

Cascaded Non-local Neural Network for Point Cloud Semantic Segmentation

Mingmei Cheng, Le Hui, Jin Xie*, Jian Yang and Hui Kong

Abstract—In this paper, we propose a cascaded non-local neural network for point cloud segmentation. The proposed network aims to build the long-range dependencies of point clouds for the accurate segmentation. Specifically, we develop a novel cascaded non-local module, which consists of the neighborhood-level, superpoint-level and global-level non-local blocks. First, in the neighborhood-level block, we extract the local features of the centroid points of point clouds by assigning different weights to the neighboring points. The extracted local features of the centroid points are then used to encode the superpoint-level block with the non-local operation. Finally, the global-level block aggregates the non-local features of the superpoints for semantic segmentation in an encoder-decoder framework. Benefiting from the cascaded structure, geometric structure information of different neighborhoods with the same label can be propagated. In addition, the cascaded structure can largely reduce the computational cost of the original non-local operation on point clouds. Experiments on different indoor and outdoor datasets show that our method achieves state-of-the-art performance and effectively reduces the time consumption and memory occupation.

I. INTRODUCTION

In recent years, 3D imaging sensors have been greatly developed to facilitate the acquisition of 3D point cloud data. With the explosive growth of 3D point cloud data, point cloud semantic segmentation has received more and more attention [1]–[3] in 3D scene understanding. Point cloud semantic segmentation aims to classify each point into a category. Due to the unordered and irregular structure of 3D point clouds, how to exploit context information of point clouds for semantic segmentation is very challenging.

Recently, various efforts have been made on point cloud semantic segmentation. PointNet [4] directly employs the multi-layer perception (MLP) to extract the local feature of a single point for point cloud segmentation. Based on PointNet, PointNet++ [5] and DGCNN [6] aggregate different local features of point clouds with the max pooling operation for segmentation. Based on the PointNet++ framework, PointWeb [7] learns the weights between each pair of points in a local region to extract the local features. Although these methods can capture the geometric structures of different neighborhoods well, the relationships between long-range neighborhoods of point clouds are ignored.

The authors are with PCA Lab, Key Lab of Intelligent Perception and Systems for High-Dimensional Information of Ministry of Education, and Jiangsu Key Lab of Image and Video Understanding for Social Security, School of Computer Science and Engineering, Nanjing University of Science and Technology, Nanjing, China. 210094. Email: {chengmm, le.hui, csjie, konghui, csjyang}@njjust.edu.cn

*The author responsible for the correspondence of this paper

This work was supported by the National Science Fund of China (Grant Nos. U1713208, 61876084), Program for Changjiang Scholars.

In fact, the context from long-range neighboring points is important for point cloud semantic segmentation. It is desirable to exploit the long-range dependencies of different neighborhoods to characterize discriminative geometric structures of point clouds. To this end, we propose a novel cascaded non-local neural network for segmentation. In our method, the non-local operation is performed on three levels, i.e., the neighborhood level, superpoint level and global level, corresponding to different scales of areas in point clouds. The non-local operation at the neighborhood level is applied to the neighboring points with the k -nearest neighbors (k -NN) algorithm. The non-local operation at the superpoint-level is conducted in the superpoint area, which is a set of points with the isotropic geometric characteristics, while the global-level operation is implemented in the global point clouds composed of all the superpoints. By cascading the three non-local operations, geometric structure information of neighboring points can be propagated level by level. Once the non-local features at the global level are obtained, we employ the encoder-decoder framework in PointNet++ to predict the label of each point for semantic segmentation. Experimental results on the S3DIS, ScanNet, vKITTI and SemanticKITTI datasets demonstrate the effectiveness of our proposed method.

In the cascaded non-local neural network, the non-local operation, as an attention mechanism, can learn the attention weights of neighboring points at each level. Thus, different weights can be highlighted for different contexts from long-range neighboring points. In addition, the original non-local operation performed on the whole point clouds is time-consuming. By partitioning the whole point clouds into the cascaded areas, the resulting cascaded non-local operation can largely reduce the computational cost of the original non-local module.

In summary, the main contributions of this paper are two-fold. On one hand, we develop a novel cascaded non-local module where the non-local features of the centroid points at different levels can be extracted. Thus, context information from long-range neighboring points can be propagated level by level. On the other hand, the cascaded non-local module can largely reduce the computational complexity of the original non-local operation.

II. RELATED WORK

A. Deep learning on point clouds

PointNet [4] is the pioneering algorithm to extract deep features of unordered point clouds. PointNet++ [5] pays more attention to extract the local feature of point clouds with

the hierarchical structure. PointWeb [7] learns the weights of pairwise points to enhance the local features of point clouds. LatticeNet [8] adopts a sparse permutohedral lattice to characterize the local features of point clouds.

Graph-based methods mainly focus on depicting the edge relationships between points, and thus boost the local feature embedding. DGCNN [6] constructs the k -nn graph to characterize the local geometric structures of point clouds so that the local features of the points can be extracted. GACNet [9] introduces a graph attention convolution to assign weights to neighboring points and extracts the local features of point clouds through an attention mechanism. Models based on the superpoint graph (SPG) framework [10], [11] partition point clouds into superpoints, and then conduct feature embedding through a graph neural network built upon SPG.

SqueezeSegV3 [12] adopts spherical projection to generate LiDAR images and proposes the spatially-adaptive convolution for point cloud feature embedding, where the spatial priors in the LiDAR images can be exploited. SalsaNet [13] is an encoder-decoder network for point cloud segmentation, where the bird-eye view image is used as the input and an auto-labeling process is employed to transfer the labels from the camera to the LiDAR. Due to the imbalanced spatial distribution, PolarNet [14] proposes the polar birds-eye-view representation so that the nearest-neighbor-free segmentation method is used for point cloud semantic segmentation.

B. Non-local neural networks

The idea of the non-local operation was originally introduced in image denoising [15]. Wang *et al.* [16] propose a neural network combining the non-local blocks with CNNs to extract the image features. However, the vast computational consumption and massive memory occupation hinder its application. In [17], Zhu *et al.* propose an asymmetric non-local neural network to reduce the computational cost through a pyramid sampling module. Zhang *et al.* [18] introduce a novel graph structure to reduce the computational complexity of the non-local operation. In [19], only the pixels along the criss-cross path participate in the non-local operation. In [20], an adaptive sampling strategy is adopted to decrease the computational complexity. These methods aiming at reducing the computational cost of the non-local operation are based on the regular structure of images, which cannot be directly extended to the irregular and unordered point clouds for point cloud feature extraction.

III. OUR METHOD

In this section, we present our cascaded non-local neural network for point cloud segmentation. In Sec. III-A, we briefly revisit the non-local operation. In Sec. III-B, we describe the details of our cascaded non-local module. Finally, we present our network architecture in Sec. III-C and analyze the computational complexity of our cascaded non-local module in Sec. III-D.

A. Revisiting the non-local operation

Recently, the non-local operation is employed to construct the deep non-local network [16]. Given the feature map




Fig. 1. The non-local operation at three levels in our cascaded non-local module, including the neighborhood level, superpoint level and global level.

$\mathbf{x} \in \mathcal{R}^{H \times W \times C}$ of an image, we denote H , W , and C as the height and width of the image and the number of the channels, respectively. Suppose that in each non-local block there are three convolutions $\theta = W_\theta(\mathbf{x})$, $\phi = W_\phi(\mathbf{x})$, and $\gamma = W_\gamma(\mathbf{x})$ for embedding, respectively, where $\theta \in \mathcal{R}^{H \times W \times C'}$, $\phi \in \mathcal{R}^{H \times W \times C'}$, $\gamma \in \mathcal{R}^{H \times W \times C'}$ and C' is the number of the output channels. We then reshape these three embeddings to $\theta \in \mathcal{R}^{N \times C'}$, $\phi \in \mathcal{R}^{N \times C'}$ and $\gamma \in \mathcal{R}^{N \times C'}$, where $N = H \times W$. The final non-local feature is calculated as:

$$\mathbf{y} = \text{softmax}(\theta \times \phi^\top) \times \gamma \quad (1)$$

where $\mathbf{y} \in \mathcal{R}^{N \times C'}$ is the weighted sum of the embeddings of all pixels.

However, the non-local operation leads to the high computational cost due to large amounts of multiplications in Eq. 1. To tackle this problem, existing methods mainly focus on sampling the feature maps to reduce the computational complexity of matrix multiplications such as [17], [19], [20]. Due to the regular grid structure of images, these methods can effectively reduce the computational cost of the non-local operation. Nonetheless, since point clouds are irregular and unordered, it is difficult to regularly sample the feature maps of point clouds. Therefore, these methods cannot be directly applied to extract non-local features of point clouds.

B. Cascaded non-local module

Given two pointwise features $\mathbf{f}_i \in \mathbb{R}^C$ and $\mathbf{f}_j \in \mathbb{R}^C$, the non-local operation for point clouds is formulated as follows:

$$\mathbf{o}_i = \sum_j h(\mathbf{f}_i, \mathbf{f}_j) * g(\mathbf{f}_j) \quad (2)$$

where $h(\mathbf{f}_i, \mathbf{f}_j) = \mathbf{W}_\theta^\top(\mathbf{f}_i - \mathbf{f}_j)$ is the pairwise function embedding the difference between the two feature vectors, $g(\mathbf{f}_j) = \mathbf{W}_\phi^\top(\mathbf{f}_j)$ is a unary function for feature embedding, the operator $*$ represents the Hadamard product. $\mathbf{W}_\theta \in \mathbb{R}^{C \times D}$ and $\mathbf{W}_\phi \in \mathbb{R}^{C \times D}$ are the two weights to be learned.

In Eq. 1, the original non-local function employs the scalar value to depict the similarity between each pair of points. However, our pairwise function employs a channel-wise vector $h(\mathbf{f}_i, \mathbf{f}_j) \in \mathbb{R}^{1 \times D}$ to describe the relationship between two points. Thus our non-local feature can capture the geometry structure more accurately. Furthermore, to deal with the feature scale of different pair of points, the pairwise function h is normalized as follow:

$$h^d(\mathbf{f}_i, \mathbf{f}_j) = \frac{\exp(h^d(\mathbf{f}_i, \mathbf{f}_j))}{\sum_q \exp(h^d(\mathbf{f}_i, \mathbf{f}_q))} \quad (3)$$

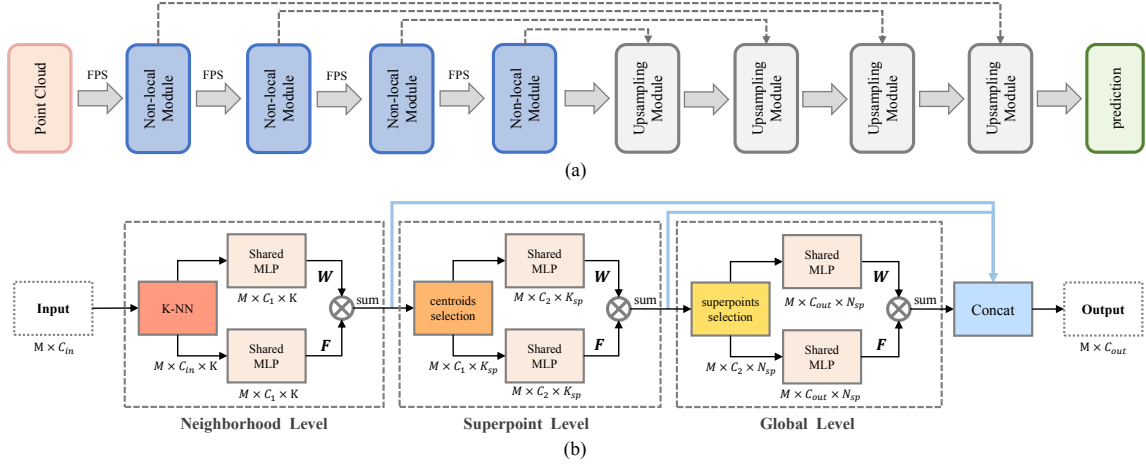


Fig. 2. (a) Framework of our cascaded non-local neural network. Four non-local modules are employed to construct the encoder network for point clouds feature embedding. The decoder is composed of four upsampling modules to predict the semantic labels of points. (b) Illustration of the proposed non-local module. The neighborhood-level non-local operation is performed to extract the local features of the neighborhood constructed by K nearest neighboring points. Then the superpoint-level non-local operation is conducted to propagate the geometric features of various neighborhoods to the superpoint. Consequently, the global-level non-local operation is performed on various superpoints, which further propagates the superpoint level non-local features in the global area. Finally, the non-local features at three levels are concatenated to generate discriminative features for point cloud semantic segmentation.

where $d = 1, 2, \dots, D$ represents the d -th channel of the feature vector.

It is noted that it is infeasible to directly perform the non-local operation on the whole point clouds since the non-local operation defined in Eq. 2 requires the huge memory occupancy. Therefore, in order to balance the memory cost of the non-local operation and accurate depiction of geometric structures of point clouds, we propose a cascaded non-local module, where the non-local operation can be performed on the point clouds at different levels. Specifically, the three-level non-local operations are conducted in three different scales of areas: neighborhood area, superpoint area and global area, respectively. It can greatly reduce the computational complexity of the non-local operation by controlling the number of points participated in each non-local operation.

Neighborhood level The non-local operation at the neighborhood level aims to extract the local features of the centroids in the corresponding neighborhoods. We first leverage the farthest point sampling (FPS) to choose M points as the centroids. For each centroid p_i , the K nearest neighboring points are used to construct the neighborhood area (Fig. 1). Consequently, we apply the non-local operation in the neighborhood area to extract the local feature of the centroid:

$$\mathbf{x}_i = \sum_{j=1}^K h(\mathbf{f}_i, \mathbf{f}_j) * g(\mathbf{f}_j) \quad (4)$$

where $i = 1, 2, \dots, M$ is the index of the centroid points. From Eq. 4, one can see that the local feature of each centroid point can be characterized by assigning different weights to the points in the neighborhood.

Superpoint level Once the local features of the centroid points are extracted, we conduct the non-local operation at the superpoint level. It is expected that geometric structure information of different neighborhoods in the superpoint can be effectively propagated. Superpoint is a set of points with isotropically geometric features. Generally, the number of

centroid points in the superpoint is different. Therefore, in order to facilitate batch processing in our neural network, we randomly sample $K_{sp} = 20$ centroid points in the superpoint. Thus, for the centroid point p_i , the non-local feature \mathbf{y}_i at the superpoint level is defined as:

$$\mathbf{y}_i = \sum_{j=1}^{K_{sp}} h(\mathbf{x}_i, \mathbf{x}_j) * g(\mathbf{x}_j) \quad (5)$$

where \mathbf{x}_i is the local feature of the centroid point p_i at the neighborhood level.

Global level In order to exploit semantic contexts from different superpoints in the point clouds, we furthermore propagate the features of the centroid points at the superpoint level. Since each superpoint contains multiple centroids, we use a max pooling operation to extract the superpoint feature \mathbf{v}_j of the j -th superpoint. As shown in Fig. 1, each small box represents the corresponding superpoint feature. For the centroid point p_i , by assigning different weights to the superpoint features, we define the following non-local operation at the global level:

$$\mathbf{z}_i = \sum_{j=1}^{N_{sp}} h(\mathbf{y}_i, \mathbf{v}_j) * g(\mathbf{v}_j) \quad (6)$$

where N_{sp} is the number of superpoints.

After propagating the non-local features of the centroids through the whole point clouds, we can obtain the fused features with a mapping function $M_\gamma: \mathbb{R}^D \rightarrow \mathbb{R}^{D^+}$. Therefore, the final feature of each centroid point is formulated as:

$$\mathcal{X}_i = M_\gamma([\mathbf{x}_i, \mathbf{y}_i, \mathbf{z}_i]) \quad (7)$$

where $[\cdot, \cdot, \cdot]$ denotes the concatenation operation and $\mathbf{x}_i, \mathbf{y}_i, \mathbf{z}_i$ represent the non-local features at three levels, respectively. With the cascaded non-local operations at three levels, the long-range dependencies across different neighboring points can be built so that we can obtain the discriminative feature of each centroid point for semantic segmentation.

C. Architecture

The overall architecture of our network is illustrated in Fig. 2. Our network is constructed based on the PointNet++ framework, combining the cascaded non-local modules to build the long-range dependencies between the points. In the encoder, we employ four non-local modules for feature embedding, while in the decoder we adopt the upsampling strategy in PointNet++ to predict the semantic labels of the points.

In PointNet++, the local features of the centroid points are extracted by performing the max pooling operation on the features of single points in a local region. Different from PointNet++, our proposed method extracts more discriminative local features by weighting the neighboring points through the non-local mechanism. In addition, we propagate the local features of the centroid points at the superpoint and global levels with the non-local operation. Thus, the discriminativeness of the local features of the centroid points can be further boosted. On the contrary, PointNet++ mainly focuses on the hierarchical local regions for feature extraction without considering non-local regions such as superpoints.

D. Computational complexity analysis

In Eq. 2, the non-local feature of each point is computed with the weighted sum of responses of all N points in the point clouds. If we ignore the feature channel, the computational complexity is $O(N^2)$. For our cascaded non-local module, the non-local operation is performed at three levels and the number of points participated in each operation is far smaller than N . Specifically, in the neighborhood level, the number of the subsampled centroid points is M ($M = 1024, 256, 64, 16$ in the four non-local modules) and each centroid point has K neighboring points. At the superpoint level, for each centroid point, the non-local operation is performed on K_{sp} centroid points in the same superpoint. At the global level, N_{sp} superpoints are involved in the non-local operation for each centroid point. In the experiment, we set $N_{sp} \leq 32$. Therefore the final computational complexity T of our hierarchical non-local operation is:

$$T = O(M \cdot K) + O(M \cdot K_{sp}) + O(M \cdot N_{sp}) \quad (8)$$

where K , K_{sp} and N_{sp} are far smaller than N and $M < N$. Thus, we can significantly reduce the computational complexity of the non-local operation on point clouds.

IV. EXPERIMENTS

In this section, we evaluate our proposed model on indoor and outdoor datasets.

A. Implementation details

To train our model, we use the SGD optimizer, where the base learning rate and mini-batch size are set to 0.05 and 16, respectively. The momentum is set to 0.9 and the weight decay is 0.0001. The approach adopted for superpoints partition follows [10]. For the S3DIS [25] dataset, the learning rate decays by 0.1 for every 25 epochs to train 100 epochs.

Method	S3DIS			ScanNet	vKITTI		
	mIoU	mAcc	OA	OA	mIoU	mAcc	OA
PointNet [4]	47.6	66.2	78.5	73.9	34.4	47.0	79.7
PointNet++ [5]	54.5	67.1	81.0	84.5	-	-	-
SPGraph [10]	62.1	73.3	85.5	-	-	-	-
ShellNet [21]	66.8	-	87.1	85.2	-	-	-
G+RCU [2]	49.7	66.4	81.1	-	35.6	57.6	79.7
G+RCU2 [22]	58.3	67.8	84.0	-	36.2	49.7	80.6
PointCNN [3]	65.4	75.6	88.1	85.1	-	-	-
PointWeb [7]	66.7	76.2	87.3	85.9	-	-	-
SSP+SPG [11]	68.4	78.3	87.9	-	52.0	67.3	84.3
HEPIN [23]	67.8	76.3	88.2	-	-	-	-
ConvPoint [24]	64.7	-	87.9	-	-	-	-
3P-RNN [1]	56.2	73.4	86.9	-	41.6	54.1	87.8
PointNL (ours)	68.4	77.5	88.2	86.7	57.6	68.1	89.8

TABLE I
SEMANTIC SEGMENTATION RESULTS ON S3DIS, SCANNET AND vKITTI DATASETS.

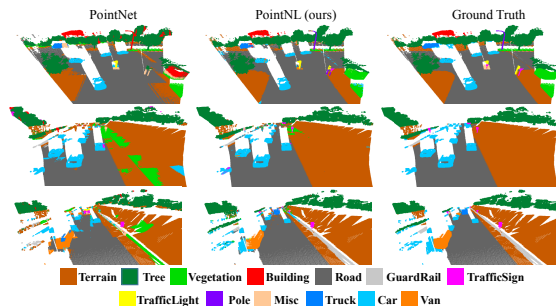


Fig. 3. Examples of semantic segmentation results on the vKITTI dataset.

For the ScanNet [26] dataset, we train the network for 200 epochs and decay the learning rate by 0.1 every 50 epochs. For the vKITTI [27] dataset, we train the network for 100 epochs and decay the learning rate by 0.1 every 25 epochs. For the outdoor dataset SemanticKITTI [28], the learning rate is 0.003 and weight decay is 0.0001. The scenes are partitioned into blocks as superpoints. Note that our method does not need to build the superpoint graphs.

B. Semantic segmentation on datasets

S3DIS. S3DIS [25] is an indoor 3D point cloud dataset. The point clouds are split into six large-scale areas, where each point is annotated with one of the semantic labels from 13 categories. Three metrics are used to quantitatively evaluate our method: mean of per-class intersection over union (mIoU), mean of per-class accuracy (mAcc), and overall accuracy (OA). For evaluation, we follow the methods [5], [10] to test the model on Area 5 and 6-fold cross validation. For training, we follow [5] to uniformly split the point clouds to blocks with an area size of $1m \times 1m$ and randomly sample 4096 points in each block. During the test, we adopt all the points for evaluation.

The quantitative results with 6-folds cross validation on S3DIS are shown in Tab. I. Our model named PointNL can achieve better or comparable performance than other methods listed in Tab. I, benefiting from the rich features provided by the non-local module. It is noted that although additional hand-craft geometry features in [10], [11] are not utilized in our model, our PointNL can still achieve

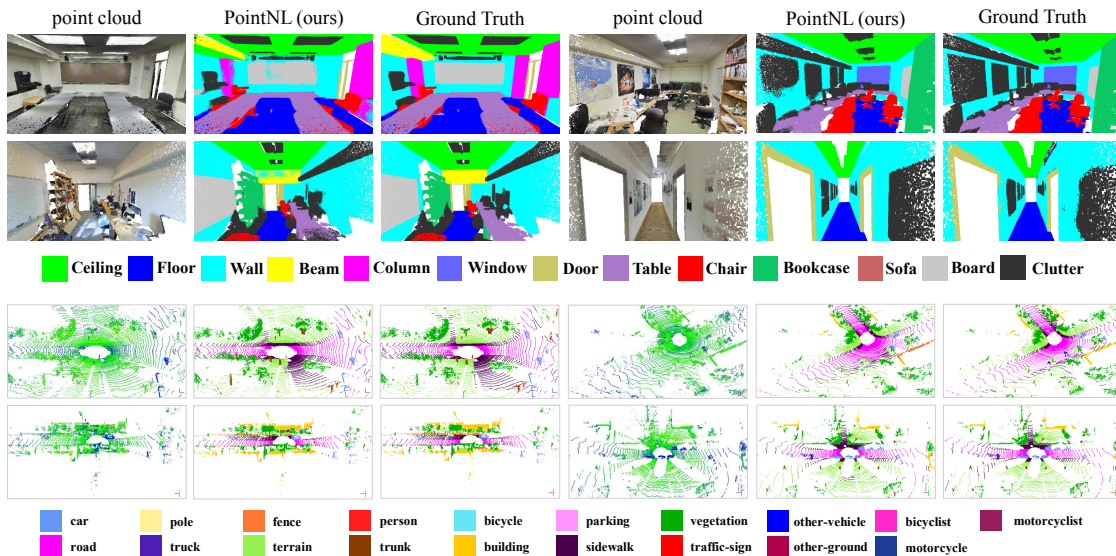


Fig. 4. Visualization results on the S3DIS and SemanticKITTI dataset, respectively. The results on S3DIS are shown in the top two rows while the results on SemanticKITTI are shown in the bottom two rows.

comparable results. The visual results are shown in Fig. 4. It can be seen that our proposed method can obtain more accurate segmentation results on the S3DIS dataset.

ScanNet. ScanNet [26] is an indoor scene dataset, which contains 1513 point clouds and points are annotated with 20 categories. Following [5], [7], the dataset is split into 1201 scenes for training and 312 scenes for testing. During the training, the points are divided into blocks of size $1.5\text{m} \times 1.5\text{m}$, and each block consists of 8192 points sampled on-the-fly. For testing, all the points in the test set are used for the evaluation. Following [7], the stride between the adjacent blocks is set to 0.5. The overall semantic voxel labeling accuracy (OA) is used to evaluate the compared segmentation methods.

We list the quantitative results on the validate dataset in Tab. I. From this table, one can see that our model outperforms the other compared methods. Due to the hierarchical non-local module, which can capture the long-range context information of point clouds, our model can obtain good segmentation results.

vKITTI. We also evaluate our method on the vKITTI [27] dataset, which mimics the real-world KITTI dataset. There are 5 sequences of synthetic outdoor scenes and 13 classes (including road, tree, terrain, car, etc.) are annotated. The 3D point clouds are obtained by projecting the 2D depth image into 3D coordinates. For evaluation, we follow the strategy adopted in [1] and split the dataset into 6 non-overlapping sub-sequences and conduct 6-fold cross validation. During the training, we split the point cloud into $4\text{m} \times 4\text{m}$ blocks and randomly sample 4096 points in each block. For evaluation, the metrics, mean of per-class intersection over union (mIoU), mean of per-class accuracy (mAcc) and overall accuracy (OA) are adopted.

The quantitative results are listed in Tab. I. From this table, one can see that our proposed PointNL can yield better

Method	mIoU	mAcc	OA	Training time	GPU memory
Non-local (baseline)	54.76	61.70	84.31	6 days	60 G
PointNL ¹ (ours)	60.03	66.35	86.56	1.4 days	5 G
PointNL ¹² (ours)	62.18	68.45	87.84	1.6 days	7 G
PointNL ¹²³ (ours)	63.50	69.24	87.90	1.8 days	8 G

TABLE II
SEMANTIC SEGMENTATION RESULTS, TRAINING TIME (DAYS) AND GPU MEMORY OCCUPATION (GBS) ON AREA 5 OF THE S3DIS DATASET.

performance on this dataset with a significant gain of 5.6% in terms of mIoU. The visual results are shown in Fig. 3, which further demonstrates the effectiveness of our method.

SemanticKITTI. SemanticKITTI [28] is a large-scale outdoor 3D dataset with 43552 LIDAR scans. Each scan contains $\sim 10^5$ points and can cover $160 \times 160 \times 20$ meters in the 3D space. The dataset is split into 21 sequences, while the sequences 00~07 and 09~10 (19130 scans) are used for training, sequence 08 (4071 scans) for validation, and the sequences 11~21 (20351 scans) for online testing. The point clouds are annotated with 19 classes without color information. For evaluation, the mean of per-class intersection over union (mIoU) is adopted following [28].

On SemanticKITTI dataset, PointNet [4], PointNet++ [5], DarkNet21Seg [28], DarkNet53Seg [28] can obtain the mIoUs of 14.6%, 20.1%, 47.4% and 49.9%, respectively, while our PointNL achieves the mIoU of 52.2%. It can be seen that our model outperforms the others. In addition, Fig. 4 visualizes the segmentation results on the SemanticKITTI dataset with our proposed PointNL.

C. Ablation study

To better demonstrate the effectiveness of our proposed method, we conduct ablation studies on Area 5 of the S3DIS dataset to analyze the effects of the non-local operations on different levels. Tab. II shows the segmentation

accuracies with the non-local operation at different levels. The neighborhood-level non-local operation (PointNL¹) can obtain the mIoU of 60.03% and is competitive with many point cloud segmentation approaches in recent years. When we further add the non-local operation at the superpoint level (PointNL¹²), the mIoU can be improved by 2.15%. When we cascade the non-local operations at three levels (PointNL¹²³), the mIoU can be further improved by 1.32%. In addition, we also implement the original non-local operation, which directly applies the non-local operation on the whole points. In Tab. II, the mIoU of the original non-local operation (baseline) is only 54.76%, which is lower than that of PointNL. It implies that our hierarchical non-local model can effectively build semantic long-range dependencies between the points to obtain the discriminative local features of point clouds.

D. Computational cost

In terms of training time and GPU memory, we compare our method to PointWeb and the original non-local operation (baseline) on the Pytorch platform. For a fair comparison, on Area 5 in the S3DIS dataset, both codes are run on Tesla P40 GPUs for 100 epochs. The computational cost, GPU memory occupancy and segmentation accuracy of the compared methods are shown in Tab. II.

For PointWeb [7], it needs at least 4.2 days, 24GB GPU memory and can obtain the mIoU of 60.28%. The original non-local operation (baseline) costs about 6 days and 60G GPU memory, and obtain the mIoU of 54.76%. While our method costs about 1.8 days, 8G GPU memory and can achieve mIoU of 63.50%.

The performance, computational cost and the GPU memory occupancy of the corresponding methods are shown in Tab. II. The total number of parameters of our model is 3.5M. For testing, the inference time of each batch is 0.35s, while PointWeb requires 1.5s. Therefore, our method not only improves the performance of semantic segmentation but also greatly reduces the time cost and memory consumption.

V. CONCLUSION

We proposed a novel cascaded non-local neural network for point cloud semantic segmentation. In our proposed non-local neural network, we developed a new cascaded non-local module to capture the neighborhood-level, superpoint-level and global-level geometric structures of point clouds. By stacking the cascaded non-local modules, semantic context information of point clouds is propagated level by level so that the discriminativeness of local features of point clouds can be boosted. Experimental results on the benchmark point cloud segmentation datasets demonstrate the effectiveness of our proposed PointNL in terms of the segmentation accuracy and computational cost.

REFERENCES

- [1] X. Ye, J. Li, H. Huang, L. Du, and X. Zhang, "3d recurrent neural networks with context fusion for point cloud semantic segmentation," in *ECCV*, 2018.
- [2] F. Engelmann, T. Kontogianni, A. Hermans, and B. Leibe, "Exploring spatial context for 3d semantic segmentation of point clouds," in *ICCV*, 2017.
- [3] Y. Li, R. Bu, M. Sun, W. Wu, X. Di, and B. Chen, "Pointcnn: Convolution on x-transformed points," in *NeurIPS*, 2018.
- [4] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, "Pointnet: Deep learning on point sets for 3d classification and segmentation," in *CVPR*, 2017.
- [5] C. R. Qi, L. Yi, H. Su, and L. J. Guibas, "Pointnet++: Deep hierarchical feature learning on point sets in a metric space," in *NeurIPS*, 2017.
- [6] Y. Wang, Y. Sun, Z. Liu, S. E. Sarma, M. M. Bronstein, and J. M. Solomon, "Dynamic graph cnn for learning on point clouds," *ACM Transactions on Graphics (TOG)*, 2019.
- [7] H. Zhao, L. Jiang, C.-W. Fu, and J. Jia, "Pointweb: Enhancing local neighborhood features for point cloud processing," in *CVPR*, 2019.
- [8] R. A. Rosu, P. Schütt, J. Quenzel, and S. Behnke, "Latticenet: Fast point cloud segmentation using permutohedral lattices," *arXiv preprint arXiv:1912.05905*, 2019.
- [9] L. Wang, Y. Huang, Y. Hou, S. Zhang, and J. Shan, "Graph attention convolution for point cloud semantic segmentation," in *CVPR*, 2019.
- [10] L. Landrieu and M. Simonovsky, "Large-scale point cloud semantic segmentation with superpoint graphs," in *CVPR*, 2018.
- [11] L. Landrieu and M. Boussaha, "Point cloud oversegmentation with graph-structured deep metric learning," *arXiv preprint arXiv:1904.02113*, 2019.
- [12] C. Xu, B. Wu, Z. Wang, W. Zhan, P. Vajda, K. Keutzer, and M. Tomizuka, "Squeezesegv3: Spatially-adaptive convolution for efficient point-cloud segmentation," *arXiv preprint arXiv:2004.01803*, 2020.
- [13] E. E. Aksoy, S. Baci, and S. Cavdar, "Salsanet: Fast road and vehicle segmentation in lidar point clouds for autonomous driving," *arXiv preprint arXiv:1909.08291*, 2019.
- [14] Y. Zhang, Z. Zhou, P. David, X. Yue, Z. Xi, B. Gong, and H. Foroosh, "Polarnet: An improved grid representation for online lidar point clouds semantic segmentation," in *CVPR*, 2020.
- [15] A. Buades, B. Coll, and J.-M. Morel, "A non-local algorithm for image denoising," in *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*. IEEE, 2005.
- [16] X. Wang, R. Girshick, A. Gupta, and K. He, "Non-local neural networks," in *CVPR*, 2018.
- [17] Z. Zhu, M. Xu, S. Bai, T. Huang, and X. Bai, "Asymmetric non-local neural networks for semantic segmentation," *arXiv preprint arXiv:1908.07678*, 2019.
- [18] S. Zhang, S. Yan, and X. He, "Latentgcn: Learning efficient non-local relations for visual recognition," *arXiv preprint arXiv:1905.11634*, 2019.
- [19] Z. Huang, X. Wang, L. Huang, C. Huang, Y. Wei, and W. Liu, "Ccnnet: Criss-cross attention for semantic segmentation," in *ICCV*, 2019.
- [20] M. Kranning, E. Chu, J. Lavaei, S. Boyd, et al., "Dynamic network energy management via proximal message passing," in *CVPR*, 2014.
- [21] Z. Zhang, B.-S. Hua, and S.-K. Yeung, "Shellnet: Efficient point cloud convolutional neural networks using concentric shells statistics," in *ICCV*, 2019.
- [22] F. Engelmann, T. Kontogianni, J. Schult, and B. Leibe, "Know what your neighbors do: 3d semantic segmentation of point clouds," in *ECCV*, 2018.
- [23] L. Jiang, H. Zhao, S. Liu, X. Shen, C.-W. Fu, and J. Jia, "Hierarchical point-edge interaction network for point cloud semantic segmentation," in *ICCV*, 2019.
- [24] A. Boulch, "Generalizing discrete convolutions for unstructured point clouds," *arXiv preprint arXiv:1904.02375*, 2019.
- [25] I. Armeni, O. Sener, A. R. Zamir, H. Jiang, I. Brilakis, M. Fischer, and S. Savarese, "3d semantic parsing of large-scale indoor spaces," in *CVPR*, 2016.
- [26] A. Dai, A. X. Chang, M. Savva, M. Halber, T. Funkhouser, and M. Nießner, "ScanNet: Richly-annotated 3d reconstructions of indoor scenes," in *CVPR*, 2017.
- [27] A. Gaidon, Q. Wang, Y. Cabon, and E. Vig, "Virtual worlds as proxy for multi-object tracking analysis," in *CVPR*, 2016.
- [28] 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 *ICCV*, 2019.