

Sparse Cross-Scale Attention Network for Efficient LiDAR Panoptic Segmentation

Shuangjie Xu^{1, 2, ‡}, Rui Wan², Maosheng Ye^{1, 2}, Xiaoyi Zou², Tongyi Cao²

¹The Hong Kong University of Science and Technology

²DEEPROUTE.AI

{sxubj, myeag}@connect.ust.hk, {ruiwan, xiaoyizou, tongyicao}@deeproute.ai

Abstract

Two major challenges of 3D LiDAR Panoptic Segmentation (PS) are that point clouds of an object are surface-aggregated and thus hard to model the long-range dependency especially for large instances, and that objects are too close to separate each other. Recent literature addresses these problems by time-consuming grouping processes such as dual-clustering, mean-shift offsets, etc., or by bird-eye-view (BEV) dense centroid representation that downplays geometry. However, the long-range geometry relationship has not been sufficiently modeled by local feature learning from the above methods. To this end, we present *SCAN*, a novel sparse cross-scale attention network to first align multi-scale sparse features with global voxel-encoded attention to capture the long-range relationship of instance context, which can boost the regression accuracy of the over-segmented large objects. For the surface-aggregated points, *SCAN* adopts a novel sparse class-agnostic representation of instance centroids, which can not only maintain the sparsity of aligned features to solve the under-segmentation on small objects, but also reduce the computation amount of the network through sparse convolution. Our method outperforms previous methods by a large margin in the SemanticKITTI dataset for the challenging 3D PS task, achieving 1st place with a real-time inference speed.

Introduction

3D scene understanding using point clouds has been an essential and challenging task for many robotics applications, including autonomous driving systems. One of the key tasks in 3D scene understanding is 3D Panoptic segmentation (PS), with two sub-tasks, semantic and instance segmentation. The semantic segmentation task aims to attach semantic information at the point level, while the instance segmentation task intends to identify individual countable objects. Point clouds have sparse, unordered, and irregular-sampled natures and aggregate only on the surface of objects. Such natures pose two main challenges for 3D PS: (1) The surface-aggregated points are far from their object centroids, leading to false segmentation of big objects (Fig. 2 a, b); (2) Closely-distributed small objects are falsely merged

Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

[‡]The work was done during an internship at DEEPROUTE.AI.

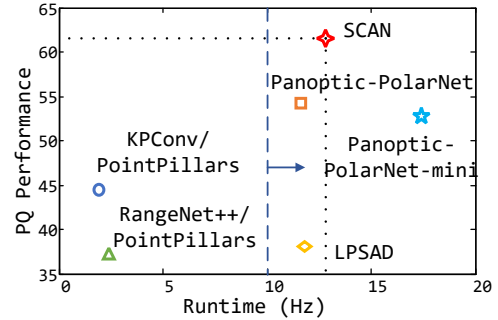


Figure 1: PQ performance vs. inference speed (Hz) in SemanticKITTI panoptic task. The blue vertical line represents the boundary of real-time runtime. The proposed *SCAN* achieves the state-of-the-art performance.

due to similarity in both physical metric space and feature space (Fig. 2 c). Therefore, how to correctly group points that belong to individual objects through efficient and effective feature learning becomes a crucial problem.

To group the surface-aggregated point clouds, many recent literature (Yang et al. 2019; Liu et al. 2020; Engelmann et al. 2020) adopts the two-stage framework to first propose bounding boxes and then segment instances. To kick over the traces of bounding boxes, most approaches (Wang et al. 2019; Lahoud et al. 2019) adopt clustering algorithms. VoteNet (Qi et al. 2019) employs voting to offer broader coverage of “good” seed points. PointGroup (Jiang et al. 2020) proposes a dual-clustering method to refine inaccurate offset predictions around object boundaries. DS-Net (Hong et al. 2020) utilizes Mean Shift with kernel functions of learned bandwidths to cluster shifted points. Such a clustering process requires high time consumption. More importantly, these methods use the center offset as the only guide to instance clustering, making it hard to capture fine-grained long-range geometry relationships and thus resulting in potential over-segmentation of large objects.

Due to the increasing demand for real-time deployment in the industry and academia, plenty of research works focus on efficient LiDAR 3D PS with range or BEV representations. LPSAD (Milioto et al. 2020) and EfficientLPS (Sirohi

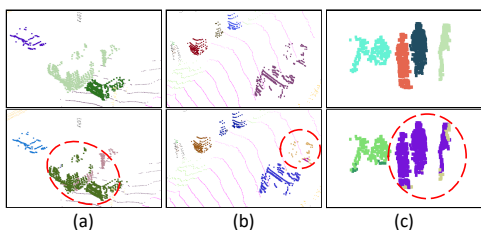


Figure 2: Challenging cases. (a) Two close cars are under-segmented. (b) Surface-aggregated points of a truck are over-segmented. (c) Three pedestrians are merged to one.

et al. 2021) transfer 3D point clouds into range images and achieve high inference speed through 2D convolutional networks. Panoptic-PolarNet (Zhou, Zhang, and Foroosh 2021) uses polar BEV grids to circumvent the issue of occlusion among instances. The above two methods compress the data from 3D to dense 2D representations, leading to the loss of original 3D correlation and the waste of most computation in the empty grids. Besides, convolution layers applied on dense representations spread information to the invalid grids, making false centroid predictions on occlusion instances and causing under-segmented issues.

Recent works have made great achievements with the sparse voxel architecture (Cheng et al. 2021; Tang et al. 2020), showing the significance of sparsity to point clouds. The widely-used sparse convolutions (Graham, Engelcke, and Van Der Maaten 2018) focus on the valid voxels to reduce computation and to avoid dilating the sparse features to invalid voxels. However, the internal association among voxels is difficult to capture with sparse convolutions, especially when the kernel size fails to cover voxels with long-range intervals. Such a “long-range” effect is not obvious in the semantic task while critical to the instance task. This can be mitigated by projecting sparse features to the 2D dense representation with 2D convolutions diffusing information, which is, however, neither efficient nor effective. To this end, we propose our efficient sparse cross-scale attention network (*SCAN*) that directly models the long-range relationship by a *cross-scale global attention* module. This bottom-up attention mechanism aggregates low-scale, geometrically strong features with high-scale, geometrically weak features, tackling the surface-aggregated problem with multi-scale internal voxel dependency. Besides, we propose *multi-scale sparse supervision* to provide fine-grained features for the attention module.

We also explore the under-segmented issue for occlusion among instances. Instead of timing-consuming clustering-based approach, we follow recent literature (Zhou, Zhang, and Foroosh 2021; Cheng et al. 2020) to use the BEV centroid representation. Compared with the existing 3D sparse and 2D dense centroid representations, we propose the *BEV sparse distribution* as our instance centroid prediction for the first time, which keeps the sparsity and the aligned long-range geometry relationships while ensures efficiency at the same time. Experimental results validate that our *SCAN* is effective and efficient: our method achieves the best perfor-

mance among the published papers with a real-time speed. Our main contributions are summarized as follows:

- We present *SCAN*, a novel sparse cross-scale attention network to first address the surface-aggregated problem by our *cross-scale global attention* module that directly models long-range dependencies of sparse voxels.
- We propose *multi-scale sparse supervision* to obtain fine-grained features for the cross-scale attention.
- We propose *BEV sparse distribution* for centroid prediction for the first time, which boosts performance for instance occlusion and ensures time efficiency.
- Our method achieves the best performance among the published papers in the SemanticKITTI dataset for the challenging 3D PS task with a real-time inference speed.

Related Work

This section briefly summarizes recent research related to our work, including 3D semantic segmentation, 3D panoptic segmentation, an attention mechanism.

3D Semantic Segmentation. Point-based methods (Qi et al. 2017a,b; Li et al. 2018; Hu et al. 2020; Thomas et al. 2019) take raw point clouds as input. They usually sample key points and rely on set abstraction to aggregate local features. However, sampling leads to information loss, and set abstraction is computationally costly. To save computational cost, regular representations, including 3d voxels and 2d grids, polar and cylinder grids, and range images (Zhou and Tuzel 2018; Zhang et al. 2020; Zhu et al. 2020b; Milioto et al. 2019; Xu et al. 2020) are used to organize sparse points. Recently, hybrid methods (Tang et al. 2020; Xu et al. 2021; Ye et al. 2021) that combine multiple representations are proposed to integrate the advantages of both fine-grained point-wise features and effective feature aggregation of regular representations. Sparse convolution (Graham 2015; Graham, Engelcke, and Van Der Maaten 2018) is also widely used to restrict convolution output only in the active regions, accelerating the volumetric convolution and enabling larger model size.

3D Panoptic Segmentation. Current 3D panoptic segmentation methods usually consist of a semantic branch and an instance branch. LPSAD (Milioto et al. 2020) obtains instances by clustering shifted points in range images. DS-Net (Hong et al. 2021) adaptively and iteratively clusters the learned point-wise centers. Based on KPConv (Thomas et al. 2019), PanosterK (Gasparini et al. 2021) introduces the impurity loss and fragmentation loss to train semantic and instance branches jointly, and outputs instance ids straightway from the network. 4D Panoptic (Aygün et al. 2021) proposes a density-based clustering as the initialization and refines it based on the temporal-spatial consistency. Panoptic-PolarNet (Zhou, Zhang, and Foroosh 2021) proposes the 2d dense center heatmap and instance offsets heads for proposal-free instance regression. Major voting is widely adopted as the post-processing to unify the final predictions. **Attention Mechanism.** Attention is defined as the weighted sum of features at multiple positions. SENet (Hu, Shen, and Sun 2018), CBAM (Woo et al. 2018), and non-local

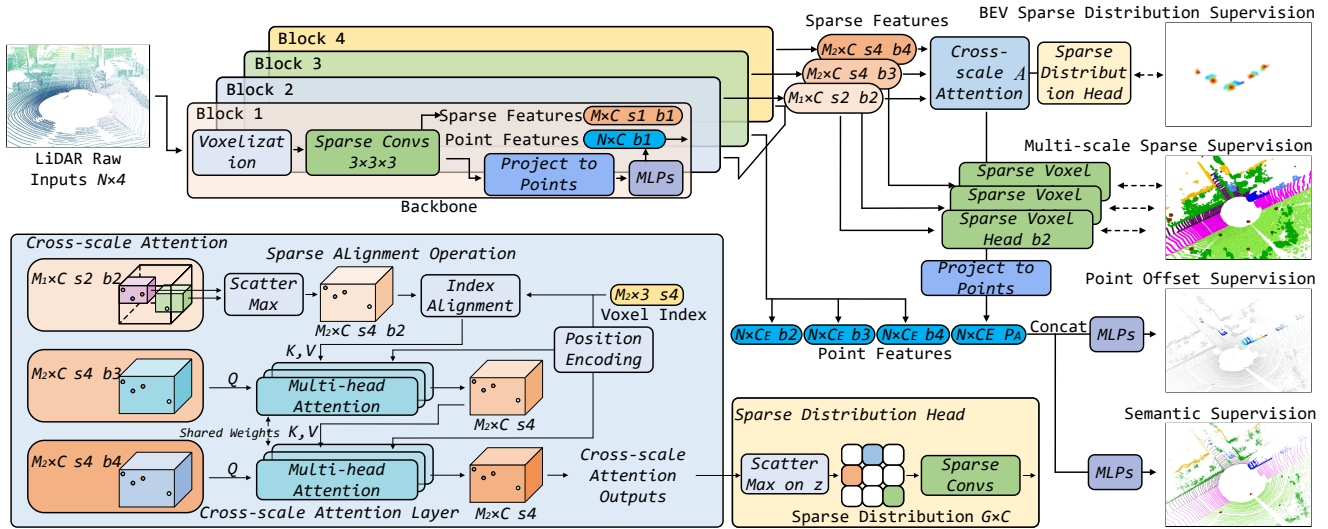


Figure 3: The overall network. Each block of network backbone encodes a sparse voxel feature and a point-wise feature. The point-wise feature is propagated into the next block. Voxel features from the last three blocks are aggregated by the proposed *cross-scale attention* module to acquire the *BEV sparse centroid distribution*. Besides, we apply *multi-scale sparse supervision* on voxel features directly for superior feature learning. Point features from the last three blocks and attention features are concatenated for point-wise offsets from centroids and semantic predictions.

operation (Wang et al. 2018) have been proposed to exploit channel-wise and spatial attention to adaptively refine features and capture long-range dependencies, which are effective plug-ins for various computer vision tasks, including classification, detection, and segmentation. Transformer (Vaswani et al. 2017) and DETR (Carion et al. 2020) are the pioneers that rely entirely on attention mechanisms to draw global dependencies between inputs and outputs by stacking self-attention and cross-attention modules. Transformer architectures have also been applied to some indoor point cloud tasks, including classification and segmentation (Engel, Belagiannis, and Dietmayer 2020; Zhao et al. 2020; Pan et al. 2021; Guo et al. 2021). However, their memory and computation complexity boost at vast key element numbers and thus hinder the model scalability. Therefore, the variants including deformable attention modules (Zhu et al. 2020a) and linear attention modules (Katharopoulos et al. 2020; Choromanski et al. 2020) have been proposed to reduce the computation by utilizing deformable convolutions and matrix properties, respectively.

Approach

The overall network is illustrated in Fig. 3. The raw point clouds are fed to our network backbone that is composed of four blocks. Each block takes a point-wise feature as the input, and outputs a point-wise feature and sparse voxel feature. The point-wise feature is propagated into the next block. Voxel features from the last three blocks are aggregated by the proposed *cross-scale attention* module (Sec. 1) to acquire the *BEV sparse centroid distribution* (Sec. 2). Moreover, we apply *multi-scale sparse supervision* (Sec. 2) on voxel features directly for superior feature learning. Addition network details are described in Sec. 3. Besides, point

features from the last three blocks and attention features are concatenated for point-wise offsets from centroids and semantic predictions.

Backbone Architecture. The raw LiDAR inputs $P \in \mathbb{R}^{N \times 4}$ (xyz and intensity) are first fed into the network backbone that is composed of four blocks. Each block $b_n, n = 1, 2, 3, 4$ first voxelizes the input point features P^{b_n} under voxel size s and encodes them to sparse features $R^{b_n} = \{I^{b_n} \in \mathbb{R}^{M \times 3}, F^{b_n} \in \mathbb{R}^{M \times C}\}$ by Submanifold Sparse Convolutional (SSC) layers (Graham, Engelcke, and Van Der Maaten 2018), where I^{b_n} is the coordinate indexes of valid voxels in the form of $x' = \lfloor x/s \rfloor, y' = \lfloor y/s \rfloor, z' = \lfloor z/s \rfloor$ with floor operation $\lfloor \cdot \rfloor$, and F^{b_n} denotes feature tensors corresponding to I^{b_n} . M denotes valid voxel numbers under current voxel size and C is the channel number of the tensor. Then sparse features are projected back to point features that flow to the next block by Multi-layer Perceptions (MLPs). The voxel size of block $b_{1 \rightarrow 4}$ is set to $s, 2s, 4s$ and $4s$ respectively, where s denotes the size of voxels in x, y, z -axis measured in meters. Inspired by great progress achieved by (Ye et al. 2021; Tang et al. 2020), we recognize the importance of multiple representation learning that helps extract better context information. Based on DRINet (Ye et al. 2021), we utilize both point-wise features and sparse voxel features. The point-wise features from multiple blocks are fused to generate point-wise offset with its corresponding semantic prediction, and voxel-wise features can provide instance-level prediction and semantic prediction in sparse format with our proposed cross scale attention.

Task Abstract. The 3D panoptic segmentation task is abstracted into three sub-tasks inspired by (Zhou, Zhang, and Foroosh 2021): the point-wise semantic predictions, BEV centroid distribution and centroid-related point offsets,

which is a preferred pipeline to take advantage of the voxel-wise features for the long-range relationship acquisition. In this work, we abstract the pipeline with four heads: 1) BEV Sparse Distribution head for instance heatmap prediction; 2) Multi-scale Sparse head for auxiliary supervision; 3) Point Offset head and 4) Point-wise Semantic head. For the first two heads, we use sparse voxel features from multi-block R^{b_n} with an attention feature A , and for the last two heads, we use point-wise features P^{b_n} , P^A projected from corresponding sparse features, as shown in Fig. 3.

The *cross-scale attention* module takes R^{b_n} under multiple scales $2s, 4s$ and multiple levels b_2, b_3, b_4 as inputs, outputting aligned and fused sparse features A with internal association among voxels. Furthermore, the additional sparse voxel semantic predictions from $b_{2 \rightarrow 4}$ are supervised only in training to instruct superior features for feeding into the *cross-scale attention*. To keep the internal association, the *BEV sparse centroid distribution* is adopted instead of the dense version in previous literature.

Cross-scale Global Attention

In this module, we first align the sparse features from different scales by the proposed *sparse alignment operation*. Then the multi-scale features with aligned coordinate encoding are fed into the *cross-scale global attention layer* that scores the relevance between each pair of voxels and captures the global geometry relationships between the source and the target sparse features. Furthermore, we propose the *multi-scale sparse supervision* to enhance the sparse features as the inputs to the cross-scale attention.

Sparse Alignment Operation. Shown as Fig. 3, sparse features from different voxel scales have different voxel coordinates. Therefore, we propose the sparse alignment operation (SA) to align R^{b_2} at scale $2s$ with R^{b_3}, R^{b_4} at scale $4s$:

$$\begin{aligned} \bar{I}^{b_2} &= \lfloor I^{b_2} / (4s/2s) \rfloor, I^a = \text{Unique}(\bar{I}^{b_2}) \\ F^a &= \max_j (F^{b_2}), \forall j \in I^a \\ R^a &= \text{SA}(R_{b_2}) = \{I^a, F^a\} \end{aligned} \quad (1)$$

, where the function Unique returns the unique coordinates of input I and \max_j is over F^{b_2} whose corresponding coordinates in \bar{I}^{b_2} are j . The coordinate I^a of aligned sparse voxel feature R^a is obtained by first downscaled from $2s$ to $4s$ and then Unique is applied to merge duplicated coordinates. The aligned tensor F^a is aggregated over same voxel indexes under scale $4s$. However, the aggregated feature R^a still doesn't align with R^{b_3} voxel by voxel, because the order of valid voxels may differ between I^a and I^{b_3} . Therefore, the additional operation Rearrange named *Voxel-wise Rearrangement* is proposed to reorder I^a and F^a with the queried mask on hashed coordinates:

$$\begin{aligned} H^{b_3} &= \text{Hash}(I^{b_3}), H^a = \text{Hash}(I^a) \\ E &= \text{HashQuery}(H^{b_3}, H^a) \\ \bar{R}^a &= \{\text{Index}(I^a, E), \text{Index}(F^a, E)\} \end{aligned} \quad (2)$$

, where the j th element $E_j = i$ of index mask E indicates the index i of I^a whose hashed value $H_i^a = H_j^{b_3}$. Then we

use E to index I^a and F^a by Index function to obtain the rearranged feature \bar{R}^a whose order is the same as R^{b_3} .

Cross-scale Global Attention Layer. We propose *cross-scale global attention layer* to exploit the inherent multi-scale property, which is proved to be crucial to sparse voxel representation (Ye et al. 2021; Xu et al. 2021). Global attention has been applied for 2D computer vision and indoor point cloud tasks (Vaswani et al. 2017; Pan et al. 2021). However, the core issue of applying global attention for long-range relationships on large-scale point clouds is that it would look over all valid voxels. When the voxel scale is low, the number of sparse valid voxels is usually on the order of ten or hundred thousands, making the computation and memory unbearable. Hence, we apply the attention layer on the high-level sparse voxel features by aligning them to the same scale $4s$. To further reduce the computation, we adopt the *Generalized Kernelizable Attention* (GKA) (Choromanski et al. 2020) as our implementation.

Given an input sparse feature R^{b_n} from current block b_n , let the aligned sparse context feature R^a from $R^{b_{n-1}}$ be encoded as key and value, the cross-scale attention feature A is calculated by:

$$\begin{aligned} \bar{R}^{b_n}, \bar{R}^a &= \text{PosEnc}(R^{b_n}), \text{PosEnc}(R^a) \\ A &= \text{GKA}(\bar{R}^{b_n}, \text{MLP}(\bar{R}^a), \text{MLP}'(R^a)) \end{aligned} \quad (3)$$

, where the three inputs of GKA are *Query*, *Key* and *Value* features respectively. To model the voxel position information into *Query* and *Key* features, we employ a *3D position embedding* function PosEnc. For dimension x in I , we use commonly used embedding functions (Vaswani et al. 2017): $PE_{2i}^x = \sin(x/10000^{2i/d_{\text{model}}})$ and $PE_{2i+1}^x = \cos(x/10000^{2i/d_{\text{model}}})$, where $d_{\text{model}} = C_A/3$ and C_A denotes the attention embedding length, i denotes the i -th position along the feature channel. The same embedding function is applied on y and z . The final position embedding $PE \in \mathbb{R}^{M_{s_i} \times C_A}$ from input voxel coordinates $I \in \mathbb{R}^{M_{s_i} \times 3}$ is obtained by setting the first d_{model} channels to PE^x , the middle d_{model} channels to PE^y and the left channels to PE^z . Finally, PE is applied to *Query* and *Key*.

Shown as Fig. 3, the bottom-up cross-scale attention starts from aggregating R^{b_2} and R^{b_3} to sparse attention feature A_1 , and then A_1 under scale $4s$ is aggregated with R^{b_4} for the attention output A . After our attention module, A now contains the internal voxel relationship cross multi-scale, and is used for the following BEV sparse centroid distribution.

Multi-scale Sparse Supervision. Existing works make hard semantic labels to supervise dense voxels, which not only costs a large memory footprint but also ignores the possible diversity of point labels within the voxels. Therefore, we supervise multi-scale sparse features from block $b_{2 \rightarrow 4}$ directly with soft voxelized semantic labels, shown as ‘‘Multi-scale Sparse Supervision’’ in Fig. 3. By making statistics of semantic labels of points in each voxel, the proportion of each category is taken as the semantic label of the voxel, which constructs the sparse voxel labels S^{b_n} . After a few SSC layers on R^{b_n} , we use the above-mentioned Rearrange to align the coordinates from sparse voxel prediction to the sparse

label S^{b_n} . We employ the L1 loss to obtain the loss of the semantic voxel:

$$L_v = \sum_{n=2 \rightarrow 4} L_{L1} (\text{Rearrange} (\text{SSC} (R^{b_n})), S^{b_n}). \quad (4)$$

A ‘‘hard’’ method takes the major vote category within the voxel as the voxel’s category. On the contrary, our ‘‘soft’’ method calculates statistics of the point number for each category, where the target to regress is N class ratio values for each voxel, making it a regression task. Therefore, we chose the L1 loss instead of a classification loss.

BEV Sparse Centroid Distribution

Many previous methods model instance segmentation based on points or centroids. Since instances are spatially separable, the discretized centroid representation (Zhou, Zhang, and Foroosh 2021) is highly suitable for LiDAR instances. Therefore, we choose to use the efficient BEV sparse centroid distribution. In this section, we rethink several possible representations and their pros and cons:

BEV Dense Distribution. As early adopted in the 2D panoptic task, some works apply the dense distribution to the 3D tasks under BEV (Zhou, Zhang, and Foroosh 2021; Ge et al. 2020). The discretized BEV centroid distribution removes the z -axis degree of freedom (DOF) and can naturally apply 2D convolutions. However, the dense distribution wastes computation on invalid positions that occupy the majority of the BEV map, especially for heavy network heads. Besides, the 2D convolutions diffuse the captured sparse relationship, which is harmful to our network.

3D Sparse Distribution. The core issues of 3D sparse distribution are twofold: 1) the load of computation/memory is heavy; 2) the z -axis DOF makes the task more difficult compared with the BEV distribution. The benefit is that it can make better use of the geometry information.

BEV Sparse Distribution. According to the above rethinking, we propose to use the BEV sparse distribution to model the instance centroids in point clouds, which can maintain the sparsity and internal relationship in voxel features while keep efficiency by only computing valid BEV positions through SSC. With the attention output A , we first set all $z \rightarrow 0$ of I^A to flatten the z -axis. Then we apply Unique on the new coordinates and obtain the max features over BEV unique voxels by \max_j operator in Equ. 1, which generates the BEV sparse feature \bar{A} . With several SSC layers, the final BEV sparse distribution is obtained as $D \in \mathbb{R}^{w \times h}$ where w and h denote grid sizes in x -axis and y -axis.

Network Details

Supervision. According to our task abstract, we divide 3D panoptic task into three sub-tasks: 1) BEV sparse centroid distribution prediction D supervised with Focal loss (Lin et al. 2017) L_d ; 2) point-wise offset prediction $O \in \mathbb{R}^{N \times 2}$ that denotes the distances on x, y axes respectively between points and the corresponding centroids; 3) point-wise semantic prediction $S \in \mathbb{R}^{N \times 1}$. We supervise point-wise O with L1 loss as L_o and S with the sum of the Lovász loss (Berman, Rannen Triki, and Blaschko 2018) and Focal loss as L_s . Besides, we mask out the background points

during calculating L_o . Furthermore, we supervise the multi-scale sparse semantic prediction with L1 loss as an auxiliary loss L_v . The total loss is the sum of the above losses.

Panoptic Inference. During inference, to further obtain the centroid prediction, we first apply sparse max pooling on D and then keep the voxel coordinates with unchanged features before and after this pooling. We keep K centroids with top confidence scores as the final centroid predictions.

By the point-wise semantic predictions S , we get the *thing* points. By the predicted K centroids and point-wise offsets O , we shift each *thing* point and then assign each shifted point to its closest centroid to get clustering results. Since K is set to cover the max number of instances, some predicted centroids are assigned with no points, which are removed during inference. To further refine panoptic results, we obtain the semantic label of each centroid by majority voting within the semantic predictions of its associated points, then we relabel the outlier points in each voxel. Besides, the instance IDs of *stuff* points are set to 0.

Experiments

In this section, we investigate our method’s performance on the standard benchmark dataset SemanticKITTI (Behley, Milioto, and Stachniss 2020) and Nuscenes (Caesar et al. 2019). We compare our model with state-of-the-art methods and perform an ablation study to demonstrate the advantage of each module in *SCAN*.

SemanticKITTI. SemanticKITTI (Behley et al. 2019; Behley, Milioto, and Stachniss 2020) is a challenging dataset, proposed to provide full 360-degree point-wise labels for the large-scale LiDAR data of the KITTI Odometry Benchmark (Geiger, Lenz, and Urtasun 2012). It contains 23201 scans with 3D semantic and instance annotations for training and 20351 for testing. The test evaluation is on the official server with 11 *stuff* classes and 8 *thing* classes.

Nuscenes. The large-scale Nuscenes dataset (Caesar et al. 2019) has newly released the *panoptic segmentation* challenge. The annotations include 10 *thing* classes and 6 *stuff* classes out of total 16 semantic classes. The dataset contains 1000 scenes, including 850 scenes for training and validation and 150 scenes for testing. Since the leaderboard has not opened until this paper is submitted, we only use the training and validation set in the experiment, which have 28130 and 6019 frames, respectively.

Evaluation. To assess the semantic segmentation, we rely on the commonly-used mean intersection-over-union (mIoU) metric (Behley et al. 2019) over all classes. To measure the quality of point cloud panoptic segmentation, we adopt the standard convention (Behley, Milioto, and Stachniss 2020) with PQ, SQ and RQ.

Implementation Details

Training. We fix the voxelization space to be limited in $[[\pm 48], [\pm 48], [-3, 1.5]]$. We do global rotation along z axis in range of $[-\pi, \pi]$ degrees and flip the points along x, y , and $x + y$ axes. Each augmentation is applied independently with a probability of 50%. In addition, we set the default scale $s = [0.2, 0.2, 0.1]$ measured in metres, thus the

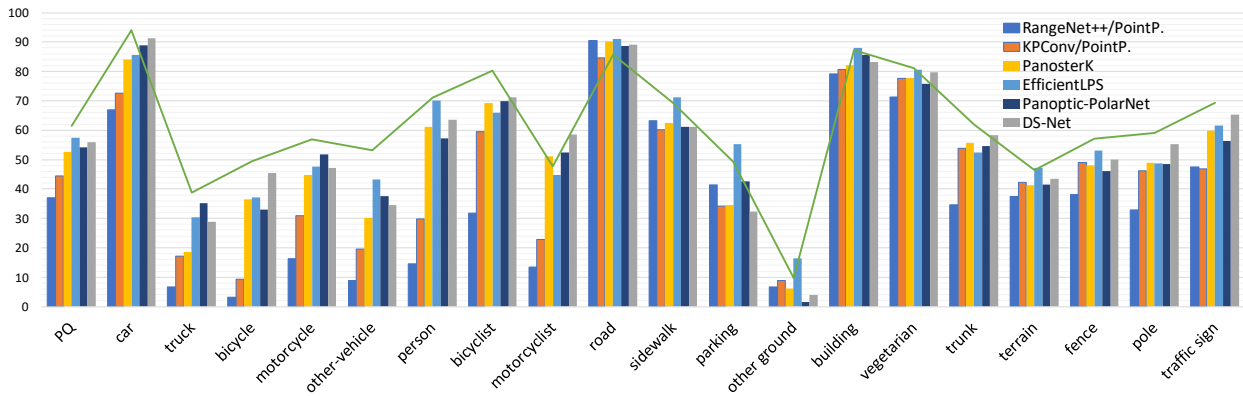


Figure 4: Per-class results of metric PQ on the SemanticKITTI test set.

Method	PQ	PQ [†]	SQ	RQ	PQ th	SQ th	RQ th	PQ st	SQ st	RQ st	mIoU	FPS
RangeNet++/PointPillars	37.1	45.9	75.9	47.0	20.2	75.2	25.2	49.3	76.5	62.8	52.4	2.4
KPConv/PointPillars	44.5	52.5	80.0	54.4	32.7	81.5	38.7	53.1	79.0	65.9	58.8	1.9
LPSAD	38.0	47.0	76.5	48.2	25.6	76.8	31.8	47.1	76.2	60.1	50.9	11.8
Panoster	52.7	59.9	80.7	64.1	49.4	83.3	58.5	55.1	78.8	68.2	59.9	-
DS-Net	55.9	62.5	82.3	66.7	55.1	87.2	62.8	56.5	78.7	69.5	61.6	-
Panoptic-PolarNet	54.1	60.7	81.4	65.0	53.3	87.2	60.6	54.8	77.2	68.1	59.5	11.6
EfficientLPS	57.4	63.2	83.0	68.7	53.1	87.8	60.5	60.5	79.5	74.6	61.4	-
GP-S3Net	60.0	69.0	82.0	72.1	65.0	86.6	74.5	56.4	78.7	70.4	70.8	-
SCAN	61.5	67.5	84.5	72.1	61.4	88.1	69.3	61.5	81.8	74.1	67.7	12.8

Table 1: Comparison of test set results on SemanticKITTI using stuff (st) and thing(th) classes. All results in [%].

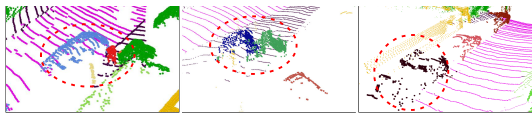


Figure 5: Results of SCAN. The first two samples demonstrate its effectiveness for challenging under-segmentation, and the last shows the robustness on over-segmentation.

$w = 120, h = 120$ for the BEV sparse centroid distribution. The feature channels are set to $C = 64$ in the network, and we configure the GKA attention by setting the number of heads 8, attention depth 2, channels of each head 16 and disabling the causality inference. We implement the network in PyTorch (Paszke et al. 2017) and train the SCAN model on 8 NVIDIA 3090 GPUs for 40 epochs with Adam (Kingma and Ba 2014) and 1-Cycle Schedule (Smith 2017). We set the batch size per GPU as 4 and the initial learning rate 0.003. The learning rate first raises tenfold before the 16-th epoch and then decays.

Inference. During inference, the auxiliary voxel semantic prediction is cut off to save the computation. The max centroid number K is set to 100 and the centroid score threshold is set to 0.1. For points out of the fixed voxelization space, we set both the semantic and instance predictions to 0.

Experiment Results

Comparison to the State-of-the-arts. Tab. 1 shows the quantitative comparison on the SemanticKITTI test set sub-

mitted to the official test server. We find that SCAN outperforms GP-S3Net (Razani et al. 2021) by 1.5% while ensures real-time inference speed, setting a new state-of-the-art performance for the LiDAR panoptic segmentation. We present our Nuscenes validation results in Table 2, and compare results with models that use the same settings of the dataset.

Inference Speed. Tab. 1 also shows the frames per second (FPS) for each approach. The reported speed for our approach is the full time including all pre- and post-processing. We investigate that our model can achieve a real-time inference speed for autonomous driving. Compared with other models that reported speed (containing preprocessing), our approach maintains a faster speed while improves accuracy.

Improvement. Many previous works are troubled by over- or under- segmentation problems. However, as shown in Fig. 4, our method handles these challenges well. In the case of large objects like *trucks* which are often over-segmented, our method outperforms other approaches by a large margin. Also, we investigate that our method is robust enough to deal with the cases of under-segmentation and small objects, especially for *person*, *other-vehicle* and *motorcyclist* classes. Such improvements demonstrate the delicate localization and effective object grouping of SCAN.

Qualitative Results. The qualitative performance on the challenging scans can be illustrated in Fig. 4. Besides, shown as Fig. 5, our approach can handle the challenging samples with over- and under- segmentation.

Method	PQ	PQ [†]	SQ	RQ	PQ th	SQ th	RQ th	PQ st	SQ st	RQ st	mIoU
DS-Net	42.5	51.0	50.3	83.6	32.5	83.1	38.3	59.2	84.4	70.3	70.7
EfficientLPS	59.2	62.8	82.9	70.7	51.8	80.6	62.7	71.5	84.3	84.1	69.4
SCAN(ours)	65.1	68.9	85.7	75.3	60.6	85.7	70.2	72.5	85.7	83.8	77.4

Table 2: Comparison of validation set results on Nuscenes using stuff (st) and thing(th) classes. All results in [%].

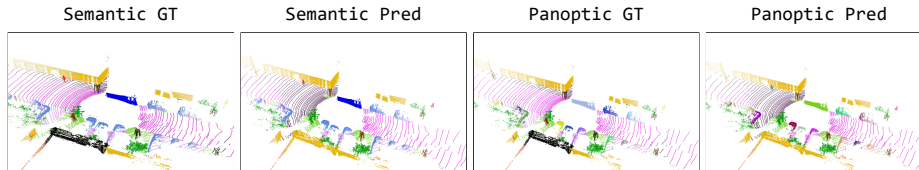


Figure 6: Qualitative results of SCAN from the SemanticKITTI validation sets.

MSS	CGA	DD	3SD	SD	PQ	mIoU	Runtime
		✓			53.3	65.8	66ms
✓		✓			53.7	66.1	66ms
	✓	✓			56.0	67.9	81ms
✓	✓	✓			56.7	68.5	81ms
✓	✓		✓		56.5	68.4	85ms
✓	✓			✓	57.2	68.9	78ms

Table 3: Ablation study of SCAN on the validation split of SemanticKITTI. CGA, MSS, SD, DD, 3SD represent for cross-scale global attention (CGA), multi-scale sparse supervision (MSS), the BEV sparse centroid distribution (SD) with its dense version (DD) and 3d sparse version (3SD).

Ablation Study

To analyze the effectiveness of different modules in SCAN, adequate ablation studies are conducted on the SemanticKITTI valid set. Starting from a baseline whose network backbone is the same with SCAN, we switch on proposed modules respectively, as shown in Tab. 3, which indicates the metrics and inference Runtime (batch size is 1) of models with NVIDIA 3090 GPU. Instead of directly summing multi-scale sparse features, the proposed cross-scale global attention CGA improves by the largest margin of 2.7% compared to the baseline in PQ with small computation increment, demonstrating the validity. The multi-scale sparse supervision MSS achieves a 0.4% promotion separately but gives 0.7% gain together with CGA, which indicates the proposed CGA can be stimulated by fine-grained features. Notice that MSS brings no influence in speed because that MSS will be cut off during inference. We also conduct experiments on different representations of centroid distribution. The performance of the default used dense version DD is 0.2% higher than the 3D sparse version 3SD. Our BEV sparse centroid distribution SD achieves the best performance with an even smaller computation.

We also conduct experiments on the cross-scale global attention module, as shown in Tab. 4 to investigate the module that brings the most gain. We first adopt various input combinations of sparse features from different blocks $b_{1 \rightarrow 4}$. Results indicate that attention on $b_{2 \rightarrow 4}$ outperforms other combinations. Only taking b_4 as input makes the module tune into the self-attention. The biggest improvement occurs be-

b_1	b_2	b_3	b_4	PE	SW	PQ	mIoU
			✓	✓		54.5	67.4
		✓	✓	✓		55.9	68.1
	✓	✓	✓	✓		56.8	68.7
✓	✓	✓	✓	✓		56.6	68.6
	✓	✓	✓			55.6	68.1
	✓	✓	✓	✓	✓	57.2	68.9

Table 4: Ablation study of the cross-scale global attention module. $b_{1 \rightarrow 4}$ stands features from which block are as inputs to the module. PE represents the 3D position encoding and SW denotes sharing weight for attention layers.

tween b_3, b_4 cross-attention and b_4 self-attention, where the former provides a foundation for the acquisition of the internal voxel relationship across sparse features. Including the b_1 feature degrades performance, for which the reason may be the insufficient learning of sparse features in the first block. Moreover, we try to remove the 3D position encoding PE and the performance has a sharp drop by 1.2%, demonstrating the importance of 3D voxel coordinates towards global attention learning. Forcing the same attention layer SW to learn the long-range relationship gives another 0.6% improvement. By sharing weights among attention layers, our cross-scale global attention module concentrates on the learning of cross-scale attention patterns.

Conclusion

We present efficient SCAN, a novel sparse cross-scale attention network to first address the surface-aggregated problem in the 3D panoptic segmentation task by modeling the long-range dependency among sparse voxel representation. The proposed cross-scale attention module introduces the attention mechanism to align and fuse multi-level and multi-scale sparse features in global instead of only stacking sparse convolution layers for local context information. Moreover, the multi-scale sparse voxel supervision is proposed to obtain fine-grained features for the cross-scale attention. In addition, we rethink centroid distributions and finally choose the BEV sparse distribution for better performance with lower computation and memory footprint. Our method achieves the state-of-the-art among published work and 1st in the SemanticKITTI challenge with a real-time runtime speed.

References

- Aygun, M.; Osep, A.; Weber, M.; Maximov, M.; Stachniss, C.; Behley, J.; and Leal-Taixé, L. 2021. 4D Panoptic LiDAR Segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 5527–5537.
- Behley, J.; Garbade, M.; Milioto, A.; Quenzel, J.; Behnke, S.; Stachniss, C.; and Gall, J. 2019. SemanticKITTI: A dataset for semantic scene understanding of lidar sequences. In *Int. Conf. Comput. Vis.*, 9297–9307.
- Behley, J.; Milioto, A.; and Stachniss, C. 2020. A Benchmark for LiDAR-based Panoptic Segmentation based on KITTI. *arXiv preprint arXiv:2003.02371*.
- Berman, M.; Rannen Triki, A.; and Blaschko, M. B. 2018. The lovasz-softmax loss: A tractable surrogate for the optimization of the intersection-over-union measure in neural networks. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 4413–4421.
- Caesar, H.; Bankiti, V.; Lang, A. H.; Vora, S.; Liong, V. E.; Xu, Q.; Krishnan, A.; Pan, Y.; Baldan, G.; and Beijbom, O. 2019. nuScenes: A multimodal dataset for autonomous driving. *arXiv preprint arXiv:1903.11027*.
- Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-end object detection with transformers. In *Eur. Conf. Comput. Vis.*, 213–229. Springer.
- Cheng, B.; Collins, M. D.; Zhu, Y.; Liu, T.; Huang, T. S.; Adam, H.; and Chen, L.-C. 2020. Panoptic-deeplab: A simple, strong, and fast baseline for bottom-up panoptic segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 12475–12485.
- Cheng, R.; Razani, R.; Taghavi, E.; Li, E.; and Liu, B. 2021. 2-S3Net: Attentive feature fusion with adaptive feature selection for sparse semantic segmentation network. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 12547–12556.
- Choromanski, K.; Likhoshesterov, V.; Dohan, D.; Song, X.; Gane, A.; Sarlos, T.; Hawkins, P.; Davis, J.; Mohiuddin, A.; Kaiser, L.; et al. 2020. Rethinking attention with performers. *arXiv preprint arXiv:2009.14794*.
- Engel, N.; Belagiannis, V.; and Dietmayer, K. 2020. Point transformer. *arXiv preprint arXiv:2011.00931*.
- Engelmann, F.; Bokeloh, M.; Fathi, A.; Leibe, B.; and Nießner, M. 2020. 3D-MPA: Multi-Proposal Aggregation for 3D Semantic Instance Segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 9031–9040.
- Gasperini, S.; Mahani, M.-A. N.; Marcos-Ramiro, A.; Navab, N.; and Tombari, F. 2021. Panoster: End-to-end Panoptic Segmentation of LiDAR Point Clouds. *IEEE Robotics and Automation Letters*.
- Ge, R.; Ding, Z.; Hu, Y.; Wang, Y.; Chen, S.; Huang, L.; and Li, Y. 2020. Afdet: Anchor free one stage 3d object detection. *arXiv preprint arXiv:2006.12671*.
- Geiger, A.; Lenz, P.; and Urtasun, R. 2012. Are we ready for autonomous driving? the kitti vision benchmark suite. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 3354–3361. IEEE.
- Graham, B. 2015. Sparse 3D convolutional neural networks. *Brit. Mach. Vis. Conf.*
- Graham, B.; Engelcke, M.; and Van Der Maaten, L. 2018. 3d semantic segmentation with submanifold sparse convolutional networks. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 9224–9232.
- Guo, M.-H.; Cai, J.-X.; Liu, Z.-N.; Mu, T.-J.; Martin, R. R.; and Hu, S.-M. 2021. PCT: Point cloud transformer. *Computational Visual Media*, 7(2): 187–199.
- Hong, F.; Zhou, H.; Zhu, X.; Li, H.; and Liu, Z. 2020. LiDAR-based Panoptic Segmentation via Dynamic Shifting Network. *arXiv preprint arXiv:2011.11964*.
- Hong, F.; Zhou, H.; Zhu, X.; Li, H.; and Liu, Z. 2021. Lidar-based panoptic segmentation via dynamic shifting network. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 13090–13099.
- Hu, J.; Shen, L.; and Sun, G. 2018. Squeeze-and-excitation networks. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 7132–7141.
- Hu, Q.; Yang, B.; Xie, L.; Rosa, S.; Guo, Y.; Wang, Z.; Trigoni, N.; and Markham, A. 2020. RandLA-Net: Efficient semantic segmentation of large-scale point clouds. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 11108–11117.
- Jiang, L.; Zhao, H.; Shi, S.; Liu, S.; Fu, C.-W.; and Jia, J. 2020. PointGroup: Dual-Set Point Grouping for 3D Instance Segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 4867–4876.
- Katharopoulos, A.; Vyas, A.; Pappas, N.; and Fleuret, F. 2020. Transformers are rnns: Fast autoregressive transformers with linear attention. In *ICML*, 5156–5165. PMLR.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Lahoud, J.; Ghanem, B.; Pollefeys, M.; and Oswald, M. R. 2019. 3d instance segmentation via multi-task metric learning. In *Int. Conf. Comput. Vis.*, 9256–9266.
- Li, Y.; Bu, R.; Sun, M.; Wu, W.; Di, X.; and Chen, B. 2018. Pointcnn: Convolution on x-transformed points. In *Adv. Neural Inform. Process. Syst.*, 820–830.
- Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollár, P. 2017. Focal loss for dense object detection. In *Int. Conf. Comput. Vis.*, 2980–2988.
- Liu, S.-H.; Yu, S.-Y.; Wu, S.-C.; Chen, H.-T.; and Liu, T.-L. 2020. Learning Gaussian Instance Segmentation in Point Clouds. *arXiv preprint arXiv:2007.09860*.
- Milioto, A.; Behley, J.; McCool, C.; and Stachniss, C. 2020. LiDAR Panoptic Segmentation for Autonomous Driving. In *IROS*.
- Milioto, A.; Vizzo, I.; Behley, J.; and Stachniss, C. 2019. RangeNet++: Fast and accurate LiDAR semantic segmentation. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 4213–4220. IEEE.
- Pan, X.; Xia, Z.; Song, S.; Li, L. E.; and Huang, G. 2021. 3d object detection with pointformer. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 7463–7472.
- Paszke, A.; Gross, S.; Chintala, S.; Chanan, G.; Yang, E.; DeVito, Z.; Lin, Z.; Desmaison, A.; Antiga, L.; and Lerer, A. 2017. Automatic differentiation in PyTorch. In *NIPS-W*.

- Qi, C. R.; Litany, O.; He, K.; and Guibas, L. J. 2019. Deep hough voting for 3d object detection in point clouds. In *Int. Conf. Comput. Vis.*, 9277–9286.
- Qi, C. R.; Su, H.; Mo, K.; and Guibas, L. J. 2017a. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 652–660.
- Qi, C. R.; Yi, L.; Su, H.; and Guibas, L. J. 2017b. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *Adv. Neural Inform. Process. Syst.*, 5099–5108.
- Razani, R.; Cheng, R.; Li, E.; Taghavi, E.; Ren, Y.; and Bingbing, L. 2021. GP-S3Net: Graph-based Panoptic Sparse Semantic Segmentation Network. *arXiv preprint arXiv:2108.08401*.
- Sirohi, K.; Mohan, R.; Büscher, D.; Burgard, W.; and Valada, A. 2021. EfficientLPS: Efficient LiDAR Panoptic Segmentation. *arXiv preprint arXiv:2102.08009*.
- Smith, L. N. 2017. Cyclical learning rates for training neural networks. In *2017 IEEE winter conference on applications of computer vision (WACV)*, 464–472. IEEE.
- Tang, H.; Liu, Z.; Zhao, S.; Lin, Y.; Lin, J.; Wang, H.; and Han, S. 2020. Searching Efficient 3D Architectures with Sparse Point-Voxel Convolution. In *Eur. Conf. Comput. Vis.*
- Thomas, H.; Qi, C. R.; Deschaud, J.-E.; Marcotegui, B.; Goulette, F.; and Guibas, L. J. 2019. Kpconv: Flexible and deformable convolution for point clouds. In *Int. Conf. Comput. Vis.*, 6411–6420.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Adv. Neural Inform. Process. Syst.*, 5998–6008.
- Wang, X.; Girshick, R.; Gupta, A.; and He, K. 2018. Non-local neural networks. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 7794–7803.
- Wang, X.; Liu, S.; Shen, X.; Shen, C.; and Jia, J. 2019. Associatively segmenting instances and semantics in point clouds. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 4096–4105.
- Woo, S.; Park, J.; Lee, J.-Y.; and Kweon, I. S. 2018. Cbam: Convolutional block attention module. In *Eur. Conf. Comput. Vis.*, 3–19.
- Xu, C.; Wu, B.; Wang, Z.; Zhan, W.; Vajda, P.; Keutzer, K.; and Tomizuka, M. 2020. Squeezesegv3: Spatially-adaptive convolution for efficient point-cloud segmentation. In *Eur. Conf. Comput. Vis.*, 1–19. Springer.
- Xu, J.; Zhang, R.; Dou, J.; Zhu, Y.; Sun, J.; and Pu, S. 2021. RPNNet: A Deep and Efficient Range-Point-Voxel Fusion Network for LiDAR Point Cloud Segmentation. *arXiv preprint arXiv:2103.12978*.
- Yang, B.; Wang, J.; Clark, R.; Hu, Q.; Wang, S.; Markham, A.; and Trigoni, N. 2019. Learning object bounding boxes for 3d instance segmentation on point clouds. In *Adv. Neural Inform. Process. Syst.*, 6740–6749.
- Ye, M.; Xu, S.; Cao, T.; and Chen, Q. 2021. DRINet: A Dual-Representation Iterative Learning Network for Point Cloud Segmentation. *arXiv:2108.04023*.
- Zhang, Y.; Zhou, Z.; David, P.; Yue, X.; Xi, Z.; Gong, B.; and Foroosh, H. 2020. PolarNet: An Improved Grid Representation for Online LiDAR Point Clouds Semantic Segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 9601–9610.
- Zhao, H.; Jiang, L.; Jia, J.; Torr, P.; and Koltun, V. 2020. Point transformer. *arXiv preprint arXiv:2012.09164*.
- Zhou, Y.; and Tuzel, O. 2018. Voxnet: End-to-end learning for point cloud based 3d object detection. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 4490–4499.
- Zhou, Z.; Zhang, Y.; and Foroosh, H. 2021. Panoptic-PolarNet: Proposal-free LiDAR Point Cloud Panoptic Segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 13194–13203.
- Zhu, X.; Su, W.; Lu, L.; Li, B.; Wang, X.; and Dai, J. 2020a. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*.
- Zhu, X.; Zhou, H.; Wang, T.; Hong, F.; Ma, Y.; Li, W.; Li, H.; and Lin, D. 2020b. Cylindrical and Asymmetrical 3D Convolution Networks for LiDAR Segmentation. *arXiv preprint arXiv:2011.10033*.