

Hierarchical Point-Edge Interaction Network for Point Cloud Semantic Segmentation

Li Jiang¹ Hengshuang Zhao¹ Shu Liu² Xiaoyong Shen² Chi-Wing Fu¹ Jiaya Jia^{1,2}

¹The Chinese University of Hong Kong ²Tencent YouTu Lab

{lijiang, hszhao, cwfu, leojia}@cse.cuhk.edu.hk {shawnshuliu, dylanshen}@tencent.com

Abstract

We achieve 3D semantic scene labeling by exploring semantic relation between each point and its contextual neighbors through edges. Besides an encoder-decoder branch for predicting point labels, we construct an edge branch to hierarchically integrate point features and generate edge features. To incorporate point features in the edge branch, we establish a hierarchical graph framework, where the graph is initialized from a coarse layer and gradually enriched along the point decoding process. For each edge in the final graph, we predict a label to indicate the semantic consistency of the two connected points to enhance point prediction. At different layers, edge features are also fed into the corresponding point module to integrate contextual information for message passing enhancement in local regions. The two branches interact with each other and cooperate in segmentation. Decent experimental results on several 3D semantic labeling datasets demonstrate the effectiveness of our work.

1. Introduction

With increasing capability of 3D sensing hardware, it is now easy to capture 3D data in many scenarios. Compared with 2D images, 3D data provides richer information about the environment. 3D data is in general view-independent and captures 3D structure, making it possible to incorporate geometry information in scene understanding tasks.

Learning-based approaches [28, 3, 10, 12, 25, 30, 15] were proposed to solve various 3D vision problems, e.g., shape classification, scene semantic/instance segmentation, and 3D object detection. Unlike 2D images, in which pixel grids are regular with object color information, 3D object data scatters, with most space actually not occupied. Therefore, directly voxelizing 3D scenes and extending deep neural network operations from 2D to 3D is inefficient. Several voxel-based methods, such as Submanifold Sparse Convolution [3] and O-CNN [23], improve the 3D convolution efficiency. However, since voxelization is accom-




Figure 1. Simple illustration of our framework. The point and edge branches work together to predict the semantic labels. Self-connected edges and edge directions are omitted.

panied by loss of information, high-resolution 3D models are needed to uphold the data precision, even though it unavoidably costs large memory and computation resource.

From another perspective, PointNet [10] directly processes 3D points in a network, only considering regions covered by the 3D points. PointNet++ [12] further adopts a hierarchical encoder-decoder structure to consider local regions, which downsamples point clouds in layers first and gradually interpolates them to the original resolution. This framework just utilizes weak connection between each point and its local context, since point features are extracted independently by the multi-layer perceptrons (MLP). In segmentation tasks, it is commonly known that local context is crucial for labeling the semantic categories. This motivates us to further explore the semantic relation between points and their local contextual neighbors to extract more discriminative features for 3D semantic scene labeling.

Our Contributions To explore the semantic relation between points in a local region and utilize the contextual information, we explicitly build edges between points and their contextual neighbors and establish a hierarchical edge branch with an auxiliary edge loss, as shown in Figure 1.

Specifically, besides the encoder-decoder point branch as in PointNet++, our new edge branch accepts point features from different layers and progressively produces edge features, which are then fed to point branch for fusing information in local graphs. For each point, the corresponding edge features provide local intrinsic geometric and regional semantic information to enhance point representation.

Instead of building isolated graphs for points in each

layer, we design a hierarchical graph construction process to gradually take point features at different layers into the edge branch. Edge features of adjacent layers are connected by an operator, named “edge upsample”. Consequently, edges on full-resolution point cloud encode multi-layer features, providing comprehensive data for final prediction.

We regularize the final edge features considering semantic consistency of the two connected points, which helps increase the discrimination ability between inter- and intra-category feature pairs, implicitly pulling points with the same semantic label closer in the feature space.

The decent performance of our method compared with all existing point-based neural networks on the large-scale scene labeling datasets, *i.e.*, Stanford Large Scale 3D Indoor Space (S3DIS) [1] and ScanNet [2], manifests the effectiveness of our framework.

2. Related Work

2.1. 3D Representation

To process 3D data, one typical approach is to store the data in volume grids and adopt 3D convolutions [28, 8, 11]. Since most voxels are unoccupied, Submanifold Sparse Convolution Network [3] defines a sparse convolution operation to process spatially-sparse 3D data. OctNet [13], on the other hand, represents the data using unbalanced octrees and defines network operations on these octrees to enable deeper neural networks without sacrificing the precision. Similarly, O-CNN [23] uses an octree to enable 3D CNN on high-resolution 3D data.

Another approach is to use multi-view 2D images, to which 2D convolutions [18, 11] can be directly applied. However, these approaches overlook the geometric structure in objects and scenes, especially the view-occluded 3D structures. Other methods [19, 9] consider 3D object surface and apply convolutions on it for semantic analysis.

2.2. Point-based Deep Neural Network

PointNet [10] is the first deep neural network to directly process 3D point coordinates, with MLPs and max-pooling for extracting features. Since max-pooling is a global operation on all the points, PointNet lacks local region understanding. PointNet++ [12] further applies a hierarchical structure and uses k -NN followed by max-pooling to capture regional information. Since it aggregates local features simply via a max-pooling, regional information is not yet fully utilized.

Recently, much effort has been made for effective local feature aggregation. SPLATNet [17] maps points into a high-dimensional sparse lattice and performs convolution on it. RSNet [4] projects features of unordered points into an ordered sequence of feature vectors and applies Recurrent Neural Network layers to model local depen-

dency. PointCNN [7] explores convolution on point clouds and addresses the point ordering issue by permuting and weighting input points and features with the \mathcal{X} -Conv operator. Besides, methods of [16, 22, 27, 26, 24] explore local context based on graphs.

Graph-based Methods ECC [16] organizes point clouds as graphs and uses graph convolutions to dynamically learn weights to combine local features. DGCNN [26] proposes the EdgeConv module to generate edge features that describe the connection between a point and its nearest neighbors. PointWeb [29] further connects every point pairs in a local region to obtain more representative region features. KCNet [14] creates k -nearest neighbor graphs and applies kernel correlation to learn local structures over point neighborhood. PCCN [24] and PointConv [27] connect each point with its k -nearest neighbors and extend the convolution operation from regular grids to irregular point clouds by adaptively projecting the relative position of two points to a convolution weight. Compared to PCCN, PointConv additionally considers point distribution density. Spectral Graph Convolution [22] performs graph convolution after a graph Fourier transform. Superpoint Graph (SPG) [6] splits the point cloud into geometrically-homogeneous partitions and builds a super-point graph, followed by a graph neural network to produce semantic labels.

In our work, we also propose a graph for point cloud processing, and yet focus particularly on exploring the semantic relation between points and their contextual neighbors for semantic segmentation through explicit edges. The key distinction of our method from other graph-based frameworks is that instead of fixing the graph and point resolution (e.g., PCCN [24] and KCNet [14]) or building independent graphs at each scale (e.g., PointConv [27], PointWeb [29] and ECC [16]), our graph is *hierarchically constructed*. We construct an edge branch, in which we fuse multi-scale point features and propagate edge features over multiple scales to enable longer distances of message passing hierarchically over edges without large memory overhead. Moreover, we propose edge loss aiming to encode the edges with exact semantic consistency information and increase the discrimination power among point features with different categories.

With meaningful edge features, we further feed edge features into each scale of the point branch to offer contextual information. To pass messages via edges, PointConv [27] and PCCN [24] adaptively learn weights from edges to fuse point features, while KCNet [14] defines a point-set kernel and kernel correlation to aggregate local features along edges. Different from these methods, our approach concatenates each point feature with the max-pooled corresponding edge features. Our approach requires less parameters to learn and preserves the distinctiveness of individual point features (Section 4.4 provides more discussions).

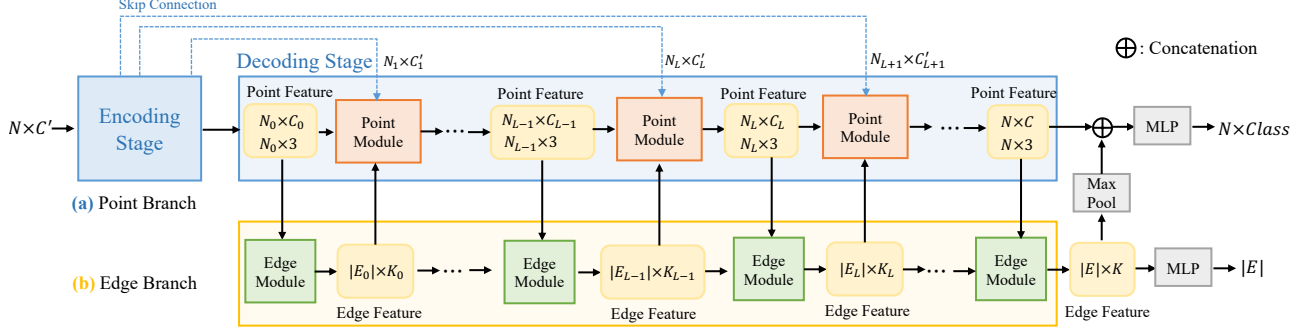


Figure 2. Overall architecture. N denotes the number of points in the original point cloud. The subscript of N is the layer index. Larger indices indicate layers with more points. C denotes the number of point feature channels. K denotes the number of edge feature channels. E denotes the edge set. The edge feature is encoded from the coarsest layer 0, and is gradually refined with the point features from later layers. Edge features in different layers also participate in the corresponding point modules to provide contextual information.

3. Our Approach

We design a hierarchical edge branch collaborating with the point prediction branch for point cloud semantic segmentation, as shown in Fig. 2. We progressively enlarge the graph, upsample edge features, and accept point features in different layers to refine the edge features. Edge features in different layers then provide extra contextual information for point feature learning. The final edge features are regularized with semantic consistency of their two-end points, which serve as auxiliary supervision for point features.

In this section, we first introduce the new edge branch, covering especially the interaction between point and edge branches, in Section 3.1. Then the hierarchical graph construction framework, which enables integration of different-layer information for edge prediction is described in Section 3.2. Section 3.3 depicts the loss regularizing both category prediction of each point and semantic-consistency prediction of each edge.

3.1. Edge Branch

Given a point cloud with N points $\mathcal{P} = \{p_1, p_2, \dots, p_N\}$, we construct a directed graph $G = (V, E)$, where $V = \mathcal{P}$ and E includes the edges that connect each point to its contextual points. Here, G is hierarchically constructed in a coarse to fine manner. We denote the graph in layer L as G_L . A larger L indicates a layer with more points, and layer 0 is the coarsest layer with the least points. The detailed graph construction process is depicted later in Section 3.2. Here, we first introduce the constitution of edge branch and how it interacts with the point branch.

As shown in Fig. 2, for the point branch, we follow PointNet++ [12] to create a hierarchical encoder-decoder structure with previous features in point encoder connected to the corresponding point decoder layers through skip-connection, thus passing detailed low-level information. The point cloud is downsampled and then upsampled in

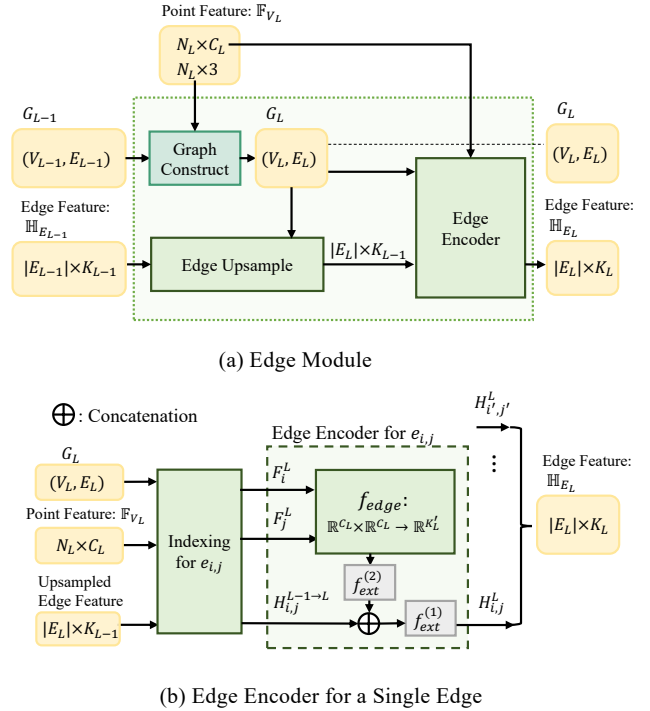


Figure 3. (a) Architecture of the *Edge Module*. (b) *Edge Encoder* block in (a). K_L and C_L represent the channel numbers in edge and point features in layer L , respectively. For simplicity, we only illustrate the edge encoding process for a single edge in (b). Edge features for all the edges in E_L constitute \mathbb{H}_{E_L} .

the process. Meanwhile, we construct an edge branch with consecutive edge modules, taking both features from the corresponding point module and the previous edge module.

The procedure is to extract edge features from the coarsest layer to grab high-level information with the largest receptive field, and progressively fuse point features from finer layers into edges, in parallel with the point decoding stage. Point features from the encoder layers are also used in the process, along with skip-connection to the corre-

sponding decoder layers.

Although both abstract global features from the coarser layers and detailed information from finer layers are important, the most essential data for edge prediction is from the last layer with the most refined point features. With this consideration, edge features are encoded in a coarse-to-fine manner, making point features in the finest layer fused at last. The hierarchical edge features are also fed to the corresponding point modules to provide additional contextual information.

3.1.1 Edge Module

At the decoding stage, for layer L , we denote the graph as $G_L = (V_L, E_L)$ and the number of points as N_L . The edge module accepts the L -layer point features \mathbb{F}_{V_L} and $(L - 1)$ -layer edge features $\mathbb{H}_{E_{L-1}}$ as arguments and returns the edge features in layer L . As shown in Fig. 3(a), the edge module is expressed as

$$\mathbb{H}_{E_L} = M_{encoder}(\mathbb{F}_{V_L}, M_{upsample}(\mathbb{H}_{E_{L-1}})), \quad (1)$$

where $M_{encoder}$ denotes the edge encoder and $M_{upsample}$ is the edge upsampling module, which maps edge features in graph G_{L-1} to graph G_L . The graph construction and edge upsampling process will be described in Section 3.2.

For each edge $e_{i,j} = (p_i, p_j) \in E_L$, its edge feature at layer L is written as

$$H_{i,j}^L = M_{encoder}(F_i^L, F_j^L, H_{i,j}^{L-1 \rightarrow L}), \quad (2)$$

where F_i^L and F_j^L are the point features of p_i and p_j , respectively. $H_{i,j}^{L-1 \rightarrow L}$ is the edge feature upsampled from layer $L - 1$ to layer L .

As illustrated in Fig. 3(b), $M_{encoder}$ for a single edge can be expanded as

$$H_{i,j}^L = f_{ext}^{(1)}([f_{ext}^{(2)}(f_{edge}(F_i^L, F_j^L)), H_{i,j}^{L-1 \rightarrow L}]), \quad (3)$$

where $[\cdot, \cdot]$ represents concatenation. The feature extractor $f_{ext} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ can be any differentiable function. In our implementation, we apply MLP as f_{ext} . The edge function f_{edge} takes the two point features it connects as input and outputs a feature for the edge. We formulate f_{edge} as

$$f_{edge}(F_i^L, F_j^L) = [(p_j - p_i), F_j^L, F_i^L], \quad (4)$$

where $[\cdot, \cdot, \cdot]$ concatenates the three elements, and p_i, p_j here represent 3D point coordinates. The two point features are concatenated for completely preserving information of the two points. Also, we provide $(p_j - p_i)$ to indicate the relative position between the two points. Other implementations of f_{edge} are discussed in the experiment part.




Figure 4. Architecture of the *Point Module*. K_L denotes the channel number of the L -layer edge features, while C_L denotes the channel number of the L -layer point feature.

3.1.2 Incorporation of Edges in Point Prediction

For layer L , every point in graph G_L links to other contextual points. So corresponding edges are expected to pass the contextual information back to the point. To this end, the edge features with respect to point p_i are operated by max-pooling as a region guidance. Let $E_L(p_i)$ denote the set containing all edges starting from p_i , the corresponding set of edge features is

$$\mathbb{H}_{E_L(p_i)} = \{H_{i,j}^L | (p_i, p_j) \in E_L(p_i)\}. \quad (5)$$

The point feature F_i^L is then updated by

$$(F_i^L)_{new} = [F_i^L, \text{MaxPool}(\mathbb{H}_{E_L(p_i)})]. \quad (6)$$

Fig. 4 gives an illustration of the process.

By incorporating edge information in point features, we enlarge the message passing range. The local region feature provided by the edges allows the point feature extractor to see farther in each layer. Additional contextual information including intrinsic geometry and semantic relation in the local region is incorporated in the region feature to benefit segmentation. We experiment with other schemes for message passing. Section 4.4 gives more discussions.

By helping feature extraction in the other branch, point and edge features become more powerful in final prediction.

3.2. Hierarchical Graph Construction

Instead of building graphs separately at each layer, we build the graph hierarchically, as shown in Fig. 5. By designing the “edge upsample” operation with each edge aware of associated edges in previous layer, we enlarge the receptive field and enable longer-range message passing for edges.

3.2.1 Graph Initialization

As shown in Fig. 5, the graph is initialized in the coarsest layer (layer 0). The initial graph G_0 is constructed by




Figure 5. Hierarchical Graph Construction. The graph is initialized in the coarsest layer and is progressively enlarged by considering both point coordinates in the current layer and the graph in previous layer.

connecting each point with its nearest k_0 points. Mathematically, $G_0 = (V_0, E_0)$ is formulated as

$$\begin{cases} V_0 = \mathcal{P}_0, \\ E_0 = \{(p_i, p_j) \mid p_i \in \mathcal{P}_0, p_j \in N_{k_0}(p_i)\}, \end{cases} \quad (7)$$

where \mathcal{P}_0 is the point set in layer 0, which is downsampled from the original point set with farthest point sampling (FPS) in encoding layers. $N_{k_0}(p_i)$ is the set of the k_0 -nearest neighbors of point p_i , including itself.

3.2.2 Hierarchical Architecture

Along with the decoding process of point features, we gradually enlarge the graph and enrich the edge features with more details. The process is illustrated in Fig. 5.

Graph Construction of Layer L Consider two adjacent layers $L-1$ and L with vertices V_{L-1} and V_L as the point set in that layer, respectively. The graph G_L is constructed by first finding the k_L nearest neighbors for each point in V_L . Let $G_L^{(0)} = (V_L, E_L^{(0)})$ denote such initial L -layer graph. For each edge $e_{i,j} = (p_i, p_j) \in E_L^{(0)}$, we consider the set consisting of possible neighboring edges in layer $L-1$ as

$$E_{ne}^{L-1}(e_{i,j}) = \{(p'_i, p'_j) \mid p'_i \in N_k^{L-1}(p_i), p'_j \in N_k^{L-1}(p_j)\},$$

where $N_k^{L-1}(p_i) \subseteq V_{L-1}$ is the k -nearest neighbors of $p_i \in V_L$ in layer $L-1$. p_i is included in $N_k^{L-1}(p_i)$ if $p_i \in V_{L-1}$.

We then check whether edges in $E_{ne}^{L-1}(e_{i,j})$ exist in E_{L-1} – the edge set of G_{L-1} . If edge $e_{i,j}$ connects two distant points, for which even in the coarser layer $L-1$ there is no connection between the two corresponding regions, we do not take the edge into consideration in layer L . Hence, if $E_{ne}^{L-1}(e_{i,j}) \cap E_{L-1} = \emptyset$, edge $e_{i,j}$ is discarded from $E_L^{(0)}$. Following this principle, the final graph $G_L = (V_L, E_L)$ has an edge set of

$$E_L = \bigcup_{p_i \in V_L} E_L(p_i),$$

where $E_L(p_i)$ (edges starting from p_i) is expressed as

$$E_L(p_i) = \{(p_i, p_j) \mid p_j \in N_{k_L}(p_i), E_{ne}^{L-1}(e_{i,j}) \cap E_{L-1} \neq \emptyset\}.$$

Note that at least $e_{i,i}$ is reserved in $E_L(p_i)$ in some extreme cases.

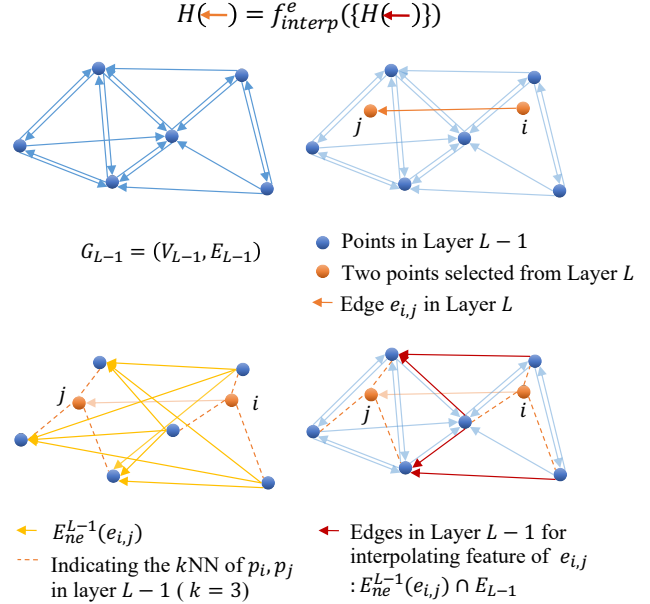


Figure 6. Demonstration of *edge upsampling*. Points in layer $L-1$ (blue ones) also exist in layer L . Self-connected edges are omitted. For edge $e_{i,j}$ in layer L , we propagate edge features in layer $L-1$ by finding its neighboring edges in E_{L-1} and interpolating features of these edges. Red arrows represent edges in G_{L-1} for interpolation, which denote intersection of E_{L-1} (blue arrows) and $E_{ne}^{L-1}(e_{i,j})$ (yellow arrows).

Edge Upsampling In PointNet++ [12], point feature of p_i in layer L is propagated from layer $L-1$ by interpolating feature values of its k nearest neighbors in layer $L-1$ as

$$F_i^{L-1 \rightarrow L} = f_{interp}^p(\{F_j^{L-1} \mid p_j \in N_k^{L-1}(p_i)\}). \quad (8)$$

We similarly propagate edge features in layer $L-1$ to layer L as

$$H_{i,j}^{L-1 \rightarrow L} = f_{interp}^e(\{H_{i',j'}^{L-1} \mid (p_{i'}, p_{j'}) \in E_{ne}^{L-1}(e_{i,j}) \cap E_{L-1}\}).$$

A demonstration is given in Fig. 6.

The interpolation weights are based on the inverse distance of the two pairs of end points. For $H_{i',j'}^{L-1}$, the weight is formulated as

$$w_{i',j'} = \frac{1}{(\|p_i - p_{i'}\|^t + \epsilon) \cdot (\|p_j - p_{j'}\|^t + \epsilon)}, \quad (9)$$

where $p_{i'}, p_{j'} \in V_{L-1}$, $p_i, p_j \in V_L$ represent point coordinates, $\epsilon = 1e - 8$ and t is set to 2. The weights are then normalized as

$$w_{i',j'}^n = \frac{w_{i',j'}}{\sum_{(p_{i' }, p_{j' }) \in E_{n\epsilon}^{L-1}(e_{i,j}) \cap E_{L-1}} w_{i',j'}}. \quad (10)$$

3.3. Loss Function

We optimize the point and edge branches jointly with the combined loss on the two branches as

$$L = \lambda_1 L_{point} + \lambda_2 L_{edge}, \quad (11)$$

where λ_1 and λ_2 adjust the ratio of the two losses.

Point Loss The final point features are followed by an MLP to produce point-wise semantic predictions. We further use the final edge predictions as weights to aggregate point scores and get refined point predictions. Cross entropy loss is applied to constrain the point predictions.

Edge Loss The edge features in the final graph G are regularized by the edge labels, which represent whether the two-end points of the edge are in the same category or not. The label for edge $e_{i,j} = (p_i, p_j) \in E$ is set as

$$l_{i,j}^e = \begin{cases} 1, & \text{if } l_i^p = l_j^p \\ 0, & \text{if } l_i^p \neq l_j^p \end{cases}. \quad (12)$$

where l_i^p and l_j^p are the point semantic labels of p_i and p_j . An MLP is adopted to produce the per-edge prediction. Binary cross entropy loss is chosen for the edge loss as

$$L_{edge} = -\frac{1}{|E|} \sum_{e_{i,j} \in E} (l_{i,j}^e \log(pred_{i,j}^e) + \alpha(1-l_{i,j}^e) \log(1-pred_{i,j}^e)), \quad (13)$$

where $pred_{i,j}^e$ is the edge prediction for $e_{i,j}$, and α balances the two kinds of edges, as there are more intra-class edges than inter-class ones considering the local neighborhood.

The final edge feature for each edge can be deemed as a function on features of the two regions centered at the two-end points. Information from different layers are taken into account. More details are preserved by encoding at last. Hence, the edge loss guides the edge encoder to seek difference between the intra- and inter-class feature pairs, and implicitly serves as auxiliary supervision for point features. It increases the discrimination power among point features in different categories. Also, with the edge supervision, more exact contextual information is passed to points via edges to enhance point features.

4. Experiments

We conducted experiments on two representative and challenging large-scale scene labeling datasets, *i.e.*, S3DIS [1] and ScanNet v2 [2], with ablation analysis presented on the ScanNet v2 val set and S3DIS Area 5.

4.1. Implementation Details

The point branch contains an encoder with four down-sampling layers and a decoder with four upsampling layers. The numbers of points, $N_0, N_1, N_2, N_3, N_4 = N$, in the decoder are 16, 64, 256, 1,024, and 4,096, respectively. The edge branch has five blocks with k (number of nearest neighbors) set to 4, 6, 10, 14, 16 from layer 0 to 4. k is chosen as 3 for point and edge feature interpolation.

The whole network was trained in an end-to-end manner using the SGD optimizer with batch size 16 and base learning rate 0.05. For S3DIS, we train the network for 100 epochs and decay the rate by 0.1 for every 25 epochs. For ScanNet, we train the network for 120 epochs and decay the rate by 0.1 for every 30 epochs. The momentum and weight decay are set to 0.9 and 0.0001 respectively.

4.2. Datasets

S3DIS The dataset [1] has 6 areas with a total of 271 rooms. Each room is provided as points with RGB information. Each point has a semantic label from 13 categories of floor, window, door, etc. In each training iteration, we randomly sample blocks in the training areas, with 4,096 points randomly selected per block. We set the block size as $0.8m \times 0.8m$ with $0.1m$ padding. Also, we represent each point as a 9D vector with XYZ , RGB , and normalized position in room. All points in the test areas are used in evaluation. Two settings are adopted [20, 6, 7]: (i) splitting Area 5 as the test set and using others for training; and (ii) adopting 6-fold cross validation, with each of the 6 areas taking as the test set once.

ScanNet v2 The dataset has 1,613 scans with a train/validation/test split of 1,201/312/100. Excluding the ‘unannotated’ points, each point in the scans has a label from 20 categories of wall, shower curtain, etc. To prepare the input data, we follow previous work [12] to randomly sample blocks in rooms and sample 4,096 points per block. Again, we use $0.8m \times 0.8m$ block size and $0.1m$ padding. Here, each input point feature is a 6D vector (XYZ & RGB). We evaluated on both the validation and test sets. Since the semantic annotation for the test sets is not publicly available, we submitted our predictions to the official server to obtain the evaluation results.

Evaluation Metric It includes the class-wise mean of intersection over union (mIoU), class-wise mean of accuracy (mAcc) and point-wise overall accuracy (OA).

4.3. Main Results

Table 1 lists quantitative results of different methods on S3DIS Area 5. Compared to previous approaches, ours yields the highest scores in terms of all the three metrics. Specifically, our model yields mIoU 61.85%, exceeding the former best by 3.58%. Table 2 shows the comparison

Methods	OA	mAcc	mIoU	ceiling	floor	wall	beam	column	window	door	table	chair	sofa	bookcase	board	clutter
PointNet [10]	-	48.98	41.09	88.80	97.33	69.80	0.05	3.92	46.26	10.76	58.93	52.61	5.85	40.28	26.38	33.22
SegCloud [20]	-	57.35	48.92	90.06	96.05	69.86	0.00	18.37	38.35	23.12	70.40	75.89	40.88	58.42	12.96	41.60
PointCNN [7]	85.91	63.86	57.26	92.31	98.24	79.41	0.00	17.60	22.77	62.09	74.39	80.59	31.67	66.67	62.05	56.74
SPGraph [6]	86.38	66.50	58.04	89.35	96.87	78.12	0.00	42.81	48.93	61.58	84.66	75.41	69.84	52.60	2.10	52.22
PCCN [24]	-	67.01	58.27	92.26	96.20	75.89	0.27	5.98	69.49	63.45	66.87	65.63	47.28	68.91	59.10	46.22
Our Method	87.18	68.30	61.85	91.47	98.16	81.38	0.00	23.34	65.30	40.02	75.46	87.70	58.45	67.78	65.61	49.36

Table 1. Semantic segmentation results evaluated on S3DIS Area 5. Most methods do not perform well on the “beam” category, which has few points (0.029%).




Figure 7. Visualization of the semantic segmentation results on the S3DIS dataset.

Methods	OA	mAcc	mIoU
PointNet [10]	78.5	66.2	47.6
RSNet [4]	-	66.45	56.47
SPGraph [6]	85.5	73.0	62.1
PointCNN [7]	88.14	75.61	65.39
Our Method	88.20	76.26	67.83

Table 2. Semantic segmentation results on the S3DIS dataset with 6-fold cross validation.

among different architectures on 6-fold cross validation. Ours also reaches the first place for all the three items.

Table 3 lists results of our framework and other point-based methods on ScanNet v2 test set. All methods use only point clouds with RGB color as input without voxelization. Our approach outperforms others by a large margin: 6.2% higher in absolute mIoU and 11.2% better relatively. Visual results are shown in Figs. 7 and 8. Our method segments objects even in complex scenes. It is notable that several detailed structures are classified and segmented from the surroundings, manifesting the effectiveness of our method.

4.4. Ablation Study

For ScanNet v2, the models are trained on training set and evaluated on validation set. For S3DIS, the models are

Method	mIoU
PointNet++ [12]	33.9
SPLATNet [17]	39.3
PointCNN [7]	45.8
PointConv [27]	55.6
Our Method	61.8

Table 3. Semantic segmentation results on ScanNet v2 test set.

trained on Areas 1-4 & 6 and evaluated on Area 5.

Edge Function We explore different ways of incorporating point information into edges, including Subtraction, Summation, Hadamard product, ‘ConcatSub’, and Concatenation. Here ‘ConcatSub’ is defined as

$$f_{edge}(F_i^L, F_j^L) = [(p_j - p_i), F_j^L, (F_j^L - F_i^L)]. \quad (14)$$

Table 4 shows comparison of the results. Overall, concatenation yields the best result due to preservation of most point information. Summation, Subtraction, and Hadamard Product all cause information loss in the level of point features. ‘ConcatSub’ achieves similar performance with Concatenation, since the two-point features can be restored in this type of operations.



Figure 8. Visualization of the semantic segmentation results on the ScanNet v2 dataset.

Methods	mIoU	mAcc	OA
Subtraction	58.31 / 58.85	67.95 / 65.66	84.02 / 86.44
Summation	57.86 / 58.96	67.25 / 65.87	83.69 / 86.48
Hadamard Product	59.07 / 58.79	68.02 / 65.27	85.31 / 86.16
ConcatSub	63.09 / 59.37	71.82 / 66.19	86.12 / 86.53
Concatenation	63.36 / 61.85	72.61 / 68.30	86.13 / 87.18

Table 4. Ablation study results for edge function f_{edge} on ScanNet v2 and S3DIS. The results are shown in format of ScanNet v2 / S3DIS. The ablation on two datasets share similar observation.

Methods	mIoU	mAcc	OA
AdaAggre (w. softmax)	56.44	66.17	83.06
AdaAggre (w.o. softmax)	55.01	64.12	82.67
MaxPool + Concat	63.36	72.61	86.13

Table 5. Ablation results for message passing by edges.

Message Passing by Edges Besides the approach described in Section 3.1.2, we also experimented with another scheme which is inspired by graph convolution [5, 21], where the edge features are further encoded to form weights for the linked points. The point features are then updated as a weighted sum of the adjacent point features. We denote this scheme as adaptive aggregation (*AdaAggre*) and test the two settings, with and without softmax, for the weights. Table 5 lists the experimental results on ScanNet v2 validation set.

The performance gain for the graph-convolution-style methods is lower than max-pooling followed by concatenation. It may be because during the point decoding, it is not very helpful to mix point features in each local neighborhood. Instead, the combined contextual feature reveals the relation of a point with its neighborhood. It can better preserve the point’s own distinctiveness.

Hierarchical Graph Construction and Edge Upsampling

We build connection between edge features of adjacent layers by “edge upsample”. We also experimented on ScanNet dataset with removing hierarchical graph construction and building the graph of each layer separately without edge upsampling.

The mIoU/mAcc/OA (%) results are 57.01/66.52/83.57 respectively, much lower than our full framework with 63.36/72.61/86.13. The connected edge branch optimally incorporates the point features in different layers, enabling effective learning for the edge features.

5. Conclusion

We have designed a hierarchical point-edge interaction network, in which an edge branch is proposed to work with the encoder-decoder point branch for point cloud semantic segmentation. The proposed hierarchical graph framework enables the edge branch to progressively integrate different-layer point features. Also, the generated edge features are incorporated into the point branch to provide contextual information. The final edge features are supervised by the semantic consistency of related points to implicitly regularize the point features. All these steps make semantic relationship with local context well utilized via edges.

With the high-quality point prediction results and generality of the framework applicable to different datasets, we believe the proposed method will broadly benefit 3D understanding in the community. In the future, we will explore multi-range edge construction to gather both close-range and long-distance contextual information.

Acknowledgments This project is supported in part by the Research Grants Council of the Hong Kong Special Administrative Region (CUHK 14203416 & 14201918).

References

- [1] Iro Armeni, Ozan Sener, Amir R Zamir, Helen Jiang, Ioannis Brilakis, Martin Fischer, and Silvio Savarese. 3D semantic parsing of large-scale indoor spaces. In *CVPR*, 2016.
- [2] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. ScanNet: Richly-annotated 3D reconstructions of indoor scenes. In *CVPR*, 2017.
- [3] Benjamin Graham, Martin Engelcke, and Laurens van der Maaten. 3D semantic segmentation with submanifold sparse convolutional networks. *CVPR*, 2018.
- [4] Qiangui Huang, Weiyue Wang, and Ulrich Neumann. Recurrent slice networks for 3D segmentation of point clouds. In *CVPR*, 2018.
- [5] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *ICLR*, 2017.
- [6] Loic Landrieu and Martin Simonovsky. Large-scale point cloud semantic segmentation with superpoint graphs. In *CVPR*, 2018.
- [7] Yangyan Li, Rui Bu, Mingchao Sun, Wei Wu, Xinhan Di, and Baoquan Chen. PointCNN: Convolution on \mathcal{X} -transformed points. In *NIPS*, 2018.
- [8] Daniel Maturana and Sebastian Scherer. VoxNet: A 3D convolutional neural network for real-time object recognition. In *IROS*, 2015.
- [9] Hao Pan, Shilin Liu, Yang Liu, and Xin Tong. Convolutional neural networks on 3D surfaces using parallel frames. *arXiv:1808.04952*, 2018.
- [10] Charles Ruizhongtai Qi, Hao Su, Kaichun Mo, and Leonidas J. Guibas. PointNet: Deep learning on point sets for 3D classification and segmentation. In *CVPR*, 2017.
- [11] Charles Ruizhongtai Qi, Hao Su, Matthias Nießner, Angela Dai, Mengyuan Yan, and Leonidas Guibas. Volumetric and multi-view CNNs for object classification on 3D data. In *CVPR*, 2016.
- [12] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J. Guibas. PointNet++: Deep hierarchical feature learning on point sets in a metric space. In *NIPS*, 2017.
- [13] Gernot Riegler, Ali Osman Ulusoy, and Andreas Geiger. OctNet: Learning deep 3D representations at high resolutions. In *CVPR*, 2017.
- [14] Yiru Shen, Chen Feng, Yaoqing Yang, and Dong Tian. Mining point cloud local structures by kernel correlation and graph pooling. In *CVPR*, 2018.
- [15] Shaoshuai Shi, Xiaogang Wang, and Hongsheng Li. PointRCNN: 3D object proposal generation and detection from point cloud. In *CVPR*, 2019.
- [16] Martin Simonovsky and Nikos Komodakis. Dynamic edge-conditioned filters in convolutional neural networks on graphs. In *CVPR*, 2017.
- [17] Hang Su, Varun Jampani, Deqing Sun, Subhransu Maji, Evangelos Kalogerakis, Ming-Hsuan Yang, and Jan Kautz. SPLATNet: Sparse lattice networks for point cloud processing. In *CVPR*, 2018.
- [18] Hang Su, Subhransu Maji, Evangelos Kalogerakis, and Erik G. Learned-Miller. Multi-view convolutional neural networks for 3D shape recognition. In *ICCV*, 2015.
- [19] Maxim Tatarchenko, Jaesik Park, Vladlen Koltun, and Qian-Yi Zhou. Tangent convolutions for dense prediction in 3D. In *CVPR*, 2018.
- [20] Lyne P. Tchammi, Christopher B. Choy, Iro Armeni, JunYoung Gwak, and Silvio Savarese. SEGCloud: Semantic segmentation of 3D point clouds. In *3DV*, 2017.
- [21] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. In *ICLR*, 2018.
- [22] Chu Wang, Babak Samari, and Kaleem Siddiqi. Local spectral graph convolution for point set feature learning. In *ECCV*, 2018.
- [23] Peng-Shuai Wang, Yang Liu, Yu-Xiao Guo, Chun-Yu Sun, and Xin Tong. O-CNN: Octree-based convolutional neural networks for 3D shape analysis. *ACM TOG (SIGGRAPH)*, 36(4):72:1–11, 2017.
- [24] Shenlong Wang, Simon Suo, Wei-Chiu Ma, Andrei Pokrovsky, and Raquel Urtasun. Deep parametric continuous convolutional neural networks. In *CVPR*, 2018.
- [25] Weiyue Wang, Ronald Yu, Qiangui Huang, and Ulrich Neumann. SGPN: Similarity group proposal network for 3D point cloud instance segmentation. In *CVPR*, 2018.
- [26] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E Sarma, Michael M Bronstein, and Justin M Solomon. Dynamic graph CNN for learning on point clouds. *ACM TOG*, 2019. to appear.
- [27] Wenxuan Wu, Zhongang Qi, and Li Fuxin. PointConv: Deep convolutional networks on 3D point clouds. In *CVPR*, 2019.
- [28] Zhirong Wu, Shuran Song, Aditya Khosla, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3D ShapeNets: A deep representation for volumetric shape modeling. In *CVPR*, 2015.
- [29] Hengshuang Zhao, Li Jiang, Chi-Wing Fu, and Jiaya Jia. PointWeb: Enhancing local neighborhood features for point cloud processing. In *CVPR*, 2019.
- [30] Yin Zhou and Oncel Tuzel. VoxelNet: End-to-end learning for point cloud based 3D object detection. In *CVPR*, 2018.