

PU-Net: Point Cloud Upsampling Network

Lequan Yu^{*1,3} Xianzhi Li^{*1} Chi-Wing Fu^{1,3} Daniel Cohen-Or² Pheng-Ann Heng^{1,3}

¹The Chinese University of Hong Kong ²Tel Aviv University

³Guangdong Provincial Key Laboratory of Computer Vision and Virtual Reality Technology,
Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences, China

{lqyu, xzli, cwfu, pheng}@cse.cuhk.edu.hk dcor@mail.tau.ac.il

Abstract

Learning and analyzing 3D point clouds with deep networks is challenging due to the sparseness and irregularity of the data. In this paper, we present a data-driven point cloud upsampling technique. The key idea is to learn multi-level features per point and expand the point set via a multi-branch convolution unit implicitly in feature space. The expanded feature is then split to a multitude of features, which are then reconstructed to an upsampled point set. Our network is applied at a patch-level, with a joint loss function that encourages the upsampled points to remain on the underlying surface with a uniform distribution. We conduct various experiments using synthesis and scan data to evaluate our method and demonstrate its superiority over some baseline methods and an optimization-based method. Results show that our upsampled points have better uniformity and are located closer to the underlying surfaces.

1. Introduction

Point cloud is a fundamental 3D representation that has drawn increasing attention due to the popularity of various depth scanning devices. Recently, pioneering works [29, 30, 18] began to explore the possibility of reasoning point clouds by means of deep networks for understanding geometry and recognizing 3D structures. In these works, the deep networks directly extract features from the raw 3D point coordinates without using traditional features, *e.g.*, normal and curvature. These works present impressive results for 3D object classification and semantic scene segmentation.

In this work we are interested in an upsampling problem: given a set of points, generate a denser set of points to describe the underlying geometry by learning the geometry of a training dataset. This upsampling problem is similar in spirit to the image super-resolution problem [33, 20]; however, dealing with 3D points rather than a 2D grid of pixels

poses new challenges. First, unlike the image space, which is represented by a regular grid, point clouds do not have any spatial order and regular structure. Second, the generated points should describe the underlying geometry of a latent target object, meaning that they should roughly lie on the target object surface. Third, the generated points should be informative and should not clutter together. Having said that, the generated output point set should be more uniform on the target object surface. Thus, simple interpolation between input points cannot produce satisfactory results.

To meet the above challenges, we present a data-driven point cloud upsampling network. Our network is applied at a patch-level, with a joint loss function that encourages the upsampled points to remain on the underlying surface with a uniform distribution. The key idea is to learn multi-level features per point, and then *expand the point set via a multi-branch convolution unit implicitly in feature space*. The expanded feature is then split to a multitude of features, which are then reconstructed to an upsampled point set.

Our network, namely PU-Net, learns geometry semantics of point-based patches from 3D models, and applies the learned knowledge to upsample a given point cloud. It should be noted that unlike previous network-based methods designed for 3D point sets [29, 30, 18], the number of input and output points in our network are *not the same*.

We formulate two metrics, distribution uniformity and distance deviation from underlying surfaces, to quantitatively evaluate the upsampled point set, and test our method using variety of synthetic and real-scanned data. We also evaluate the performance of our method, and compare it to baseline and state-of-the-art optimization-based methods. Results show that our upsampled points have better uniformity, and are located closer to the underlying surfaces.

Related work: optimization-based methods. An early work by Alexa *et al.* [2] upsamples a point set by interpolating points at vertices of a Voronoi diagram in the local tangent space. Lipman *et al.* [24] present a locally optimal projection (LOP) operator for points resampling and surface

*Equal contribution.

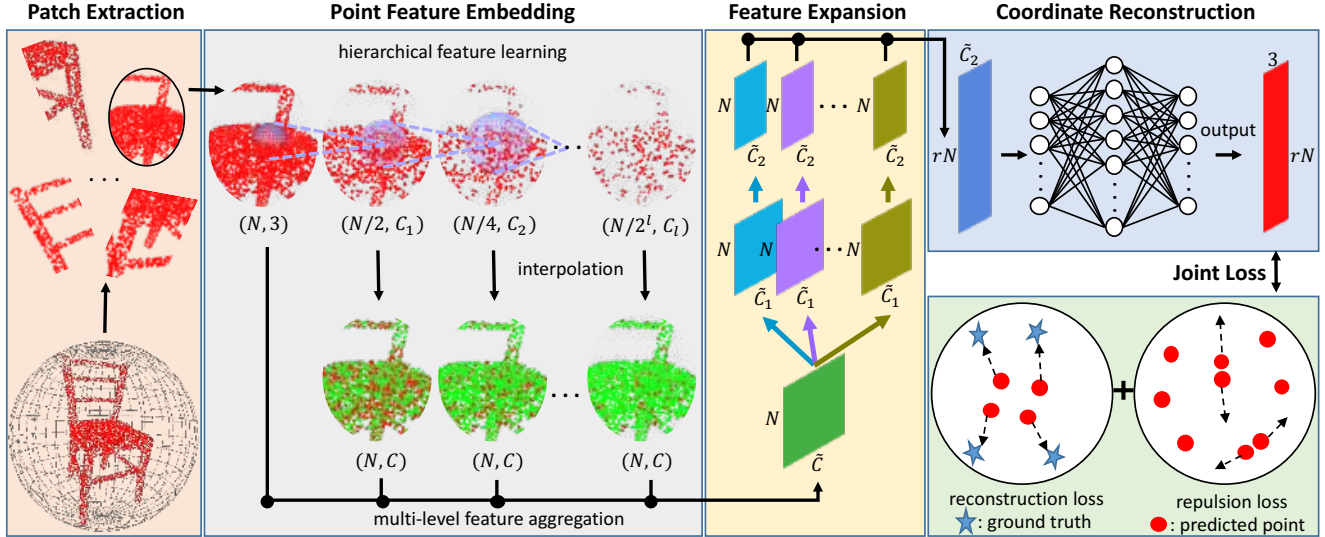


Figure 1. The architecture of PU-Net (better view in color). The input has N points, while the output has rN points, where r is the upsampling rate. C_i , \tilde{C} and \tilde{C}_i represent the feature channel number. We restore different level features for the original N points with interpolation, and reduce all level features to a fixed dimension C with a convolution. The red color in the point feature embedding component shows the original and progressively subsampled points in hierarchical feature learning, while the green color indicates the restored features. We jointly adopt the reconstruction loss and repulsion loss in the end-to-end training of PU-Net.

reconstruction based on an L_1 median. The operator works well even when the input point set contains noise and outliers. Successively, Huang *et al.* [14] propose an improved weighted LOP to address the point set density problem.

Although these works have demonstrated good results, they make a strong assumption that the underlying surface is smooth, thus restricting the method’s scope. Then, Huang *et al.* [15] introduce an edge-aware point set resampling method by first resampling away from edges and then progressively approaching edges and corners. However, the quality of their results heavily relies on the accuracy of the normals at given points and careful parameter tuning. It is worth mentioning that Wu *et al.* [35] propose a deep points representation method to fuse consolidation and completion in one coherent step. Since its main focus is on filling large holes, global smoothness is, however, not enforced, so the method is sensitive to large noise. Overall, the above methods are not data-driven, thus heavily relying on priors.

Related work: deep-learning-based methods. Points in a point cloud do not have any specific order nor follow any regular grid structure, so only a few recent works adopt a deep learning model to directly process point clouds. Most existing works convert a point cloud into some other 3D representations such as the volumetric grids [27, 36, 31, 6] and geometric graphs [3, 26] for processing. Qi *et al.* [29, 30] firstly introduced a deep learning network for point cloud classification and segmentation; in particular, the PointNet++ uses a hierarchical feature learning architecture to capture both local and global geometry context. Subsequently, many other networks were proposed for high-level

analysis problems with point clouds [18, 13, 21, 34, 28]. However, they all focus on global or mid-level attributes of point clouds. In another work, Guerrero *et al.* [10] developed a network to estimate the local shape properties in point clouds, including normal and curvature. Other relevant networks focus on 3D reconstruction from 2D images [8, 23, 9]. To the best of our knowledge, there are no prior works focusing on point cloud upsampling.

2. Network Architecture

Given a 3D point cloud with point coordinates in nonuniform distributions, our network aims to output a denser point cloud that follows the underlying surface of the target object while being uniform in distribution. Our network architecture (see Fig. 1) has four components: *patch extraction*, *point feature embedding*, *feature expansion*, and *coordinate reconstruction*. First, we extract patches of points in varying scales and distributions from a given set of prior 3D models (Sec. 2.1). Then, the *point feature embedding* component maps the raw 3D coordinates to a feature space by hierarchical feature learning and multi-level feature aggregation (Sec. 2.2). After that, we expand the number of features using the *feature expansion* component (Sec. 2.3) and reconstruct the 3D coordinates of the output point cloud via a series of fully connected layers in the *coordinate reconstruction* component (Sec. 2.4).

2.1. Patch Extraction

We collect a set of 3D objects as prior information for training. These objects cover a rich variety of shapes, from

smooth surface to shapes with sharp edges and corners. Essentially, for our network to upsample a point cloud, it should learn local geometry patterns from the objects. This motivates us to take a patch-based approach to train the network and to learn the geometry semantics.

In detail, we randomly select M points on the surface of these objects. From each selected point, we grow a surface patch on the object, such that any point on the patch is within a certain geodesic distance (d) from the selected point over the surface. Then, we use Poisson disk sampling to randomly generate \hat{N} points on each patch as the referenced ground truth point distribution on the patch. In our upsampling task, both local and global context contribute to a smooth and uniform output. Hence, we set d with varying sizes, so that we can extract patches of points on the prior objects with varying scale and density.

2.2. Point Feature Embedding

To learn both local and global geometry context from the patches, we consider the following two feature learning strategies, whose benefits complement each other:

Hierarchical feature learning. Progressively capturing features of growing scales in a hierarchy has been proved to be an effective strategy for extracting local and global features. Hence, we adopt the recently proposed hierarchical feature learning mechanism in PointNet++ [30] as the very frontal part in our network. To adopt hierarchical feature learning for point cloud upsampling, we specifically use a relatively small grouping radius in each level, since generating new points usually involves more of the local context than the high-level recognition tasks in [30].

Multi-level feature aggregation. Lower layers in a network generally correspond to local features in smaller scales, and vice versa. For better upsampling results, we should optimally aggregate features in different levels. Some previous works adopt skip-connections for cascaded multi-level feature aggregation [25, 32, 30]. However, we found by experiments that such top-down propagation is not very efficient for aggregating features in our upsampling problem. Therefore, we propose to directly combine features from different levels and let the network learn the importance of each level [11, 38, 12].

Since the input point set on each patch (see point feature embedding in Fig. 1) is subsampled gradually in hierarchical feature extraction, we concatenate the point features from each level by first restoring the features of all original points from the downsampled point features by the interpolation method in PointNet++ [30]. Specifically, the features of an interpolated point x in level ℓ is calculated by:

$$f^{(\ell)}(x) = \frac{\sum_{i=1}^3 w_i(x) f^{(\ell)}(x_i)}{\sum_{i=1}^3 w_i(x)}, \quad (1)$$

where $w_i(x)=1/d(x, x_i)$, which is an inverse distance weight, and x_i, x_2, x_3 are the three nearest neighbors of x in level ℓ . We then use a 1×1 convolution to reduce the interpolated feature in different level to the same dimension, i.e., C . Finally, we concatenate the features from each level as the embedded point feature f .

2.3. Feature Expansion

After the point feature embedding component, we expand the number of features in the feature space. This is equivalent to expanding the number of points, since *points and features are interchangeable*. Suppose the dimension of f is $N \times \tilde{C}$, N is the number of input points and \tilde{C} is the feature dimension of the concatenated embedded feature. The *feature expansion* operation would output a feature f' with dimension $rN \times \tilde{C}_2$, where r is the upsampling rate and \tilde{C}_2 is the new feature dimension. Essentially, this is similar to feature upsampling in image-related tasks, which can be done by deconvolution [25] (also known as transposed convolution) or interpolation [7]. However, it is non-trivial to apply these operations to point clouds due to the non-regularity and unordered properties of points.

We therefore propose an efficient *feature expansion* operation based on the sub-pixel convolution layer [33]. This operation can be represented as:

$$f' = \mathcal{RS}([C_1^2(C_1^1(f)), \dots, C_r^2(C_r^1(f))]), \quad (2)$$

where $C_i^1(\cdot)$ and $C_i^2(\cdot)$ are two sets of separate 1×1 convolutions, and $\mathcal{RS}(\cdot)$ is a reshape operation to convert an $N \times r\tilde{C}_2$ tensor to a tensor of size $rN \times \tilde{C}_2$. We emphasize that the feature in the embedding space has already encapsulated the relative spatial information from the neighborhood points via the efficient multi-level feature aggregation, so we do not need to explicitly consider the spatial information when performing this feature expansion operation.

It is worth mentioning that the r feature sets generated from the first convolution $C_i^1(\cdot)$ in each set have a high correlation, and this would cause the final reconstructed 3D points to be located too close to one another. Hence, we further add another convolution (with separate weights) for each feature set. Since we train the network to learn the r different convolutions for the r feature sets, these new features can include more diverse information, thus reducing their correlations. This *feature expansion* operation can be implemented by applying separated convolutions to the r feature sets; see Fig. 1. It can also be implemented by more computation efficient grouped convolution [19, 37, 39].

2.4. Coordinate Reconstruction

In this part, we reconstruct the 3D coordinates of output points from the expanded feature with the size of $rN \times \tilde{C}_2$. Specifically, we regress the 3D coordinates via a series of

fully connected layers on the feature of each point, and the final output is the upsampled point coordinates $rN \times 3$.

3. End-to-End Network Training

3.1. Training Data Generation

Point cloud upsampling is an ill-posed problem due to the uncertainty or ambiguity of upsampled point clouds. Given a sparse input point cloud, there are many feasible output point distributions. Therefore, we do not have the notion of “correct pairs” of input and ground truth. To alleviate this problem, we propose an on-the-fly input generation scheme. Specifically, the referenced ground truth point distribution of a training patch is fixed, whereas *the input points are randomly sampled from the ground truth point set with a downsampling rate of r at each training epoch*. Intuitively, this scheme is equivalent to simulating many feasible output point distributions for a given sparse input point distribution. Additionally, this scheme can further enlarge the training dataset, allowing us to depend on a relatively small dataset for training.

3.2. Joint Loss Function

We propose a novel joint loss function to train the network in an end-to-end fashion. As we mentioned earlier, the function should encourage the generated points to be located on the underlying object surfaces in a more uniform distribution. Therefore, we design a joint loss function that combines the *reconstruction loss* and *repulsion loss*.

Reconstruction loss. To put points on the underlying object surfaces, we propose to use the Earth Mover’s distance (EMD) [8] as our reconstruction loss to evaluate the similarity between the predicted point cloud $S_p \subseteq \mathbb{R}^3$ and the referenced ground truth point cloud $S_{gt} \subseteq \mathbb{R}^3$:

$$L_{rec} = d_{EMD}(S_p, S_{gt}) = \min_{\phi: S_p \rightarrow S_{gt}} \sum_{x_i \in S_p} \|x_i - \phi(x_i)\|_2, \quad (3)$$

where $\phi: S_p \rightarrow S_{gt}$ indicates the bijection mapping.

Actually, Chamfer Distance (CD) is another candidate for evaluating the similarity between two point sets. However, compared with CD, EMD can better capture the shape (see [8] for more details) to encourage the output points to be located close to the underlying object surfaces. Hence, we choose to use EMD in our reconstruction loss.

Repulsion loss. Although training with the reconstruction loss can generate points on the underlying object surfaces, the generated points tend to be located near the original points. To distribute the generated points more uniformly, we design the repulsion loss, which is represented as:

$$L_{rep} = \sum_{i=0}^{\hat{N}} \sum_{i' \in K(i)} \eta(\|x_{i'} - x_i\|) w(\|x_{i'} - x_i\|), \quad (4)$$

where $\hat{N} = rN$ is the number of output points, $K(i)$ is the index set of the k -nearest neighbors of point x_i , and $\|\cdot\|$ is the L2-norm. $\eta(r) = -r$ is called the *repulsion term*, which is a decreasing function to penalize x_i if x_i is located too close to other points in $K(i)$. To penalize x_i only when it is too close to its neighboring points, we add two restrictions: (i) we only consider points $x_{i'}$ in the k -nearest neighborhood of x_i ; and (ii) we use the fast-decaying weight function $w(r)$ in the repulsion loss; that is, we follow [24, 14] to set $w(r) = e^{-r^2/h^2}$ in our experiments.

Altogether, we train the network in an end-to-end manner by minimizing the following joint loss function:

$$L(\theta) = L_{rec} + \alpha L_{rep} + \beta \|\theta\|^2, \quad (5)$$

where θ indicates the parameters in our network, α balances the reconstruction loss and repulsion loss, and β denotes the multiplier of weight decay. For simplicity, we ignore the index of each training sample.

4. Experiments

4.1. Datasets

Since there are no public benchmarks for point cloud upsampling, we collect a dataset of 60 different models from the Visionair repository [1], ranging from smooth non-rigid objects (*e.g.*, Bunny) to steep rigid objects (*e.g.*, Chair). Among them, we randomly select 40 for training, and use the rest for testing¹. We crop 100 patches for each training object, and we use $M=4000$ patches to train the network in total. For testing objects, we use Monte-Carlo random sampling approach to sample 5000 points on each object as input. To further demonstrate the generalization ability of our network, we directly test our well-trained network on the SHREC15 [22] dataset, which contains 1200 shapes from 50 categories. In detail, we randomly choose one model from each category for testing, considering that each category contains 24 similar objects in various poses. As for ModelNet40 [36] and ShapeNet [4], we found it hard to extract patches from those objects due to the low mesh quality (*e.g.*, holes, self-intersection, etc.). Therefore, we use them for testing; see the supplementary material for the results.

4.2. Implementation Details

The default point number \hat{N} of each patch is 4096, and the upsampling rate r is 4. Therefore, each input patch has 1024 points. To avoid overfitting, we augment the data by randomly rotating, shifting and scaling the data. We use 4 levels with grouping radii 0.05, 0.1, 0.2 and 0.3 in the point feature embedding component, and the dimension C of the restored feature is 64. For details on other network architecture parameters, please see our supplementary material.

¹The complete object list can be found in the supplementary material.




Figure 2. Example point distributions with corresponding normalized uniformity coefficients (NUC) computed with $p = 0.2\%$.

Parameters k and h in repulsion loss are set as 5 and 0.03, respectively. The balancing weights α and β are set as 0.01 and $\beta = 10^{-5}$, respectively. The implementation is based on TensorFlow². For the optimization, we train the network for 120 epoch using the Adam [17] algorithm with a mini-batch size of 28 and a learning rate of 0.001. Generally, the training took about 4.5h on the NVIDIA TITAN Xp GPU.

4.3. Evaluation Metric

To quantitatively evaluate the quality of the output point sets, we formulate two metrics to measure the deviation between the output points and the ground truth meshes, as well as the distribution uniformity of the output points. For surface deviation, we find the closest point \hat{x}_i on the mesh for each predicted point x_i , and calculate the distance between them. Then we compute the mean and standard deviation over all the points as one of our metrics.

As for the uniformity metric, we randomly put D equal-size disks on the object surface ($D = 9000$ in our experiments) and calculate the standard deviation of the number of points inside the disks. We further normalize the density of each object and then compute the overall uniformity of the point sets over all the objects in the testing dataset. Therefore, we define the *normalized uniformity coefficient* (NUC) with disk area percentage p as:

$$avg = \frac{1}{K * D} \sum_{k=1}^K \sum_{i=1}^D \frac{n_i^k}{N^k * p},$$

$$NUC = \sqrt{\frac{1}{K * D} \sum_{k=1}^K \sum_{i=1}^D \left(\frac{n_i^k}{N^k * p} - avg \right)^2},$$
(6)

where n_i^k is the number of points within the i -th disk of the k -th object, N^k is the total number of points on the k -th object, K is the total number of test objects, and p is the percentage of the disk area over the total object surface area. Note that we use geodesic distance rather than Euclidean distance to form the disks. Fig. 2 shows three different point distributions with their corresponding NUC values. As we

² <https://github.com/yulequan/PU-Net>

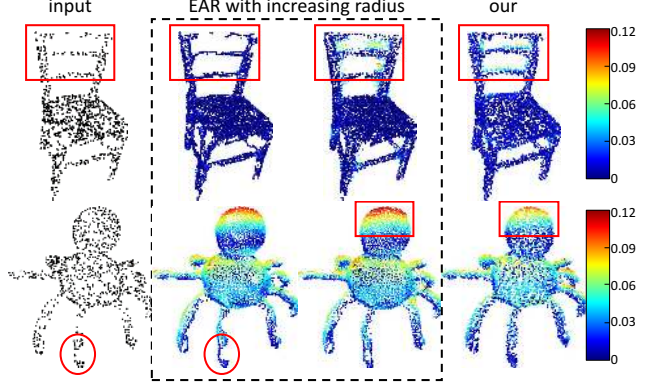


Figure 3. Comparison with the EAR method [15]. We color-code all point clouds to show the deviation from the ground truth mesh. For the EAR method, the radius of the Chair model is 0.1 and 0.182, while the radius of the Spider model is 0.106 and 0.155.

can see, the proposed NUC metric can effectively reveal the point set uniformity: the lower the UNC value, the more uniform the point set distribution is.

4.4. Comparisons with Other Methods

Comparison with an optimization-based method. We compare our method with the Edge Aware Resampling (EAR) method [15], which is a state-of-art method for point cloud upsampling. The results are shown in Fig. 3, where the Chair is from our collected testing dataset and the Spider is from SHREC15. We color-code the point clouds to show the deviation from the ground truth meshes. There are 1024 points in the input and we do a 4X upsampling. Since EAR relies on the normal information, to be fair, we calculate the normal according to the ground truth mesh. We show two results of EAR with increasing radius, while setting other parameters to their default values. As we can see, the radius parameter has a great influence on EAR’s performance. For relatively small radius, the output has low surface deviation but the added points are not uniform, while more outliers are introduced if the radius is large. In contrast, our method can better balance the deviation and uniformity without the need to carefully tune the parameters.

Comparison with deep learning-based methods. As far as we know, we are not aware of any deep learning-based method for point cloud upsampling, so we design some baseline methods for comparison. Since PointNet [29] and PointNet++ [30] are pioneers for 3D point cloud reasoning with deep learning techniques, we design the baselines based on them. Specifically, we adopt the semantic segmentation network architecture for point feature embedding and use one set of convolutions for feature expansion. Note that we consider two versions of PointNet++: basic PointNet++ and PointNet++ with multi-scale grouping (MSG) for handling non-uniform sampling density; hence, we have three

Table 1. Quantitative comparison on our collected dataset.

Method	NUC with different p						Deviation (10^{-2})		Time (s)
	0.2%	0.4%	0.6%	0.8%	1.0%	1.2%	mean	std	
Input	0.316	0.224	0.185	0.164	0.150	0.142	-	-	-
PointNet [29]	0.409	0.334	0.295	0.270	0.252	0.239	2.27	3.18	0.05
PointNet++ [30]	0.215	0.178	0.160	0.150	0.143	0.139	1.01	0.83	0.16
PointNet++(MSG) [30]	0.208	0.169	0.152	0.143	0.137	0.134	0.78	0.61	0.45
PU-Net (Ours)	0.174	0.138	0.122	0.115	0.112	0.110	0.63	0.53	0.15

Table 2. Quantitative comparison on SHREC15 dataset.

Method	NUC with different p						Deviation (10^{-2})		Time (s)
	0.2%	0.4%	0.6%	0.8%	1.0%	1.2%	mean	std	
Input	0.315	0.222	0.184	0.165	0.153	0.146	-	-	-
PointNet [29]	0.490	0.405	0.360	0.330	0.309	0.293	2.03	2.94	0.05
PointNet++ [30]	0.259	0.218	0.197	0.185	0.176	0.170	0.88	0.75	0.16
PointNet++(MSG) [30]	0.250	0.199	0.175	0.161	0.153	0.148	0.69	0.59	0.45
PU-Net (Ours)	0.219	0.173	0.154	0.144	0.140	0.137	0.52	0.45	0.15

baselines in total, and we train them only with the reconstruction loss. Please refer to the supplemental material for details of the baseline network architectures. Although we modify the PointNet, PointNet++, and PointNet++(MSG) architectures for the upsampling problem, for convenience, we still call the baselines by their original names.

Tables 1 and 2 list the quantitative comparison results on our collected dataset and the SHREC15 dataset, respectively. Note that measuring NUC with small p shows local distribution uniformity in small regions, while measuring NUC with large p shows more global uniformity. Among the baselines, PointNet performs the worst, since it cannot capture local structure information. Compared with PointNet++, PointNet++(MSG) can slightly improve the uniformity due to the explicit multi-scale information grouping. However, it involves more parameters, thus significantly prolonging the training and testing time. Overall, our PU-Net achieves the best performance with the lowest deviation from surface and the best distribution uniformity compared to the baselines, especially on the local uniformity.

Fig. 4 shows results for visual comparison, where points are colored by their distance deviations from surface. As we can see, the point clouds predicted by our method better match the underlying surface with lower deviations.

4.5. Architecture Design Analysis

Analyzing the feature expansion. We compare our proposed feature expansion scheme with two interpolation-like schemes. The first one is similar to a naive point interpolation (denoted as *interp1*). After extracting the point feature of each point, we combine its own features and the features from the nearest neighboring points to generate the upsampled features. The second one introduces more randomness

Table 3. Architecture design analysis on our collect dataset.

Methods	NUC with different p				Deviation (10^{-2})	
	0.4%	0.6%	0.8%	1.0%	mean	std
<i>interp1</i> +EMD	0.153	0.136	0.126	0.121	0.67	0.54
<i>interp2</i> +EMD	0.144	0.129	0.122	0.118	0.71	0.57
CD	0.185	0.167	0.154	0.147	0.87	0.69
EMD	0.140	0.126	0.119	0.116	0.68	0.58
Ours	0.138	0.122	0.115	0.112	0.63	0.53

(denoted as *interp2*). Instead of using the features from the nearest neighbors, we use a radius based ball query to find the neighborhood and combine the features from these points to generate the upsampled features. We train these two networks with the reconstruction loss (also named as the EMD loss) and the results are listed in the top two rows of Table 3. For fair comparison, we also train our network only with the EMD loss. The results are shown in the fourth row. Comparing with these two interpolation-like schemes, our proposed scheme can generate more uniform outputs with comparable surface distance deviation.

Comparing different loss functions. As mentioned above, the EMD can better capture the object shape than CD. Comparing their performance in Table 3, we can see that the EMD loss improves the output uniformity with low surface distance deviation when comparing with the CD loss, meaning that the EMD loss can better encourage the output points to lie on the underlying surface. Furthermore, by comparing the last two rows in Table 3, we can see that the repulsion loss can further improve the uniformity of the output.

4.6. More Experiments

Surface reconstruction from upsampled point sets. An important application of point cloud upsampling is to improve the surface reconstruction quality. Hence, we com-

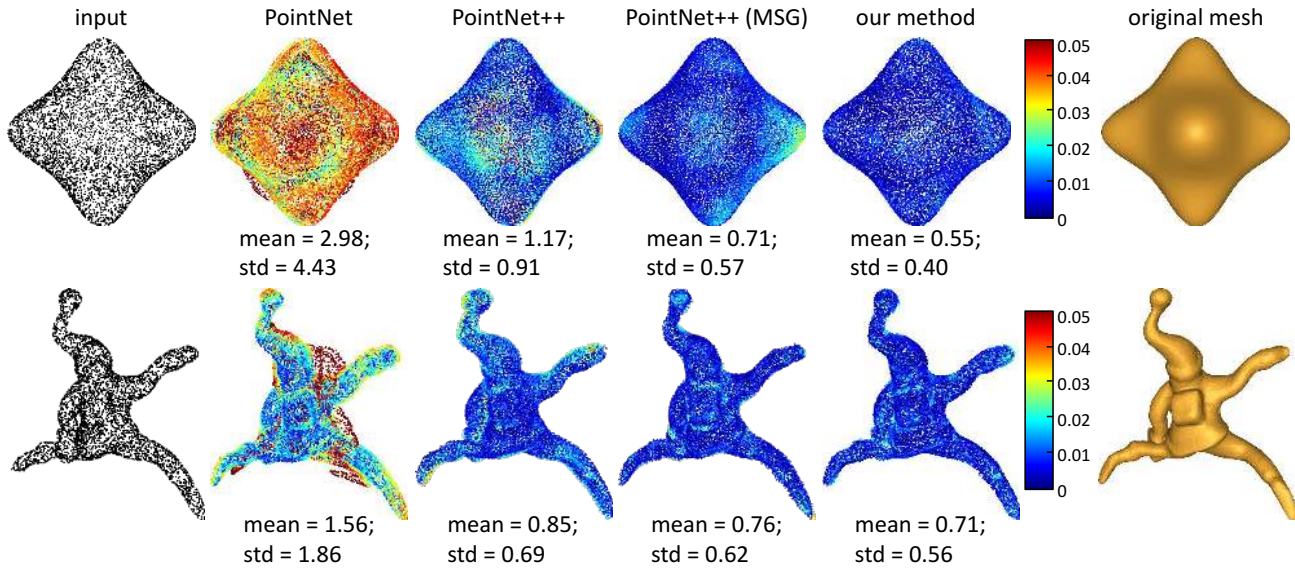


Figure 4. Visual comparison on samples from our collected dataset (top row) and SHREC (bottom row). The colors on points (see color map) reveal the surface distance errors, where blue indicates low error and red indicates high error.

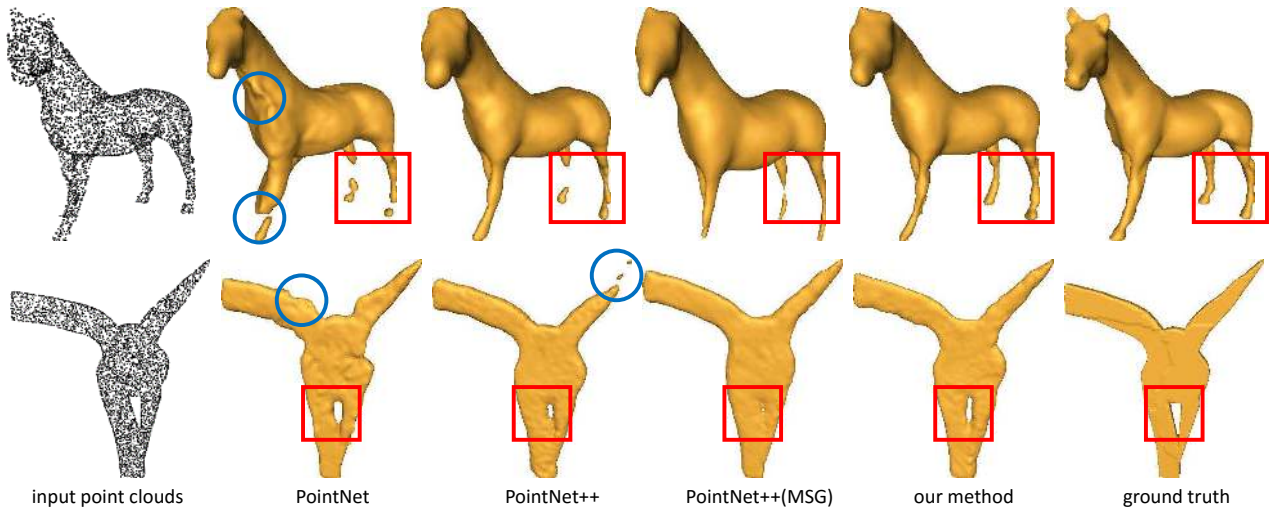


Figure 5. Surface reconstruction results from the upsampled point clouds.




Figure 6. Results of iterative upsampling. We color-code points by the depth information. Blue points are closer to us.

pare the reconstruction results of different methods with the direct Poisson surface reconstruction method [16] provided in MeshLab [5]; see Fig. 5. We can observe that the reconstruction result from our method is the closest to the ground truth, while other methods either miss certain structures (e.g., the leg of the Horse) or overfill the hole.

Results of iterative upsampling. To study the ability of our network to handle varying number of input points, we design an iterative upsampling experiment, which takes the output of the previous iteration as the input of the next iteration. Fig. 6 shows the results. The initial input point cloud has 1024 points and we increase fourfold in each iteration. From the results, we can see that our network can produce reasonable results for different number of input points. Furthermore, this iterative upsampling experiment also shows

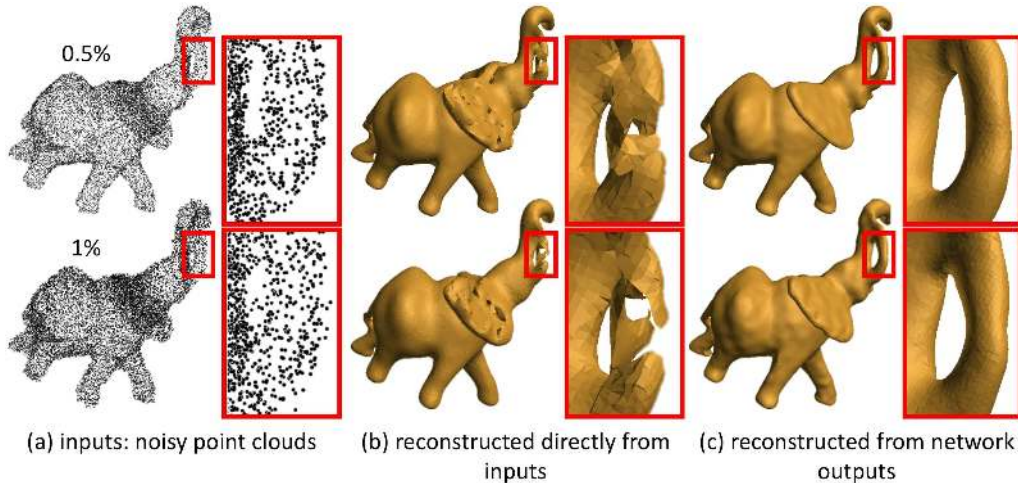


Figure 7. Surface reconstruction results from noisy input points.

the anti-noise ability of our network to resist the accumulated errors introduced in the iterative upsampling.

Results from noisy input point sets. Fig. 7 shows the surface reconstruction results from noisy point clouds (Gaussian noise of 0.5% and 1% of object bounding box diagonal), which demonstrate that our network facilitates the production of better surfaces even with noisy inputs.

Results on real-scanned point clouds. Lastly, we evaluated the ability of our network to upsample real-scanned point clouds, which were downloaded from Visionair [1]. In Fig. 8, the left-most column presents the real-scanned point clouds. Even though each real-scanned point cloud contains millions of points, the phenomenon of inhomogeneity still exists. For better visualization, we cut small patches from the original point clouds and show the patches in the middle column. We can observe that the real-scanned points tend to have line structures, while our network still has the ability to uniformly add points in the sparse regions.

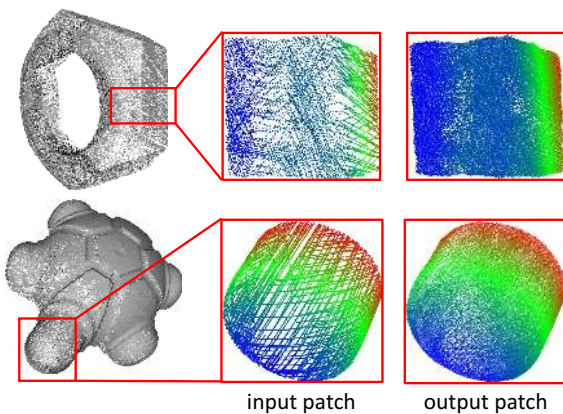


Figure 8. Results on real-scanned point clouds (Screw nut & Turtle). We color-code input patches and upsampling results to show the depth information. Blue points are closer to viewpoint.

5. Conclusion

In this paper, we present a deep network for point cloud upsampling, with the goal of generating a denser and uniform set of points from a sparser set of points. Our network is trained at a patch-level using a multi-level feature aggregation manner, thus capturing both local and global information. The design of our network bypasses the need for a prescribed order among the points, by operating on individual features that contain non-local geometry to allow a context-aware upsampling. Our experiments demonstrate the effectiveness of our method. As the first attempt using deep networks, our method still has a number of limitations. Firstly, it is not designed for completion, so our network can not fill large holes and missing parts. Besides, our network may not be able to add meaningful points for tiny structures that are severely undersampled.

In the future, we would like to investigate and develop more means to handle irregular and sparse data, both for regression purposes and for synthesis. One immediate step is to develop a downsampling method. Although, downsampling seems like a simpler problem, there is room to devise proper losses and architecture that maximize the preservation of information in the decimated point set. We believe that in general, the development of deep learning methods for irregular structures is a viable research direction.

Acknowledgments. We thank anonymous reviewers for the comments and suggestions. The work is supported in part by the National Basic Program of China, the 973 Program (Project No. 2015CB351706), the Research Grants Council of the Hong Kong Special Administrative Region (Project no. CUHK 14225616), the Shenzhen Science and Technology Program (No. JCYJ20170413162617606), and the CUHK strategic recruitment fund.

References

- [1] Visionair. [Online; accessed on 14-November-2017].
- [2] M. Alexa, J. Behr, D. Cohen-Or, S. Fleishman, D. Levin, and C. T. Silva. Computing and rendering point set surfaces. *IEEE Trans. Vis. & Comp. Graphics*, 9(1):3–15, 2003.
- [3] J. Bruna, W. Zaremba, A. Szlam, and Y. LeCun. Spectral networks and locally connected networks on graphs. In *Int. Conf. on Learning Representations (ICLR)*, 2014.
- [4] A. X. Chang, T. Funkhouser, L. J. Guibas, P. Hanrahan, Q. Huang, Z. Li, S. Savarese, M. Savva, S. Song, H. Su, et al. ShapeNet: An information-rich 3D model repository. *arXiv preprint arXiv:1512.03012*, 2015.
- [5] P. Cignoni, M. Callieri, M. Corsini, M. Dellepiane, F. Ganovelli, and G. Ranzuglia. MeshLab: an open-source mesh processing tool. In *Eurographics Italian Chapter Conference*, 2008.
- [6] A. Dai, C. R. Qi, and M. Niessner. Shape completion using 3D-Encoder-Predictor CNNs and shape synthesis. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [7] C. Dong, C. C. Loy, K. He, and X. Tang. Image super-resolution using deep convolutional networks. *IEEE Trans. Pattern Anal. & Mach. Intell.*, 38(2):295–307, 2016.
- [8] H. Fan, H. Su, and L. J. Guibas. A point set generation network for 3D object reconstruction from a single image. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [9] T. Groueix, M. Fisher, V. G. Kim, B. C. Russell, and M. Aubry. AtlasNet: A papier-mâché approach to learning 3D surface generation. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2018. to appear.
- [10] P. Guerrero, Y. Kleiman, M. Ovsjanikov, and N. J. Mitra. PCPNet: Learning local shape properties from raw point clouds. *Computer Graphics Forum (Eurographics)*, 2018. to appear.
- [11] B. Hariharan, P. Arbeláez, R. Girshick, and J. Malik. Hypercolumns for object segmentation and fine-grained localization. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [12] Q. Hou, M. Cheng, X. Hu, A. Borji, Z. Tu, and P. Torr. Deeply supervised salient object detection with short connections. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [13] B.-S. Hua, M.-K. Tran, and S.-K. Yeung. Pointwise convolutional neural networks. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2018. to appear.
- [14] H. Huang, D. Li, H. Zhang, U. Ascher, and D. Cohen-Or. Consolidation of unorganized point clouds for surface reconstruction. *ACM Trans. on Graphics (SIGGRAPH Asia)*, 28(5):176:1–8, 2009.
- [15] H. Huang, S. Wu, M. Gong, D. Cohen-Or, U. Ascher, and H. R. Zhang. Edge-aware point set resampling. *ACM Trans. on Graphics*, 32(9):1–12, 2013.
- [16] M. Kazhdan, M. Bolitho, and H. Hoppe. Poisson surface reconstruction. In *Eurographics Symposium on Geometry Processing (SGP)*, 2006.
- [17] D. Kingma and J. Ba. Adam: A method for stochastic optimization. In *Int. Conf. on Learning Representations (ICLR)*, 2015.
- [18] R. Klokov and V. Lempitsky. Escape from cells: deep Kd-Networks for the recognition of 3D point cloud models. In *IEEE Int. Conf. on Computer Vision (ICCV)*, 2017.
- [19] A. Krizhevsky, I. Sutskever, and G. E. Hinton. ImageNet classification with deep convolutional neural networks. In *Int. Conf. on Neural Information Processing Systems (NIPS)*, 2012.
- [20] C. Ledig, L. Theis, F. Huszar, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [21] Y. Li, R. Bu, M. Sun, and B. Chen. PointCNN. *arXiv preprint arXiv:1801.07791*, 2018.
- [22] Z. Lian, J. Zhang, S. Choi, H. ElNaghy, J. El-Sana, T. Furuya, A. Giachetti, R. A. Guler, L. Lai, C. Li, H. Li, F. A. Limberger, R. Martin, R. U. Nakanishi, A. P. Neto, L. G. Nonato, R. Ohbuchi, K. Pevzner, D. Pickup, P. Rosin, A. Sharf, L. Sun, X. Sun, S. Tari, G. Unal, and R. C. Wilson. Non-rigid 3D shape retrieval. In *Proc. of the 2015 Eurographics Workshop on 3D Object Retrieval*, 2015.
- [23] C.-H. Lin, C. Kong, and S. Lucey. Learning efficient point cloud generation for dense 3D object reconstruction. In *AAAI Conf. on Artificial Intell. (AAAI)*, 2017.
- [24] Y. Lipman, D. Cohen-Or, D. Levin, and H. Tal-Ezer. Parameterization-free projection for geometry reconstruction. *ACM Trans. on Graphics (SIGGRAPH)*, 26(3):22:1–5, 2007.
- [25] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [26] J. Masci, D. Boscaini, M. Bronstein, and P. Vandergheynst. Geodesic convolutional neural networks on Riemannian manifolds. In *IEEE Int. Conf. on Computer Vision (ICCV) workshops*, 2015.
- [27] D. Maturana and S. Scherer. Voxnet: A 3D convolutional neural network for real-time object recognition. In *Int. Conf. on Intell. Robots and Systems (IROS)*, 2015.
- [28] C. R. Qi, W. Liu, C. Wu, H. Su, and L. J. Guibas. Frustum PointNets for 3D object detection from RGB-D data. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2018. to appear.
- [29] C. R. Qi, H. Su, K. Mo, and L. J. Guibas. PointNet: Deep learning on point sets for 3D classification and segmentation. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [30] C. R. Qi, L. Yi, H. Su, and L. J. Guibas. PointNet++: Deep hierarchical feature learning on point sets in a metric space. In *Int. Conf. on Neural Information Processing Systems (NIPS)*, 2017.
- [31] G. Riegler, A. O. Ulusoy, and A. Geiger. OctNet: Learning deep 3D representations at high resolutions. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2017.

- [32] O. Ronneberger, P. Fischer, and T. Brox. U-Net: Convolutional networks for biomedical image segmentation. In *Int. Conf. on MICCAI*, 2015.
- [33] W. Shi, J. Caballero, F. Huszár, J. Totz, A. P. Aitken, R. Bishop, D. Rueckert, and Z. Wang. Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [34] W. Wang, R. Yu, Q. Huang, and U. Neumann. SGPN: Similarity group proposal network for 3D point cloud instance segmentation. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2018. to appear.
- [35] S. Wu, H. Huang, M. Gong, M. Zwicker, and D. Cohen-Or. Deep points consolidation. *ACM Trans. on Graphics (SIGGRAPH Asia)*, 34(6):176:1–13, 2015.
- [36] Z. Wu, S. Song, A. Khosla, F. Yu, L. Zhang, X. Tang, and J. Xiao. 3D ShapeNets: A deep representation for volumetric shapes. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [37] S. Xie, R. Girshick, P. Dollár, Z. Tu, and K. He. Aggregated residual transformations for deep neural networks. *arXiv preprint arXiv:1611.05431*, 2016.
- [38] S. Xie and Z. Tu. Holistically-nested edge detection. In *IEEE Int. Conf. on Computer Vision (ICCV)*, 2015.
- [39] X. Zhang, X. Zhou, M. Lin, and J. Sun. ShuffleNet: An extremely efficient convolutional neural network for mobile devices. *arXiv preprint arXiv:1707.01083*, 2017.