 6
- [2] Mehmet Aygün, Aljosa Osep, Mark Weber, Maxim Maximov, Cyrill Stachniss, Jens Behley, and Laura Leal-Taixé. 4d panoptic lidar segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5527–5537, 2021. 2, 3
- [3] Jens Behley, Martin Garbade, Andres Milioto, Jan Quenzel, Sven Behnke, Cyrill Stachniss, and Jurgen Gall. Semantickitti: A dataset for semantic scene understanding of lidar sequences. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9297–9307, 2019. 5
- [4] Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multi-modal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020. 5
- [5] Xinyuan Chang, Maixuan Xue, Xinran Liu, Zheng Pan, and Xing Wei. Driving by the rules: A benchmark for integrating traffic sign regulations into vectorized hd map. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2025. 1
- [6] Yujun Chen, Xin Tan, Zhizhong Zhang, Yanyun Qu, and Yuan Xie. Beyond the label itself: Latent labels enhance semi-supervised point cloud panoptic segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1245–1253, 2024. 2, 6
- [7] Hui-Xian Cheng, Xian-Feng Han, and Guo-Qiang Xiao. Cenet: Toward concise and efficient lidar semantic segmentation for autonomous driving. In *2022 IEEE international conference on multimedia and expo (ICME)*, pages 01–06. IEEE, 2022. 8
- [8] Mingmei Cheng, Le Hui, Jin Xie, and Jian Yang. Spsc-net: Semi-supervised semantic 3d point cloud segmentation network. In *Proceedings of the AAAI conference on artificial intelligence*, pages 1140–1147, 2021. 1, 2
- [9] Christopher Choy, JunYoung Gwak, and Silvio Savarese. 4d spatio-temporal convnets: Minkowski convolutional neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3075–3084, 2019. 2, 3, 5
- [10] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. 4
- [11] Juncong Fei, Kunyu Peng, Philipp Heidenreich, Frank Bieder, and Christoph Stiller. Pillarsegnet: Pillar-based semantic grid map estimation using sparse lidar data. In *2021 IEEE intelligent vehicles symposium (IV)*, pages 838–844. IEEE, 2021. 1
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 3
- [13] Siyuan Huang, Yichen Xie, Song-Chun Zhu, and Yixin Zhu. Spatio-temporal self-supervised representation learning for 3d point clouds. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6535–6545, 2021. 1
- [14] Alok Jhaldiyal and Navendu Chaudhary. Semantic segmentation of 3d lidar data using deep learning: a review of projection-based methods. *Applied Intelligence*, 53(6): 6844–6855, 2023. 1
- [15] Li Jiang, Shaoshuai Shi, Zhuotao Tian, Xin Lai, Shu Liu, Chi-Wing Fu, and Jiaya Jia. Guided point contrastive learning for semi-supervised point cloud semantic segmentation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6423–6432, 2021. 2, 6
- [16] Lingdong Kong, Youquan Liu, Runnan Chen, Yuexin Ma, Xinge Zhu, Yikang Li, Yuenan Hou, Yu Qiao, and Ziwei Liu. Rethinking range view representation for lidar segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 228–240, 2023. 2
- [17] Lingdong Kong, Youquan Liu, Xin Li, Runnan Chen, Wenwei Zhang, Jiawei Ren, Liang Pan, Kai Chen, and Ziwei Liu. Robo3d: Towards robust and reliable 3d perception against corruptions. *arXiv preprint arXiv:2303.17597*, 2023. 8
- [18] Lingdong Kong, Jiawei Ren, Liang Pan, and Ziwei Liu. Lasermix for semi-supervised lidar semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21705–21715, 2023. 2, 5, 6, 7, 8
- [19] Lingdong Kong, Xiang Xu, Jiawei Ren, Wenwei Zhang, Liang Pan, Kai Chen, Wei Tsang Ooi, and Ziwei Liu. Multi-modal data-efficient 3d scene understanding for autonomous driving. *arXiv preprint arXiv:2405.05258*, 2024. 2, 6
- [20] Xin Lai, Yukang Chen, Fanbin Lu, Jianhui Liu, and Jiaya Jia. Spherical transformer for lidar-based 3d recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17545–17555, 2023. 2, 4, 6, 8
- [21] Bo Lan, Fei Wang, Lekun Xia, Fan Nai, Shiqiang Nie, Han Ding, and Jinsong Han. Bullydetect: Detecting school physical bullying with wi-fi and deep wavelet transformer. *IEEE Internet of Things Journal*, 2024. 4
- [22] Alex H Lang, Sourabh Vora, Holger Caesar, Lubing Zhou, Jiong Yang, and Oscar Beijbom. Pointpillars: Fast encoders for object detection from point clouds. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12697–12705, 2019. 2, 8

- [23] Jianan Li and Qiulei Dong. Density-guided semi-supervised 3d semantic segmentation with dual-space hardness sampling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3260–3269, 2024. 1, 2, 5, 6
- [24] Jiale Li, Hang Dai, Hao Han, and Yong Ding. Mseg3d: Multi-modal 3d semantic segmentation for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 21694–21704, 2023. 1, 2
- [25] Jinyu Li, Chenxu Luo, and Xiaodong Yang. Pillarnext: Rethinking network designs for 3d object detection in lidar point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17567–17576, 2023. 2
- [26] Li Li, Hubert PH Shum, and Toby 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*, pages 9361–9371, 2023. 2, 5, 6, 8
- [27] Mengtian Li, Yuan Xie, Yunhang Shen, Bo Ke, Ruizhi Qiao, Bo Ren, Shaohui Lin, and Lizhuang Ma. Hybridrc: Weakly-supervised 3d point cloud semantic segmentation via hybrid contrastive regularization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14930–14939, 2022. 1
- [28] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017. 5
- [29] Jiahui Liu, Chirui Chang, Jianhui Liu, Xiaoyang Wu, Lan Ma, and Xiaojuan Qi. Mars3d: A plug-and-play motion-aware model for semantic segmentation on multi-scan 3d point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9372–9381, 2023. 2, 6
- [30] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019. 5
- [31] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 652–660, 2017. 2, 3, 8
- [32] Laurenz Reichardt, Nikolas Ebert, and Oliver Wasenmüller. 360deg from a single camera: a few-shot approach for lidar segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1075–1083, 2023. 6
- [33] Hanyu Shi, Jiacheng Wei, Ruibo Li, Fayao Liu, and Guosheng Lin. Weakly supervised segmentation on outdoor 4d point clouds with temporal matching and spatial graph propagation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11840–11849, 2022. 2, 3
- [34] Jingang Shi, Yusi Wang, Zitong Yu, Guanxin Li, Xiaopeng Hong, Fei Wang, and Yihong Gong. Exploiting multi-scale parallel self-attention and local variation via dual-branch transformer-cnn structure for face super-resolution. *IEEE Transactions on Multimedia*, 26:2608–2620, 2023. 4
- [35] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *Advances in neural information processing systems*, 30, 2017. 2, 5
- [36] Hugues Thomas, Charles R Qi, Jean-Emmanuel Deschaud, Beatriz Marcotegui, François Goulette, and Leonidas J Guibas. Kpconv: Flexible and deformable convolution for point clouds. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6411–6420, 2019. 2, 8
- [37] Ozan Unal, Dengxin Dai, Lukas Hoyer, Yigit Baran Can, and Luc Van Gool. 2d feature distillation for weakly-and semi-supervised 3d semantic segmentation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 7336–7345, 2024. 2, 6
- [38] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 4
- [39] Aoran Xiao, Jiaxing Huang, Dayan Guan, Kaiwen Cui, Shijian Lu, and Ling Shao. Polarmix: A general data augmentation technique for lidar point clouds. *Advances in Neural Information Processing Systems*, 35:11035–11048, 2022. 6
- [40] Xiang Xu, Lingdong Kong, Hui Shuai, and Qingshan Liu. Frnet: Frustum-range networks for scalable lidar segmentation. *IEEE Transactions on Image Processing*, 2025. 6, 8
- [41] Zongyi Xu, Bo Yuan, Shanshan Zhao, Qianni Zhang, and Xinbo Gao. Hierarchical point-based active learning for semi-supervised point cloud semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 18098–18108, 2023. 2
- [42] Kangwei Yan, Fei Wang, Bo Qian, Han Ding, Jinsong Han, and Xing Wei. Person-in-wifi 3d: End-to-end multi-person 3d pose estimation with wi-fi. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 969–978, 2024. 4
- [43] Xu Yan, Chaoda Zheng, Ying Xue, Zhen Li, Shuguang Cui, and Dengxin Dai. Benchmarking the robustness of lidar semantic segmentation models. *International Journal of Computer Vision*, 132(7):2674–2697, 2024. 8
- [44] Dongqiangzi Ye, Zixiang Zhou, Weijia Chen, Yufei Xie, Yu Wang, Panqu Wang, and Hassan Foroosh. Lidarmultinet: Towards a unified multi-task network for lidar perception. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 3231–3240, 2023. 2
- [45] Yin Zhou and Oncel Tuzel. Voxelnet: End-to-end learning for point cloud based 3d object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4490–4499, 2018. 1, 8
- [46] Xinge Zhu, Hui Zhou, Tai Wang, Fangzhou Hong, Yuexin Ma, Wei Li, Hongsheng Li, and Dahua Lin. Cylindrical and asymmetrical 3d convolution networks for lidar segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9939–9948, 2021. 1, 2, 3, 5, 6