

Human Vision Based 3D Point Cloud Semantic Segmentation of Large-Scale Outdoor Scenes

Sunghwan Yoo, Yeonjeong Jeong, Maryam Jameela, Gunho Sohn
Department of Earth and Space Science and Engineering, York University
Toronto, ON M3J 1P3, Canada

(jacobyyoo, yjjeong, maryumja, gsohn)@yorku.ca

Abstract

This paper proposes EyeNet, a novel semantic segmentation network for point clouds that addresses the critical yet often overlooked parameter of coverage area size. Inspired by human peripheral vision, EyeNet overcomes the limitations of conventional networks by introducing a simple but efficient multi-scale input and a parallel processing network with connection blocks between parallel streams. The proposed approach effectively addresses the challenges of dense point clouds, as demonstrated by our ablation studies and state-of-the-art performance on Large-Scale Outdoor datasets.

1. Introduction

Recently, there has been growing interest in developing digital twins of the three-dimensional world, driven by their various applications. With advancements in LiDAR devices and survey techniques, point cloud datasets have become more accurate, dense, and spatially extensive, both on the ground level [3, 14, 29, 30, 34] and in the airborne level [10, 16, 28, 35].

However, the functional coverage area of an input batch is critical for effective feature learning in semantic segmentation networks, especially as data sets become denser [16, 35]. Traditional approaches to semantic segmentation, such as increasing the number of input points [17, 33, 40] and downsampling the point cloud [25, 26] to extract features from high-density outdoor scenes, have limitations. For instance, increasing the number of input points is restricted by the memory capacity of GPUs, while further downsampling can result in information loss.

In this work, we propose a 3D semantic segmentation network called EyeNet, inspired by the human visual system, peripheral vision. The network takes input from multi-scale regions, including the central and peripheral regions. The high-density central region is advantageous for extract-




Figure 1. The design of EyeNet is based on the concept of the human vision system. We drew inspiration from the peripheral vision system, which has a low density of receptors but covers a larger region than the central vision. As a result, we added the peripheral region around the central region to achieve a similar effect in our network.

ing precise geometric features, while the low-density peripheral region with a larger coverage area is beneficial for extracting features from objects that require such assets. To process these regions, the study uses a parallel feature processing design with connection layers for exchanging features between parallel streams. The proposed network outperforms state-of-the-art semantic segmentation methods on the SensatUrban benchmark [16] and Toronto3D [30]. The design philosophy of the EyeNet is presented in Figure 1.

Our key contributions are:

- Our approach introduces a memory-efficient human vision based input that draws inspiration from human vision, in order to address the limitations of traditional semantic segmentation networks.
- To facilitate feature learning using the human vision based input, we proposed an efficient parallel processing architecture.
- To ensure that complex features from both the central and peripheral regions are preserved and effectively

merged for improved learning, we proposed an effective deep learning connection layer.

2. Related Work

Semantic segmentation of outdoor 3D point clouds is a challenging task that has received increasing attention recently. Numerous deep learning-based approaches have been proposed to address this problem. [13, 17, 25, 26, 33, 36, 40]

Deep learning-based methods have shown promising results in semantic segmentation of outdoor 3D point clouds. However, achieving optimal coverage area of the deep learning network is one of crucial factors. Two commonly used methods to achieve an effective coverage area in outdoor 3D point clouds are downsampling the input point cloud and increasing the number of input points.

2.1. Downsampling Input Point Cloud

One of the most effective ways to manage computational costs and ensure sufficient coverage area of the input in outdoor 3D point cloud semantic segmentation is downsampling. Early networks, such as PointNet [25, 26] often utilized the constant-volume downsampling method, which collects points in a fixed volume and downsamples them to a fixed number of points. More recent networks such as KPConv [17, 33] adopted the grid sampling method where one point per voxel in a 3D grid was kept. Although this approach is computationally efficient, it disregards differences in point distribution within the point cloud, resulting in a loss of geometric information and a decrease in performance [16].

2.2. Increasing the Number of Inputs

Recently, the trend has shifted away from constant-volume downsampling towards increasing the number of input points through grid sampling [6, 17, 33, 40]. RandLA [17] is a network that processes large-scale point clouds by utilizing random sampling based on k-nearest neighbors (KNN), thus avoiding the loss of geometric information. However, a critical limitation of this method is that the ability to increase the number of input points can be restricted by the memory capacity of GPUs.

2.3. Peripheral Vision in Computer Vision

The integration of peripheral vision into deep learning has been a widely explored area in 2D computer vision research. It is primarily investigated in 2D object detection and 2D semantic segmentation [7, 9, 12, 24]. Also, recent studies have explored the combination of peripheral vision and transformer networks [23], leading to the development of the Peripheral Vision Transformer network. However, the application of peripheral vision systems in 3D computer vision tasks has not been extensively explored.

3. Methodology

3.1. Overview

Real-world 3D scene’s point clouds have tremendous points and varying densities, requiring sampling methods like grid subsampling which scale is determined by voxel size [32]. Then, unsampled points are predicted based on the nearest sampled point’s prediction usually decided by voting because random cropping is used due to GPU memory limit. When we crop the input point cloud randomly for training and inference, the optimal coverage size of the input point cloud is not well-defined.

To address this issue, we introduce a novel approach in EyeNet that utilizes both central and peripheral regions, allowing for a parallel processing network that facilitates communication between these regions to retain important features. By incorporating the human vision based input in this manner, we can maximize the benefits of these inputs. A summary of the EyeNet architecture is shown in Figure 2. Our EyeNet can basically be applied to most existing point-based networks [17, 33, 40] easily. In this paper, we select RandLA as our baseline. Therefore, any process not described in this section followed the same procedures used in the baseline RandLA implementation [1].

As depicted in Figure 2, we constructed the network by stacking two streams of RandLA-Net’s encoder and decoder in parallel. To extract fundamental features in each encoding layer, we utilized the k-NN based RandLA’s local feature aggregation (LFA) module. The LFA module outputs from both streams were randomly downsampled and fused through the connection layer. After the five encoding layers, we included five decoding layers with a bilinear interpolation operation. Skip connections were utilized to concatenate features from encoding and decoding layers to minimize information loss due to random downsampling. Finally, the outputs from both streams were merged using the Feature Merging Block and passed through fully connected layers (fc) to assign predicted labels to each point.

3.2. Human Vision Based Input

Due to the extensive number of points present in urban scale large outdoor point cloud, it is not feasible to input all points into the network simultaneously. Thus, for both training and inference purposes, it is necessary to crop the input region in the entire scene. To do this, our approach, called Human Vision Based Input, is centred around two key factors: the coverage areas and point densities of the central and peripheral regions. Let N be the number of input points for each feature processing stream. We use the K-nearest neighbours (KNN) algorithm to gather N points from the centroid C_0 , forming the input points for the central stream $P_c = \{p_c^1 \cdots p_c^n \cdots p_c^N\}$. Here, we define R as the distance between C_0 and the farthest point p_c^f in P_c , and

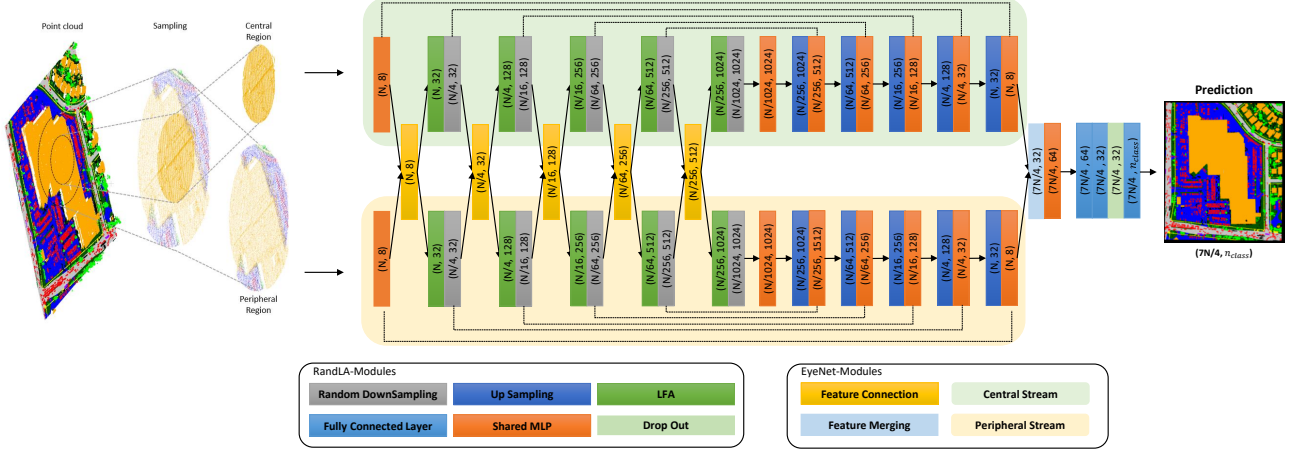


Figure 2. The figure displays the architecture of EyeNet, which utilizes a parallel processing approach to handle the human vision based input. The network also includes feature connection and feature merging blocks to facilitate information exchange between the central and peripheral streams. N and d are used to represent the number of points and feature dimension, respectively. The central and peripheral regions share $N/4$ points as inputs, resulting in a total of $7N/4$ input points. MLP represents the multilayer perceptron.

we set the radius of the central region to R and the peripheral region to $2R$.

We randomly sample $N/4$ points from the central region, and we refer to these points as the messenger points, which are represented as $P_m = \{p_m^1 \cdots p_m^n \cdots p_m^{N/4}\} \subseteq P_c$.

Moreover, we randomly select an additional $3N/4$ points from the region between R and $2R$, and we define them as $P_{ph-m} = \{p_{ph}^1 \cdots p_{ph}^n \cdots p_{ph}^{3N/4}\}$. We choose the number $3N/4$ to ensure uniform point density with P_m .

Finally, the combination of the point clouds P_{ph-m} and P_m becomes the input point cloud for the peripheral stream, which we denote as $P_{ph} = P_{ph-m} + P_m = \{p_{ph}^1 \cdots p_{ph}^n \cdots p_{ph}^{3N/4}, p_m^1 \cdots p_m^n \cdots p_m^{N/4}\}$. Therefore, the messenger points are an intersection between the points in the central and peripheral regions. An overview of the operation is shown in Figure 4.

Our proposed human vision based technique for collecting network input points allows us to cover the same input region for a deep neural network using substantially fewer points than the conventional fixed-number-based approaches like RandLA. In all experiments, we used nearly 60% fewer points than the baseline. The efficient sampling approach enabled our EyeNet to utilize GPU memory more effectively. During inference, our EyeNet was capable of employing four times more batches than the baseline due to this approach.

3.3. Parallel Processing and Receptive Field

The network comprises two streams, namely the central and peripheral streams. The central and peripheral streams

process point clouds P_c and P_{ph} , respectively, with different densities. It makes the receptive field of the peripheral region larger than that of the central region when we use the KNN-based module for feature aggregation with the same number of K for both regions. In this concept, a point can learn its feature from a variable size of the receptive field - small when in the central region, large when in the peripheral region, or both when selected as a messenger point due to the use of random cropping in large-scale 3D point cloud training and inferencing.

3.4. Connection Block

As we utilize a parallel processing network, establishing connections between two streams is crucial. We were motivated by previous works such as [18, 39] and developed the connection block that exchanges information between streams through messenger points. Since the feature information from the central and peripheral streams vary, we adopted the two-way self-attentive pooling block from [17] along with the channel enhancement block from [6] for better information exchange.

Feature Preparation: To start, we consider the features of points in the output of the central and peripheral streams, at a given layer l , denoted as $F_c^l = \{f_c^{l1} \cdots f_c^{ln} \cdots f_c^{lN^l}\}$ and $F_{ph}^l = \{f_{ph}^{l1} \cdots f_{ph}^{ln} \cdots f_{ph}^{lN^l}\}$, respectively. From these features, we extract the features of messenger points $F_{m.c}^l = \{f_{m.c}^{l1} \cdots f_{m.c}^{ln} \cdots f_{m.c}^{lN^l/4}\} \subseteq F_c^l$ and $F_{m.ph}^l = \{f_{m.ph}^{l1} \cdots f_{m.ph}^{ln} \cdots f_{m.ph}^{lN^l/4}\} \subseteq F_{ph}^l$. The remaining points' features are referred to as $F_{c-m}^l =$

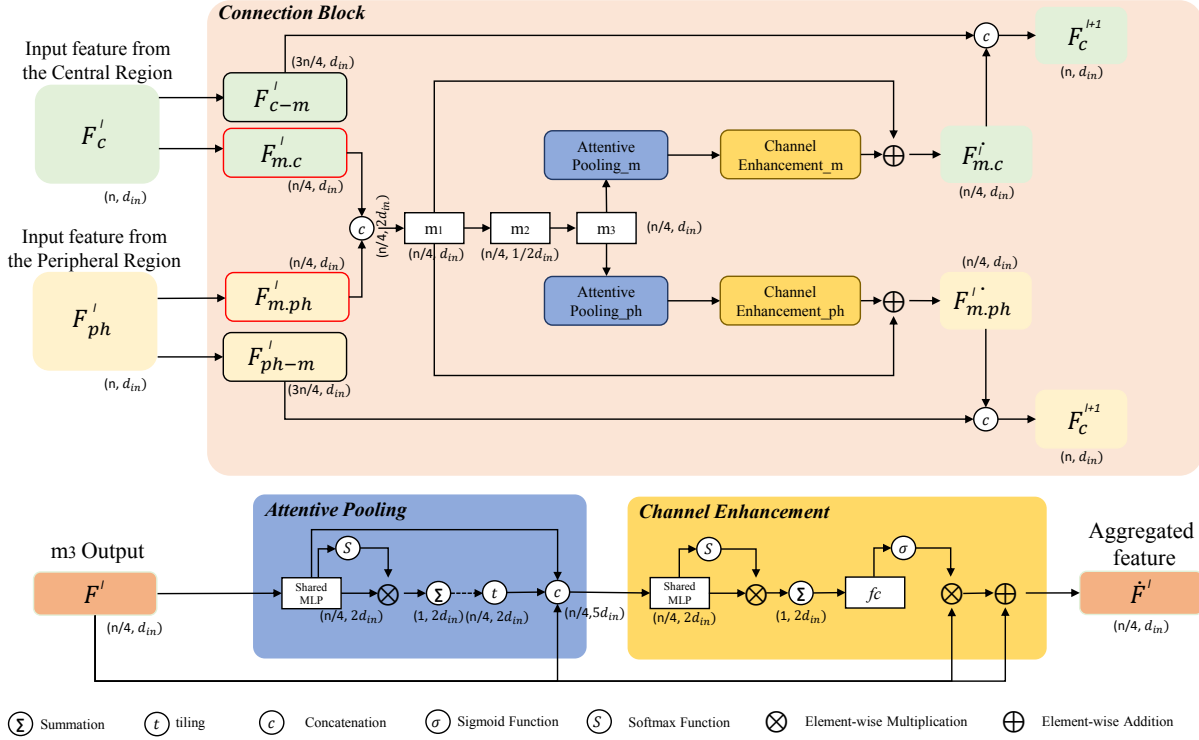


Figure 3. The proposed connection block. The top panel shows the overview of the connection block that is developed to facilitate information exchange between the central and peripheral streams. The bottom panel shows the self-attentive pooling block that weighs the most important features and the channel enhancement block that further enhances the features from the attentive pooling block. f_c : the fully connected layer.




Figure 4. The overview of our human vision based input. Points in the central region are generated based on number-based sampling, and points in the peripheral region are generated based on volume-based down sampling.

$\{f_{c-m}^{l1} \dots f_{c-m}^{ln} \dots f_{c-m}^{l3N^l/4}\} \subseteq F_c^l$ and $F_{ph-m}^l = \{f_{ph-m}^{l1} \dots f_{ph-m}^{ln} \dots f_{ph-m}^{l3N^l/4}\} \subseteq F_{ph}^l$. Afterward, we concatenate the channel-wise features of $F_{m,c}^l$ and $F_{m,ph}^l$ into F_m^l and process the resulting output through shared MLP function $m()$ three times. The shared MLPs' operations are denoted as follows:

$$f^{li} = m_3(m_2(m_1(f^{li}, W_1), W_2), W_3) \quad (1)$$

where $F^l = \{f^{l1} \dots f^{ln} \dots f^{lN^l/4}\}$ and W_s are the learn-

able weights of shared MLPs.

Attentive Pooling: Inspired by self-attentive pooling block in RandLA [17], we used the powerful attention mechanism to exchange and fuse important features automatically, and it consists of the following steps.

We utilized a shared MLP followed by softmax to learn a unique attention score for each feature from F^l . It is formally defined as follows:

$$f_{ap}^{li} = L(m_a(f^{li}, W_a)) \quad (2)$$

where W is the learnable weights of shared MLP and L is the Relu activation function.

$$s_{ap}^{li} = S(f_{ap}^{li}) \quad (3)$$

where S is softmax function. We can consider the attention scores that have been learned as a softmax function that automatically picks out significant characteristics. In technical terms, these traits are combined through a weighted summation process, as described below:

$$f^{li'} = \sum (f_{ap}^{li} \cdot s_{ap}^{li}) \quad (4)$$

Channel Enhancement: In addition, we improved the features by creating more prominent global features while minimizing the loss of information during the self-attention encoding process. This method was initially presented in

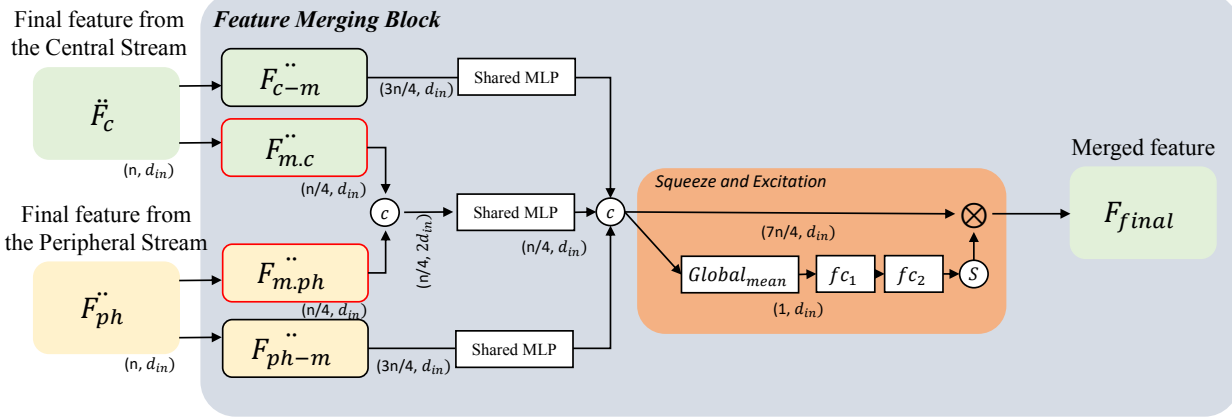


Figure 5. The overview of our feature merging block. First, the features are processed, and shared MLPs are applied to smooth the feature spaces. Then, the point-wise concatenation of the features occurs, followed by the squeeze and excitation that weighs the most important features for predictions.

RandLA++ [6] and involved the following procedures.

To start, intermediate feature sets F^l , $F^{l'}$, and F_{ap}^l are concatenated into F_{cat}^l . Later, the concatenated features are operated by shared MLP function $m_{ce}()$, followed by ReLU L .

$$f_{ce}^{li} = L(m_{ce}(f_{cat}^{li}, W_{ce})) \quad (5)$$

Subsequently, the output feature f_{ce}^{li} is improved channel-wise by attention pooling using Eq. (6) and Eq. (7) into f_{ceat}^{li} . Ultimately, a fully connected layer yields the combined feature f^{li} . It is formally defined as follows:

$$s^{li}_{ce} = S(f_{ce}^{li}) \quad (6)$$

$$f^{li}_{ceat} = \sum (f_{ce}^{li} \cdot s^{li}_{ce}) \quad (7)$$

$$f^{li}_{ce} = \sigma(fc(f^{li}_{ceat}, W_{fc})) \quad (8)$$

where σ is a sigmoid function.

Finally, element-wise multiplication between F^l and F_{ce}^{li} is applied, which is then added to F^l to generate channel enhanced feature \hat{F}^l .

Lastly, the outputs of both branches, f_c^{li} and f_{ph}^{li} respectively, are added with the corresponding output of m_1 operation of the concatenated messenger features. Then point-wise concatenation is performed with the corresponding remaining points' features. It is defined as follows:

$$f_{m.c}^{li} = f_c^{li} + m_1(f_m^{li}, W_1) \quad (9)$$

$$f_{m.ph}^{li} = f_{ph}^{li} + m_1(f_m^{li}, W_1) \quad (10)$$

In summary, we begin by preparing features through our feature preparation process. These features are then divided into two streams of attention pooling and channel enhancing blocks. This is done to allow for the effective exchange and merging of features needed for both the central and pe-

ripheral streams. Please see the diagram of our connection block in Figure 3 for further illustration.

3.5. Feature Merging

The output features \hat{F}_c and \hat{F}_{ph} from the central and peripheral streams, respectively, have to be merged back to match the input size of the human vision based input. We utilize squeeze and excitation method [15] to perform feature merging.

First, we extract $F_{m.c}$ and $F_{m.ph}$, features of messenger points, from \hat{F}_c and \hat{F}_{ph} respectively, resulting in corresponding remained features F_{c-m} and F_{ph-m} . Next, we perform channel-wise concatenation between $F_{m.c}$ and $F_{m.ph}$ to obtain F_m . We then use shared MLPs to smooth the feature spaces of \hat{F}_m , F_{c-m} , and F_{ph-m} , followed by point-wise concatenation. Subsequently, we apply a squeeze and excitation module to enhance the salient features. The feature merging operation is illustrated in Figure 5.

4. Experiments

4.1. Implementation details

We implemented EyeNet in Tensorflow, employing the grid sampling strategy as input pre-processing, followed by KPConv [33] and RandLA [17]. The Adam optimizer is employed, with an initial learning rate of 0.005, which decays by 5% after each epoch. For each layer, the number of nearest points K is set to [16, 21, 21, 21, 16]. We sample a fixed number of 28672 points from each point cloud as input, and use a batch size of 16. Our choice of loss function is lovasz-softmax loss [4]. We followed the same voxel size and input features with ones used in RandLA [17] except YUTO dataset. All experiments are performed on an

Method	OA(%)	mIoU(%)	ground	veg.	building	wall	bridge	parking	rail	traffic	street	car	footpath	bike	water
PointNet [25]	80.78	23.71	67.96	89.52	80.05	0.00	0.00	3.95	0.00	31.55	0.00	35.14	0.00	0.00	0.00
PointNet++ [26]	84.30	32.92	72.46	94.24	84.77	2.72	2.09	25.79	0.00	31.54	11.42	38.84	7.12	0.00	56.93
TangentConv [31]	76.97	33.30	71.54	91.38	75.90	35.22	0.00	45.34	0.00	26.69	19.24	67.58	0.01	0.00	0.00
SPGraph [19]	85.27	37.29	69.93	94.55	88.87	32.83	12.58	15.77	15.48	30.63	22.96	56.42	0.54	0.00	44.24
SparseConv [13]	88.66	42.66	74.10	97.90	94.20	63.30	7.50	24.20	0.00	30.10	34.00	74.40	0.00	0.00	54.80
KPConv [33]	93.20	57.58	87.10	98.31	95.33	74.40	28.69	41.38	0.00	55.99	54.43	85.67	40.39	0.00	86.30
RandLA [17]	89.78	52.69	80.11	98.07	91.58	48.88	40.75	51.62	0.00	56.67	33.23	80.14	32.63	0.00	71.31
RandLA++ [6]	91.90	57.10	84.10	98.20	94.80	58.60	59.80	53.40	0.00	54.60	42.60	78.20	38.20	0.00	69.70
EyeNet(ours)	93.70	62.30	86.60	98.60	96.20	65.80	59.20	64.80	17.90	64.80	49.80	83.10	46.20	11.10	65.40

Table 1. Sensat Urban Performance Comparison. Performances of other works except RandLA++ are taken from the SensatUrban paper [16] RandLA++ performance is taken from its paper. [6]

Method	OA(%)	mIoU(%)	road	road mark.	natural	building	utility line	pole	car	fence
PointNet++ [26]	92.56	59.47	92.90	0.00	86.13	82.15	60.96	62.81	76.41	14.43
DGCNN [37]	94.24	61.79	93.88	0.00	91.25	80.39	62.40	62.32	88.26	15.81
KPConv [33]	95.39	69.11	94.62	0.06	96.07	91.51	87.68	81.56	85.66	15.72
MS-PCNN [22]	90.03	65.89	93.84	3.83	93.46	82.59	67.80	71.95	91.12	22.50
TGNet [20]	94.08	61.34	93.54	0.00	90.83	81.57	65.26	62.98	88.73	7.85
MS-TGNet [30]	95.71	70.50	94.41	17.19	95.72	88.83	76.01	73.97	94.24	23.64
RandLA [17]	92.95	77.71	94.61	42.62	96.89	93.01	86.51	78.07	92.85	37.12
Rim et al. [27]	72.55	66.87	92.74	14.75	88.66	93.52	81.03	67.71	39.65	56.90
MappingConvSeg [38]	93.17	77.57	95.02	39.27	96.77	93.32	86.37	79.11	89.81	40.89
EyeNet(ours)	94.63	81.13	96.98	65.02	97.83	93.51	86.77	84.86	94.02	30.01

Table 2. Toronto3D Performance Comparison. RGB features are not used to obtain results. Performances of other works are taken from the Toronto3D paper [30]

NVIDIA Quadro RTX 6000 GPU.

4.2. Semantic Segmentation on Benchmark datasets

To evaluate the performance of EyeNet, we conducted experiments on the semantic segmentation task. We selected a range of outdoor benchmark datasets, including SensatUrban [16], Toronto3D [30], DALES [35], as well as our own YUTO dataset.

Sensat Urban: The SensatUrban [16] dataset for the semantic segmentation task consists of large areas of three UK cities with a coverage area of 7.64 km^2 . 2847 million points were collected using a UAV Photogrammetry and have an average point density of 372 points per square meter. Each point in this dataset is assigned a semantic label of 13 classes (ground, vegetation, building, etc.). The point clouds are divided into 34 tiles, and we follow their work’s same train, validation, and test split. For evaluation, we upload the test labels on the online test server [2] and acquire mean class-wise intersection over union (mIoU), overall point-wise accuracy (OA), and per class IoU. We set the voxel size of the grid sampling to 0.2m, and use the provided 3D coordinates and color information for training and testing our network.

The performance comparison results for SensatUrban

are presented in Table 1. EyeNet has set a new state-of-the-art for both OA and mIoU metrics, surpassing all existing methods by a considerable margin. Remarkably, EyeNet has shown the best performance in seven out of thirteen semantic classes. The SensatUrban dataset is notorious for having severe class imbalance, and none of the previous approaches were able to predict the rail and bike classes, except for SPGraph [19]. In contrast, EyeNet successfully predicted these classes without any significant data engineering efforts. Furthermore, recent works such as RandLA [17], KPConv [33], Point Transformers [40], Omni-Supervised [11], FG-Net [21], and SCF-Net [8] have increased the number of input point clouds up to 100K for expanding the input coverage area. Despite using an input size of approximately 28K, which is far smaller than other approaches, EyeNet demonstrated superior performance in both small and large object classes. An illustration of segmented scenes on Sensat Urban is shown in Fig. 6.

Toronto3D: Toronto3D [30] is a publicly available dataset that contains 3D point clouds and associated images of urban environments in Toronto, Canada. The dataset covers a total length of 1 kilometer and was collected using a mobile lidar scanner. The dataset contains 78.3 million points with an average point density of 391 points per square me-

Method	OA(%)	mIoU(%)	ground	vegetation	car	truck	power lines	poles	fences	building
KPCConv [33]	97.8	81.1	97.1	96.6	85.3	41.9	75.0	95.5	63.5	94.1
RandLA [17]	97.09	79.27	96.99	93.16	82.22	38.6	94.96	58.01	73.56	96.64
PointNet++ [26]	95.7	68.3	94.1	89.1	75.4	30.3	40.0	79.9	46.2	91.2
ConvPoint [5]	97.2	67.4	96.9	96.3	75.5	21.7	40.3	86.7	29.6	91.9
SPGraph [19]	95.5	60.6	94.7	93.4	62.9	18.7	28.5	65.2	33.6	87.9
PointCNN [30]	97.2	58.4	97.5	95.7	40.6	4.8	57.6	26.7	52.6	91.7
ShellNet [17]	96.4	57.4	96.0	95.4	32.2	39.6	20.0	27.4	60.0	88.4
EyeNet(ours)	97.18	79.6	97.15	93.37	83.53	41.3	94.65	57.38	72.54	96.87

Table 3. DALES Performance Comparison. RandLA results are directly obtained from the author of RandLA. All other results are taken from DALES paper. [35]

Method	mIoU(%)	ground	vegetation	building	water	car	truck	traffic road	sidewalk	parking
RandLA [17]	58.37	80.61	94.44	95.39	3.34	74.59	13.87	78.10	23.43	61.56
KPCConv [33]	56.14	86.94	96.25	94.01	0.00	84.02	0.00	79.93	3.26	60.83
EyeNet(ours)	63.44	86.26	95.94	96.78	13.61	83.02	14.26	84.65	31.08	65.34

Table 4. YUTO Performance Comparison. Results of RandLA and KPCConv are taken from internal experiments.

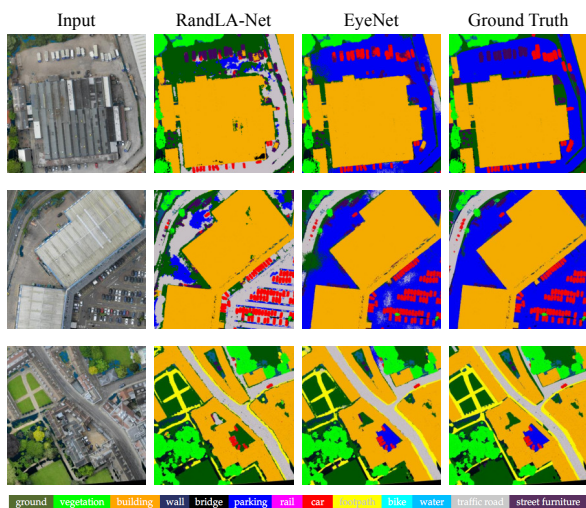


Figure 6. Visualization comparison of SensatUrban between RandLA and EyeNet. From left to right: the input point cloud with RGB colours, semantic segmentation results of RandLA, semantic segmentation results of EyeNet, the ground truth of the input pointcloud

ter. Each point in the dataset is labelled with one of 8 semantic classes such as road, road markings, or natural objects. The dataset is divided into four different sections, and the L002 section was selected as the test split for this study. The evaluation was done by calculating mIoU, OA, and per class IoU for the test split. A voxel size of 0.06 meters was used for grid sampling, and the provided 3D coordinates and intensity information were used for training and testing the network.

The results of the performance comparison for Toronto3D are presented in Table 2. The EyeNet model achieved a new state-of-the-art performance for mIoU metric, surpassing all existing methods by a considerable margin. Notably, EyeNet demonstrated the best performance in four out of eight semantic classes. The Toronto3D dataset includes a unique class called road marking, which is challenging to classify without RGB features, and previous approaches have struggled to accurately classify it. However, EyeNet successfully predicted road markings, highlighting its superior performance.

DALES: The Dayton Annotated LiDAR Earth Scan (DALES) [35] dataset is a publicly available dataset of LiDAR point clouds covering an area of approximately 10 square kilometers in Dayton, Ohio, USA and was collected using an Aerial Laser Scanner (ALS). The dataset includes over 505 million points and has an average point density of 25.5 points per square meter. Each point in the dataset is assigned a semantic label from one of 8 classes, including ground, building, vegetation, water, and others. The dataset is divided into 40 sections and split into training, validation, and test sets for semantic segmentation tasks. The evaluation was done by calculating mIoU, OA, and per class IoU for the test split. A voxel size of 0.32 meters was used for grid sampling, and the provided 3D coordinates was used for training and testing the network.

The performance comparison results for DALES are presented in Table 3. The EyeNet model achieved the second-best performance for mIoU metrics. However, since DALES is the most sparse dataset among the tested benchmarks, the benefits of using the human vision based input in EyeNet are minimal.

YUTO: The York University Toronto Ontario (YUTO) dataset is a privately owned dataset of LiDAR point clouds covering an area of approximately 9.5 square kilometers in the York University campus. The dataset contains 738 million points collected using an ALS and has an average point density of 78 points per square meter. Each point in the dataset is assigned a semantic label from one of nine classes, including ground, traffic road, sidewalk, water, and others. The dataset is divided into 41 tiles, with 11 tiles used for evaluation. The evaluation was performed by calculating mIoU and per-class IoU for the test split. The performance of two popular semantic segmentation networks, RandLA [17] and KPConv [33], was also evaluated for comparison. A voxel size of 0.20 meters was used for grid sampling, and the provided 3D coordinates and intensity were used for training and testing the network.

The results of the performance comparison for YUTO are presented in Table 4. The EyeNet model outperformed the two tested networks by a significant margin in terms of the mIoU metric. It is noteworthy that EyeNet exhibited the highest performance in six out of the nine semantic categories. The YUTO dataset encompasses classes related to the terrain, including traffic road, ground, sidewalk, parking, and water, which is a distinctive feature of ALS datasets. It poses a challenge to classify these classes without the RGB channel accurately. Nevertheless, EyeNet achieved superior results in recognizing these terrain-related classes.

4.3. Ablation Study

An ablation study was conducted to evaluate the effectiveness of the introduced features in the Sensat Urban test set through an online evaluation server.

In order to demonstrate the efficacy of our parallel processing network architecture, we constructed and compared three different types of network structures: Baseline, Sequential, and Parallel. The Baseline structure utilized the pre-existing RandLA-Net architecture. The Sequential structure consisted of two RandLA-Net structures connected sequentially. The Sequential structure was constructed to demonstrate that the performance gain from EyeNet was not solely due to increased parameters. The Parallel structure stacked two RandLA-Net structures in parallel like Fig. 2.

The study compared the performance of four network structures, including the baseline network RandLA, sequentially stacked RandLAs, EyeNet with only Feature Merging Block, and EyeNet. The results, presented in Table 5, indicated that the sequential structure resulted in a 6.7% decrease in mIoU, implying that increasing parameter size alone did not improve performance. In contrast, our EyeNet with only Feature Merging Block yielded improvements of 3.2% and 8.8% in OA and mIoU, respec-

Peri. Input	Structure	CB	OA	mIoU
✗	Baseline	✗	89.8	52.7
✗	Sequential	✗	89.4	46.0
✓	Parallel	✗	93.0	61.5
✓	Parallel	✓	93.7	62.3

Table 5. Ablation results of the peripheral input (Peri. Input), network structure, and connection block (CB).

tively, over the baseline. These results suggest that our human vision-inspired approach contributed to performance enhancement. Additionally, the connection block led to further improvements of 0.7% and 0.8% in OA and mIoU, respectively, indicating that the connection between the two streams improved the network’s feature understanding and led to a performance boost.

5. Conclusion

Determining an optimal coverage area for the input batch is crucial for effective feature learning in 3D point cloud semantic segmentation, but it is often neglected. To address this issue, we have developed a human vision based input processing semantic segmentation network architecture, which overcomes the limitations of traditional semantic segmentation networks. We have also introduced a merging block, connection block and parallel stack design for effective feature learning. As a result, our approach has achieved state-of-the-art performance on large scale outdoor benchmark datasets such as SensatUrban and Toronto3D. However, it is possible that the performance of our network may be limited by the baseline network. Therefore, in future work, we plan to investigate the application of our method to different network architectures since our method is model-agnostic.

6. Acknowledgment

This research project has been supported by the Natural Sciences and Engineering Research Council of Canada (NSERC)’s Collaborative Research and Development Grant (CRD) - 3D Mobility Mapping Artificial Intelligence (3DMMAI) and Teledyne Geospatial Inc. We’d like to thank Alvin Poernomo (Machine Learning Developer), Hamdy Elsayed (Innovation Manager) and Chris Verheggen (SVP R&D).

References

- [1] Randla-net github repository. <https://github.com/QingyongHu/RandLA-Net>. 2

- [2] Sensaturban evaluation server. <https://github.com/QingyongHu/SensatUrban>. Accessed: 2022-11-11. 6
- [3] Jens Behley, Martin Garbade, Andres Milioto, Jan Quenzel, Sven Behnke, Cyrill Stachniss, and Jurgen Gall. Semantickitti: A dataset for semantic scene understanding of lidar sequences. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9297–9307, 2019. 1
- [4] Maxim Berman, Amal Rannen Triki, and Matthew B Blaschko. The lovász-softmax loss: A tractable surrogate for the optimization of the intersection-over-union measure in neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4413–4421, 2018. 5
- [5] Alexandre Boulch. Convpoint: Continuous convolutions for point cloud processing. *Computers & Graphics*, 88:24–34, 2020. 7
- [6] Jiaqing Chen, Yindi Zhao, Congtang Meng, and Yang Liu. Multi-feature aggregation for semantic segmentation of an urban scene point cloud. *Remote Sensing*, 14(20):5134, 2022. 2, 3, 5, 6
- [7] Arturo Deza and Talia Konkle. Emergent properties of foveated perceptual systems. *arXiv preprint arXiv:2006.07991*, 2020. 2
- [8] Siqi Fan, Qiulei Dong, Fenghua Zhu, Yisheng Lv, Peijun Ye, and Fei-Yue Wang. Scf-net: Learning spatial contextual features for large-scale point cloud segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14504–14513, 2021. 6
- [9] Lex Fridman, Benedikt Jenik, Shaiyan Keshvari, Bryan Reimer, Christoph Zetsche, and Ruth Rosenholtz. Sideeye: A generative neural network based simulator of human peripheral vision. *arXiv preprint arXiv:1706.04568*, 2017. 2
- [10] Weixiao Gao, Liangliang Nan, Bas Boom, and Hugo Ledoux. Sum: A benchmark dataset of semantic urban meshes. *ISPRS Journal of Photogrammetry and Remote Sensing*, 179:108–120, 2021. 1
- [11] Jingyu Gong, Jiachen Xu, Xin Tan, Haichuan Song, Yanyun Qu, Yuan Xie, and Lizhuang Ma. Omni-supervised point cloud segmentation via gradual receptive field component reasoning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11673–11682, 2021. 6
- [12] Stephen Gould, Joakin Arfvidsson, Adrian Kaehler, Benjamin Sapp, Marius Messner, Gary Bradski, Paul Baumstarck, Sukwon Chung, Andrew Y Ng, et al. Peripheral-foveal vision for real-time object recognition and tracking in video. 2007. 2
- [13] Benjamin Graham, Martin Engelcke, and Laurens Van Der Maaten. 3d semantic segmentation with submanifold sparse convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9224–9232, 2018. 2, 6
- [14] Timo Hackel, Nikolay Savinov, Lubor Ladicky, Jan D Wegner, Konrad Schindler, and Marc Pollefeys. Semantic3d. net: A new large-scale point cloud classification benchmark. *arXiv preprint arXiv:1704.03847*, 2017. 1
- [15] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7132–7141, 2018. 5
- [16] Qingyong Hu, Bo Yang, Sheikh Khalid, Wen Xiao, Niki Trigoni, and Andrew Markham. Towards semantic segmentation of urban-scale 3d point clouds: A dataset, benchmarks and challenges. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4977–4987, 2021. 1, 2, 6
- [17] Qingyong Hu, Bo Yang, Linhai Xie, Stefano Rosa, Yulan Guo, Zhihua Wang, Niki Trigoni, and Andrew Markham. Randla-net: Efficient semantic segmentation of large-scale point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11108–11117, 2020. 1, 2, 3, 4, 5, 6, 7, 8
- [18] Tengting Huang, Zhe Liu, Xiwu Chen, and Xiang Bai. Ep-net: Enhancing point features with image semantics for 3d object detection. *CoRR*, abs/2007.08856, 2020. 3
- [19] Loic Landrieu and Martin Simonovsky. Large-scale point cloud semantic segmentation with superpoint graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4558–4567, 2018. 6, 7
- [20] Ying Li, Lingfei Ma, Zilong Zhong, Dongpu Cao, and Jonathan Li. Tgnet: Geometric graph cnn on 3-d point cloud segmentation. *IEEE Transactions on Geoscience and Remote Sensing*, 58(5):3588–3600, 2019. 6
- [21] Kangcheng Liu, Zhi Gao, Feng Lin, and Ben M Chen. Fg-net: Fast large-scale lidar point clouds understanding network leveraging correlated feature mining and geometric-aware modelling. *arXiv preprint arXiv:2012.09439*, 2020. 6
- [22] Lingfei Ma, Ying Li, Jonathan Li, Weikai Tan, Yongtao Yu, and Michael A Chapman. Multi-scale point-wise convolutional neural networks for 3d object segmentation from lidar point clouds in large-scale environments. *IEEE Transactions on Intelligent Transportation Systems*, 22(2):821–836, 2019. 6
- [23] Juhong Min, Yucheng Zhao, Chong Luo, and Minsu Cho. Peripheral vision transformer. *arXiv preprint arXiv:2206.06801*, 2022. 2
- [24] M Hamed Mozaffari and Won-Sook Lee. Semantic segmentation with peripheral vision. In *Advances in Visual Computing: 15th International Symposium, ISVC 2020, San Diego, CA, USA, October 5–7, 2020, Proceedings, Part II 15*, pages 421–429. Springer, 2020. 2
- [25] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 652–660, 2017. 1, 2, 6
- [26] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in neural information processing systems*, 30, 2017. 1, 2, 6, 7
- [27] Beanbonyka Rim, Ahyoung Lee, and Min Hong. Semantic segmentation of large-scale outdoor point clouds by encoder–decoder shared mlps with multiple losses. *Remote Sensing*, 13(16):3121, 2021. 6

- [28] Franz Rottensteiner, Gunho Sohn, Jaewook Jung, Markus Gerke, Caroline Baillard, Sebastien Benitez, and Uwe Breitkopf. The isprs benchmark on urban object classification and 3d building reconstruction. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences I-3 (2012)*, Nr. 1, 1(1):293–298, 2012. [1](#)
- [29] Andrés Serna, Beatriz Marcotegui, François Goulette, and Jean-Emmanuel Deschaud. Paris-rue-madame database: a 3d mobile laser scanner dataset for benchmarking urban detection, segmentation and classification methods. In *4th international conference on pattern recognition, applications and methods ICPRAM 2014*, 2014. [1](#)
- [30] Weikai Tan, Nannan Qin, Lingfei Ma, Ying Li, Jing Du, Guorong Cai, Ke Yang, and Jonathan Li. Toronto-3d: A large-scale mobile lidar dataset for semantic segmentation of urban roadways. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 202–203, 2020. [1](#), [6](#), [7](#)
- [31] Maxim Tatarchenko, Jaesik Park, Vladlen Koltun, and Qian-Yi Zhou. Tangent convolutions for dense prediction in 3d. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3887–3896, 2018. [6](#)
- [32] Hugues Thomas. *Learning new representations for 3D point cloud semantic segmentation*. PhD thesis, Université Paris sciences et lettres, 2019. [2](#)
- [33] Hugues Thomas, Charles R Qi, Jean-Emmanuel Deschaud, Beatriz Marcotegui, François Goulette, and Leonidas J Guibas. Kpconv: Flexible and deformable convolution for point clouds. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6411–6420, 2019. [1](#), [2](#), [5](#), [6](#), [7](#), [8](#)
- [34] Bruno Vallet, Mathieu Brédif, Andrés Serna, Beatriz Marcotegui, and Nicolas Paparoditis. Terramobilita/iqmulus urban point cloud analysis benchmark. *Computers & Graphics*, 49:126–133, 2015. [1](#)
- [35] Nina Varney, Vijayan K Asari, and Quinn Graehling. Dales: A large-scale aerial lidar data set for semantic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 186–187, 2020. [1](#), [6](#), [7](#)
- [36] Lei Wang, Yuchun Huang, Yaolin Hou, Shenman Zhang, and Jie Shan. Graph attention convolution for point cloud semantic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10296–10305, 2019. [2](#)
- [37] 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 Transactions On Graphics (tog)*, 38(5):1–12, 2019. [6](#)
- [38] Kunping Yan, Qingyong Hu, Hanyun Wang, Xiaohong Huang, Li Li, and Song Ji. Continuous mapping convolution for large-scale point clouds semantic segmentation. *IEEE Geoscience and Remote Sensing Letters*, 19:1–5, 2021. [6](#)
- [39] Jin Hyeok Yoo, Yecheol Kim, Ji Song Kim, and Jun Won Choi. 3d-cvf: Generating joint camera and lidar features using cross-view spatial feature fusion for 3d object detection. *CoRR*, abs/2004.12636, 2020. [3](#)
- [40] Hengshuang Zhao, Li Jiang, Jiaya Jia, Philip Torr, and Vladlen Koltun. Point transformer, 2020. [1](#), [2](#), [6](#)