

Point Cloud Upsampling via Cascaded Refinement Network

Hang Du*, Xuejun Yan*, Jingjing Wang, Di Xie, and Shiliang Pu^(✉)

Hikvision Research Institute, Hangzhou, China
{duhang, yanxuejun, wangjingjing9, xiedi,
pushiliang.hri}@hikvision.com

Abstract. Point cloud upsampling focuses on generating a dense, uniform and proximity-to-surface point set. Most previous approaches accomplish these objectives by carefully designing a single-stage network, which makes it still challenging to generate a high-fidelity point distribution. Instead, upsampling point cloud in a coarse-to-fine manner is a decent solution. However, existing coarse-to-fine upsampling methods require extra training strategies, which are complicated and time-consuming during the training. In this paper, we propose a simple yet effective cascaded refinement network, consisting of three generation stages that have the same network architecture but achieve different objectives. Specifically, the first two upsampling stages generate the dense but coarse points progressively, while the last refinement stage further adjust the coarse points to a better position. To mitigate the learning conflicts between multiple stages and decrease the difficulty of regressing new points, we encourage each stage to predict the point offsets with respect to the input shape. In this manner, the proposed cascaded refinement network can be easily optimized without extra learning strategies. Moreover, we design a transformer-based feature extraction module to learn the informative global and local shape context. In inference phase, we can dynamically adjust the model efficiency and effectiveness, depending on the available computational resources. Extensive experiments on both synthetic and real-scanned datasets demonstrate that the proposed approach outperforms the existing state-of-the-art methods. The code is publicly available at <https://github.com/hikvision-research/3DVision>.

1 Introduction

Point clouds have been widely adopted in many 3D computer vision studies [18,19,37,13,28,39,15] in recent years. However, in real-world scenarios, the raw point clouds produced by the 3D sensors are often sparse, noisy, and non-uniform, which have negative impact on the performance of the point cloud analysis and processing tasks. Therefore, in order to facilitate the downstream point cloud tasks, it is necessary to upsample sparse point clouds to a dense, uniform and high-fidelity point set.

* Equal contribution.




Fig. 1: Upsampling results ($\times 4$) on real-scanned point clouds. The top line is the upsampled point clouds, and the bottom line shows the corresponding 3D surface reconstruction results using the ball-pivoting algorithm [2]. Compared with the recent state-of-the-art methods, we can generate high-fidelity point clouds with uniform distribution, resulting in a more smooth reconstruction surface.

In recent years, many deep learning-based methods have been proposed for point cloud upsampling. Compared with the traditional optimization-based methods [1,14,7,29], deep learning-based methods [34,35,21,11,20] are able to handle more complex geometric structures of 3D objects since they can effectively extract deep features from point clouds and learn to generate new points in a data-driven manner. Among them, as a pioneering point cloud upsampling method, Yu *et al.* [34] propose a classic upsampling framework, named PU-Net, which develops a PointNet-based network to learn multi-scale point features and expand the number of points by multi-branch Multi-layer Perceptrons (MLP). Based on the framework of PU-Net, the recent advanced upsampling methods [11,20,16] have made remarkable progress by designing more effective point feature extractor and upsampling unit. However, these methods directly generate the dense and high-fidelity points via a single-stage network, which makes the network challenging to meet multiple objectives and cannot obtain an optimal result through one-stage structure (Fig. 1).

Coarse-to-fine is a widely-used scheme in many point cloud generative models [38,26,31,30,32]. Certain methods [27,12] also have applied such scheme for point cloud upsampling. Specifically, MPU [27] contains a series of upsampling sub-networks that focus on different level of details. During the training, it needs

to extract the local patches from the previous stage and activate multiple upsampling units progressively, which is complicated and time-consuming, especially on a large upsampling rate. More recently, Li *et al.* [12] propose Dis-PU to disentangle the multiple objectives of upsampling task, which consists of a dense generator for coarse generation and a spatial refiner for high-quality generation. In order to make the results more reliable, Dis-PU utilizes a warm-up training strategy that gradually increases the learning rate of the spatial refiner, which requires a longer training epoch to ensure full convergence of the networks.

In the view of these limitations, we argue that a better coarse-to-fine point upsampling framework should be flexible and easily optimized during the training. To this end, we propose to accomplish coarse-to-fine point cloud upsampling via three cascaded generation stages, each of which has the same network architecture but focuses on different purposes. The overview of the proposed network is shown in Fig. 2. Among each generation stage, we leverage the local self-attention mechanism by designing a transformer-based feature extraction module that can learn both the global and local shape context. As for feature expansion, we employ a two-branch based module to exploit the geometric information and learn shape variations rather than directly duplicating the initial point-wise features.

In addition, the learning conflicts between multiple stages is a major challenge to the coarse-to-fine framework. To tackle this problem, we simply adopt a residual learning scheme for point coordinate reconstruction, which predicts the offset of each point and further adjusts the initial point position. In this manner, the entire cascaded refinement network enables to be optimized without extra training strategies.

Resorting to the above practices, we can build a flexible and effective point cloud upsampling framework. Compared with the existing methods, the proposed framework can obtain significant performance improvements. According to the available computational resources, the feed-forward process of refinement stage is optional. Thus, we can dynamically adjust the model efficiency and effectiveness in inference phase. Extensive experiments on both synthetic and real-scanned datasets demonstrate the effectiveness of our proposed method.

The contributions of this paper are summarized as follows:

- We propose a cascaded refinement network for point cloud upsampling. The proposed network can be easily optimized without extra training strategies. Besides, our method has better scalability to dynamically adjust the model efficiency and effectiveness in inference phase.
- We adopt the residual learning scheme in both point upsampling and refinement stages, and develop a transformer-based feature extraction module to learn both global and local shape context.
- We conduct comprehensive experiments on both synthetic and real-scanned point cloud datasets, and achieve the leading performance among the state-of-the-art methods.

2 Related Work

Our work is related to the point cloud upsampling and the networks in point cloud processing. In this section, we provide a brief review of recent advances in these areas.

2.1 Point Cloud Upsampling

Existing point cloud upsampling methods can be roughly divided into traditional optimization-based approaches [1,14,7,29] and the deep learning-based approaches [34,35,27,11,21,33,20,40,16,22,12]. The former generally requires the geometric prior information of point clouds (*e.g.*, edges and normal), and achieves poor results on the complex shapes. In contrast, the deep learning-based approaches have been prevailing and dominated the recent state-of-the-arts. Particularly, PU-Net [34] proposes a classic learning-based upsampling framework that involves three components, including feature extraction, feature expansion (upsampling unit), and point coordinate reconstruction. Later on, many works [11,21,20,40,16,22] follow this classic framework and focus on improving it from many aspects, such as feature extractor, upsampling scheme, training supervision. For example, PU-GAN [11] introduces a upsampling adversarial network with a uniform loss to ensure the point uniformity, and PU-GCN [20] leverages graph convolutional networks for both feature extraction and expansion. The above methods employ a single-stage network to accomplish the multiple objectives of point cloud upsampling. However, it is challenging to achieve all the objectives at the same time, and coarse-to-fine manner is a more appropriate solution. To this end, MPU [27] and Dis-PU [12] propose to divide the upsampling process into multiple steps, however, they require progressive training or warm-up strategy to ensure a better optimization, which is complicated and time-consuming during the training. In this work, instead of using extra training strategies, we develop a more flexible and effective coarse-to-fine framework.

2.2 Point Cloud Processing

There are three major categories of point cloud processing networks, including projection-based, voxel-based, and point-based networks. Since the point clouds have no spatial order and regular structure, projection-based methods [24,9] and voxel-based methods [31,41,5] transform irregular point cloud to 2D pixel or 3D voxel, and then apply 2D/3D convolution to extract the regular representations. In contrast, point-based networks are designed to process the irregular point cloud directly, which is able to avoid the loss of shape context during the projection or voxelization. Among them, some methods [18,19] employ MLP-based structure to extract point-wise features, and others [28,10] utilize graph-based convolutional networks to aggregate the point neighbourhood information. Besides, certain methods [13,23] develop continuous convolutions that can be directly applied to the point cloud. More recently, the success of vision transformer has inspired many 3D point cloud processing approaches [6,3,17,39,36]. As far

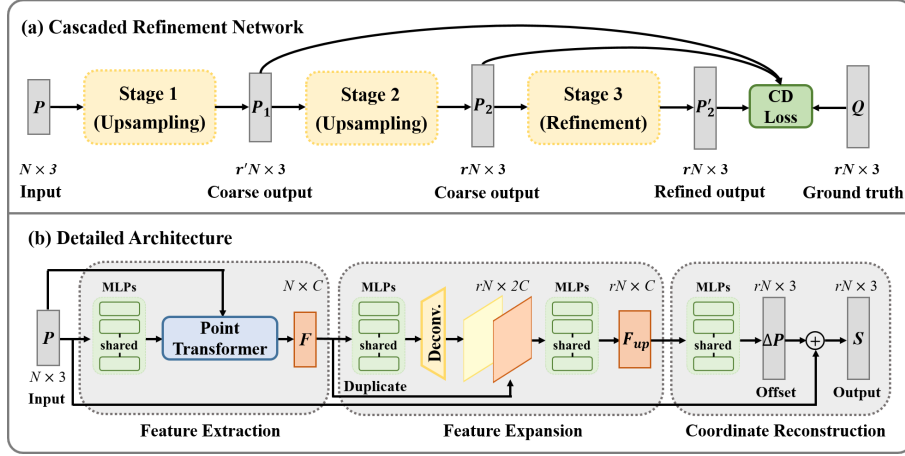


Fig. 2: (a) Overview of our cascaded refinement network for point cloud upsampling, which consists of two upsampling stages and one refinement stage. (b) For each upsampling or refinement stage, the detailed network architecture contains a transformer-based feature extraction module, a two-branch feature expansion module, and a coordinate reconstruction module.

as we can see, the transformer-based structure is under-explored in point cloud upsampling. Thus, based on the architecture of point transformer [39], we build a transformer-based feature extraction module in our upsampling network.

3 Proposed Method

Given a sparse point cloud set $\mathcal{P} = \{p_i\}_{i=1}^N$, where N is the number of points and p_i is the 3D coordinates, the objective of point clouds upsampling is to generate a dense and high-fidelity point cloud set $\mathcal{S} = \{s_i\}_{i=1}^{rN}$, where r is the upsampling rate. Fig. 2 shows the overview of the proposed cascaded refinement network, where the first two upsampling stages generate a coarse dense point cloud set progressively and the last refinement stage serves as a refiner to adjust the coarse points to a better position. During the training, we train the entire framework in an end-to-end manner, and adopt three Chamfer Distance (CD) losses to constrain each stage, simultaneously. In the following sections, we elaborate the detailed network architecture of our framework and the training loss function.

3.1 Network Architecture

As shown in the bottom of Fig. 2, the network architecture of upsampling or refinement stage follows the commonly-used pipeline, which consists of a feature extraction module, a feature expansion module and a coordinate reconstruction

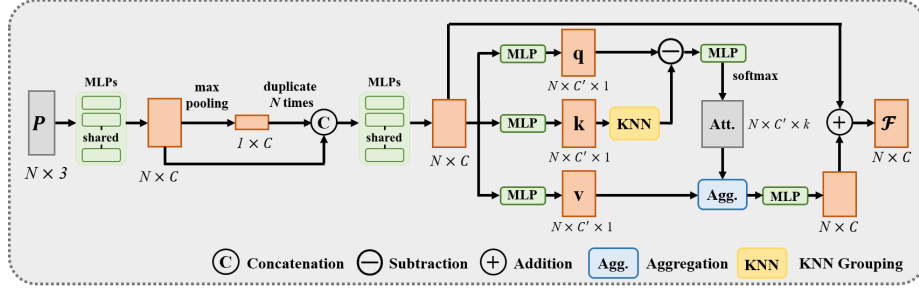


Fig. 3: The detailed network architecture of transform-based feature extraction module. The attention map is calculated locally within the k-nearest neighbors of the current query. We omit the position encoding in value vectors and attention map for the clearness.

module. In the following, we will provide a detailed introduction of each module, respectively.

Feature extraction. Given a input point cloud $\mathcal{P} = \{p_i\}_{i=1}^N$, the feature extraction module aims to encode point features $\mathcal{F} = \{f_i\}_{i=1}^N$ of C channel dimensions. Most existing point cloud upsampling methods [34,11,12] fulfill this objective using a MLP-based structure. However, the local details are limited by the insufficient feature representation capability of MLPs. To solve this problem, some upsampling approaches [28,10,20] employ a GCN-based structure to capture the local neighborhood information for point feature extraction.

In the view of the success of transformer-based networks [17,39,6] in the field of 3D computer vision, we explore a transformer-based feature extraction module for the task of point cloud upsampling. Fig. 3 illustrates the details on our proposed network architecture. Based on the network architecture in [19], we first extract the point-wise features from the input point clouds via a set of MLPs, and then apply a max pooling operation to obtain the global features. The global features are duplicated N times and concatenated with the initial features, subsequently. After that, a point transformer layer [39] is utilized to refine the local shape context and obtain the output point features. By applying self-attention locally, the learned features can incorporate the point neighborhood information between the point feature vectors. Finally, the encoded point features $\mathcal{F} = \{f_i\}_{i=1}^N, f_i \in \mathbb{R}^C$ are added with the input features through a residual connection. The locally self-attention operations can be formulated as:

$$\mathbf{f}_i = \mathbf{x}_i \oplus \sum_{j \in \mathcal{N}(i)} \text{softmax}(\mathbf{q}_i - \mathbf{k}_j) \odot \mathbf{v}_j, \quad (1)$$

where \mathbf{x}_i is i th input point feature vector, \mathbf{q}_i is the corresponding query vector, \mathbf{k}_j is the j th key vector from the k-nearest neighbors of \mathbf{x}_i , and \mathbf{v}_j is the j th value vector. The attention weights are calculated between the query and its k-nearest neighbors by softmax function. Note that we omit the position encoding

in Equation 1. One can refer to [39] for more details of this part. In this way, we are able to encode the point features by combining the local and global information.

Feature expansion. Subsequently, the extracted point features are feed into a feature expansion module that produces the expanded point feature $\mathcal{F}_{\text{up}} = \{f_i\}_{i=1}^{rN}$, where r is the upsampling rate. Most previous upsampling approaches apply duplicate-based [27,11,12] or shuffle [20] operation to increase the numbers of the input points features. In this work, we employ a two-branch based scheme for feature expansion, which combines the advantages of duplicate-based [27,11] and learning-based methods [8,31,30]. Specifically, we first duplicate the input features $\mathcal{F} = \{f_i\}_{i=1}^N$ with r copies. Meanwhile, a transposed convolution branch is used for feature expansion, which is utilized to learn new point features through a learnable manner. Certain methods [8,30] have discussed the advantages of learning-based feature interpolation for generative models, which can exploit more local geometric structure information with respect to the input shape. Finally, we concatenate the point features from two branches and feed them to a set of MLPs which produces the expanded features as follows:

$$\mathcal{F}_{\text{up}} = \text{MLP}(\text{Concat}[\text{Dup}(\mathcal{F}, r); \text{Deconv}(\mathcal{F})]). \quad (2)$$

The duplicated-based branch is able to preserve the initial shape, and the learnable transposed convolution branch can provide local geometric variations on the input point clouds. The combination of them enables to produce a more expressive upsampled features for reconstructing the point coordinate reconstruction subsequently.

Coordinate reconstruction. The objective of coordinate reconstruction is to generate a new point set $\mathcal{S} = \{s_i\}_{i=1}^{rN}$ from the expanded feature vectors $\mathcal{F}_{\text{up}} = \{f_i\}_{i=1}^{rN}$. To accomplish this objective, a common way is to regress the 3D point coordinates directly. However, it is difficult to generate high-fidelity points from latent space without noises [12]. In order to solve this problem, several methods [27,12] apply a residual learning strategy that predicts the offset $\Delta\mathcal{P}$ of each point to adjust the initial position, which can be formulated as:

$$\mathcal{P}_{\text{output}} = \Delta\mathcal{P} + \text{Dup}(\mathcal{P}_{\text{input}}, r), \quad (3)$$

where r is the upsampling rate. Moreover, residual learning is a simple yet effective scheme to mitigate the learning conflicts between multiple stages in a coarse-to-fine framework. Thus, we consider the advantages of such scheme are two folds: (i) learning per-point offset can greatly decrease the learning difficulty of a cascaded network and mitigate the conflicts between multiple stages, which results in a stable and better optimization; (ii) residual connection enables to preserve the geometric information of initial shape and provide variations for generating new points. Hence, we adopt residual learning scheme for coordinate reconstruction, which contains two MLP layers to generate the per-point offset $\Delta\mathcal{P}$ from \mathcal{F}_{up} ($\Delta\mathcal{P} = \text{MLP}(\mathcal{F}_{\text{up}})$), and then adds it on the r times duplicated input points. In this manner, the entire cascaded refinement network can be

easily optimized together without extra training strategies (*e.g.*, progressive or warm-up training).

3.2 Compared with Previous Coarse-to-Fine Approaches

The coarse-to-fine scheme has been studied in the previous point cloud upsampling methods [27,12]. Nonetheless, the proposed approach is more efficient and flexible during the training and inference. Specifically, MPU [27] needs to extract the local training patches from the previous stage, and activates the training of each upsampling stage progressively. Dis-PU [12] adopts a warm-up training strategy to control the importance of dense generator and spatial refiner in different training stage. Both of them are complex during the training and cost a longer time to ensure the full convergence of the networks. Moreover, the complicated training strategies rely on the experience of hyper-parameters tuning, which may be sensitive on the new datasets. In contrast, our advantages are two folds. Firstly, the proposed cascaded refinement network can be easily trained without progressive or warm-up training strategies, which enables to require less training consumption. Secondly, our framework is more flexible and scalable. We can dynamically adjust the model efficiency and effectiveness at the stage of inference. For example, we can pursue a faster inference time only using two-stage upsampling structure, and still achieve a leading performance among other competitors (Table 3). Therefore, the proposed coarse-to-fine upsampling framework can be more effective and practical in real-world applications.

3.3 Training Loss Function

During the training, our approach is optimized in an end-to-end manner. To encourage the upsampled points more distributed over the target surface, we employ Chamfer Distance (CD) [4] \mathcal{L}_{CD} for training loss function, which is formulated as:

$$\mathcal{L}_{CD}(\mathcal{P}, \mathcal{Q}) = \frac{1}{|\mathcal{P}|} \sum_{p \in \mathcal{P}} \min_{q \in \mathcal{Q}} \|p - q\|_2^2 + \frac{1}{|\mathcal{Q}|} \sum_{q \in \mathcal{Q}} \min_{p \in \mathcal{P}} \|p - q\|_2^2, \quad (4)$$

where \mathcal{P} and \mathcal{Q} denote the generated point clouds and their corresponding ground truth, respectively. CD loss aims to calculate the average closest point distance between the two sets of point cloud.

In order to explicitly control the generation process of point clouds, we employ the CD loss on three generation stages, simultaneously. Therefore, the total training loss function \mathcal{L}_{total} can be defined as:

$$\mathcal{L}_{total} = \mathcal{L}_{CD}(\mathcal{P}_1, \mathcal{Q}) + \mathcal{L}_{CD}(\mathcal{P}_2, \mathcal{Q}) + \mathcal{L}_{CD}(\mathcal{P}'_2, \mathcal{Q}), \quad (5)$$

where \mathcal{P}_1 , \mathcal{P}_2 and \mathcal{P}'_2 are the output point clouds from three generation stages, respectively. We validate the effectiveness of the penalty on three stages in the experiment (please refer to supplementary materials). The results show that the simultaneous supervision on three generation stages results in a better performance. Therefore, we adopt CD loss on three generation stages simultaneously.

4 Experiments

In this section, we conduct extensive experiments to verify the effectiveness of the proposed point upsampling network. The section is organized as follows. Section 4.1 introduces the detail experimental settings. Section 4.2 demonstrates the obvious improvements by our approach on two synthetic datasets. Section 4.3 shows the qualitative comparisons on real-scanned data. Section 4.4 provides model complexity analysis with other counterparts. Section 4.5 studies the effect of major components in our approach.

4.1 Experimental Settings

Datasets. We employ two point cloud upsampling datasets, including PU-GAN [11] and PU1K [20]. Among them, PU-GAN dataset contains 24,000 training patches collected from 120 3D models and 27 3D models for testing. In contrast, PU1K dataset is a large-scale point cloud upsampling dataset which is newly released by PU-GCN [20]. PU1K dataset consists of 69,000 training patches collected from 1,020 3D models and 127 3D models for testing. In addition, we also utilize a real-scanned point cloud dataset ScanObjectNN [25] for qualitative evaluation.

Training and evaluation. Our models are trained with 100 epochs. The batch size is set as 64 for PU1K and 32 for PU-GAN, respectively. The learning rate begins at 0.001 and drops by a decay rate of 0.7 every 50k iterations. The ground truth of each training patch contains 1,024 points and the input contains 256 points that are randomly sampled from the ground truth. In practical implementation, we set the upsampling rate $r = 2$ in the upsampling stage and $r = 1$ in the refinement stage, to achieve $\times 4$ upsampling. More detailed settings of our network can be found in the supplementary material. For inference, we follow the common settings [11,20] to divide the input point clouds into multiple patches based on the seed points. Then, the patches are upsampled with r times. After that, FPS algorithm is used to combine all the upsampled patches as the output point clouds. For quantitative evaluation, we apply three commonly used metrics, including Chamfer Distance (CD), Hausdorff Distance (HD), and Point-to-Surface Distance (P2F). A smaller value of these metrics denotes a better performance.

Comparison methods. In the following, we will compare the proposed method with five existing point cloud upsampling methods, including PU-Net [34], MPU [27], PU-GAN [11], Dis-PU [12], PU-GCN [20], and PU-EVA [16]. For a fair comparison, we reproduce these methods by their officially released codes and recommended settings on the same experimental environment.

4.2 Results on Synthetic Dataset

PU1K dataset. As shown in Table 1, we conduct experiments on three different input sizes of point, including sparse (512), medium (1,024), and dense (2,048). From the results, we can observe the proposed method can achieve the

Table 1: Quantitative comparison ($\times 4$ upsampling) on PU1K dataset with different input sizes of point cloud. The values of CD, HD, and P2F are multiplied by 10^3 . A smaller value denotes a better performance.

Methods	Sparse (512) input			Medium (1,024) input			Dense (2,048) input		
	CD	HD	P2F	CD	HD	P2F	CD	HD	P2F
PU-Net [34]	2.990	35.403	11.272	1.920	24.181	7.468	1.157	15.297	4.924
MPU [27]	2.727	30.471	8.190	1.268	16.088	4.777	0.861	11.799	3.181
PU-GAN [11]	2.089	22.716	6.957	1.151	14.781	4.490	0.661	9.238	2.892
Dis-PU [12]	2.130	25.471	7.299	1.210	16.518	4.606	0.731	9.505	2.719
PU-GCN [20]	1.975	22.527	6.338	1.142	14.565	4.082	0.635	9.152	2.590
PU-EVA [16]	1.942	20.980	6.366	1.065	13.376	4.172	0.649	8.870	2.715
Ours	1.594	17.733	4.892	0.808	10.750	3.061	0.471	7.123	1.925

Table 2: Quantitative comparison on PU-GAN [11] dataset. The input size of point cloud is 1,024. The values of CD, HD, and P2F are multiplied by 10^3 .

Methods	$\times 4$ Upsampling			$\times 16$ Upsampling		
	CD	HD	P2F	CD	HD	P2F
PU-Net [34]	0.883	7.132	8.127	0.845	11.225	10.114
MPU [27]	0.589	6.206	4.568	0.365	8.113	5.181
PU-GAN [11]	0.566	6.932	4.281	0.390	8.920	4.988
Dis-PU [12]	0.527	5.706	3.378	0.302	6.939	4.146
PU-GCN [20]	0.584	5.257	3.987	0.320	6.567	4.381
PU-EVA [16]	0.571	5.840	3.937	0.342	8.140	4.473
Ours	0.520	6.102	3.165	0.284	7.143	3.724

leading performance in all metrics and obvious improvements on the sparse input. Besides, compared the results of left three columns (sparse input) with the right three columns (dense input), we can observe that the sparse input size is a more challenging scenario. Nonetheless, our method can consistently yield evident improvements over the existing state-of-the-art methods, which verifies the robustness of our method to different input sizes. Moreover, the qualitative results in Fig. 4 can also demonstrate that our cascaded refinement network enables to generate more uniform point clouds with better fine-grained details and less outliers, such as the arms of the chair and the wings of the plane.

PU-GAN dataset. Following the previous settings [12], we also conduct experiments under two different upsampling rates at the input size of 1,024 on PU-GAN dataset. Table 2 reports the quantitative results based on $\times 4$ and $\times 16$ upsampling rates. Since PU-GAN dataset only provides training patches based on $\times 4$ upsampling rate, we apply the model twice for accomplishing $\times 16$ upsampling rates in this experiment. Overall, our method can achieve the best performance in most metrics. For Hausdorff Distance metric, some counterparts




Fig. 4: Point cloud upsampling ($\times 4$) results on PU1K dataset. The point clouds are colored by the nearest distance between the ground truth and the prediction of each method. The blue denotes the small errors. One can zoom in for more details.

slightly outperform our method. We consider the most possible reason is that our transformer-based feature extraction module require more variations during the training. Compared with the 1,020 3D meshes in PU1K, PU-GAN is a much smaller dataset that only contains 120 meshes for training, which limits the capacity of the proposed method. Nonetheless, as for the results of Chamfer Distance and Point-to-Surface Distance, our method gains evident improvement over other counterparts.

4.3 Results on Real-scanned Dataset

To verify the effectiveness of our method in real-world scenarios, we utilize the models trained on PU1K and conduct the qualitative experiments on ScanObjectNN [25] dataset. We only visualize the input sparse points and upsampled results, since ScanObjectNN dataset does not provide the corresponding dense points. For clarity, we only choose three recent advanced methods (Dis-PU [12], PU-GCN [20], and PU-EVA [16]) for visualization. As shown in Fig. 5, we can observe upsampling the real-scanned data is more challenging, since the input




Fig. 5: Point cloud upsampling ($\times 4$ and $\times 16$) results on real-scanned sparse inputs. Compared with the other counterparts, our method can generate more uniform with detailed structures. The 3D meshes are reconstructed by ball-pivoting algorithm [2].

sparse points are non-uniform and incomplete. Comparing with other counterparts, our method can generate more uniform and high-fidelity point clouds. The visualization of 3D meshes also shows the proposed is able to obtain a better surface reconstruction result. More visualization results on real-scanned data can be found in the supplementary material.

4.4 Model Complexity Analysis

In this section, we investigate the model complexity of our method compared with existing point cloud upsampling methods. The experiments are conducted on a Titan X Pascal GPU. As shown in Table 3, we can find that the proposed method is comparable with other counterparts. In point of training speed, our method is more efficiency than PU-GAN [11], Dis-PU [12], and PU-EVA [16]. Note that Dis-PU [12] employs a warm-up strategy during the training and requires more epoch to ensure the convergence of networks. Thus, the total training time of Dis-PU costs about 48h on PU-GAN dataset, which is much longer than ours (8h). As for inference time, we obtain a comparable and affordable inference speed compared with existing methods. In terms of model parameters, the proposed method can also be regarded as a lightweight network for point cloud upsampling. In addition, our method is scalable to achieve a trade-off between model effectiveness and efficiency by removing the refinement stage during the inference. In this manner, we can achieve leading performance as well as the model

Table 3: Comparison of model complexity with existing point cloud upsampling methods. The metrics of CD, HD, and P2F are calculated on PU1K dataset.

Methods	CD	HD	P2F	Training Speed (s/batch)	Inference Speed (ms/sample)	Params. (Kb)
PU-Net [34]	1.157	15.297	4.924	0.13	9.9	812
MPU [27]	0.861	11.799	3.181	0.23	8.8	76
PU-GAN [11]	0.661	9.238	2.892	0.39	13.1	542
PU-GCN [20]	0.635	9.152	2.590	0.13	8.6	76
Dis-PU [12]	0.731	9.505	2.719	0.58	11.2	1047
PU-EVA [16]	0.649	8.870	2.715	0.57	12.8	2869
Ours (w/o refiner)	0.492	7.337	1.984	0.37	7.3	567
Ours (full model)	0.471	7.123	1.925	0.37	10.8	847

efficiency among other counterparts (the first bottom line). Overall, our cascaded refinement network can not only obtain the obvious performance improvements but also achieve comparable model complexity with other counterparts.

4.5 Ablation Study

To verify the effect of major components in our approach, we quantitatively evaluate the contribution of them. In the following, we provide the detailed discussions respectively.

Ablation of generation stages. Our full cascaded refinement network consists of two upsampling stage and one refinement stage. In order to analyze the effectiveness of each sub-network, we investigate the setting of the number of generation stages. As shown in Table 4, the “two-stage” refers to remove the refinement stage, and the “four-stage” denotes adding an extra upsampling stage. All the models are trained from scratch. Thus, the “two-stage” is different with merely removing the refiner in inference phase. We can find that the performance consistently increases with the number of generation stages. Our “two-stage” still outperforms the existing SOTA methods with a faster inference speed. In the view of the balance between effectiveness and efficiency, we adopt two upsampling stage and a refinement stage in this work.

Effect of feature extraction module. Feature extraction module plays an important role in point cloud upsampling. As shown in Table 5, we investigate the effect of different network architecture on feature extraction. Specifically, “MLP-based” denotes the same feature extraction unit in MPU [27], and “DenseGCN” refers to replace the point transformer layer with a DenseGCN block in [10]. From the results, we can find that the proposed transform-based feature extraction module outperforms other counterparts, indicating its effectiveness of the locally self-attention operator on shape context learning. Moreover, the comparison of DenseGCN and MLP-based networks proves that an effective feature extraction network can promote obvious improvements on point cloud upsampling.

Study of residual learning. For a coarse-to-fine upsampling framework, it is crucial to mitigate the learning conflicts between multiple generation units. As mentioned in Section 3.2, MPU needs to progressively activate the training

Table 4: Ablation of generation stages. The values of CD, HD, and P2F are multiplied by 10^3 .

Model	Dense (2048) input			Inference Speed (ms/sample)	Params. (Kb)
	CD	HD	P2F		
two-stage	0.513	7.808	2.078	7.3	567
three-stage	0.471	7.123	1.925	10.8	847
four-stage	0.466	6.793	1.821	14.7	1127

Table 5: Effect of feature extraction module.

Model	Dense (2,048) input		
	CD	HD	P2F
MLP-based	0.628	9.687	2.721
DenseGCN	0.531	8.462	2.254
Ours	0.471	7.123	1.925

Table 6: Effect of residual learning scheme.

Warm-up	Residual learning	Dense (2,048) input		
		CD	HD	P2F
✓	-	0.511	7.599	2.986
✓	✓	0.472	7.329	1.933
-	-	0.733	9.667	3.380
-	✓	0.471	7.123	1.925

of each upsampling unit, and Dis-PU employs a warm-up strategy that gradually increases the learning rate of spatial refiner. In contrast, we only adopt the residual learning in each generation stage without progressive or warm-up training strategy. From the results in Table 6, we can find that residual learning is more effective than warm-up training strategy in our framework. The reason behind is, that the residual connection makes each stage only regress the point offsets of the initial shape and thus mitigate the learning conflicts between multiple stages, resulting in a better optimization. Therefore, we equip the residual learning for point reconstruction, which makes the proposed method be easily optimized without extra training strategies.

5 Conclusion

In this paper, we propose a cascaded refinement network for point cloud upsampling. The proposed method consists of three generation stages, which progressively generates the coarse dense points and refine the coarse points to a better position in the end. We adopt a residual learning scheme for point coordinate reconstruction, which enables to decrease the difficulty for regressing the new points and mitigate the conflicts between multiple stages. Moreover, a transformer-based feature extraction module is designed to aggregate the global and local shape context, and a two-branch based feature expansion module enables to learn a more expressive upsampled features. Compared with the existing coarse-to-fine point upsampling methods, the proposed network can be easily optimized without extra training strategies. Moreover, we can also dynamically adjust the model efficiency and effectiveness at the stage of inference, depending on the available computational resources. Both quantitative and qualitative results demonstrate the superiority of our method over the state-of-the-art methods.

References

1. Alexa, M., Behr, J., Cohen-Or, D., Fleishman, S., Levin, D., Silva, C.T.: Computing and rendering point set surfaces. *IEEE Trans. Vis. Comput. Graph.* **9**, 3–15 (2003)
2. Bernardini, F., Mittleman, J., Rushmeier, H., Silva, C., Taubin, G.: The ball-pivoting algorithm for surface reconstruction. *IEEE transactions on visualization and computer graphics* **5**(4), 349–359 (1999)
3. Engel, N., Belagiannis, V., Dietmayer, K.C.J.: Point transformer. *IEEE Access* **9**, 134826–134840 (2021)
4. Fan, H., Su, H., Guibas, L.J.: A point set generation network for 3d object reconstruction from a single image. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 605–613 (2017)
5. Graham, B., Engelcke, M., van der Maaten, L.: 3d semantic segmentation with sub-manifold sparse convolutional networks. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition* pp. 9224–9232 (2018)
6. Guo, M.H., Cai, J., Liu, Z.N., Mu, T.J., Martin, R.R., Hu, S.: Pct: Point cloud transformer. *Comput. Vis. Media* **7**, 187–199 (2021)
7. Huang, H., Wu, S., Gong, M., Cohen-Or, D., Ascher, U.M., Zhang, H.: Edge-aware point set resampling. *ACM Transactions on Graphics (TOG)* **32**, 1 – 12 (2013)
8. Hui, L., Xu, R., Xie, J., Qian, J., Yang, J.: Progressive point cloud deconvolution generation network. In: *ECCV* (2020)
9. Lang, A.H., Vora, S., Caesar, H., Zhou, L., Yang, J., Beijbom, O.: Pointpillars: Fast encoders for object detection from point clouds. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* pp. 12689–12697 (2019)
10. Li, G., Müller, M., Thabet, A.K., Ghanem, B.: Deepgcns: Can gcns go as deep as cnns? In: *ICCV*. pp. 9266–9275 (2019)
11. Li, R., Li, X., Fu, C.W., Cohen-Or, D., Heng, P.A.: Pu-gan: A point cloud upsampling adversarial network. In: *ICCV*. pp. 7202–7211 (2019)
12. Li, R., Li, X., Heng, P.A., Fu, C.W.: Point cloud upsampling via disentangled refinement. In: *CVPR*. pp. 344–353 (2021)
13. Li, Y., Bu, R., Sun, M., Wu, W., Di, X., Chen, B.: Pointcnn: Convolution on x-transformed points. In: *NeurIPS* (2018)
14. Lipman, Y., Cohen-Or, D., Levin, D., Tal-Ezer, H.: Parameterization-free projection for geometry reconstruction. In: *SIGGRAPH 2007* (2007)
15. Long, C., Zhang, W., Li, R., Wang, H., Dong, Z., Yang, B.: Pc2-pu: Patch correlation and position correction for effective point cloud upsampling. *ArXiv abs/2109.09337* (2021)
16. Luo, L., Tang, L., Zhou, W., Wang, S., Yang, Z.X.: Pu-eva: An edge-vector based approximation solution for flexible-scale point cloud upsampling. In: *ICCV* (2021)
17. Pan, X., Xia, Z., Song, S., Li, L.E., Huang, G.: 3d object detection with pointformer. In: *CVPR*. pp. 7459–7468 (2021)
18. Qi, C., Su, H., Mo, K., Guibas, L.J.: Pointnet: Deep learning on point sets for 3d classification and segmentation. In: *CVPR*. pp. 77–85 (2017)
19. Qi, C., Yi, L., Su, H., Guibas, L.J.: Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In: *NIPS* (2017)
20. Qian, G., Abualshour, A., Li, G., Thabet, A.K., Ghanem, B.: Pu-gcn: Point cloud upsampling using graph convolutional networks. In: *CVPR*. pp. 11678–11687 (2021)
21. Qian, Y., Hou, J., Kwong, S.T.W., He, Y.: Pugeo-net: A geometry-centric network for 3d point cloud upsampling. In: *ECCV* (2020)

22. Qian, Y., Hou, J., Kwong, S.T.W., He, Y.: Deep magnification-flexible upsampling over 3d point clouds. *IEEE Transactions on Image Processing* **30**, 8354–8367 (2021)
23. Shi, S., Wang, X., Li, H.: Pointcnn: 3d object proposal generation and detection from point cloud. 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) pp. 770–779 (2019)
24. Tatarchenko, M., Park, J., Koltun, V., Zhou, Q.Y.: Tangent convolutions for dense prediction in 3d. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 3887–3896 (2018)
25. Uy, M.A., Pham, Q.H., Hua, B.S., Nguyen, D.T., Yeung, S.K.: Revisiting point cloud classification: A new benchmark dataset and classification model on real-world data. In: *ICCV*. pp. 1588–1597 (2019)
26. Wang, X., Ang Jr, M.H., Lee, G.H.: Cascaded refinement network for point cloud completion. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 790–799 (2020)
27. Wang, Y., Wu, S., Huang, H., Cohen-Or, D., Sorkine-Hornung, O.: Patch-based progressive 3d point set upsampling. In: *CVPR*. pp. 5951–5960 (2019)
28. Wang, Y., Sun, Y., Liu, Z., Sarma, S.E., Bronstein, M.M., Solomon, J.M.: Dynamic graph cnn for learning on point clouds. *ACM Transactions on Graphics (TOG)* **38**, 1 – 12 (2019)
29. WuShihao, Huanghui, GongMinglun, ZwickerMatthias, Cohen-OrDaniel: Deep points consolidation. *ACM Transactions on Graphics* (2015)
30. Xiang, P., Wen, X., Liu, Y.S., Cao, Y.P., Wan, P., Zheng, W., Han, Z.: SnowflakeNet: Point cloud completion by snowflake point deconvolution with skip-transformer. In: *ICCV* (2021)
31. Xie, H., Yao, H., Zhou, S., Mao, J., Zhang, S., Sun, W.: Grnet: Gridding residual network for dense point cloud completion. In: *ECCV* (2020)
32. Yan, X., Yan, H., Wang, J., Du, H., Wu, Z., Xie, D., Pu, S., Lu, L.: Fbnet: Feedback network for point cloud completion. In: *European Conference on Computer Vision* (2022)
33. Ye, S., Chen, D., Han, S., Wan, Z., Liao, J.: Meta-pu: An arbitrary-scale upsampling network for point cloud. *IEEE transactions on visualization and computer graphics* **PP** (2021)
34. Yu, L., Li, X., Fu, C.W., Cohen-Or, D., Heng, P.: Pu-net: Point cloud upsampling network. In: *CVPR*. pp. 2790–2799 (2018)
35. Yu, L., Li, X., Fu, C.W., Cohen-Or, D., Heng, P.A.: Ec-net: an edge-aware point set consolidation network. In: *ECCV* (2018)
36. Yu, X., Rao, Y., Wang, Z., Liu, Z., Lu, J., Zhou, J.: Pointtr: Diverse point cloud completion with geometry-aware transformers. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 12498–12507 (2021)
37. Yuan, W., Khot, T., Held, D., Mertz, C., Hebert, M.: Pcn: Point completion network. *International Conference on 3D Vision (3DV)* pp. 728–737 (2018)
38. Yuan, W., Khot, T., Held, D., Mertz, C., Hebert, M.: Pcn: Point completion network. In: *2018 International Conference on 3D Vision (3DV)*. pp. 728–737. *IEEE* (2018)
39. Zhao, H., Jiang, L., Jia, J., Torr, P.H.S., Koltun, V.: Point transformer. In: *ICCV* (2021)
40. Zhao, Y., Hui, L., Xie, J.: Sspu-net: Self-supervised point cloud upsampling via differentiable rendering. In: *ACMMM* (2021)
41. Zhou, Y., Tuzel, O.: Voxnet: End-to-end learning for point cloud based 3d object detection. 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition pp. 4490–4499 (2018)