PU-EDGEFORMER: EDGE TRANSFORMER FOR DENSE PREDICTION IN POINT CLOUD UPSAMPLING

Dohoon Kim[1], Minwoo Shin[1], and Joonki Paik[1,2]

[1]Department of Image, Chung-Ang University, Seoul, Korea
[2]Department of Artificial Intelligence, Chung-Ang University, Seoul, Korea

ABSTRACT

Despite the recent development of deep learning-based point cloud upsampling, most MLP-based point cloud upsampling methods have limitations in that it is difficult to train the local and global structure of the point cloud at the same time. To solve this problem, we present a combined graph convolution and transformer for point cloud upsampling, denoted by PU-EdgeFormer. The proposed method constructs EdgeFormer unit that consists of graph convolution and multi-head self-attention modules. We employ graph convolution using Edge-Conv, which learns the local geometry and global structure of point cloud better than existing point-to-feature method. Through in-depth experiments, we confirmed that the proposed method has better point cloud upsampling performance than the existing state-of-the-art method in both subjective and objective aspects. The code is available at `https://github.com/dohoon2045/PU-EdgeFormer`.

Index Terms— Point cloud upsampling, neural networks, graph convolution, vision transformer

1. INTRODUCTION

Point clouds are the most widely used representation of three-dimensional (3D) data acquired with 3D sensors because they use less memory and computational cost than meshes or voxels. In recent research, point cloud is used as input to various 3D applications such as autonomous driving, 3D reconstruction, virtual/augmented reality, and robotics. However, the acquired point cloud is spatially sparse, and the coordinates are noisy and non-uniform because of various 3D acquisition problems including: i) occlusion, ii) light reflection, iii) limited hardware, and iv) computational cost. When this point cloud is applied to 3D applications related to Neural Network as it is, information cannot be clustered due to sparsity and non-uniform characteristics. Therefore, it is very important

to pre-process the raw point cloud because it causes performance degradation by failing to learn the contextual manifold of the local structure. Successful point cloud upsampling is an important preprocessing process that can improve the performance of various 3D tasks because it generates dense, uniform, and noise-free point clouds.

Early point cloud upsampling used optimization methods [1, 2]. After Qi *et al*. proposed PointNet, which successfully applied the deep learning method to the point cloud [3], Yu *et al*. first proposed deep learning-based point cloud upsampling method called PU-Net [4]. Based on [3, 4], Yu *et al*. proposed EC-Net to perform upsampling while preserving the edge of the point cloud [5]. Yifan *et al*. proposed MPU upsampling progressively for each point patch [6], and Li *et al*. proposed PU-GAN to which GAN and attention unit were applied [7]. Qian *et al*. proposed PUGeo-Net, which converts samples in 2D domain into 3D by applying linear transformation [8], and Qian *et al*. suggested NodeShuffle to the point shuffle process of the upsampling module to improve the performance of PU-GCN [9].

Although the performance of deep learning-based point cloud upsampling is continuously improving, the 2D convolution operation used for feature extraction in most upsampling methods does not reflect the geometric relationship of the input point cloud because each point is learned independently. To solve this problem, MPU and PU-GCN have proposed methods of reflecting the local geometry of the input point cloud as a point feature, but there is still a problem that the global structure information of the input point cloud is not reflected as a point feature.

In this paper, we propose a novel unit called EdgeFormer that extracts features from point cloud for successful upsampling. Using EdgeFormer we calculate the attention score from the input point through the multi-head self-attention of transformer [10], and designed the network to reflect the global structure in the point feature and also the local geometric structure through EdgeConv [11].

2. PROPOSED METHOD

Given a point cloud of N points, the point cloud upsampling process with an upsampling ratio r produces a dense
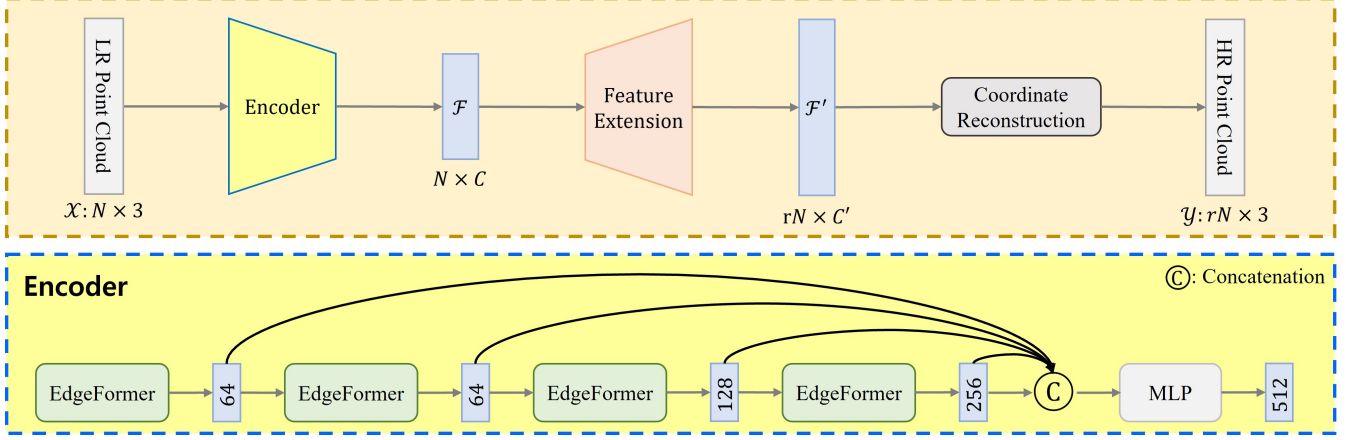
Fig. 1: Architecture of PU-EdgeFormer (top) and the encoder module (bottom)

point cloud $\mathcal{Y} \in \mathbb{R}^{rN \times 3}$ based on the geometric information of the sparse point cloud $\mathcal{X} \in \mathbb{R}^{N \times 3}$. A well-generated high-resolution (HR) point cloud $\mathcal{Y}$ should not simply have increased number of points, but should have uniformly distributed points to have similar edge surface to the ground truth. Since this upsampling process corresponds to a very challenging ill-posed problem, in order to solve it with a deep learning method, the point feature must well represent the geometry of the low-resolution (LR) point cloud $\mathcal{X}$. Unlike the existing upsampling method that extracts point features by applying MLP to LR point cloud $\mathcal{X}$, we calculate the point-to-point score using the multi-head self-attention proposed by transformer [10] with the point feature to make the point cloud feature represent the global structure of the point cloud well. However, the point features extracted using transformer do not have the local geometric structure information of the point cloud. To solve this problem, we applied EdgeConv proposed in DGCNN [11] to multi-head self-attention operation. EdgeConv composes the local neighborhood graph of the point cloud, performs a convolution operation on it, and aggregates it so that the point feature represents the geometrical relationship between points in the point cloud. As a result, the proposed EdgeFormer calculates the entire relationship of points through multi-head self-attention, so that the global structure of $\mathcal{X}$ can be reflected in the point features. Also, it is designed to reflect the local geometric information of $\mathcal{X}$ to point features by applying EdgeConv.

2.1. Network Architecture

The proposed point cloud upsampling EdgeFormer (PU-EdgeFormer) consists of: i) encoder, ii) feature extension, and iii) coordinate reconstruction as shown in Fig. 1. Each step of PU-EdgeFormer is described as follows:

Encoder Since an LR point cloud $\mathcal{X}$ is a structure in which the coordinate values x, y, z of the three-dimensional Euclidean space are simply arranged as row vectors, it is nec-

essary to convert it into a point feature, which is a point set of the latent space. In this paper, we construct the module to extract the point feature $\mathcal{F} \in \mathbb{R}^{N \times C}$ through MLP operation by concatenating the LR point cloud through four EdgeFormers. We will discuss about the EdgeFormer in section 2.2.

Feature Extension To upscale the extracted point feature $\mathcal{F}$, the feature extension module reshapes $\mathcal{F}$ to have a dimension $rN \times C/r$ using a shuffling operator, and then applies the MLP operator, $\Phi_i(\cdot; \theta_i)$, for $i = 1, \cdots, n$. A mathematical expression of this process is given as

$$\mathcal{F}_i = \Phi_i(\mathcal{F}_{i-1}; \theta_i), \tag{1}$$

where θ_i denotes a learnable parameter that operated by input point feature $\mathcal{F}_{i-1}$. In this paper, we set $n = 2$, and used 256 and 32 as the output dimension of each MLP. $\mathcal{F}_n$ in (1) becomes feature extension result $\mathcal{F}' \in \mathbb{R}^{rN \times C'}$.

Coordinate Reconstruction This module generates HR point cloud $\mathcal{Y}$ through the result of feature extension $\mathcal{F}'$ and LR point cloud $\mathcal{X}$. The module makes HR point cloud $\mathcal{Y}$ through summation between final Point feature that reshaped by MLP operator $\Phi(\cdot; \theta)$ on $\mathcal{F}'$ and $\tilde{\mathcal{X}} \in \mathbb{R}^{rN \times 3}$ that duplicated r times on N points:

$$\mathcal{Y} = \tilde{\mathcal{X}} + \Phi(\mathcal{F}'; \theta). \tag{2}$$

As a result, point feature $\mathcal{F}'$ that added with duplicated LR point cloud $\tilde{\mathcal{X}}$ are important to upsample point cloud, where $\mathcal{F}'$ is a point set of latent space obtained by the encoder module. Because good upsampling results can be derived when the geometry of the LR point cloud is well reflected as a feature, the encoder plays a very important role in the point cloud upsampling work.

2.2. EdgeFormer

The encoder shown in the bottom of Fig. 1 consists of four EdgeFormer units. An EdgeFormer unit is shown in Fig. 2.
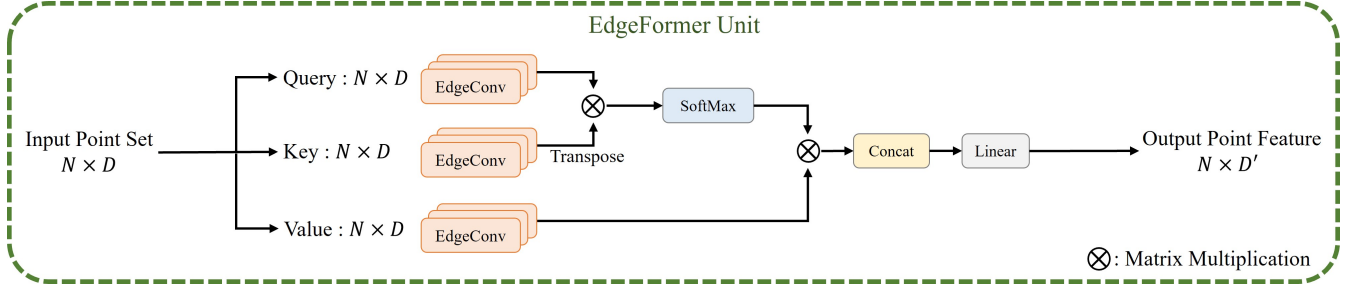
Fig. 2: Proposed EdgeFormer unit.

Unlike an MLP applied to the existing 2D image transformer, in the point cloud, each row vector appears as a point, which is permutation invariance. So the input point set becomes a query, key, and value directly without going through the positional encoding process. The EdgeConv operation is applied to query, key, and value, and the results are divided by the number of heads h. The query and key divided into each head are matrix multiplied, and the attention score for the current point set is calculated through SoftMax operation. By multiplication of this value and matrix, point features are derived in the direction where the value of attention score is high. This operation is repeated for each head and all the calculated results are concatenated. Finally, the output point feature is derived through linear operation.

The proposed EdgeFormer replaces the linear operation for query, key, and value with EdgeConv. After grouping points using k-nearest neighbor (k-NN) operation, $\mathcal{N}(\cdot; k)$, EdgeConv performs convolution, $\psi(\cdot; \theta_f)$, of the grouping results and the learnable parameter θ_f, and then aggregates the results using max pooling.

$$f' = \text{MaxPool}(\psi(\mathcal{N}(f; k); \theta_f)), \qquad (3)$$

where f denotes each input of query, key, and value, and f' the result of EdgeConv. The EdgeConv operation in (3) can solve the problem of losing local features when MLP is used. We designed the EdgeFormer unit to not only represent the local geometric structure of the input by applying the EdgeConv operation to each input, but also represent the global structure through the multi-head self-attention operation.

3. EXPERIMENTS

3.1. Implementation Details

Datasets. We used PU1K dataset that includes large-scale objects with complex shapes provided with PU-GCN [9]. The PU1K dataset was collected by PU-GAN [7] and ShapeNet-Core [12]. We set the number of points in the LR point cloud of the training dataset to 256 and the number of points in the ground truth to 1024. We set the upsampling ratio as $\times 4$ just like the other existing methods.

Experimental Settings. Our experiments were run on the tensorflow platform. As hyper-parameters of the experiment, the batch size was set to 64, the training epochs were set to 100, the learning rate of Adam optimization was set to 0.001, $k = 16$ of EdgeConv in the EdgeFormer unit, and the number of heads h was set to 8. In addition, we performed data augmentation such as rotation, scaling, random perturbations on the training data to avoid overfitting.

Loss Functions To train our network, we used Chamfer distance [4] that computes the distance between the two nearest points in two point sets:

$$
\begin{aligned}
\mathcal{L}(\mathcal{Y}^{Pred}, \mathcal{Y}^{GT}) = {} & \frac{1}{|\mathcal{Y}^{Pred}|} \sum_{p \in \mathcal{Y}^{Pred}} \min_{q \in \mathcal{Y}^{GT}} \|p - q\|_2^2 \\
& + \frac{1}{|\mathcal{Y}^{GT}|} \sum_{p \in \mathcal{Y}^{GT}} \min_{q \in \mathcal{Y}^{Pred}} \|p - q\|_2^2.
\end{aligned} \qquad (4)
$$

Evaluation Metrics. According to the quantitative experimental results of recent papers on point cloud upsampling, we use three metrics: Chamfer distance (CD) [4], Hausdorff distance (HD) [13], and point-to-surface distance (P2F) [14]. For all of these metrics, lower values indicate better performance.

3.2. Comparative Evaluation

Qualitative Results. Fig. 3 shows the subjective quality of four different methods including ours. Compared with the existing methods, we show that our method derives less noise and generates HR point clouds with robust edges.

Quantitative Results. Table 1 compares the quantitative experimental results. Our method shows good performance in CD and HD, similar to state-of-the-art, but shows poor performance in P2F. This is because the edge is excessively smoothed by applying EdgeConv every time in EdgeFormer of Encoder.

3.3. Robustness Test

A robustness test was performed to show that our method is robust to sparse and noisy inputs.

(a) Input	(b) GT	(c) PU-Net	(d) MPU	(e) PU-GCN	(f) Ours
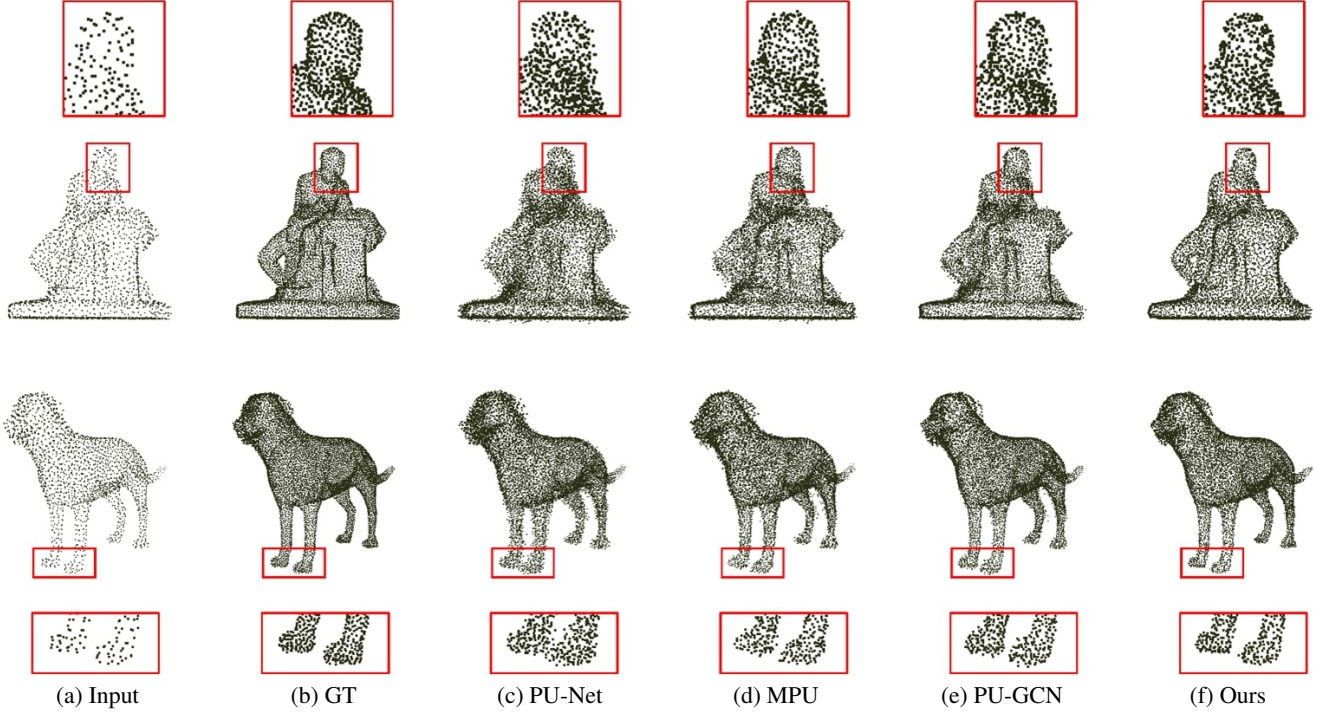
Fig. 3: Qualitative upsampling results. We apply x4 upsampling in (a) input that consists of 2048 points. For precise comparison with other methods, we zoom-in on local parts of point clouds in red boxes.

Table 1: Quantitative upsampling results. ↓ means lower value denotes better performance.

Methods	CD↓ ($\times 10^{-3}$)	HD↓ ($\times 10^{-3}$)	P2F↓ ($\times 10^{-3}$)
PU-Net	1.155	15.170	4.834
MPU	0.935	13.327	3.551
PU-GCN	0.585	7.577	**2.499**
Ours	**0.462**	**3.813**	2.869

Table 2: Quantitative results on noisy input. All values denote CD ($\times 10^{-3}$).

Methods	$\sigma = 0.1$	$\sigma = 0.5$	$\sigma = 1$	$\sigma = 2$
PU-Net	1.124	1.054	1.768	3.901
MPU	0.954	0.911	1.464	3.539
PU-GCN	0.632	0.809	1.416	3.410
Ours	**0.589**	**0.720**	**1.202**	**2.880**

0.5, 1, and 2 on the test dataset. Table 2 demonstrates that our method yields robust results against noisy inputs.

4. CONCLUSION

We proposed EdgeFormer, a new method for extracting features of point clouds. EdgeFormer allows the feature to have a global structure through transformer and a local geometric structure through EdgeConv instead of the MLP used when extracting features from the existing point cloud upsampling. In the experimental results, PU-EdgeFormer showed higher performance than the existing state-of-the-art, but the edge was excessively smoothed. To solve this problem, we will study the method of estimating the manifold in the point cloud and refine it based on this in the future.

(a) Input	(b) × 16	(c) × 256	(d) GT

Fig. 4: Sparsity robustness test.

Sparsity Robustness Test. For a sparse input consisting of 256 points, we use upsampling 16 times and 256 times. Fig. 4 is the visual experimental result, and it can be seen that our method predicts the GT manifold well even when an input is sparse.

Noise Robustness Test. We conducted a robustness test by setting the Gaussian noise with standard deviation of 0.1,

5. REFERENCES

[1] Marc Alexa, Johannes Behr, Daniel Cohen-Or, Shachar Fleishman, David Levin, and Claudio T. Silva, "Computing and rendering point set surfaces," *IEEE Transactions on visualization and computer graphics*, vol. 9, no. 1, pp. 3–15, 2003.

[2] Yaron Lipman, Daniel Cohen-Or, David Levin, and Hillel Tal-Ezer, "Parameterization-free projection for geometry reconstruction," *ACM Transactions on Graphics (TOG)*, vol. 26, no. 3, pp. 22–es, 2007.

[3] 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*, 2017, pp. 652–660.

[4] Lequan Yu, Xianzhi Li, Chi-Wing Fu, Daniel Cohen-Or, and Pheng-Ann Heng, "Pu-net: Point cloud upsampling network," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 2790–2799.

[5] Lequan Yu, Xianzhi Li, Chi-Wing Fu, Daniel Cohen-Or, and Pheng-Ann Heng, "Ec-net: an edge-aware point set consolidation network," in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 386–402.

[6] Wang Yifan, Shihao Wu, Hui Huang, Daniel Cohen-Or, and Olga Sorkine-Hornung, "Patch-based progressive 3d point set upsampling," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 5958–5967.

[7] Ruihui Li, Xianzhi Li, Chi-Wing Fu, Daniel Cohen-Or, and Pheng-Ann Heng, "Pu-gan: a point cloud upsampling adversarial network," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 7203–7212.

[8] Yue Qian, Junhui Hou, Sam Kwong, and Ying He, "Pugeo-net: A geometry-centric network for 3d point cloud upsampling," in *European conference on computer vision*. Springer, 2020, pp. 752–769.

[9] Guocheng Qian, Abdulellah Abualshour, Guohao Li, Ali Thabet, and Bernard Ghanem, "Pu-gcn: Point cloud upsampling using graph convolutional networks," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 11683–11692.

[10] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.

[11] 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)*, vol. 38, no. 5, pp. 1–12, 2019.

[12] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al., "Shapenet: An information-rich 3d model repository," *arXiv preprint arXiv:1512.03012*, 2015.

[13] Matthew Berger, Joshua A Levine, Luis Gustavo Nonato, Gabriel Taubin, and Claudio T Silva, "A benchmark for surface reconstruction," *ACM Transactions on Graphics (TOG)*, vol. 32, no. 2, pp. 1–17, 2013.

[14] Kok-Lim Low, "Linear least-squares optimization for point-to-plane icp surface registration," *Chapel Hill, University of North Carolina*, vol. 4, no. 10, pp. 1–3, 2004.