

CPSeg: Cluster-free Panoptic Segmentation of 3D LiDAR Point Clouds

Enxu Li^{*,2}, Ryan Razani^{*,1}, Yixuan Xu², and Bingbing Liu¹

¹Huawei Noah’s Ark Lab ²University of Toronto

{thomas.enxu.li, ryan.razani, richard.xu2 liu.bingbing}@huawei.com

Abstract—A fast and accurate panoptic segmentation system for LiDAR point clouds is crucial for autonomous driving vehicles to understand the surrounding objects and scenes. Existing approaches usually rely on proposals or clustering to segment foreground instances. As a result, they struggle to achieve real-time performance. In this paper, we propose a novel real-time end-to-end panoptic segmentation network for LiDAR point clouds, called CPseg. In particular, CPseg comprises a shared encoder, a dual-decoder, and a cluster-free instance segmentation head, which is able to dynamically pillarize foreground points according to the learned embedding. Then, it acquires instance labels by finding connected pillars with a pairwise embedding comparison. Thus, the conventional proposal-based or clustering-based instance segmentation is transformed into a binary segmentation problem on the pairwise embedding comparison matrix. To help the network regress instance embedding, a fast and deterministic depth completion algorithm is proposed to calculate the surface normal of each point cloud in real-time. The proposed method is benchmarked on two large-scale autonomous driving datasets: SemanticKITTI and nuScenes. Notably, extensive experimental results show that CPseg achieves state-of-the-art results among real-time approaches on both datasets.

I. INTRODUCTION

Recently, panoptic segmentation systems start to draw the attention of the autonomous driving community, since both foreground dynamic objects (i.e. the *thing*) and background static scenes (i.e. the *stuff*) can be perceived simultaneously. However, even though LiDAR is a well-concurred primary perception sensor for autonomous driving for its active sensing nature with high resolution of sensor readings, panoptic segmentation systems using LiDAR point cloud still lack sufficient research compared to image-based approaches.

GP-S3Net [1] is the current state of the art in LiDAR panoptic segmentation task. The authors proposed using a 3D sparse convolution-based UNet as a semantic backbone and a combination of HDBSCAN and GCNN to segment instances. However, GP-S3Net is computationally intensive and is thus not a real-time method. Among all published works for LiDAR panoptic segmentation, only few of them [2], [3], [4] are capable of operating in real-time (see Figure 1). A performance gap still exists when comparing real-time methods with the current state of the art. This poses a question: is it possible to build a panoptic segmentation system with a competitive PQ yet still with a fast runtime?

Given the runtime constraint for real-time systems, proposal-free methods are favorable as they are more compu-

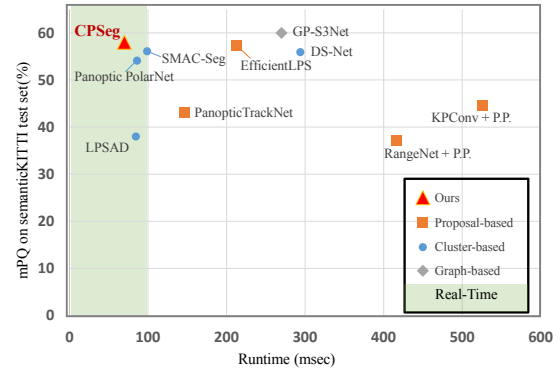


Fig. 1. PQ vs runtime on SemanticKITTI [5] test set. Methods that run in 100 ms or less are considered to be real-time (green area). Our proposed CPseg outperforms all other available real-time approaches.

tationally efficient. Current proposal-free approaches usually rely on clustering [4], [3] or graphs [1] to segment foreground objects. These methods mainly originate from 2D image processing tasks. Designing an effective 3D panoptic segmentation system based on unique characteristics of LiDAR point clouds is still an open research problem. This motivates us to find a more suitable and unique design targeting the panoptic segmentation task in the LiDAR domain. In this work, we take advantage of geometric patterns in LiDAR point clouds and present a new proposal-free and cluster-free approach to segment foreground objects in real-time. In particular, we propose a network to predict object centroid as the embedding of each point and dynamically group points with similar embedding as pillars in the sparse 2D space. Then, objects are formed by building connections of pillars.

Our main contributions can be summarized as, 1) a real-time panoptic segmentation network that is end-to-end (i.e., not relying on deterministic clustering or proposals to segment instances) and achieves state-of-the-art results without extensive post-processing, 2) a fast surface normal calculation module to aid the process of regressing foreground instance embedding with a novel deterministic depth completion algorithm, 3) a comprehensive qualitative and quantitative comparison between proposed method and existing methods on both large-scale datasets of SemanticKITTI and nuScenes, and 4) a thorough ablation analysis of how each proposed component contributes to the overall performance.

II. RELATED WORK

The panoptic segmentation task jointly optimizes semantic and instance segmentation. LiDAR-based semantic segmen-

* Indicates equal contribution.

This work was done by all authors while at Huawei Noah’s Ark Lab.

tation can be categorized into either a projection-based, a voxel-based, or a point-based method depending on the format of data being processed. Projection-based methods project a 3D point cloud into a 2D image plane either in spherical Range-View (RV) [6], Bird-Eye-View (BEV) [3], [7], or multi-view representations [8]. Voxel-based methods transform a point cloud into volumetric grids to be processed using 3D convolutions. Processing these 3D grids using 3D convolution is computationally expensive. Therefore, some methods leverage sparse convolutions to alleviate this limitation and to fully exploit sparsity of point clouds [9], [10], [11]. Point-based methods [12], however, process the unordered point cloud directly. Despite having high accuracy, these methods are inefficient and require large memory consumption. Similar to instance segmentation, panoptic segmentation can be divided into top-down (proposal-based) or bottom-up (proposal-free) methods, as elaborated below.

Proposal-based panoptic segmentation Top-down panoptic segmentation is a two-stage approach. First, foreground object proposals are generated, and subsequently, they are further processed to extract instance information that is fused with background semantic information. Mask R-CNN [13] is commonly used for instance segmentation with a light-weight stuff branch segmentation. To resolve the overlapping instance predictions by Mask R-CNN and the conflict between instance and semantic predictions, EfficientPS [14] proposes to fuse according to their confidence. Inspired by image-based methods, MOPT [15] and EfficientLPS [14] attach a semantic head to Mask R-CNN to generate panoptic segmentation. [16] proposes to use center features in regressing 3D bounding boxes and is later adopted to provide instance segmentation results by [17]. However, these top-down methods contain multiple slow sequential processes.

Proposal-free panoptic segmentation In contrast to proposal-based methods, bottom-up panoptic segmentation predicts semantic segmentation and groups the *thing* points into clusters to achieve instance segmentation. LPSAD [2] presents a shared encoder with a dual-decoder, followed by a clustering algorithm to segment instances based on the predicted semantic embedding and object centroids. Panoster [18] proposes a learnable clustering module to assign instance class labels to every point. It requires extensive post-processing steps (e.g. using DBSCAN [19] to merge nearby object predictions) to refine the predictions. DS-Net [20], however, offers a learnable dynamic shifting module to shift points in 3D space towards the object centroids. GP-S3Net [1] proposes a graph-based instance segmentation network for LiDAR-based panoptic segmentation. It uses HDBSCAN [21] to cluster raw point cloud into graph nodes, which can cause confusion in crowded scenes with multiple close-range instances. In contrast, we shift points towards their object centroids using a fully learnable method, improving runtime and performing well even in crowded scenes. Moreover, SMAC-Seg [4] introduces a Sparse Multi-directional Attention Clustering module with a repel loss to better supervise the network separating the instances. However, its usage of large kernels in SMAC introduces additional computation

cost. Their method works best with using 0.5m grid size, potentially mixing close-range instances into one grid cell. On the other hand, our dynamic approach of finding connected pillars provides a significant runtime advantage, allowing us to use much finer resolution (e.g. 0.15m grid size) during sub-sampling while retaining real-time performance.

III. PROPOSED METHOD

A. Problem Formulation

Let $(\mathbf{P}_{set}, \mathbf{L}) = \{\mathbf{p}_i, (c_i, o_i)\}_{i=1}^N$ be N unordered points of a point cloud where $\mathbf{p}_i \in \mathbb{R}^{C_0}$ is the input feature for point i , tuple $(c_i, o_i) \in \mathbb{C} \times \mathbb{O}$ is the semantic class label and instance ID label for point i . \mathbb{C} is a set of semantic class labels and \mathbb{O} is a set of instance IDs. \mathbb{C} can be further divided into \mathbb{C}_{thing} and \mathbb{C}_{stuff} , representing a set of countable foreground thing classes and a set of background stuff classes, respectively. Note that instance label o_i is only valid if $c_i \in \mathbb{C}_{thing}$. The goal is to learn a function $\mathcal{F}(\cdot, \Theta)$, parameterized by Θ , that takes input feature \mathbf{p}_i and assigns a semantic label for each point and an instance label if it is part of the foreground.

B. Network Architecture

The overview of our panoptic segmentation framework is depicted in Figure 2. We first transform the LiDAR point cloud $\mathbf{P} \in \mathbb{R}^{N \times C_0}$ with C_0 as input features (Cartesian coordinates, remission and depth) into a 2D range image with spatial dimension $H \times W$ using spherical projection similar to [2]. At the same time, we build a dense depth map, which will be utilized as a guidance for the depth completion algorithm to extract surface normal features in the following stage. Then, CPseg takes both inputs and predicts semantic and instance segmentation results in the range view (RV). When re-projecting the results to the 3D point cloud, KNN-based post-processing is utilized to refine the output, as introduced in [22]. Lastly, we fuse the results to obtain panoptic labels and use majority voting to refine the semantic segmentation results where different semantics are predicted in the same instance.

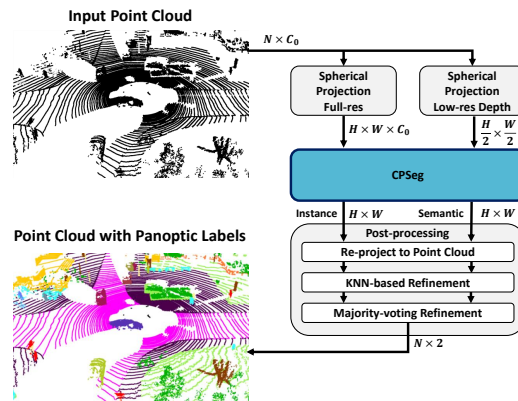


Fig. 2. Our method first projects all points into range images to predict instance embedding and semantics. Then, it uses a post-processing to aggregate consistent instances.

Our proposed model is summarized in Figure 3. It consists of three main components: (A) a dual-decoder U-Net, (B) a

surface normal calculation module, which takes the depth maps and computes normal vectors to benefit instance embedding regression, (C) a cluster-free instance segmentation head, which segments the foreground instance embedding into objects. A 2D RV representation of a LiDAR point cloud is fed into CPSEg. The output \hat{G}_1 , \hat{G}_2 of the dual-decoder U-Net are the semantic prediction and instance embedding of the projected point cloud, respectively. In particular, the semantic decoder generates $\hat{G}_1 \in \mathbb{R}^{H \times W \times C_{class}}$, where C_{class} is the number of semantic classes. With the Cartesian xy coordinates added as a prior, the instance decoder outputs $\hat{G}_2 \in \mathbb{R}^{H \times W \times 2} = P_{xy} + O_{xy}$ where P_{xy} is the xy coordinates of the point cloud in RV and O_{xy} is the output from the last block of the instance decoder. Essentially, the instance embedding in 2D space, \hat{G}_2 , is the predicted 2D location of the object centroids. We then filter using ground truth semantic labels during training or semantic predictions (i.e. $\text{argmax}(\hat{G}_1)$) during inference to obtain foreground embedding, denoted as $\hat{G}_{2,th} \in \mathbb{R}^{N \times 2}$ where N is the number of foreground points and 2 refers to the embedding in 2D space.

Basic architecture We adopt CLSA module from [4] to extract contextual features. CLSA block learns to recover local geometry in the neighbourhood, which is beneficial for RV-based methods to learn contextual information. The output of the CLSA module is then fed to a shared encoder with five residual blocks, similarly to [23], where we obtain multi-scale feature maps, J_1, J_2, J_4, J_8 (subscript indicates the stride with respect to the full resolution downsampled by AvgPool layer at the end of each encoder block). Detailed architecture of encoder and decoder blocks are shown on top left of Figure 3. Note that each convolution layer is followed by a BatchNorm and a LeakyReLU layer, and that the pooling layer and transposed convolution layer are removed in the last encoder and decoder blocks, respectively.

C. Surface Normal

Surface normal vectors provide additional geometric cues to regress instance embedding, the shifted 2D location of the object centroid. We use a Diamond Inception Module, adopted from [8], to extract geometric features from surface normals and directly fuse them with the features in the instance decoder using concatenation operation followed by convolutional layers to obtain instance embedding.

In this section, we describe a deterministic way of calculating surface normal features of the point cloud with a novel depth completion algorithm. The inputs to this module are (A) \mathbf{r}_s , sparse 2D depth map with a scale of $H \times W$, and (B) \mathbf{r}_d , dense 2D depth map with a scale of $\frac{H}{2} \times \frac{W}{2}$. The depth map is obtained by projecting the LiDAR point cloud onto a 2D map with a specified size using discretized indices from spherical transformation, as introduced by [6]. We obtain \mathbf{r}_{row} , a completed depth map with a weighted row fill using k row neighbours for every entry, as given by,

$$r_{row}[i, j] = \frac{1}{n_{row}[i, j]} \sum_{v=-\lfloor k/2 \rfloor}^{\lfloor k/2 \rfloor} r_s[i, j+v] \cdot o[i, j+v] \cdot w(v) \quad (1)$$

$$n_{row}[i, j] = \sum_{v=-\lfloor k/2 \rfloor}^{\lfloor k/2 \rfloor} o[i, j+v] \cdot w(v) \quad (2)$$

where $r_s[i, j]$ and $o[i, j]$ are the depth value and binary occupancy at i^{th} row and j^{th} column respectively. The operation $\lfloor \cdot \rfloor$ denotes the floor function. The weights $w(v) = ae^{-\frac{v^2}{2b^2}}$ are sampled from a Gaussian distribution where the center point receives the largest attention and nearby points are weighted less as they deviate away from the center. Here, a and b are hyperparameters, which are set to 1 in our method. We then re-use equations 1 and 2 to obtain \mathbf{r}_{col} using k column neighbours to fill. Next, we bilinear upsample \mathbf{r}_d to $H \times W$ to obtain a coarse but dense depth map, denoting as \mathbf{r}_u . From \mathbf{r}_u , we calculate $\frac{\partial \mathbf{r}_u}{\partial \theta}$ and $\frac{\partial \mathbf{r}_u}{\partial \phi}$ using finite difference approximation along horizontal and vertical directions where θ and ϕ are the azimuth and elevation angles for each entry. Local geometry could be interpreted from the two signals. Hence, they serve as the guidance signal to adaptively select a horizontal or vertical fill for each empty entry. For instance, when the magnitude of $\frac{\partial \mathbf{r}_u}{\partial \phi}$ is small (i.e. $\|\frac{\partial \mathbf{r}_u}{\partial \phi}\| \rightarrow 0$), it indicates the point is on a pole-like or wall-like object. Therefore, a completion using weighted average of the valid column neighbours is more desired since the change in depth in the vertical direction is relatively small. In summary, each entry in the completed depth map can be expressed as,

$$r = \begin{cases} r_s & \text{if } o = 1 \\ r_{row} & \text{if } o = 0 \ \& \ \|\frac{\partial \mathbf{r}_u}{\partial \theta}\| \leq \|\frac{\partial \mathbf{r}_u}{\partial \phi}\| \\ r_{col} & \text{if } o = 0 \ \& \ \|\frac{\partial \mathbf{r}_u}{\partial \theta}\| > \|\frac{\partial \mathbf{r}_u}{\partial \phi}\| \end{cases} \quad (3)$$

where index $[i, j]$ is omitted for brevity and o denotes occupancy. From the completed depth map, \mathbf{r} , we follow [24] to calculate the gradients $\frac{\partial \mathbf{r}}{\partial \theta}$ and $\frac{\partial \mathbf{r}}{\partial \phi}$ and transform them into Cartesian frame centered at the LiDAR sensor, obtaining $(\hat{\mathbf{n}}_x, \hat{\mathbf{n}}_y, \hat{\mathbf{n}}_z)$. Note that the purpose of the depth completion algorithm above is to ensure the neighbourhood of valid entry is smooth such that gradients are not influenced by noise.

D. Cluster-free Instance Segmentation

Given the 2D embedding of the foreground $\hat{G}_{2,th}$ from the instance decoder, the goal of the cluster-free instance segmentation module is to segment them into instances. First, we dynamically group the foreground points into pillars according to $\hat{G}_{2,th}$, their location in the 2D embedding space, such that points within grid size d are inside the same pillar (see bottom right of Figure 3). The embedding of each resulting pillar is the average embedding of the points being grouped together. Pillarized foreground embedding is denoted as $\hat{G}_p \in \mathbb{R}^{M \times 2}$, where M is the number of pillars. Next, we construct a pairwise comparison matrix $\hat{G}_m \in \mathbb{R}^{M \times M}$ to find connected pillars with each entry as, $\hat{G}_{m,ij} = f(\hat{G}_{p,i}, \hat{G}_{p,j})$, which represents the connectivity probability of pillar i and j . A large probability indicates the network is confident that the points in the two pillars belong to the same object. In order for \hat{G}_m to provide meaningful connectivity indications, we need the function, f , to follow several constraints: **1).** $f(\hat{G}_{p,i}, \hat{G}_{p,i}) = 1$, a

Method	PQ	PQ [†]	RQ	SQ	PQ Th	RQ Th	SQ Th	PQ St	RQ St	SQ St	mIoU	FPS
RangeNet++ [22] + PointPillars [27]	37.1	45.9	47.0	75.9	20.2	25.2	75.2	49.3	62.8	76.5	52.4	2.4
PanopticTrackNet [15]	43.1	50.7	53.9	78.8	28.6	35.5	80.4	53.6	67.3	77.7	52.6	6.8
KPCConv [12] + PointPillars [27]	44.5	52.5	54.4	80.0	32.7	38.7	81.5	53.1	65.9	79.0	58.8	1.9
Panoster [18]	52.7	59.9	64.1	80.7	49.4	58.5	83.3	55.1	68.2	78.8	59.9	—
DS-Net [20]	55.9	62.5	66.7	82.3	55.1	62.8	87.2	56.5	69.5	78.7	61.6	3.4*
EfficientLPS [14]	57.4	63.2	68.7	83.0	53.1	60.5	87.8	60.5	74.6	79.5	61.4	4.7
GP-S3Net [1]	60.0	69.0	72.1	82.0	65.0	74.5	86.6	56.4	70.4	78.7	70.8	3.7*
LPSAD [2]	38.0	47.0	48.2	76.5	25.6	31.8	76.8	47.1	60.1	76.2	50.9	11.8
Panoptic-PolarNet [3]	54.1	60.7	65.0	81.4	53.3	60.6	87.2	54.8	68.1	77.2	59.5	11.6
SMAC-Seg [4]	56.1	62.5	67.9	82.0	53.0	61.8	85.6	58.4	72.3	79.3	63.3	10.1
CPSeg [Ours]	56.9	63.4	68.7	82.3	54.7	63.6	86.2	58.5	72.4	79.4	62.6	14.2

TABLE I

COMPARISON ON SEMANTICKITTI [5] TEST DATASET. METRICS ARE PROVIDED IN [%] AND FPS IS IN [HZ].(*: SOURCE FROM [4])

Method	PQ	PQ [†]	RQ	SQ	PQ Th	RQ Th	SQ Th	PQ St	RQ St	SQ St	mIoU	FPS
PanopticTrackNet [15]	51.6	56.1	63.3	80.4	45.9	56.1	81.4	61.0	75.4	79.0	58.9	—
EfficientLPS [14]	62.4	66.0	74.1	83.7	57.2	68.2	83.6	71.1	84.0	83.8	66.7	—
SPVNAS [11] + CenterPoint [16]	72.2	76.0	81.2	88.5	71.7	79.4	89.7	73.2	84.2	86.4	76.9	—
Cylinder3D++ [10] + CenterPoint [16]	76.5	79.4	85.0	89.6	76.8	84.0	91.1	76.0	86.6	87.2	77.3	—
(AF)2-S3Net [9] + CenterPoint [16]	76.8	80.6	85.4	89.5	79.8	86.8	91.8	71.8	83.0	85.7	78.8	—
Panoptic-PolarNet [3]	63.6	67.1	75.1	84.3	59.0	69.8	84.3	71.3	83.9	84.2	67.0	10.1
PolarStream [7]	70.9	74.4	81.7	85.9	70.3	80.3	86.7	71.7	84.2	84.4	69.7	22.0 [‡]
CPSeg [Ours]	73.2	76.3	82.7	88.1	72.9	81.3	89.2	74.0	85.0	86.3	73.7	10.2

TABLE II

COMPARISON ON NUSCENES [17] TEST DATASET. METRICS ARE PROVIDED IN [%] AND FPS IS IN [HZ]. (‡: APPROXIMATED USING FIGURE 5 IN [7])

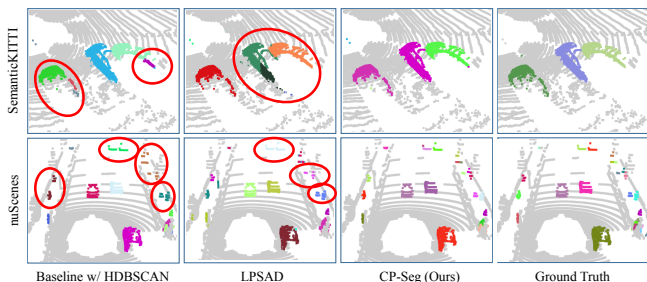


Fig. 4. Qualitative comparison of CPSEg with other methods.

150 for validation, and 150 for testing. 10 label classes are considered *things*, and 6 are considered *stuff*.

Baselines We use a dual-decoder U-Net based on SalsaNext [23] as the baseline. In particular, the two decoders generate semantic segmentation and instance embedding respectively. Then, a clustering algorithm (e.g. BFS, HDBSCAN) is added after the instance decoder to segment the objects based on the predicted embedding. To be fair in comparison, we add the CLSA Feature Extractor Module in front of the encoder to match our network design. Moreover, we implement LPSAD based on [2] as an additional baseline. Quantitative and qualitative results are compared against the proposed methods on the SemanticKITTI and nuScenes validation set.

Evaluation Metric We follow [29] to use the mean Panoptic Quality (PQ) as our main metric to evaluate and compare the results with others. In addition, we also report Recognition Quality (RQ), and Segmentation Quality (SQ). They are calculated separately on *stuff* and *thing* classes, providing PQ^{St} , SQ^{St} , RQ^{St} and PQ^{Th} , SQ^{Th} , RQ^{Th} .

A. Experimental Setup

For both datasets, we trained CPSEg end-to-end for 150 epochs using SGD optimizer and exponential-decay learning rate scheduler with initial learning rate starting at 0.01 and

a decay rate of 0.99 every epoch. A weight decay of 10^{-4} was used. The model was trained on 4 NVIDIA V100 GPUs with a batch size of 4 per GPU. The weights for the losses were set to $\beta_1 = 1.0$, $\beta_2 = 0.1$, $\beta_3 = 0.2$. We used a range image with resolution of ($H = 64$, $W = 2048$).

B. Quantitative Evaluation

In Table I and Table II, we compile the results of CPSEg compared to other models, separating the models into two groups based on their inference speed. For evaluations on SemanticKITTI test dataset (Table I), only the models in row 8-11 are known to have real-time performance, with FPS’s above 10Hz. With a PQ of 56.9% and an FPS of 14.2Hz, CPSEg achieves performances that match state-of-the-art models. More importantly, it establishes a new benchmark in PQ for real-time models, surpassing the next best real-time model, SMAC-Seg, by 0.8%. Specifically, with a 1.8% increase in RQ^{Th} and 3.1Hz improvement in FPS over SMAC-Seg, we demonstrate that CPSEg is better in recognizing foreground objects while using less computation. These improvements can be mainly attributed to the use of cluster-free instance segmentation module and the incorporation of surface normal to aid embedding regression.

On nuScenes test dataset (Table II), CPSEg again achieves competitive segmentation ability. Notably, it obtains the highest PQ for models with real-time performances, outperforming the PolarStream and Panoptic-PolarNet by 2.3% and 9.6%. Although Cylinder3D++ and (AF)2-S3Net, when combined with CenterPoint, achieve better PQ, their feature extraction based on 3D voxels is much more time-consuming compared to the 2D feature extraction in CPSEg.

C. Qualitative Evaluation

The performance of CPSEg can also be seen in Figure 4, where we compare its inference results to LPSAD, our implementation based on [2], and baseline models. For a closedup

view of a scene from SemanticKITTI dataset (row 1) where three cars are lined up closely, only CPSEg segments the instance points without errors. LPSAD identifies the car in the middle as two separate instances, whereas the baseline model produces even worse over-segmentation errors.

In a complex scene from nuScenes (row 2), with variations in instance classes and few sparse points describing each instance, correctly recognizing and distinguishing each instance is proved to be more difficult. For areas where pedestrians walk closely or where cars are positioned further away, the baseline model using HDBSCAN and LPSAD are prone to making under-segmentation errors. In such a complex scene, only CPSEg is able to segment accurately.

D. Ablation Studies

In this section, we present an extensive ablation analysis on proposed components in CPSEg. Note that all results are compared on SemanticKITTI validation set (Seq 08). First, we investigate the individual contribution of each component in the network, as shown in Table III. The cluster-free instance segmentation module is the key component, introducing 9.6% increase on the PQ (compare to the baseline with BFS) while eliminating the computation of clustering. Moreover, extracting surface normal brings another jump in PQ since the network receives guidance on regressing the embedding for each foreground object. Lastly, the model achieves the best result by incorporating binary Lovász loss in supervising the segmentation on the pairwise matrix.

Architecture	Cluster-free	3D Normal	Lovász	mPQ
Baseline w/ BFS				44.9
Baseline w/ HDBSCAN				52.7
Proposed	✓			54.5
	✓	✓		55.6
	✓		✓	55.3
	✓	✓	✓	56.2

TABLE III

ABLATION STUDY OF THE PROPOSED MODEL WITH INDIVIDUAL COMPONENTS VS BASELINE. METRICS ARE PROVIDED IN [%].

We experiment with changing α , the parameter used to map the pillar embedding to the connectivity probability. We set the threshold T to be 0.5, and pillar grid size d to be 0.15 for the experiments on α . In the first setting, α is learned from the corresponding pillar feature from the instance decoder. In particular, $\alpha = MLP(F_{ins,i} \parallel F_{ins,j})$, where $F_{ins,i}$ and $F_{ins,j}$ are the corresponding features of pillar i and j from the instance decoder, and \parallel denotes the concatenation operation. In the second setting, we set α to be various fixed values. From the results in Table IV, constant $\alpha = 2.5$ yields the best results. A fixed value works relatively better than learning from the feature; for panoptic segmentation tasks on outdoor autonomous driving dataset, the difference in the regressed 2D embedding is enough to determine the connectivity of the pillars. However, we think that a learned α could potentially work better if the scene is dense and crowded (indoor scenes) such that the network

requires more information in making connections. Also note that $\alpha = 2.5 \approx -\ln(T)/(2 \times d)$. We draw conclusions that α can be regarded as a function of the threshold, T , and pillar grid size, d . Hence, we choose α to be $-\frac{\ln(T)}{2d}$ for the rest of the experiments. This design choice ensures that adjacent pillars are considered to be connected.

	α	PQ	PQ ^{1h}	RQ ^{1h}	SQ ^{1h}
Fixed	0.1	44.7	31.5	38.7	71.9
	1.0	56.0	58.4	66.2	76.7
	2.5	56.2	58.7	66.6	76.5
	5.0	48.7	41.0	51.2	68.0
Learned	—	55.3	56.2	64.9	75.5

TABLE IV

ABLATION STUDY OF USING DIFFERENT α . METRICS PROVIDED IN [%].

One may concern about the complexity of the model as it grows quadratically with M , the number of pillars. Note that we provide the average number of pillars resulted from using different grid sizes in Table V. Typically, a SemanticKITTI LiDAR scan contains an average number of 12 instances and 6.8k number of foreground points. We find that M is proportional to the number of instances in the scan but significantly less than the number of points. As the point embedding is learned to shift together in the network, dynamically grouping the foreground points together using pillars according to their embedding significantly reduces the computation the network needs to carry. The baselines with BFS clustering (row 1 in Table III) and HDBSCAN (row 2 in Table III) run in 8.6Hz and 4.8Hz respectively. In contrast, CPSEg runs in 14.2Hz. By dynamically grouping LiDAR points with similar embedding, CPSEg only processes on average 141 pillars instead of thousands of LiDAR points.

Grid Size (m)	M	PQ	PQ ^{1h}	RQ ^{1h}	SQ ^{1h}	Runtime (ms)
0.05	430	55.1	56.7	65.6	75.0	78
0.15	141	56.2	58.7	66.6	76.5	70
0.30	103	56.0	57.6	65.7	76.2	68
0.50	91	55.8	57.6	65.4	76.6	67

TABLE V

ABLATION STUDY OF USING DIFFERENT GRID SIZES IN PILLARIZING FOREGROUND POINTS BY EMBEDDING. METRICS PROVIDED IN [%].

V. CONCLUSION

In this work, we propose a novel real-time proposal-free and cluster-free panoptic segmentation network for 3D point cloud, called CPSEg. Our method builds upon an efficient semantic segmentation network and addresses the instance segmentation by incorporating a unique cluster-free instance head where the foreground point cloud is dynamically pillarized in the sparse space according to the learned embedding and object instances are formed by building connection of pillars. CPSEg outperforms existing real-time LiDAR-based panoptic segmentation methods on both SemanticKITTI and nuScenes datasets. The thorough analysis illustrates the robustness and effectiveness of the proposed method, which could inspire the field and push the panoptic segmentation research towards a proposal-free and cluster-free direction.

REFERENCES

- [1] R. Razani, R. Cheng, E. Li, E. Taghavi, Y. Ren, and L. Bingbing, "Gp-s3net: Graph-based panoptic sparse semantic segmentation network," in *Proc. of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2021, pp. 16 076–16 085.
- [2] A. Milioto, J. Behley, C. McCool, and C. Stachniss, "Lidar panoptic segmentation for autonomous driving," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 8505–8512.
- [3] Z. Zhou, Y. Zhang, and H. Foroosh, "Panoptic-polarnet: Proposal-free lidar point cloud panoptic segmentation," in *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [4] E. Li, R. Razani, Y. Xu, and B. Liu, "Smac-seg: Lidar panoptic segmentation via sparse multi-directional attention clustering," in *2022 International Conference on Robotics and Automation (ICRA)*, 2022, pp. 9207–9213.
- [5] J. Behley, M. Garbade, A. Milioto, J. Quenzel, S. Behnke, C. Stachniss, and J. Gall, "Semantickitti: A dataset for semantic scene understanding of lidar sequences," in *2019 IEEE/CVF International Conference on Computer Vision*. IEEE, 2019, pp. 9296–9306. [Online]. Available: <https://doi.org/10.1109/ICCV.2019.00939>
- [6] B. Wu, A. Wan, X. Yue, and K. Keutzer, "Squeezeseg: Convolutional neural nets with recurrent crf for real-time road-object segmentation from 3d lidar point cloud," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 1887–1893.
- [7] Q. Chen, S. Vora, and O. Beijbom, "Polarstream: Streaming object detection and segmentation with polar pillars," in *Advances in Neural Information Processing Systems*, 2021.
- [8] M. Gerdzhev, R. Razani, E. Taghavi, and B. Liu, "Tornado-net: multi-view total variation semantic segmentation with diamond inception module," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021.
- [9] R. Cheng, R. Razani, E. Taghavi, E. Li, and B. Liu, "(af)2-s3net: Attentive feature fusion with adaptive feature selection for sparse semantic segmentation network," in *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021.
- [10] X. Zhu, H. Zhou, T. Wang, F. Hong, Y. Ma, W. Li, H. Li, and D. Lin, "Cylindrical and asymmetrical 3d convolution networks for lidar segmentation," in *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021, pp. 9939–9948.
- [11] H. Tang, Z. Liu, S. Zhao, Y. Lin, J. Lin, H. Wang, and S. Han, "Searching efficient 3d architectures with sparse point-voxel convolution," in *European Conference on Computer Vision*, 2020.
- [12] H. Thomas, C. R. Qi, J.-E. Deschaud, B. Marcotegui, F. Goulette, and L. J. Guibas, "Kpconv: Flexible and deformable convolution for point clouds," in *Proc. of the IEEE International Conference on Computer Vision*, 2019, pp. 6411–6420.
- [13] K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask r-cnn," in *Proc. of the IEEE International Conference on Computer Vision (ICCV)*, Oct 2017.
- [14] K. Sirohi, R. Mohan, D. Büscher, W. Burgard, and A. Valada, "Efficientlps: Efficient lidar panoptic segmentation," *IEEE Transactions on Robotics*, vol. 38, no. 3, pp. 1894–1914, 2022.
- [15] J. V. Hurtado, R. Mohan, W. Burgard, and A. Valada, "Mopt: Multi-object panoptic tracking," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshop on Scalability in Autonomous Driving*, 2020.
- [16] T. Yin, X. Zhou, and P. Krahenbuhl, "Center-based 3d object detection and tracking," in *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021, pp. 11 784–11 793.
- [17] W. K. Fong, R. Mohan, J. V. Hurtado, L. Zhou, H. Caesar, O. Beijbom, and A. Valada, "Panoptic nusenes: A large-scale benchmark for lidar panoptic segmentation and tracking," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 3795–3802, 2022.
- [18] S. Gasperini, M.-A. N. Mahani, A. Marcos-Ramiro, N. Navab, and F. Tombari, "Panoster: End-to-end panoptic segmentation of lidar point clouds," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 3216–3223, 2021.
- [19] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu, "A density-based algorithm for discovering clusters in large spatial databases with noise," *Kdd*, vol. 96, p. 226–231, 1996.
- [20] F. Hong, H. Zhou, X. Zhu, H. Li, and Z. Liu, "Lidar-based panoptic segmentation via dynamic shifting network," in *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, June 2021.
- [21] R. J. G. B. Campello, D. Moulavi, and J. Sander, "Density-based clustering based on hierarchical density estimates," *Advances in Knowledge Discovery and Data Mining Lecture Notes in Computer Science*, p. 160–172, 2013.
- [22] A. Milioto, I. Vizzo, J. Behley, and C. Stachniss, "Rangenet++: Fast and accurate lidar semantic segmentation," in *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2019.
- [23] T. Cortinhal, G. Tzelepis, and E. E. Aksoy, "Salsanext: Fast semantic segmentation of lidar point clouds for autonomous driving," in *2020 IEEE Intelligent Vehicles Symposium (IV)*, 2020, pp. 655–661.
- [24] H. Badino, D. D. Huber, Y. Park, and T. Kanade, "Fast and accurate computation of surface normals from range images," in *2011 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2011.
- [25] D. J. Pearce, "An improved algorithm for finding the strongly connected components of a directed graph," *Victoria University, Wellington, NZ, Tech. Rep.*, 2005.
- [26] M. Berman, A. Rannen Triki, and M. B. Blaschko, "The lovasz-softmax loss: a tractable surrogate for the optimization of the intersection-over-union measure in neural networks," in *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [27] A. H. Lang, S. Vora, H. Caesar, L. Zhou, J. Yang, and O. Beijbom, "Pointpillars: Fast encoders for object detection from point clouds," in *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [28] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "nusenes: A multimodal dataset for autonomous driving," in *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 11 621–11 631.
- [29] A. Kirillov, K. He, R. Girshick, C. Rother, and P. Dollár, "Panoptic segmentation," in *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 9404–9413.